

Taxpayer Responses in Good Times and Bad

Enda Patrick Hargaden*

April 30, 2020

Abstract

I show that the magnitude of taxpayer responsiveness, a key parameter in public finance, varies through time. Using linked administrative data and identifying responses from changes in notches, I document a marked decline in responsiveness during the Great Recession that cannot be explained by increased enforcement. I characterize which employee-employer pairs are best at reporting tax-advantaged incomes. Workers in industries with above-average responsiveness, such as construction, were disproportionately affected by the recession. I show responsiveness also declined for workers who remained matched with the same employers throughout the period.

Keywords : Taxable income; behavioural response; notches; bunching.

JEL Classification : H24, H26, E62.

*University of Tennessee, Knoxville. 702 Stokely Management Center, 916 Volunteer Blvd, Knoxville, TN 37996. Email: enda@utk.edu. An earlier draft of this paper circulated as “Taxpayer Responses over the Cycle: Evidence from Irish Notches.” Generous financial support for this project was provided by the University of Michigan's Rackham Graduate School, the Michigan Institute for Teaching and Research Excellence, and the University of Tennessee's Boyd Center for Business and Economic Research, all of which is greatly appreciated. I thank Matias Cattaneo, Bill Fox, Jim Hines, Jason Kerwin, Henrik Kleven, Daniel Reck, Barra Roantree, Georg Schaur, Joel Slemrod, Mel Stephens, Ugo Troiano, Johannes Voget, Caroline Weber, Chenyu Yang, anonymous referees, and seminar participants at the University of Michigan, University of Melbourne, ZEW-University of Mannheim, University of Vienna, Umeå University, Queen's University Belfast, Claremont McKenna College, Economic and Social Research Institute, Institute for Fiscal Studies, National Tax Association, International Institute for Public Finance, and the Royal Economic Society for very useful feedback and comments on earlier drafts. I would like to thank Ben Berstock, Berni Dunne, John Dunne, and Mervyn Ó'Lúing in the Central Statistics Office for providing access to the data, and Philip O'Connell and Karl Whelan for hosting me as a Visitor to University College Dublin.

1 Introduction

In this paper I provide evidence that responsiveness to income taxation fell considerably during the Great Recession. The magnitude of the behavioural response to taxation is a key element of optimal tax policy (Feldstein, 1999; Saez et al., 2012). The level of responsiveness determines the deadweight loss associated with a tax, and thus a large literature attempts to measure it (e.g. Gruber and Saez, 2002; Kopczuk, 2005; Weber, 2014).

This paper documents that responsiveness is time-varying. I use the extent of bunching below tax notches as a measure of responsiveness. Notches are thresholds in the tax system where marginal rates exceed 100%. They provide strong incentives to adjust reported income. For example, by earning €1 over €26,000 in 2010 Irish workers' tax liability increased by €1,040. The Irish tax system included notches before, during, and in the immediate aftermath of the Great Recession. Additionally, these thresholds changed regularly. This permits estimation that incorporates elements from both bunching and difference-in-difference designs. There is unambiguous evidence of tax-advantageous reporting behaviour in the pre-recession period. This is simply not present after 2008. The recession was particularly severe in Ireland: unemployment rose from 5% at the start of 2008 to 14% by the end of 2009. Using the estimator of Chetty et al. (2011), bunching is not statistically significant in 2009, or any subsequent years in the data. The pattern in bunching is consistent with reduced responsiveness during the recession, and cannot be explained by increased enforcement.

Though this paper relates to tax collection in Ireland, it is of broader interest on at least two dimensions. Firstly, an emerging literature has found suggestive evidence of reduced responsiveness of U.S. taxpayers during the Great Recession (Guyton et al., 2016; Buhlmann et al., 2018; Mortenson and Whitten, 2020), reason to believe the findings of this paper may be valid externally. Secondly, systematic variation in the sensitivity of taxpayers over the cycle has direct analogies to the literature on fiscal multipliers (Auerbach and Gorodnichenko, 2013).

The data source and identifying variation in this paper contribute to the literature on taxpayer responses on several grounds. Firstly, employee tax returns in Ireland are filed to the tax authorities¹ directly by their employer. The data come from these administrative records. As wages are tax-deductible for firms and filing incorrect returns is a serious criminal offence, it is unlikely that the results are driven solely by misreporting.² Secondly, the primary identification strategy is based on changes in notch thresholds. This helps disentangle tax-induced responses from learning and round-number effects. As the changes in tax liability are on average worth about a week's wages, these are large and likely salient magnitudes to workers. Lastly, the setting saw particularly volatile movements in output and unemployment. The data contain both the booming Celtic Tiger economy when GNP growth averaged 5% per annum, and also the Great Recession that immediately followed. For context, the drop in quarterly output from peak to trough was four times larger in Ireland than

¹The data come from a tax agency in a developed country responsible for the collection of a large share of national product: income taxes in 2016 amounted to €19.2bn of a total GNP of €226.7bn. Sources: Department of Finance Databank, CSO Quarterly National Accounts.

²Though misreporting is difficult to confirm, I provide some evidence of misreporting explaining much of the decline later in the paper.

in the United States. This volatility provides large variation to analyze macroeconomic effects, improving the precision of the estimates, and mitigating concerns about relevance/salience.

I supplement the bunching results with two additional empirical approaches. I test for between-years differences in the income distribution near the notch with Kolmogorov-Smirnov tests, and estimate responses using a fixed-effect Poisson regression. Both indicate a decline in responsiveness. The Kolmogorov-Smirnov tests show no statistical significance after 2009. From the Poisson regression I do find statistically significant responses, but the incidence rate ratios are three times larger in 2006–2008 than in 2009–2013.

The pattern of responsiveness warrants decomposition. I analyze the determinants of reporting tax-advantaged³ or tax-disadvantaged incomes. The characteristics of both the employee (e.g. age, nationality, full-time status) and the firm (e.g. industry, size, form of incorporation) — and thus the employer-employee pairing — are determinants of reporting behaviour. The industries most likely to report tax-advantaged incomes include construction and agriculture, while employees in state-owned utilities and public administration do not exhibit bunching in tax-advantaged regions. As cash-based industries like construction lack third-party reporting, this evidence is consistent with responses being driven by misreporting.

I use DiNardo, Fortin and Lemieux (1996) counterfactuals to demonstrate that some of the overall change in responsiveness can be explained by changes in the composition of the labour force. Attrition from the labour market during recessions is non-random. The counterfactuals estimate what the wage distribution would look like had certain covariates (e.g. employee age, firm size, sector, legal form of incorporation) remained constant. As these covariates may themselves be endogenous to the overall labour market, strong assumptions are needed to interpret the counterfactuals causally. Nonetheless the counterfactuals suggest that labour force composition must be considered with, and may be a determinant of, taxpayer responsiveness.

I then re-examine bunching behaviour using inverse probability weights to control for these changes in the characteristics of employee-employer pairings. The determinants of bunching appear to change over time, even for those who remained with the same employer. Holding the labour force fixed, responsiveness to taxation appears to have declined during the recession.

In the appendix, I provide a theoretical model showing how wage-freezes can restrict responsiveness. Section 2 discusses the institutional background and data. Section 3 provides the empirical evidence that responsiveness declined during the recession. Section 4 analyzes the determinants of reporting an income just below notch thresholds, and highlights the employee-employer pairs most likely to report tax-advantaged incomes. Section 5 provides counterfactual distributions to estimate how much of the decline in responsiveness was due to changes in the employee-employer pairings/composition of the labour force. Section 6 concludes.

³An effective tax rate exceeding 100% means reporting a pre-tax income below the notch returns a *higher* after-tax income. I refer to a pre-tax income just below the notch threshold as “tax-advantaged”, with a comparable “tax-disadvantaged” definition for pre-tax incomes just above.

2 Institutional Background and Data

2.1 Tax Collection in Ireland

As they are the individuals required to remit employees' tax payments to the Office of the Revenue Commissioners ("Revenue"), the administrative burden of income taxation in Ireland is largely borne by employers. Employers must register each employee with Revenue when they begin working, and employers are also required to calculate and deduct income tax each time payments are made to an employee.⁴ Taxpayers may elect to file individually or jointly with a spouse. Via a 'tax credit certificate'⁵ issued by Revenue when the employee is registered, the employer has a record of the correct rate of tax to deduct from an employee's pay/wages. At the end of the year, employers are required to file a P35L form to Revenue detailing the total taxable income and amount of taxes deducted for each employee.

The principal income tax in Ireland is the Pay As You Earn (PAYE) tax. There are additional social security-related taxes such as Pay-Related Social Insurance (PRSI). The focus on this paper is on three other taxes that produced notches in take-home pay: the Health Levy, the Income Levy, and the Universal Social Charge (USC). These are individualized taxes, not affected by marital or filing status.

2.2 Health Levy

Officially but less frequently called the Health "Contribution", the Health Levy was instituted by Section 4 of the Health Contributions Act, 1979. An unusual feature of the Health Levy is that its eligibility threshold created a notch in take-home pay. Income below the threshold was exempt from the Levy, but an individual whose annual earnings exceeded the threshold was liable to pay the Levy on the entirety of their income. Earmarked for health expenditure but carrying no additional benefits to payers, the eligibility threshold has a long history of change.⁶ Unfortunately, the data on the behavioural response to the Health Levy prior to the mid-2000s are not available. From the beginning of the dataset through mid-2009, passing the Health Levy threshold incurred a liability of 2% of total income.

2.3 Income Levy

With the deterioration in the public finances caused by the decline in economic activity and the recapitalization of the banking sector, the Irish government announced the introduction of the Income Levy in 2009. The Income Levy operated in much the same way as the Health Levy, but

⁴Revenue exempt firms from this requirement if the employee earns less than €8 per week.

⁵Tax credits are not easily manipulated. As of 2015 a single individual receives a tax credit of €1,650, and this is doubled for married couples. The primary sources of additional credits are those available for widows, carers of incapacitated children or relatives, and blind people.

⁶Announcing changes to the tax system in December 1999, Minister for Finance Charlie McCreevy T.D. stated "As in previous years, the threshold for the payment of the health levy will be increased by £500 from £11,250 to £11,750 per annum." (Oireachtas Debates, Wednesday 1 December 1999)

was activated at a lower income level. With a rate of 2% applicable for two-thirds of the year (1 May–31 December) on incomes of at least €15,028, this implies a liability of €200. In the subsequent year although the threshold did not change, the tax penalty associated with crossing the notch increased by 50%, implying a liability of just over €300.

2.4 Universal Social Charge

The Universal Social Charge (USC) was introduced in 2011 to replace the Health and Income Levies. The threshold was set at €4,004 in 2011 and increased to €10,036 for 2012 and 2013. A 2% liability in the relevant income range was constant over the period. Although neither the threshold nor the penalty changed in the 2013, I include the results for completeness.⁷

Table 1: Principal list of notch thresholds and liabilities

Year	Health Levy	Income Levy	USC	Liability
2005	20,800			(416)
2006	22,880			(458)
2007	24,960			(499)
2008	26,000			(520)
2009	26,000	15,028		(867) (200)
2010	26,000	15,028		(1,040) (301)
2011			4,004	(80)
2012			10,036	(201)
2013			10,036	(201)

All units are denominated in nominal euros. The right-hand column indicates the additional tax liability incurred from exceeding thresholds on the left. For example, earning €1 in excess of €26,000 in 2008 increased tax liability by €520.

Table 1 summarizes the notch thresholds and associated liabilities. Notches are quite an unusual feature in developed countries' income tax systems. It is unprecedented to have such rich variation in notch thresholds in a developed country. The focus of the empirical estimation in Section 3 will be to compare taxpayer responses to these changes conditional on the employer-employee pairing.

2.5 Data

The data for this paper come from four administrative sources provided to the Central Statistics Office. The sources are the P35L employee tax return file, the Central Business Register, the Department of Social Protection's Client Record System, and the Job Churn Statistical Product. The four sources are merged using PPSNs⁸ and constitute a panel of all registered employees and

⁷The primary result for the €4,004 threshold is that no excess bunching is observed. This is consistent with a reduced ability to manipulate income during a recession, but I am less convinced by this result because the €4,004 threshold is further away from the other thresholds.

⁸The PPSN is comparable to a Social Security Number.

employers in the state. Variables include employee identifiers, the number of weeks employed, total taxable pay, month and year of birth, sex, nationality, employer’s form of incorporation and NACE sector, the total number of employees, the number of hires, and the number of separations. All these datasets are complete for 2006–2013.⁹ A 10% random sample of employees was provided to me by CSO. It is a representative panel of the universe of workers from 2006–2013.

There are some limitations of the data, e.g. a lack of educational attainment data or the marital status of individuals. This limitation is relatively minor. Though joint-filing has a significant impact on the primary Pay As You Earn (PAYE) income tax,¹⁰ it does not change the thresholds for the taxes studied here. The levies are individual-level taxes. It is not possible for married couples to avoid these taxes by pooling/separating the incomes.

Minimal cleaning of the data was needed. I removed people earning precisely the amount provided by various welfare programs.¹¹ These schemes support thousands of people, creating additional spikes near the notch points. These are government-sponsored non-market programs, and so I exclude them to avoid contamination with more traditional forms of income.

3 Dynamics of Bunching Behaviour

3.1 Bunching in Levels

The first evidence of taxpayer responses can be seen from histograms of the income distribution about a notch. Figure 1 shows the clear bunching of people earning between €25,991 and €26,000 in 2008.¹² The green dots indicate the empirical distribution of earnings about a notch at €26,000, represented by the vertical red line. Earnings are grouped into €10 bins.

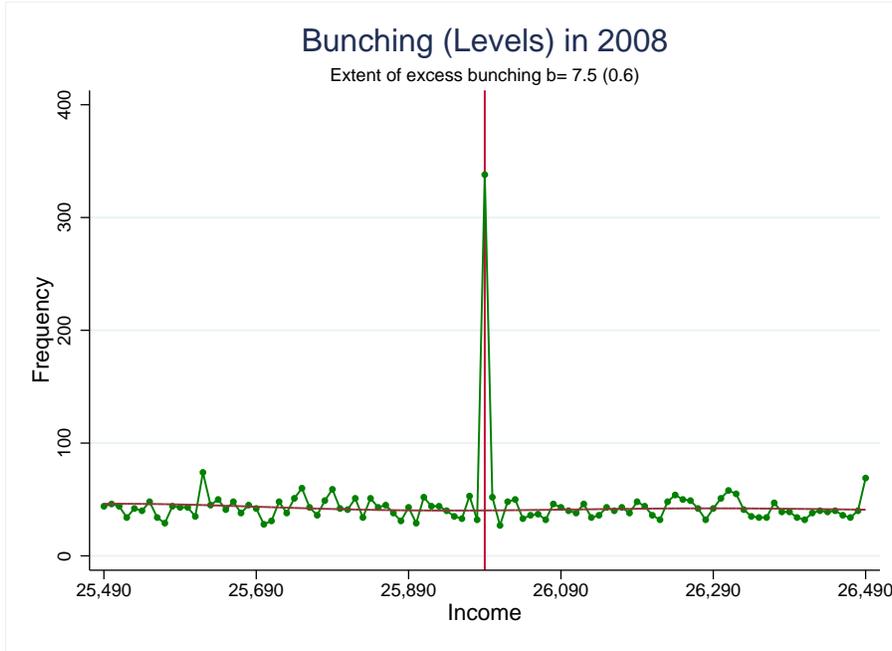
⁹Although it is a non-random sample for that year, the Job Churn data also include the 2005 taxable income for ‘job-stayers’, i.e. the 79% of workers who remained with the same employer in 2005 and 2006. This allows a basis to estimate changes in bunching between 2005 and 2006.

¹⁰The PAYE tax is progressive, so has kinks where the marginal rates change. These kinks are considerably above the notch thresholds.

¹¹I remove people who were on a Community Employment Scheme, receiving Carer’s Allowance, or people reporting earnings precisely at the Spousal Earning Thresholds. Community Employment (CE) schemes have been in effect since the mid-1990s. They are subsidized forms of employment for the long-term unemployed. Approximately 20,000 people are employed in CE Schemes. The Carer’s Allowance is a payment people who care for others (typically other family members) who for medical reasons require full-time assistance. Approximately 50,000 people are in receipt of a carer’s allowance. The Spousal Earning Thresholds generate kinks in the marginal tax rate for the lower-earning spouse.

¹²All notches relate to income exceeding the threshold. Therefore incomes in 2008 of €25,999 or €26,000 were essentially equivalent for tax purposes, but substantially different from an income of €26,001.

Figure 1: Bunching at the 2008 notch point



The vertical axis measures the total number of people in the 2008 data reporting incomes in the €10 bins defined by the horizontal axis. Note the large spike in the number of people earning an income that just avoids the tax liability triggered by passing the notch threshold.

Figure 1 also includes a high-order polynomial of best fit around the empirical distribution excluding a small window about the notch.¹³ This method is original to Chetty et al. (2011).¹⁴ By comparing the true size of the bins with the size predicted by the polynomial, we can compute the excess mass (“bunching”) at the notch. As with Chetty et al. (2011), I find that the results are not particularly sensitive to the precise specification of the counterfactual, e.g. to small changes in order of the polynomial. As the counterfactual is effectively a straight line, this is unsurprising.

We can compare the ratio of the excess number of people to the left of the notch to the number of people predicted by the polynomial. This comparison provides our estimate of the rate of excess bunching b . Bootstrapped standard errors for this estimate are included in parentheses. We can see clear evidence of excess bunching that is statistically significant at all conventional levels. Level-bunching graphs for all notches are included in Appendix B.

3.2 Bunching in Differences

Although the excess bunching is clear in the previous section, it is not apparent whether the bunching is caused by the notch. The requirement needed for parameter identification in level-bunching estimation is that the earnings distribution would be continuous absent the notch. This assumption is routinely made in the literature, e.g. Saez (2010).

However, this assumption is likely violated by several real-world features, the most prominent

¹³The provision for estimation excluding a small window permits non-precise bunching.

¹⁴I am indebted to Tore Olsen for generously sharing the code for this portion of the paper.

of which is round-number bunching. Round-number bunching occurs when there is excess mass in the income distribution at numbers such as €25,000. Casual inspection of the earnings distribution in Appendix F shows very strong evidence for round-number bunching. Furthermore, the extent of round-number bunching is not uniform across different round-numbers (for example, there is considerably more bunching at €26,000 than €27,000) and therefore dummy variables capturing round-numbering are unlikely to adequately control for this phenomenon.

The nature of the quasi-experimental variation in this paper, namely that the location of notch thresholds and/or associated penalties change over time, naturally leads towards an analysis not in levels but in differences. This approach permits a more compelling identification strategy than the levels-bunching estimation previously employed in the literature. In particular, the identification assumption for this estimator is that the change in the income distribution would have been flat absent the introduction of the notch. The difference-bunching approach eliminates concerns about round-number bunching under the assumption that the taste for round-number bunching is not changing from one year to the next.

Figure 2 is the primary figure of this paper. It portrays the change in the number of people reporting income in a particular bin between 2006 and 2012. The measure of excess bunching b is the difference between b_i and $\mathbb{E}[b_i]$. Under stationarity the expected change in the number of people in a bin is zero, and the polynomial of best fit confirms that the expected change in any particular bin for any given year is approximately zero. The red vertical line in the middle of each graph marks the notch threshold for that year. Inference in the difference-bunching estimator is conducted against a null hypothesis of $H_0 : b = 0$ and the calculated bootstrapped standard errors and the full set of individual figures are included in Appendix C.

There is clear and statistically significant bunching-in-differences in the early years. In 2006, 2007, and 2008 we see an excess of people reporting an income within €10 of the notch. This is very precise reporting of annual earnings.¹⁵

It is important to note the years studied in this paper. We have seen the evidence of tax-advantageous reporting through 2008. September 2008 saw Lehman Brothers declaring bankruptcy, one of the iconic triggers precipitating the Great Recession. In June 2008, the unemployment rate in Ireland was 6%. Within a year of the collapse of Lehman, unemployment had reached 13.9%. It would continue to rise through 2010–2011 until peaking at 16% in early 2012, finally dipping below 10% only in 2015. The 2009 tax year saw the introduction of the Income Levy notch at €15,028.¹⁶

In stark contrast to the prior years, the evidence of manipulation disappears in 2009. Note that the vertical axes in Figure 2 are all at the same scale. Bunching, to the extent there is any,¹⁷ is an order of magnitude smaller than the comparable estimates from the prior years. This is true for both the new notch in 2009 and the pre-existing notch that had an increased tax penalty.

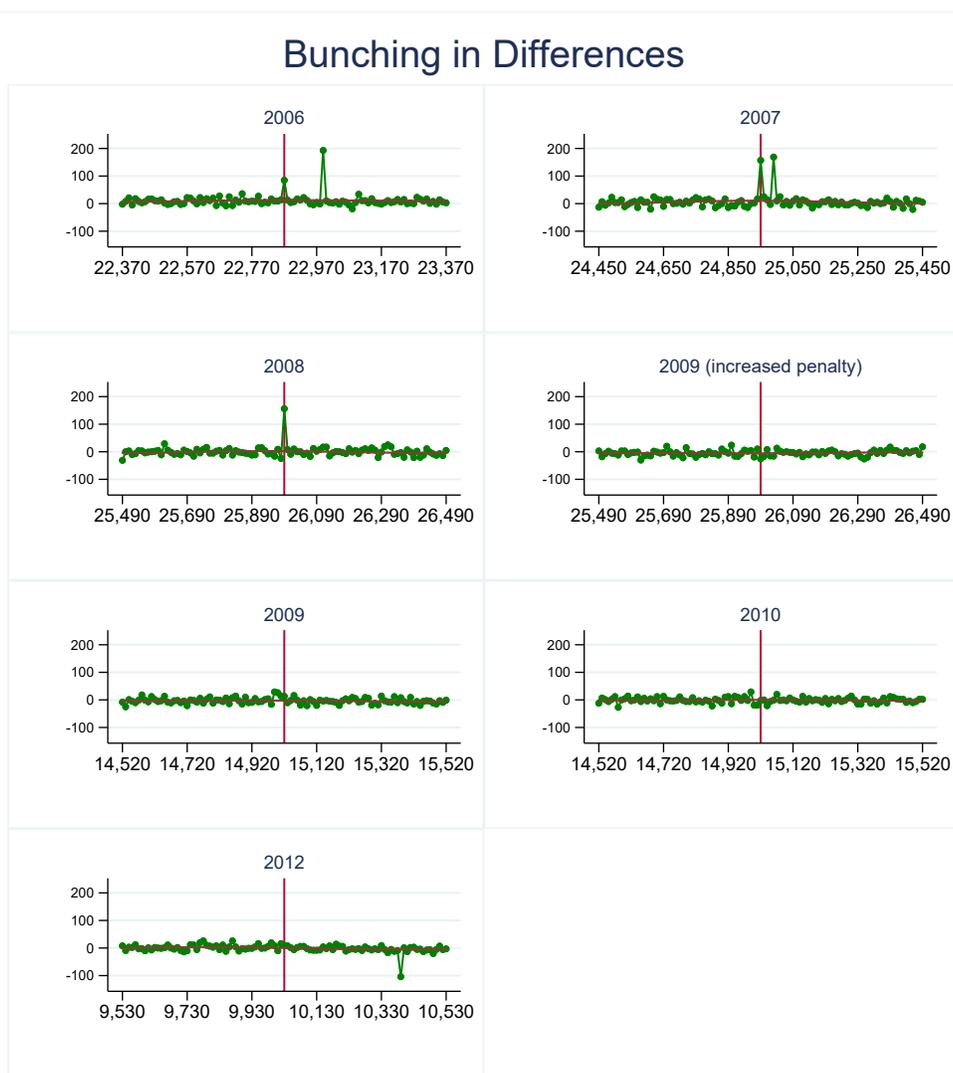
It is clear that the pattern continues in subsequent years. After 2008, there is no evidence of

¹⁵There are additional spikes in 2006, 2007, and 2012.

¹⁶And also the retention of the notch at €26,000. Although the threshold remained the same, the tax penalty for exceeding the €26,000 notch increased from €520 to €867. Due to the 67% increase in the tax penalty, a constant rate of responsiveness/elasticity would imply a substantial increase in bunching in this year.

¹⁷It is not statistically significant.

Figure 2: Difference-bunching through time



The vertical axis measures the change in the number of people reporting incomes in the €10 bins defined by the horizontal axis. Note the large spike in the number of people earning an income that just avoids the tax liability triggered by passing the notch threshold. This ceases in 2009, for both new thresholds and increased penalties on existing thresholds.

manipulation of reported incomes about the notches. A similar pattern is observed for ‘reverse bunching’, decreases in the mass of people when the tax incentives from a notch are removed. This is shown in Figure 12 in Appendix G. There is substantial evidence of reverse bunching in the pre-2009 period, but not after.

When analyzing the responsiveness of taxpayers in bunching designs, it is commonplace to report the elasticity of taxable income (ETI). Theoretically, the ETI can be a sufficient statistic for the excess burden of the tax (Feldstein, 1999; Chetty, 2009). Empirically, the extent of bunching can be mapped into the ETI (Kleven and Waseem, 2013). A more recent literature is critical of

using bunching estimates to infer ETI, suggesting that it is not identified without quite restrictive assumptions (Blomquist and Newey, 2017). Specifically, the actual elasticity cannot be identified without an assumption on the slope of the counterfactual and thus the estimate of the ETI will depend on what assumption the econometrician is willing to make (Bertanha et al., 2018).

The existing best estimate for the ETI of the relevant population in Ireland, using a difference-in-difference identification strategy, is 0.07 (Acheson et al., 2018). This is on the lower-end of estimates relative to those found in other countries, e.g. Gruber and Saez (2002). However, differences in institutional administration means comparing magnitudes across countries is not particularly informative (Slemrod and Kopczuk, 2002). The empirical focus of this paper is the change in bunching in the income distribution. While this approach cannot be directly mapped to an elasticity, it does not require strong structural assumptions and remains informative about responsiveness in one institutional setting over time.

3.3 Additional Evidence on Dynamics

In this section I include two additional measures of the dynamics of responsiveness. The first is the Kolmogorov-Smirnov test of equivalent distributions, and the second is a Poisson fixed-effect regression.

Under an assumption of stationarity, the difference-bunching estimates above provide a clean test of excess bunching at the specified notch point. This section performs a similar test but generalizes the test to differences in distributions rather than differences at a particular point. The Kolmogorov-Smirnov (KS) statistic is a non-parametric way of comparing two distributions. It is robust to transformations that do not affect the relative magnitude of different bins within the window. The KS procedure finds the largest gap between the two distributions, then tests if the largest difference is statistically significant. For this reason it is applicable to testing for bunching.

Table 2: Kolmogorov-Smirnov Test Results

Year	Test Statistic	p -value
2007	0.0236	0.001
2008	0.0230	0.001
2009	0.0180	0.146
2010	0.0158	0.299
2011	0.0101	0.705
2012	0.0086	0.894
2013	0.0078	0.946

The table shows the test statistic and p -values of differences in distributions about the notches from years 2007 through 2013.

The Kolmogorov-Smirnov results for all years are shown in Table 2. As this is a difference test, the first reported year is 2007. Interpretation of the results is straightforward. The distributions

about the notch from 2006 to 2007 return a test statistic of 0.0236, and the p -value easily rejects the null that the distributions are the same. This is because there is excess mass about the notch in 2007 relative to 2006. Similarly, we can reject the null of no excess mass moving from the 2007 to the 2008 income distributions. However, we find no evidence of excess mass for any year after 2008.

The second estimation procedure used to test for excess bunching is more parametric. The difference-bunching method employed in Section 3.2 is based on the changes in the number of people in a €10 income bin. Difference estimates have the advantage of removing any time-invariant properties of particular bins, such as round-number bunching. An alternative empirical approach is to provide each bin with a fixed effect. This will also control for all inherent properties of the bin. Indeed, when $T = 2$, first difference estimation is equivalent to fixed effect estimation.

Table 3: Poisson FE Estimation

	Incident Rate Ratio
Treatment 2006	2.162*** (0.017)
Treatment 2007	2.474*** (0.021)
Treatment 2008	1.211*** (0.010)
Treatment 2009	1.329*** (0.010)
Treatment 2010	1.315*** (0.011)
Treatment 2012	1.006 (0.009)
Treatment 2013	0.937*** (0.009)
Income Bin & Year FE	Yes
N	17,760

Table shows declining IRRs for treatments after the onset of the Great Recession. Treatment in 2009 is the Income Levy, which was introduced that year. The unit of analysis is ten-euro income bins, annually.

Particularly when $T > 2$ there are advantages to pursuing both difference- and fixed effect-based estimation procedures. The approach of this section does just that. Borrowing from the RCT literature, we can interpret a notch as a treatment and test if the treatment has a statistically significant effect on the number of people reporting incomes near the affected bin. One can estimate this model using OLS with bin fixed effects, but as the number of people in a bin is a count (i.e. a

discrete number) it is more appropriate to place restrictions on the error term.¹⁸

One such structure is provided by the Poisson fixed effect model (Cameron and Trivedi, 2005). The Poisson FE estimator is unusual in nonlinear panel models in that it does not suffer from an incidental parameters problem, and is consistent even under error term misspecification. Results are shown in Table 3.

The coefficients of a Poisson regression are non-standard: in particular $\beta = \log(c_1) - \log(c_0)$ where c_i is the expected count of people in an income bin depending on whether the dummy treatment variable is 1 or 0. For this reason I exponentiate the coefficients so they represent incidence rate ratios (IRRs), where the $IRR = e^\beta = \frac{c_1}{c_0}$. The interpretation of the IRR is the factor increase in the count of people in the particular income bin. For example, an IRR of 1.3 indicates a 30% increase in the count of people in the income bin. The null hypothesis for inference is $H_0 : IRR = 1$. Standard errors are robust to heteroskedasticity.

We again see the pattern of declining responsiveness in the latter period of the sample. For example, the coefficient on the Health Levy 2006 dummy indicates that there was more than twice as many people in the effect income bin while that treatment lasted. Note that the calculation of SE's in the Poisson FE model differs from the bootstrapped estimates of the bunching procedure, with consequent differences in statistical significance. Consistent with all previous evidence, we find that the responses are significantly larger in early years. The treatment effects are approximately three times as large prior to 2008 than they are afterwards.

3.4 Interpretation of Results

The empirical evidence on responsiveness raises a number of questions. These questions include whether the recession had a causal effect on behaviour or whether they are just correlated; whether other countries experienced similar declines; and whether the results are real behavioural responses or avoidance/evasion responses.

The principal finding of this paper is that taxpayer responsiveness in Ireland declined at the onset of the Great Recession. While it may be tempting to posit a causal relationship here, there are any number of potentially omitted variables that could violate the exclusion restriction. One natural candidate to explain the decline in responsiveness is a change in enforcement intensity during the recession. Table 4 documents the number of audits by and annual budget of the revenue authority over the period.

The table shows the decline in taxpayer responsiveness is not well-explained by an increase in enforcement activity. The number of audits fell from about 13,500 in 2008 to about 11,000 in 2010. This is consistent with the path of the budget for enforcement, which fell by over 20% over the same period.¹⁹ It is possible that *perceptions* of enforcement followed a different pattern, but as an empirical matter both the operating budget and the number of audits were declining as

¹⁸As a practical matter, the effects reported below are relatively unresponsive to changes in the estimation procedure.

¹⁹I have been unable to ascertain the cost of administration for 2013, though the 2013 Annual Report notes that staff levels were down 13% since 2008.

Table 4: Indicators of Tax Enforcement in Ireland, 2003–2013

Year	Total Audit Interventions	Cost of Administration (€m)
2003	16,029	341.8
2004	16,321	358.6
2005	14,214	378.9
2006	13,626	416.5
2007	14,308	451.5
2008	13,414	482.1
2009	12,419	449.7
2010	11,008	377.1
2011	11,066	372.4
2012	9,066	365.8
2013	8,037	—

responsiveness went down.

While this evidence means the decline in responsiveness cannot be explained by an increase in enforcement, the extent of the sharp bunching in the pre-recession period is suggestive of earnings responding through misreporting rather than real behavioural mechanisms. Other papers have found evidence of misreporting in Ireland. Hargaden and Roantree (2020) looks at bunching at different set of notches in the Irish tax code. While the focus of that paper is quite separate, an appendix studies the real labour supply (hours) response of minimum-wage workers. They find no evidence of hours adjustment in response to increased taxes. To the extent that the behaviour of minimum-wage workers can generalize to the results found in this paper, it suggests most of action is on the reporting margin. A conclusion that misreporting is a key feature of the response is consistent with the precise bunching found in this paper. For example, it is consistent with employer and employee agreeing on a set salary that just avoids the notch threshold.²⁰

The low level of responsiveness found by the bunching estimator is consistent with the low ETI for Ireland found in official estimates (Acheson et al., 2018). After all, it appears that the bunching at tax notches is only of the same magnitude as the bunching at round numbers.²¹ However, it is important to note that bunching estimators only capture precise responses to taxes. Specifically, bunching estimators only capture responses within the tight window (such as a few euro a week) below thresholds. There are surely less precise responses, including extensive margin responses, that bunching estimates are unlikely to capture.

A focus on such precise responses will thus likely be concentrated on tax avoidance. In terms of welfare, tax avoidance is a subtle source of deadweight loss. In simplistic public finance models, workers facing increased taxes substitute towards leisure. This behavioural response creates an unambiguous deadweight loss. Realistically, workers have multiple margins (e.g. spending hours in

²⁰This would represent a transfer between employer and employee, and have no impact of welfare.

²¹Dube et al. (2020) argue that round number bunching is evidence of optimization frictions. As the round number bunching is approximately the same magnitude as the tax-induced bunching, they appear to be equally important features of the data.

the library researching deductions) by which they can lower their reported income to tax authorities. Under neoclassical assumptions, these too generate costs for individuals and should be considered in deadweight loss calculations (Feldstein, 1999). An exception to this noted by Chetty (2009) is when avoidance/evasion creates transfers to other people, such as accountants. Doerrenberg et al. (2017) note that charitable-giving can also provide an exception to this, but the institutional arrangements for charitable-giving in Ireland mean that is less applicable in this setting.

With these context-specific results noted, a natural question is how generalizable these results are to contexts outside of the recession in Ireland. The Great Recession was an international event. Finding evidence of time-varying response in other countries would be suggestive of external validity. Further, it would strengthen the interpretation of the decline being caused by the recession.

Guyton et al. (2016) is a field experiment on the Earned Income Tax Credit (EITC) in the United States. The EITC is an income supplement for low-earners, but requires filing a tax return. The research question relates to eligible people’s non-participation in the program, and if reminder letters could encourage filing, i.e. an extensive margin response. While the experiment itself is not directly relevant to this paper, the institutional background of non-response is. In particular, the authors document a clear increase in non-filing of EITC claims between 2007 and 2009. The figures depicted in Appendix G are taken from that paper. An increase in non-filing for the EITC (a tax credit) is equivalent to a reduction in responsiveness.

Also using the EITC, but this time within the intensive margin at discontinuities in the Federal schedule, Mortenson and Whitten (2020) also exploit administrative data to examine the extent of bunching. The authors find “bunching at the first EITC kink drops in 2008 at the onset of the Great Recession.” Their result matches the findings in this paper. More supporting evidence from the U.S. is found in Buhlmann et al. (2018). Rather than relying on changes in the Federal schedule, Buhlmann et al. (2018) use variation in state-level EITC top-ups to estimate the extent of bunching for the EITC. Again, the authors note a decline in recession-era responsiveness: “while we document a strong relationship up until 2007, we find no effect during the Great Recession.” The evidence on both extensive and intensive margins from below-average earners in the United States is that responsiveness declined at the same time it declined in Ireland. While this is by no means conclusive proof that the declines in responsiveness are caused by economic downturns, the existence of supporting evidence in other countries mitigates concerns that the results were purely institutional specific.

With the international evidence noted, it is worthwhile considering whether we could disentangle real responses versus misreporting. Suppose the response is all evasion. Then we would expect to see earnings responses concentrated in worker-types and industries most able to avoid detection from the authorities. One mechanism used for tax enforcement is third-party reporting (Slemrod et al., 2017; Pomeranz, 2015). When a third-party (such as a credit card company) is compelled to report a transaction, it makes evasion more difficult. There are two main impediments to third-party reporting. The first impediment is industries with many cash transactions, like construction. Hiring someone to fix a gutter is likely transacted with cash and thus not reported to the tax authority.

The second impediment is the existence of self-employment. While self-employed company-owners have some incentives to employ themselves within their own firm,²² the exact amount to declare as labour income versus profit is made without third-party oversight. There is a large literature from other countries indicating that this strategy is used as a mechanism for tax avoidance, e.g. Alstadsæter et al. (2014); Harju and Matikka (2016); Le Maire and Schjerner (2013); Almunia and Lopez-Rodriguez (2018).

While decomposition of the results into real and reporting responses is not possible in the available administrative data, it is possible to investigate the determinants of responsiveness. In particular we can investigate if the characteristics of ‘bunchers’ suggest greater labour market flexibility (such as workers in the retail sector) or greater reporting ability (such as cash-based transactions, or existence of self-employment income). We unpack these determinants in Section 4.

4 Determinants of Responsiveness

Under the relatively innocuous assumptions that income is increasing in hours worked and that individuals enjoy leisure, notches create regions of the income distribution that are strictly dominated by lower pre-tax incomes.²³ By substituting labour for leisure to escape the tax penalty, individuals can increase both consumption and leisure.

One feature of the bunching estimates we saw in Section 3 is the pervasiveness of suboptimal behaviour. Even though bunching is observed in the pre-recession years, more people report income in the strictly dominated region than in an interval just below the notch threshold. This has been found in previous research (Kleven and Waseem, 2013; Chetty et al., 2011) and attributed to broadly defined optimization frictions. While it is obvious that the characteristics of the employee will determine reported income, the nature of tax system administration in Ireland means it is likely that the characteristics of the employer also play a role. One contribution of this paper is to provide insight into the employee-employer matches that display suboptimal behaviour.

Defining a particular notch threshold as x , we expect to observe excess bunching to the left of x . Consequently let us call anyone whose annual income Y is close to (within €100) but below the threshold x , i.e. $Y \in (x - 100, x]$ as a “buncher”. Who are the bunchers?²⁴ Who are the “non-bunchers”, whose income gross income lies in the strictly dominated region?

$$\text{Dominated Region} = \begin{cases} 1 & \text{if income in dominated region} \\ 0 & \text{if income} \in (\text{Threshold}-\text{€}100, \text{Threshold}] \end{cases}$$

One useful feature of the dataset is the inclusion of demographic information on the individuals

²²To enter the social insurance system, for example.

²³When crossing a notch threshold x increases an average tax rate from t_1 to t_2 then the dominated region is $\left(x, \frac{1-t_1}{1-t_2}x\right]$.

²⁴People can be bunchers in one year, non-bunchers the next, not local to the notch the following year, etc. The specific €100 threshold is arbitrary. Other papers have used comparable amounts (Hargaden and Roantree, 2020, use €3 per week=€156 per year, for example) and find similar results.

and basic characteristics on the employers. For employees, the relevant demographics are age, nationality, and sex.²⁵ On the employer side, the most important variables are the number of employees, industry (NACE code), and legal structure. NACE is the standard industry classification code in Europe, comparable to SIC and NAICS in the United States. In particular, the NACE 2 classification system is used.²⁶ Legal structure captures the form of incorporation.²⁷ Another variable of interest is the number of weeks during the year the employee worked for the firm, which I split into an indicator for 52 weeks or not.²⁸ By focusing on taxable income, we incorporate all margins (both legal and illegal) that may adjust to taxation.

The binary nature of the Dominated Region variable leads to the standard analytical approaches for limited dependent variables, namely probit and logit models. Though some instrument would be needed to ascertain causality, estimating the determinants provides insight into the characteristics of those who list incomes in a strictly dominated region relative to a tax-advantaged region. Table 5 pools all years together and shows the determinants of reporting income as tax-advantaged (=0) or in the strictly dominated region (=1) for people local to the notch. The overall picture is consistent regardless of the exact specification, and I show results for two examples: the Health Levy, and “any notch”. The odd-numbered columns in Table 5 show estimates using the probit model, and the even-numbered columns show comparable figures using logits. I report mean marginal effects. Positive coefficients indicate a greater likelihood of registering income in the strictly dominated region.

The results are largely consistent with a model that smaller firms in informal sectors are more likely to report tax-advantageous incomes, similar to results found in Feldman and Slemrod (2007). Having multiple jobs makes one less likely to report a tax-advantaged income. One possible cause is that switching employers mid-year makes precise control of annual income more difficult.²⁹ Consistent with this result, working with an employer for fifty-two weeks of the year implies one is much less likely to be in the strictly dominated region. Being one’s own employer, i.e. having self-employment income, is by far the most robust predictor of reporting a tax-advantaged income.

Different forms of legal incorporation are strongly associated with different reporting behaviours. The legal incorporation coefficients are relative to the base of ‘Sole proprietorships’. The more informal forms of business (sole proprietorships, small partnerships, etc.) increase the likelihood of reporting income in the tax-advantaged region. All coefficients are positive and significant, implying that employees of sole proprietors are the most able to adjust incomes. Sole proprietors are generally smaller companies, and this is consistent with the positive coefficient on firm size. Working for a larger firm makes a person more likely to be in the dominated region. The effect is small, although

²⁵Unfortunately no data on e.g. education are made available to researchers.

²⁶The NACE 2 classification system is a four-digit industry identifier. For example, 85 indicates “Education”, 855 indicates “Other education”, and 8553 indicates “Driving school activities”.

²⁷Individual proprietorship; Partnership; Co-operative society; Public Limited Company; Private Unlimited Company; Private Limited Company; Statutory Body; Branch of a Foreign Company; Other. There is substantial heterogeneity within and between, say, individual proprietorships and public limited companies. Unfortunately, most companies that would be generally classified as foreign-owned are registered for the purposes of taxation as domestic in Ireland, and thus few firms are registered as foreign-held.

²⁸The qualitative interpretation is the same if weeks is used in the continuous form.

²⁹It is also more likely that income trajectories (“career concerns”) are more relevant for those that actively switch employers.

Table 5: Determinants of Exceeding Notch Thresholds

		Health Levy		Any Notch	
		(1)	(2)	(3)	(4)
Personal Characteristics	Number of employers	0.019***	0.021***	0.006	0.006
	Irish	0.003	0.003	0.003	0.003
	EU 2004	0.020**	0.021**	0.020*	0.019**
	Male	0.007	0.006	-0.006	-0.006
	Fifty-two weeks	-0.037***	-0.039***	-0.063***	-0.066***
	Age (decade)	-0.004	-0.004	-0.003	-0.002
	Self-employment	-0.160***	-0.148***	-0.210***	-0.205***
Legal Incorporation	Partnership	0.037**	0.036**	0.042**	0.042***
	Co-Operative	0.053*	0.052*	0.041	0.041
	Public Limited	0.102***	0.105***	0.082***	0.082***
	Private Unlimited	0.106***	0.107***	0.065***	0.063***
	Private Limited	0.054***	0.053***	0.040***	0.039***
	Statutory Body	0.062***	0.063***	0.023*	0.020
	Foreign Branch	0.100***	0.102***	0.048**	0.046**
Other	0.098***	0.104***	0.063***	0.061***	
Industry	Agriculture	-0.035	-0.036	-0.017	-0.016
	Utilities	0.042	0.044	0.042	0.042
	Construction	-0.030***	-0.032***	-0.029**	-0.030**
	Wholesale/retail	-0.020**	-0.020**	0.007	0.008
	Transport	-0.023	-0.024	-0.005	-0.005
	Hotels/restaurants	-0.032**	-0.033**	0.029**	0.030**
	Information Technology	0.015	0.014	0.003	0.002
	Finance	0.020	0.019	0.007	0.006
	Real estate	-0.011	-0.008	-0.008	-0.006
	Professional/scientific	-0.008	-0.010	-0.012	-0.012
	Admin and support	-0.008	-0.010	0.010	0.010
	Public Administration	0.032**	0.032*	0.030**	0.030*
	Education	0.002	0.001	0.012	0.014
	Health/social work	0.004	0.002	0.011	0.012
	Arts	-0.008	-0.008	-0.006	-0.005
Other service	-0.028	-0.028	0.000	0.001	
Within house/other	0.008	0.009	0.027	0.030	
Firm Size	Employees ('000s)	0.002**	0.002*	0.002***	0.002***
	Workforce separation (%)	0.000	0.000	0.000*	0.000**
Sample Size		20,625	20,625	33,508	33,508
Wald χ^2		787.6	815.4	2955.1	2790.3
(Pseudo-)R ²		0.059	0.059	0.084	0.084

Tables show average marginal effects for Probit (odd-numbered columns) and Logit (even-numbered) regressions on reporting an income that fell in the strictly dominated tax-disadvantaged region relative to reporting income below but within €100 of the notch threshold. All regressions include year fixed effects and are clustered at the individual level. Legal incorporation effects are relative to a base value of Sole proprietorships. Industry effects are relative to a base value of Manufacturing.

it is precisely estimated.

The sectoral indicators are relative to the manufacturing industry. The sectors with negative coefficients relative to manufacturers, and therefore the sectors with the lowest tendency to report incomes in the dominated region, are construction and agriculture (though agriculture is insignificant). In both of these sectors, substantive portions of business are conducted in informal settings. The sectors with the largest positive coefficients are utilities and public administration. The utilities in Ireland are almost all state-owned enterprises. Construction firms avoid taxes, but government agencies do not. This finding is consistent with Paulus (2015).

Men are more likely to be in the dominated region for the Health Levy than women, but less likely to be in the dominated region for other notches. This result is perhaps surprising, given the evidence that women have a larger labour supply elasticity (Devereux, 2004), though statistically insignificant. Irish people are slightly more likely to be in the dominated region than non-natives, but this result is again not significant. Citizens of countries that only joined the EU in 2004, who generally have little experience with the Irish tax system, are more likely to report tax-disadvantaged incomes.

Overall we conclude that the characteristics of both employer and employee (the “employer-employee pairing”) predict reporting behaviour. The importance of this result becomes more apparent when one considers the employer-employee pairings most likely to be affected by the economic downturn in Ireland.

5 Intertemporal Comparisons of Bunching

While the observed decline in aggregate responsiveness is quite apparent, the cause is not. Why did the behaviour of taxpayers as a whole change? It seems intuitive that income effects would make tax minimization more enticing during a recession.

One potential avenue for explaining the decline in responsiveness comes from behavioural economics: salience effects. There is convincing evidence that consumers are relatively unresponsive to less salient taxes (Finkelstein, 2009; Chetty et al., 2009). While there is no direct way to measure salience and other behavioural explanations, dramatic declines in salience seem unlikely to be driving the results. The introduction of the USC was extensively covered by the media (Reilly, 2014; Feehan, 2015; Flanagan, 2016). Furthermore, Google Trends data included in the appendix indicate, if anything, an increase in salience of the taxes over time.

A larger issue seems to be the effects of the macroeconomy on the employee-employer pairings. A sufficient condition for a change in average responsiveness is non-random attrition from the labour force. If a sector-specific shock adversely impacts the employment conditions of a group, such as young semi-skilled men, then we expect average responsiveness to change.³⁰

As emphasized above, Ireland experienced severe a severe labour market shock. The high unemployment rate omits the large number of discouraged workers: in a country with a population

³⁰Unless avoidance behaviour is identical in all sectors in the economy.

Figure 3: Industry-specific shocks to the labour market

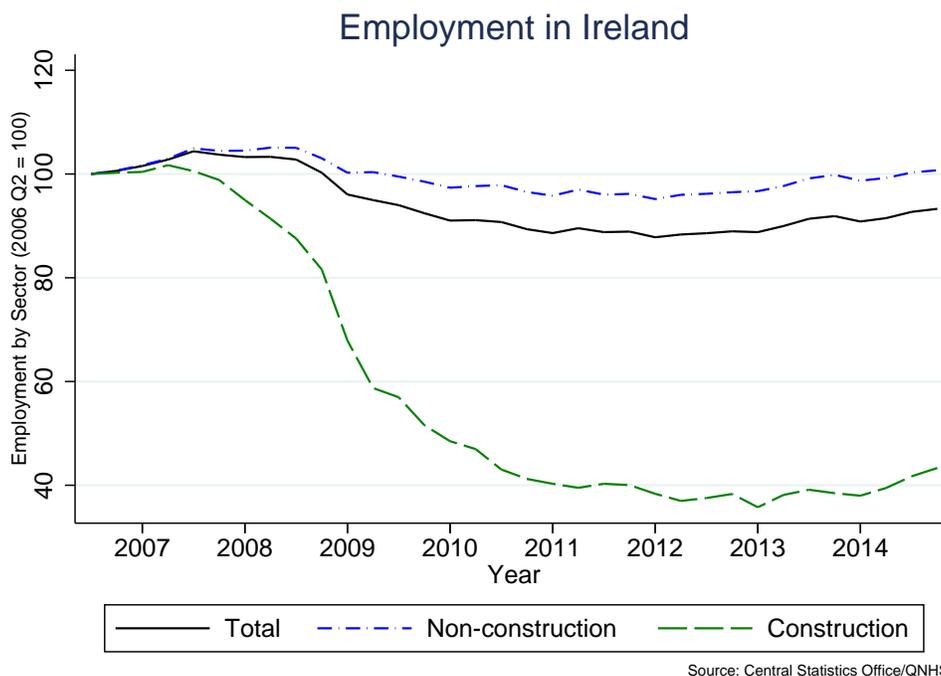


Figure shows the dynamics of employment in construction and non-construction sectors. While non-construction employment did drop by approximately 10% during the recession, more than half of construction workers lost their jobs.

of 4.5 million, the labour force declined by 340,000 people. The effect was widespread but fell disproportionately on the construction sector, an industry traditionally associated with employment of younger men. Figure 3 shows that the level of employment in all non-construction sectors was approximately constant over the decade, whereas the numbers employed in construction fell by more than half.

With such a large change in the composition of the labour force, it is instructive to generate counterfactual wage distributions. One simple approach would be to re-analyze the data without construction workers.³¹ More sophisticated approaches exist, and these can decompose results with better precision.

5.1 DFL Counterfactual Estimates

The method developed in DiNardo, Fortin and Lemieux (1996), hereinafter “DFL”, is well-suited to a decomposition of bunching at notches and has been used previously in the literature (Onji, 2009). Rather than bluntly removing certain segments of the data, the DFL technique asks what would the income distribution look like in year t had it the characteristics from year $t - 1$.

It is important to note that strong assumptions are needed to interpret these compositions

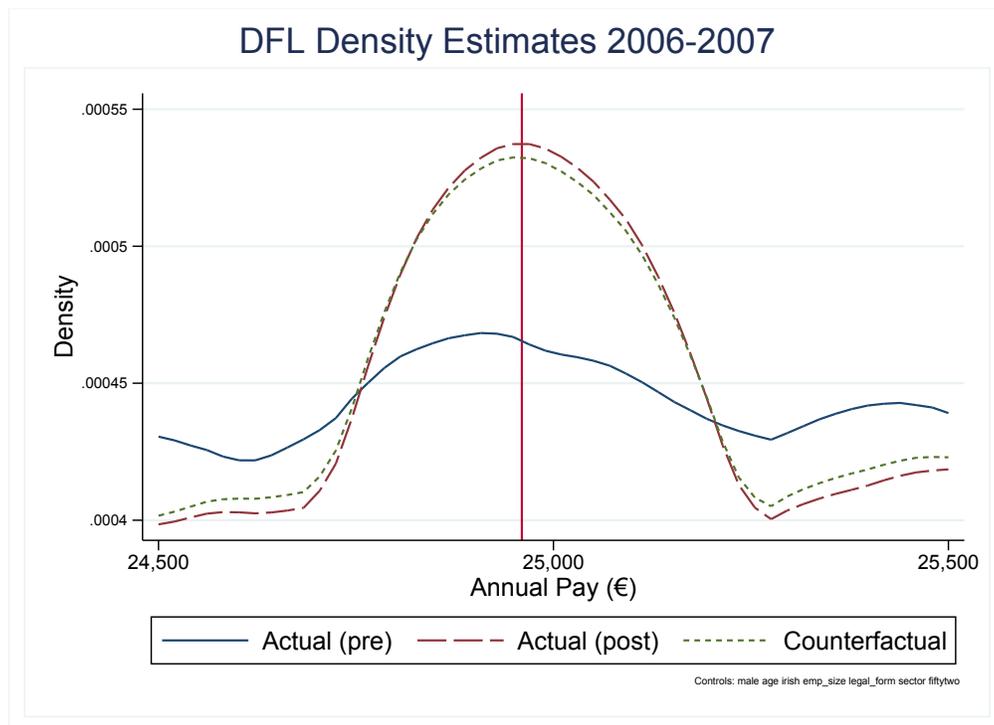
³¹And this is included in Appendix D.

causally. The technique ‘controls for’ compositional effects using a suite of variables (such as firm size) that are likely themselves endogenous. I include the decompositions to document that compositional factors can be statistically important in explaining bunching behaviour. The decompositions are motivation for Section 5.2, which will re-estimate the determinants of bunching adjusting for these compositional shifts.

The DFL decomposition is a generalization of the method of Oaxaca and Blinder. The Oaxaca-Blinder decomposition can be seen as measuring how the average pay of how one group (e.g. women) would be paid if they had another group’s (e.g. men’s) characteristics. Instead of focusing attention exclusively on the mean, the DFL technique generalizes this intuition to the entire distribution of wages.

When applied to this paper, the DFL technique constructs a counterfactual wage density after the introduction of a notch had the employer-employee attributes remained as they were before the notch. That is, it decomposes the excess bunching into a component explained by changes in the characteristics of the employer-employee match and an unexplained residual component. The component explained by differences in attributes is the gap between the actual post-notch distribution and the generated counterfactual distribution. The residual component is attributed as due to the notch.

Figure 4: Counterfactual estimates of bunching based on changes in covariates



The figure shows actual and counterfactual wage densities for 2006 to 2007.

Figure 4 depicts an example of this exercise. The figure has three lines. Two are the observed

densities before (2006) and after (2007) the introduction of the notch. From these two lines we observe excess bunching at the notch the year of its introduction. The dotted line is the counterfactual estimate. The counterfactual is constructed by rescaling the actual post-notch distribution to account for changes in the characteristics of the employer-employee match. As we can see in Figure 4, the counterfactual almost perfectly matches the actual post-notch distribution. The conclusion is that controlling for changes in characteristics has only a small effect. This is intuitively quite sensible because it occurs before the labour market shock.

Figure 5 depicts a DFL decomposition for a later notch, this time in 2009-2010. There are two takeaways from this graph. The first is that there is less observed excess bunching in the real, empirical densities. This is consistent with the evidence of reduced bunching presented earlier. The second takeaway is that, relatively speaking, the counterfactual does not match the actual post-notch density as well as in previous years. This means the changing composition of the labour force appears to be more important in this case.

Figure 5: Counterfactual estimates of bunching based on changes in covariates

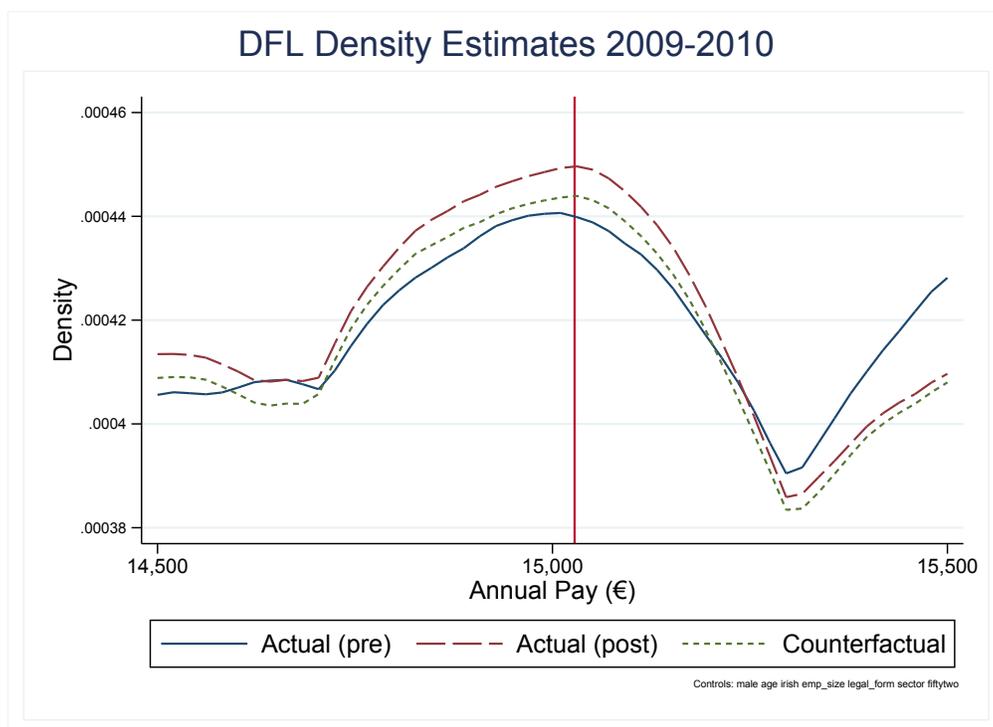


Figure 5 shows the counterfactual wage density from 2009 to 2010. Changes in the covariates do a less impressive job matching the 2010 density.

The full set of counterfactual density estimates are shown in Appendix E. Collectively, the evidence indicates that changes in labour market composition can be an important consideration to understand the observed change in bunching. Figure 5 suggests that the importance of the compositional change may be increasing over time. It is thus important to investigate how much of the change in bunching is explained by changing characteristics. One way to do this is to re-estimate

the determinants of bunching over time, holding the labour force constant.

5.2 Inverse Probability Counterfactual Estimates

The preceding analysis suggests that while the changes in the labour force can help explain the intertemporal pattern in bunching behaviour, its explanatory power is neither complete nor constant. Holding the composition of the labour force fixed, do the predictors of bunching themselves change over time? Do recessions affect not just the circumstances of those that lose their job, but also the circumstances of people that remain matched with the same employer?

In addition to the construction of DFL counterfactuals, I investigate the determinants of bunching by reweighting the post-recession labour force to match the covariates to those of their pre-recession counterparts. Conditional on the effectiveness of the reweighting procedure, this holds the composition of the labour force constant between the two periods.

Any two samples can be reweighted to look similar to each other in observed characteristics. I present 2006 and 2010, as they represent two years of the same notch (the Health Levy) either side of 2008. This reweighting is achieved through inverse probability weighting. Inverse probability weighting is a two-step procedure in which observations are given weights according to their probability of attriting from the sample based on their covariates. For example, 2010 observations who work in construction are given additional weight to compensate for the loss of jobs in construction between 2006 and 2010. I base the reweighting on covariates that are less controllable by worker alone: age, gender, Irish citizenship, EU citizenship, and an indicator for being in the construction sector.

The idea of omitting choice variables from the reweighting variables is to minimize endogenous selection, e.g. switching between firms. As a practical matter, the results are relatively insensitive to variations of this procedure, i.e. including a full set of sector dummies. Holding these variables constant, the likelihood of reporting an income in the dominated region increased by 6.3% ($p < 0.01$) between 2006 and 2010. The interpretation of this is of reduced responsiveness, even controlling for declines in the construction sector. Expanding the sectoral controls from just construction to including the full set of sectoral dummies, the point estimate changes of this recession ‘treatment effect’ falls to 5.7% ($p < 0.01$). Reweighting a large set of covariates to match the earlier sample, the probability of reporting a tax-advantaged income during the recession was still lower.

Table 6 shows how the IPW-reweighted 2010 sample compares to the IPW-reweighted 2006 sample. The reweighting procedure is based on people who earned between €20,000–€30,000. The procedure does a good job: the characteristics are quite similar. For example, recall from Figure 3 that the fraction of workers employed in construction decreased by more than half (from 12% to 5%) between 2006 and 2010. After reweighting the samples to adjust for propensity to attrit from the sample, the difference collapses to essentially zero. The only nonzero difference is for age, with the average age of the reweighted samples about one month different. These differences are not statistically significant.

Using the reweighted sample, we can analyze the determinants of bunching absent any bias due

Table 6: Sample means of selected variables before and after inverse probability weighting

	Unweighted			Weighted		
	2006	2010	Difference	2006	2010	Difference
Age	36.48	38.29	1.82	37.82	37.72	-0.11
Construction sector	0.12	0.05	0.08	0.07	0.07	0.00
Irish	0.67	0.69	0.02	0.69	0.69	0.00
Male	0.48	0.47	0.01	0.48	0.48	0.00
EU 2004	0.11	0.11	0.00	0.11	0.11	0.00

Table shows the means of variables after reweighting the 2006 and 2010 samples to have statistically indistinguishable characteristics. The primary difference after reweighting the samples is the convergence of the share of workers in construction.

Table 7: Determinants of reporting a tax-disadvantaged income after inverse probability reweighting

	IPW 2006	IPW 2010
Fifty-two weeks	-0.0233	-0.0519***
Firm Size ('000s)	-0.0006	0.0037**
Number of employers	0.0163	0.0171
Self-employment	-0.1262***	-0.1633***
Age (decade)	-0.0050	-0.0052
Irish	0.0132	-0.0011
Male	-0.0027	0.0205**
EU 2004	0.0308	0.0144
Workforce separation (%)	-0.0004*	0.0002
Sample Size	3,133	5,514
Wald χ^2	90.04	294.88
(Pseudo-)R ²	0.0325	0.0909

Table shows selected mean marginal effects from probit regression on the determinants of bunching. The first column shows the results for the notch at €22,880 in 2006. The second column shows the results for the 2010 notch, but with the results weighted so the covariates match those of the 2006 sample.

to changes in labour force composition. Table 7 presents evidence that the economic environment of those that remained in employment changed, and in particular that the determinants of bunching varied. The first column of results in Table 7 shows the mean marginal effects of a selection of covariates from a probit on the probability of reporting a tax-disadvantaged income in 2006. The right-hand column presents the results for the 2010, but reweighting the sample to ‘look like’ the 2006 sample. We can thus interpret the right-hand column of Table 7 as the determinants of bunching as if the composition of the labour force were unchanged from the 2006 column.

The overall impression is similar, but note that the coefficients do change. For example, the

marginal effect of working for the same employer 52 weeks a year more than doubles in magnitude. Men are more likely to report a tax-disadvantaged income, suggesting that women are more adept/flexible with earnings. Note also that the marginal effect of workforce separation³² switches sign (though towards statistical insignificance). This could indicate that the natures of separations vary, for example from voluntary to enforced, between 2006 and 2010. The self-employment indicator is highly significant and is the largest absolute determinant in both specifications. Frictions for regular workers appear even larger in recessions: the marginal effect of having any self-employment income increases by about a third.

It is reassuring that the results between the 2006 sample and the 2010 reweighted sample largely coincide. Nonetheless, the marginal effects move a non-trivial amount. As discussed above, workforce separation change sign in an important way (as do firm size, and sex). Some coefficients move from statistically insignificant to significant. Recall the earlier results that the evidence of bunching disappears. The regression results above demonstrate that this cannot be attributed solely to a composition effect. Holding worker characteristics constant, there is a differential tendency to report tax-advantaged earnings during the recession.

Perhaps there are additional constraints on reporting behaviour during recessions. The precise source of the additional constraints on behaviour can be speculated on. One interpretation is that it is harder to find opportunities to avoid taxes during recessions, for example by switching jobs or targeting overtime hours. Owners and managers may be less receptive to worker requests when labour is in abundant supply, or if they suspect the firm is not viable in the long run. Alternatively, social norms may constrain renegotiation of wages during a time of redundancies. These are far from the only possible explanations; many other plausible interpretations exist to could help explain the changes in bunching behaviour. Further research could shed light on this.

6 Conclusion

The extent to which taxpayers respond to changes in tax rates is a central parameter in public finance. In this paper I investigate if responsiveness is constant over time. It is not. Using administrative tax return data from Ireland, I find that responses were three times larger at the peak of Ireland’s business cycle than at the trough, despite the treatment intensity being largest at the height of the recession. The clear evidence of tax-advantageous reporting behaviour in the pre-recession period is not statistically significant in any year after 2008. Although the findings relate to Ireland, research designs using administrative data in the United States have found similar results for both nonfiling (Guyton et al., 2016) and bunching behaviour at the federal- (Mortenson and Whitten, 2020) and state-level (Buhlmann et al., 2018).

The characteristics of the employer-employee pairing are good predictors of reporting a tax-advantaged income. For example, individuals with self-employment income tend to successfully avoid exceeding notch thresholds, and construction workers report more tax-advantaged incomes

³²The fraction of employees who separate from the firm in that year.

than comparable workers in public utilities.

Sector-specific shocks can alter the overall composition of the labour force. Counterfactual wage distributions indicate that some of the decline can be explained, at least in a statistical sense, by changes in the composition of the labour force. The recession disproportionately reduced employment in sectors which had previously exhibited above-average ability to report tax-advantaged incomes. Using inverse probability weights to control for changes in the labour force, I find that the determinants of earnings responses themselves change over the cycle. This result is consistent with the presence of additional constraints on reporting tax-advantaged earnings during a recession.

Though these findings are significant contributions to the literature on the behavioural response to taxation, I suggest that they may be of interest to related fields. The importance of labour force composition (Solon et al., 1994) and frictions (Michaillat, 2012) to taxable responses, and highlighting the role of employee-employer pairings in tax minimization strategies, suggest an avenue for further research in labour economics. The potential interaction of nominal rigidity and taxpayer responsiveness (Grigsby et al., 2019) should be of interest to macroeconomists.

Indeed, an entire macroeconomics literature exists about whether the magnitude of the fiscal multiplier varies over time (Auerbach and Gorodnichenko, 2013; Ramey and Zubairy, 2018). This paper identifies an analogous effect on the revenue collection side. The evidence is that the response of taxpayers to changes in tax rates was substantially lower during the recession. A time-varying (and particularly a business cycle-dependent) taxpayer response has considerable implications for policy implementation. A decreased taxpayer response generates a lower deadweight loss of taxation, which implies the costs of taxation could be lower during a recession; alternatively a decreased ability of taxpayers to avoid taxes suggests less resources need to be allocated to enforcement during recessions.

References

- Acheson, Jean, Seán Kennedy, Edgar Morgenroth, and Brian Stanley**, “Do Marginal Tax Rates Matter?,” *Department of Finance Discussion Paper*, March 2018.
- Adam, Stuart, James Browne, David Phillips, and Barra Roantree**, “Frictions and taxpayer responses: evidence from bunching at personal tax thresholds,” *IFS Working Paper Series*, August 2017.
- Almunia, Miguel and David Lopez-Rodriguez**, “Under the Radar: The Effects of Monitoring Firms on Tax Compliance,” *American Economic Journal: Economic Policy*, February 2018, *10* (1), 1–38.
- Alstadsæter, Annette, Wojciech Kopczuk, and Kjetil Telle**, “Are closely held firms tax shelters?,” *Tax Policy and the Economy*, 2014, *28* (1), 1–32.

- Auerbach, Alan J. and Yuriy Gorodnichenko**, “Fiscal Multipliers in Recession and Expansion,” in Alberto Alesina and Francesco Giavazzi, eds., *Fiscal Policy after the Financial Crisis*, University of Chicago Press, 2013, pp. 63–98.
- Bertanha, Marinho, Andrew H McCallum, and Nathan Seegert**, “Better Bunching, Nicer Notching,” *Working paper*, 2018.
- Blomquist, Sören and Whitney Newey**, “The bunching estimator cannot identify the taxable income elasticity,” *National Bureau of Economic Research Working Papers*, 2017.
- Buhlmann, Florian, Benjamin Elsner, and Andreas Peichl**, “Tax Refunds and Income Manipulation: Evidence from the EITC,” *International Tax and Public Finance*, 2018, *25*, 1490–1518.
- Cameron, A. Colin and Pravin K. Trivedi**, *Microeconometrics: Methods and Applications*, Cambridge University Press, 2005.
- Chetty, Raj**, “Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance,” *American Economic Journal: Economic Policy*, August 2009, *1* (2), 31–52.
- , **Adam Looney, and Kory Kroft**, “Salience and taxation: Theory and evidence,” *American Economic Review*, 2009, *99* (4), 1145–77.
- , **John N. Friedman, Tore Olsen, and Luigi Pistaferri**, “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records,” *The Quarterly Journal of Economics*, May 2011, *126* (2), 749–804.
- Devereux, Paul J.**, “Changes in Relative Wages and Family Labor Supply,” *The Journal of Human Resources*, July 2004, *39* (3), 698–722.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux**, “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, September 1996, *64* (5), 1001–44.
- Doerrenberg, Philipp, Andreas Peichl, and Sebastian Siegloch**, “The elasticity of taxable income in the presence of deduction possibilities,” *Journal of Public Economics*, 2017, *151*, 41 – 55.
- Dube, Arindrajit, Alan Manning, and Suresh Naidu**, “Monopsony and Employer Misoptimization Explain Why Wages Bunch at Round Numbers,” Working Paper 24991, National Bureau of Economic Research January 2020.
- Feehan, Conor**, “Vote us in and we’ll abolish the USC, promises Lucinda Creighton,” *Evening Herald*, January 5 2015.

- Feldman, Naomi E. and Joel Slemrod**, “Estimating Tax Noncompliance with Evidence from Unaudited Tax Returns,” *The Economic Journal*, March 2007, *117* (518), 327–352.
- Feldstein, Martin**, “Tax avoidance and the deadweight loss of the income tax,” *Review of Economics and Statistics*, 1999, *81* (4), 674–680.
- Finkelstein, Amy**, “E-ztax: Tax salience and tax rates,” *The Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.
- Flanagan, Pat**, “Over 80% of taxpayers want USC abolished and believe the State taxes the poor too much,” *Irish Mirror*, June 20 2016.
- Grigsby, John, Erik Hurst, and Ahu Yildirmaz**, “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data,” *NBER Working Paper 25628*, 2019.
- Gruber, Jon and Emmanuel Saez**, “The Elasticity of Taxable Income: Evidence and Implications,” *Journal of Public Economics*, April 2002, *84* (1), 1–32.
- Guyton, John, Day Manoli, Brenda Schafer, and Michael Sebastiani**, “Reminders & Recidivism: Evidence from Tax Filing & EITC Participation among Low-Income Nonfilers,” *NBER Working Paper Series*, January 2016.
- Hargaden, Enda Patrick and Barra Roantree**, “Does Statutory Incidence Matter? Earnings Responses to Social Security Contributions,” *Working paper*, February 2020.
- Harju, Jarkko and Tuomas Matikka**, “The elasticity of taxable income and income-shifting: what is “real” and what is not?,” *International Tax and Public Finance*, 2016, *23* (4), 640–669.
- 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**, “Tax bases, tax rates and the elasticity of reported income,” *Journal of Public Economics*, 2005, *89* (11-12), 2093–2119.
- Le Maire, Daniel and Bertel Schjerning**, “Tax bunching, income shifting and self-employment,” *Journal of Public Economics*, 2013, *107*, 1–18.
- Michaillat, Pascal**, “Do Matching Frictions Explain Unemployment? Not in Bad Times,” *American Economic Review*, June 2012, *102* (4), 1721–50.
- Mortenson, Jacob A. and Andrew Whitten**, “Bunching to Maximize Tax Credits: Evidence from Kinks in the U.S. Tax Schedule,” *American Economic Journal: Economic Policy*, Forthcoming 2020.
- 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.

- Paulus, Alari**, “Tax Evasion and Measurement Error: An Econometric Analysis of Survey Data Linked with Tax Records,” ISER Working Paper Series 2015-10, Institute for Social and Economic Research May 2015.
- Pomeranz, Dina**, “No taxation without information: Deterrence and self-enforcement in the value added tax,” *American Economic Review*, 2015, *105* (8), 2539–69.
- Ramey, Valerie A and Sarah Zubairy**, “Government spending multipliers in good times and in bad: Evidence from US historical data,” *Journal of Political Economy*, 2018, *126* (2), 850–901.
- Reilly, Jerome**, “Left should unite and scrap USC,” *Irish Independent*, November 23 2014.
- Rubinstein, Ariel**, “Perfect equilibrium in a bargaining model,” *Econometrica*, 1982, pp. 97–109.
- 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.
- Slemrod, Joel and Wojciech Kopczuk**, “The Optimal Elasticity of Taxable Income,” *Journal of Public Economics*, April 2002, *84* (1), 91–112.
- , **Brett Collins, Jeffrey L Hoopes, Daniel Reck, and Michael Sebastiani**, “Does credit-card information reporting improve small-business tax compliance?,” *Journal of Public Economics*, 2017, *149*, 1–19.
- , **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.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker**, “Measuring the Cyclicity of Real Wages: How Important Is Composition Bias?,” *The Quarterly Journal of Economics*, February 1994, *109* (1), 1–25.
- Tithe an Oireachtais**, “Official Record of Dáil and Seanad Debates, 1919–2015,” 2015. Accessed on July 19, 2015.
- Weber, Caroline E**, “Toward obtaining a consistent estimate of the elasticity of taxable income using difference-in-differences,” *Journal of Public Economics*, 2014, *117*, 90–103.

A Theoretical Motivation

Though a large literature analyzing taxpayer behaviour exists, prevailing models typically omit external macroeconomic conditions. To take Saez (2010)'s EITC paper as an example, when faced with a kinked tax schedule workers with quasi-linear utility choose their labour supply to optimize their consumption-leisure decision. Exogenous macroeconomic conditions are arguably under-analyzed: for example, the job vacancy rate could be an important element in explaining taxpayer responsiveness to kinks and notches. To incorporate these factors, I model the wage decision as a bargain between employees and employers with stochastic outside options. Commencing renegotiation carries an exogenous risk of breakdown, and thus employees will occasionally not commence bargaining even in the case of mutually beneficial renegotiation.

I motivate the model's primary elements (employee-employer bargaining, and outside options) with some stylized facts. Self-employed people are more responsive to taxes than regular employees.³³ This fact suggests a role for the employer-employee match in tax reporting, and Chetty et al. (2011) find employers' posted offers do influence employees' behaviour. For outside options, the job vacancy rate fluctuates throughout the business cycle.³⁴ This fluctuation directly affects the availability of outside options.

Figure 6: Substantial increase in pay freezes for those remaining in jobs

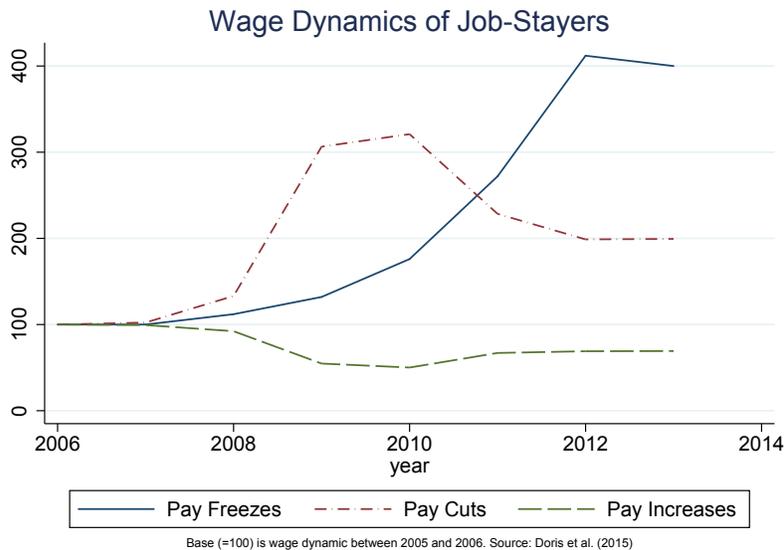


Figure shows the increasing prominence of pay freezes from 2009 onwards. The fraction of the work force accepting pay freezes was four times higher in 2012 than in 2006.

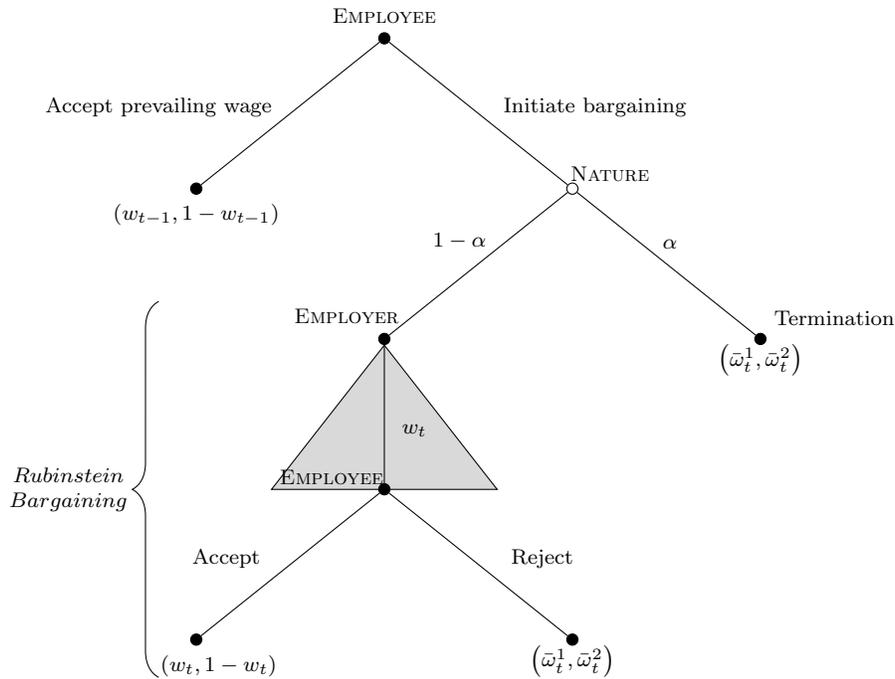
³³In an RCT on tax compliance, Slemrod et al. (2001) found treatment effects were statistically significant only for those with self-employment or farm incomes. Adam et al. (2017) finds the most responsive workers to tax incentives are company owner-managers.

³⁴The seasonally-adjusted US job vacancy rate was 240% higher in the second quarter of 2018 than the same period in 2009. The comparable figure is 320% for Ireland. Sources: BLS Job Openings and Labour Turnover Survey, CSO Earnings Hours and Employment Costs Survey.

Nominal rigidity, in the form of pay freezes, also influences employer-employee bargaining. Documenting the importance of macroeconomic factors for bargaining, Grigsby et al. (2019) state “any model with a constant fraction of wage adjustments. . . will fail to match the wage setting patterns over a business cycle.” Although the labour market in Ireland changed most dramatically for those who lost their jobs, conditions also changed for those who remained in employment. For example, within the subset of employees who remained matched with the same firm, Figure 6 shows that the dynamics of wage setting changed dramatically over the period.³⁵

Motivated by these facts, I modify the Rubinstein (1982) wage bargaining model to incorporate stochastic outside options and pay freezes. In the model, employees have the initial option to accept the prevailing wage or to initiate bargaining. Continuing with the prevailing wage can be seen as a safe option. An exogenous probability of a breakdown in negotiations induces a real tradeoff between accepting the prevailing wage and initiating bargaining. This tradeoff has the analogue to the real world that formally raising the issue of conditions is not costless and sometimes results in termination.

Figure 7: Graphical depiction of wage-bargaining model



In the model, Employee and Employer bargain over the distribution of surplus π_t .³⁶ The employee share of π_t is denoted $w_t \in [0, 1]$. Last period the players agreed a distribution of $(w_{t-1}, 1 - w_{t-1})$ going to the Employee and Employer, respectively.

The Employee chooses whether or not to initiate bargaining at the start of the game. If the

³⁵Figure 6 also documents the rise of pay cuts. This reflects reduced base pay, and also fewer bonuses and overtime opportunities.

³⁶Matching the macroeconomic framework, π_t is lower during recessions.

Employee accepts the prevailing wage, the game ends, and last year's wage persists. Assuming negotiations begin, then with probability α Nature terminates the contract and the Employee receives the outside option $\bar{\omega}_t^1 > 0$. These outside options are high or low, depending on the macroeconomic environment.

With probability $1 - \alpha$ bargaining continues, and the Employer proposes a distribution of surplus $(w_t, 1 - w_t)$. The Employee may accept this offer, or reject with a counter-offer. This delay from counter-offering is costly due to discount rates δ_1 and δ_2 for Employee and Employer respectively.

Formally, counter-offers may continue indefinitely. The key insight of Rubinstein (1982) is that although this game is notionally indefinite, under modest restrictions the unique equilibrium is that the first offer is accepted. My modification of this game continues this tradition, incorporating the addition of accepting the prevailing wage as one option.

When an Employee's discount rate is δ_1 and the Employer's is δ_2 , the unique equilibrium split of the surplus is that the Employer receives $1 - w_t = \frac{1 - \delta_1}{1 - \delta_1 \delta_2}$ and the Employee receives the complementary $w_t = \frac{\delta_1(1 - \delta_2)}{1 - \delta_1 \delta_2}$. This division is scaled up by the time-varying benefits π_t .

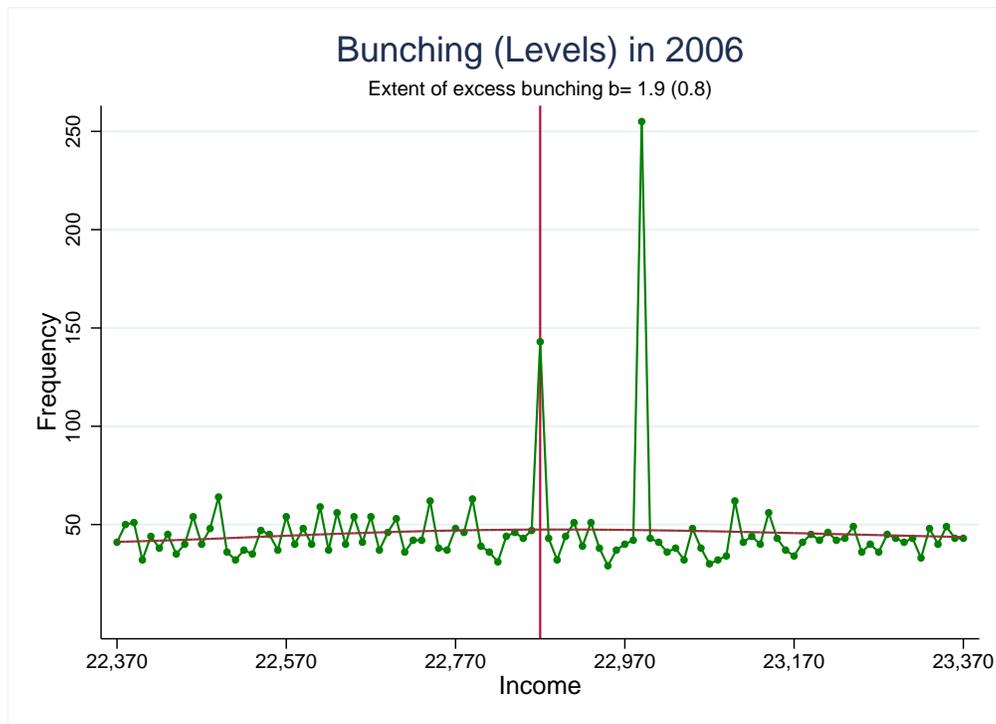
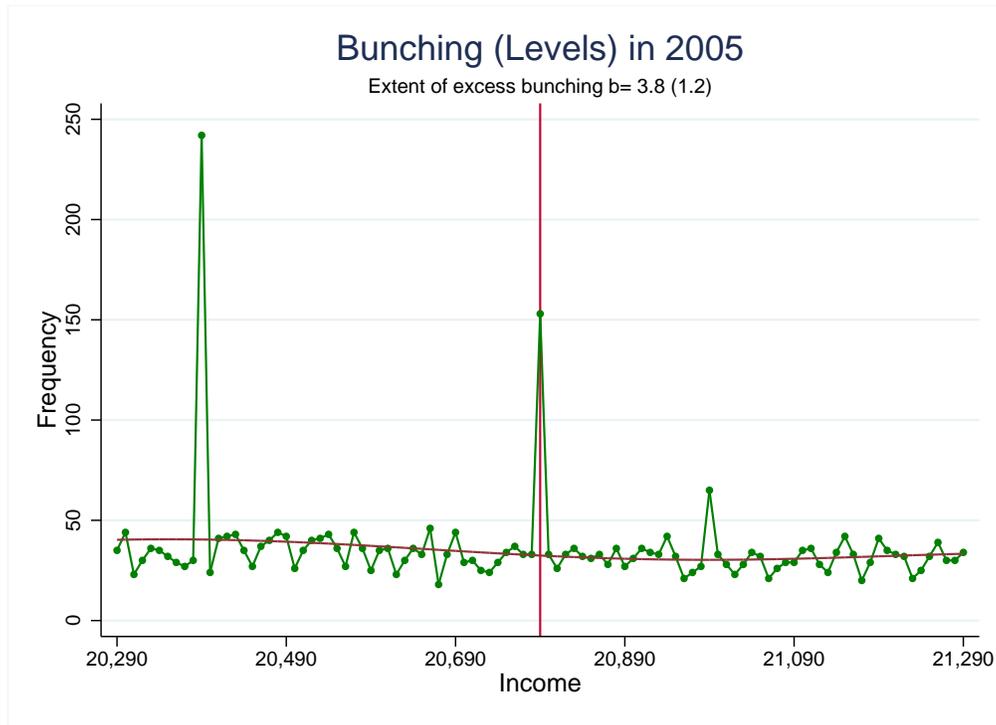
Given the payoff structure, we can use backward induction to see if the Employee will initiate negotiations. In equilibrium, initiating negotiations will with probability α breakdown and result in payoff $\bar{\omega}_t^1$ and with probability $1 - \alpha$ payoff $\frac{\pi_t \delta_1(1 - \delta_2)}{1 - \delta_1 \delta_2}$. With a utility function $u(\cdot)$ that satisfies the Inada conditions,³⁷ the expected utility from initiating bargaining is $\alpha u(\bar{\omega}_t^1) + (1 - \alpha) u\left(\frac{\pi_t \delta_1(1 - \delta_2)}{1 - \delta_1 \delta_2}\right)$. This risky option is to be contrasted with the sure wage of $\pi_{t-1} w_{t-1}$, which returns utility $u(\pi_{t-1} w_{t-1})$. It is a relatively simple decision, depending on the parameters, whether to engage in negotiations or not. If π_t is sufficiently larger than π_{t-1} , negotiations commence.

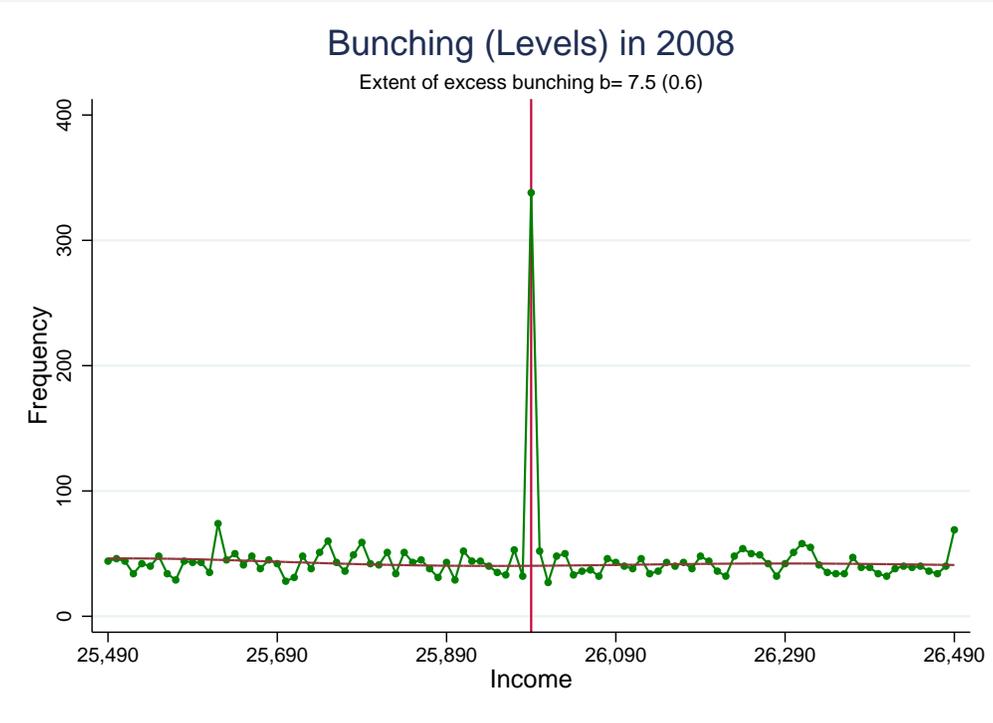
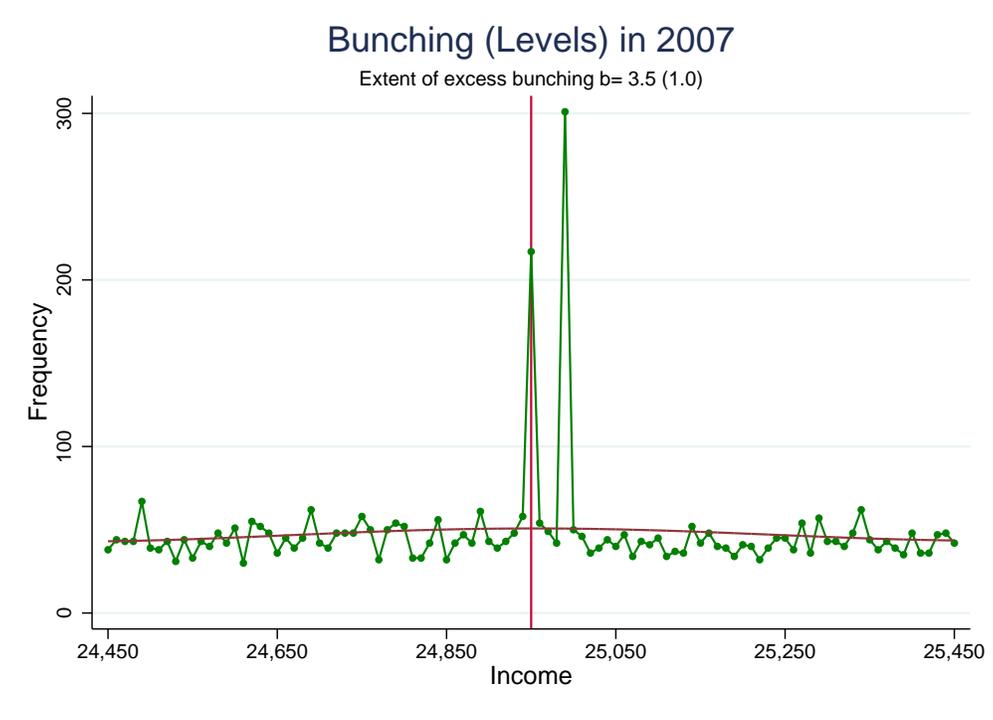
Define ρ as the fraction of employees accepting the existing wage structure, i.e. those that accept pay freezes. It is clear that both $\frac{\delta \rho}{\delta \pi_t} < 0$ and $\frac{\delta \rho}{\delta \bar{\omega}_t^1} < 0$. In a recession, both π_t and $\bar{\omega}_t^1$ are small. Consequently ρ is higher during recessions. Consistent with the pattern we observe in Figure 6, more job-stayers will accept wage freezes when outside options are poor.

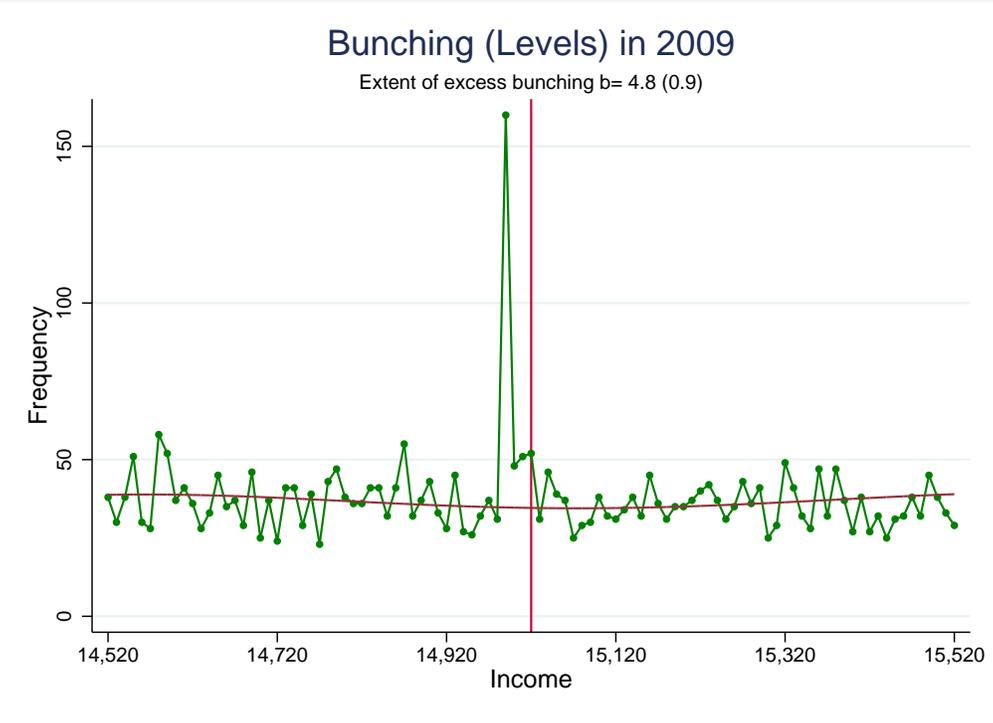
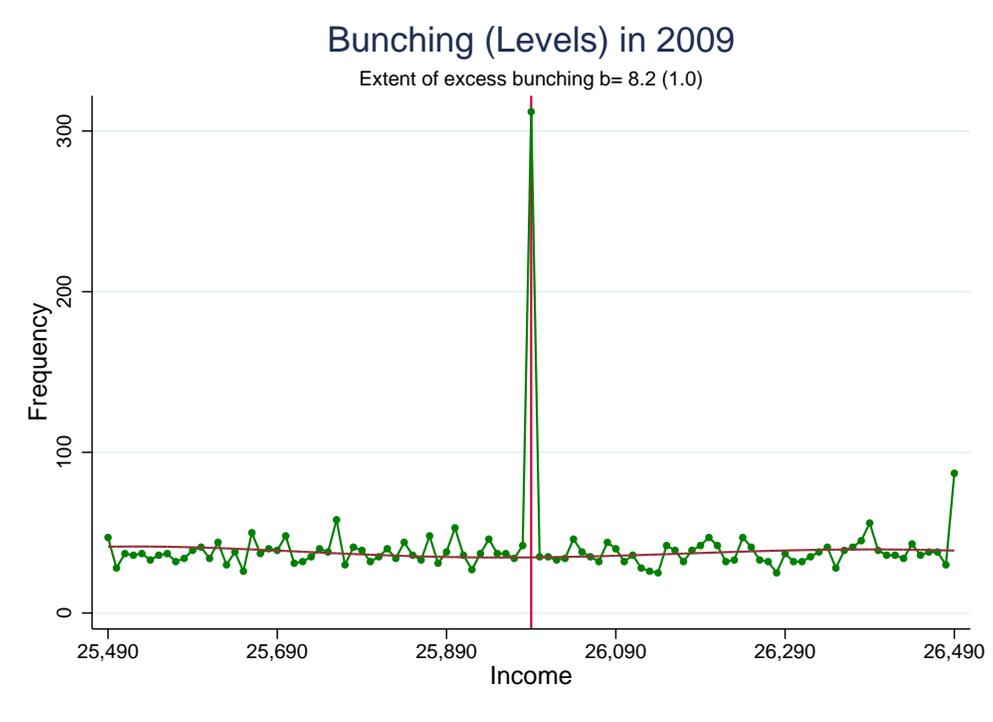
This simple model, though stylized, matches and highlights some key features of wage dynamics over the business cycle. Employers and outside options are relevant considerations for employees' earnings. Demanding renegotiation of wages is not costless, especially with an employer facing financial difficulty. It can be entirely rational to accept a wage freeze if such a strategy minimizes the likelihood of unemployment. Diminished outside options and pay freezes combine to restrict the capacity of workers to respond to changes in taxation.

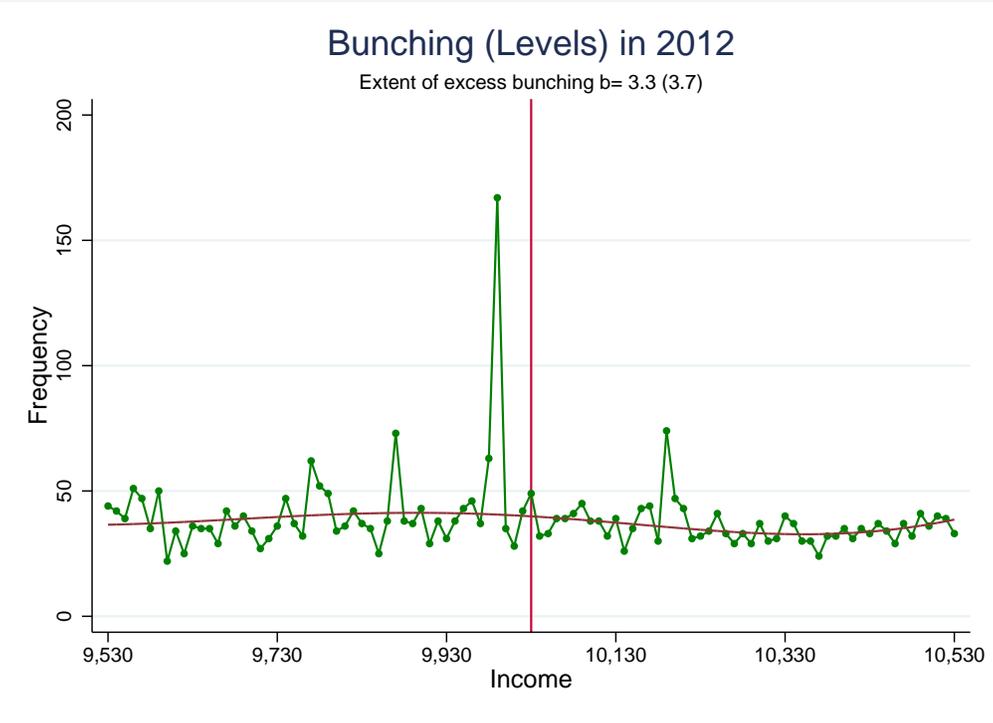
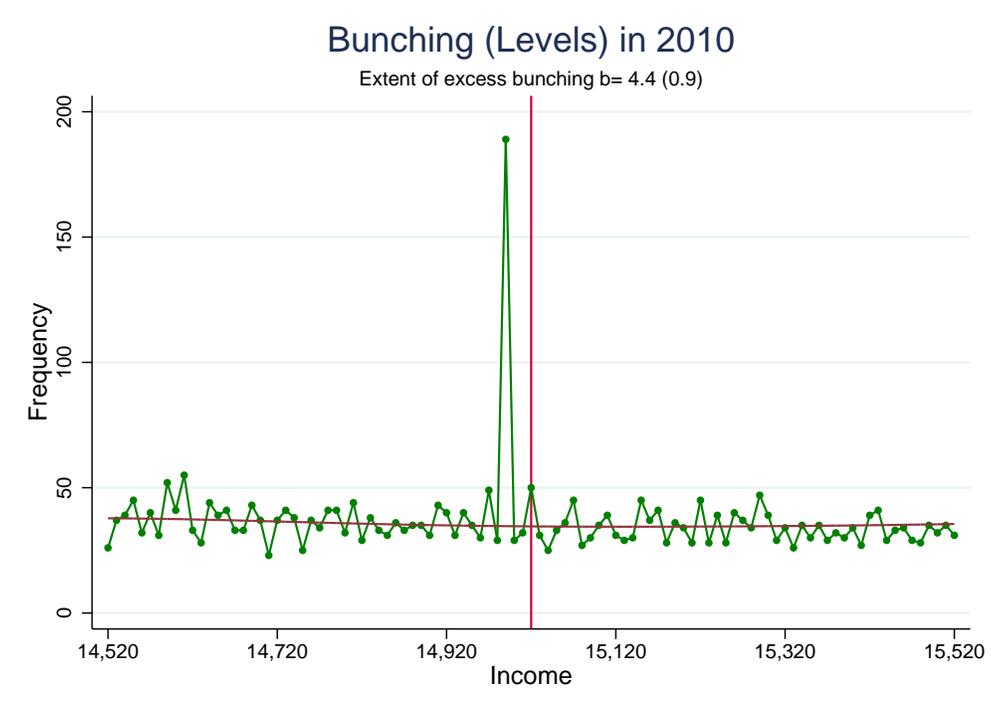
³⁷This ensures an interior solution.

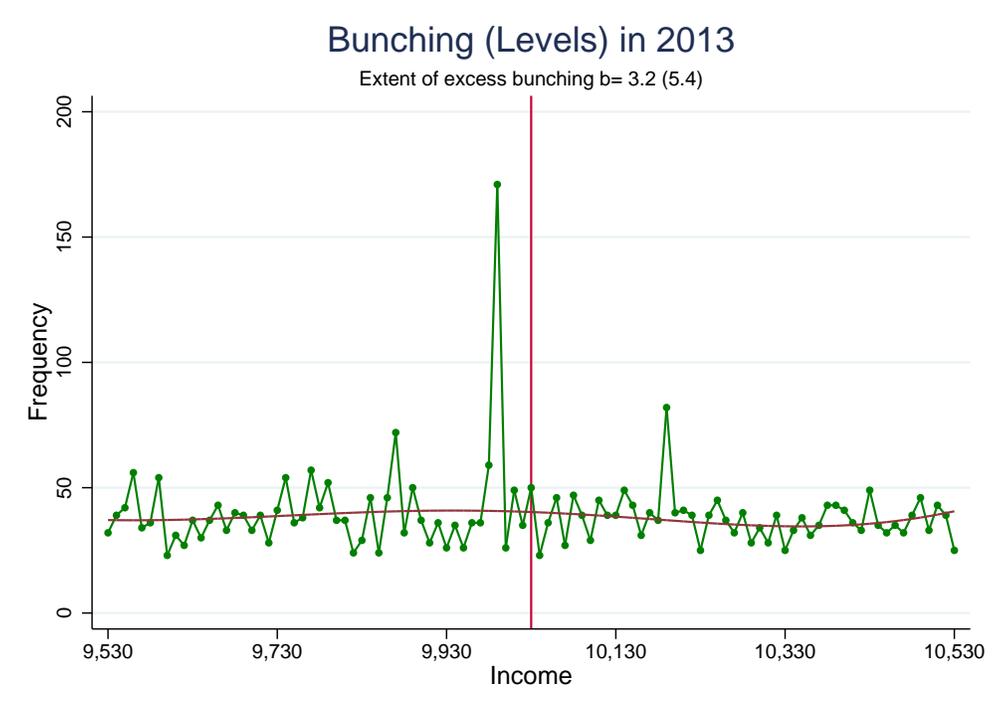
B Bunching (Levels) Figures



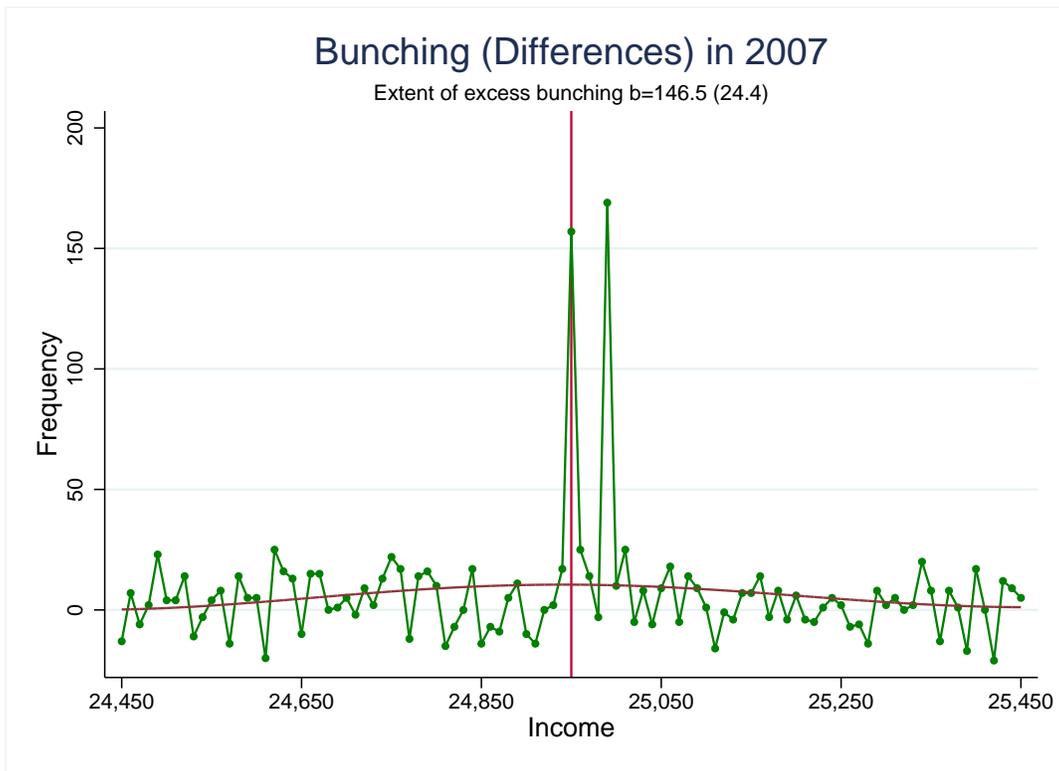
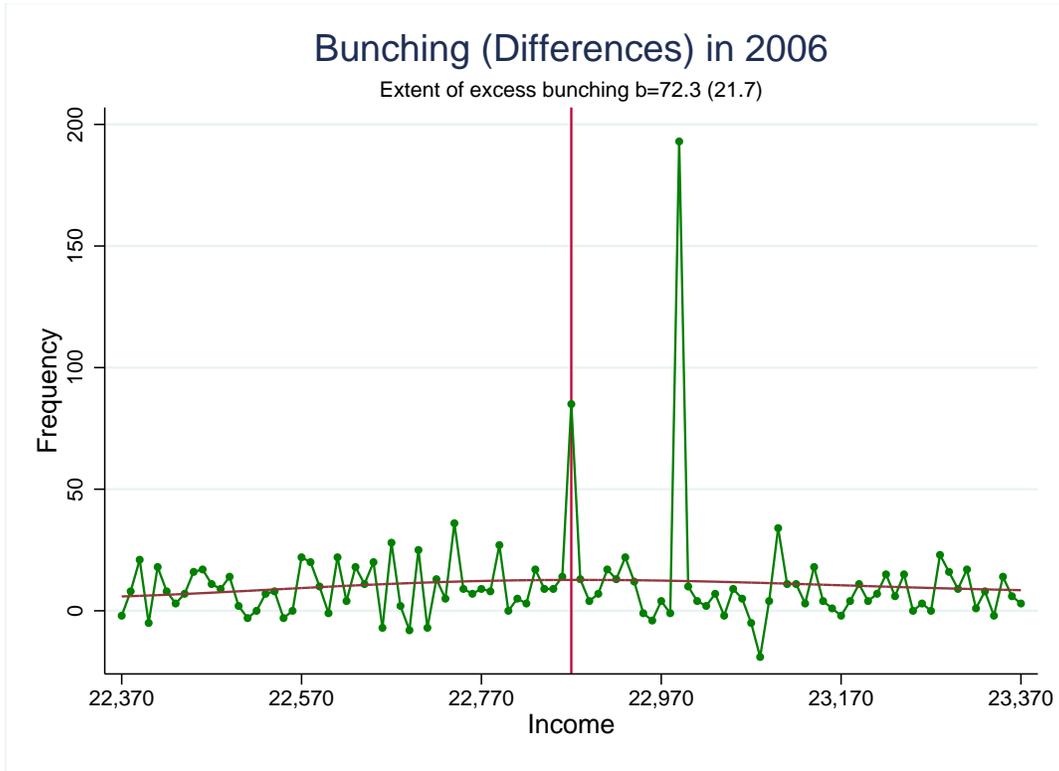


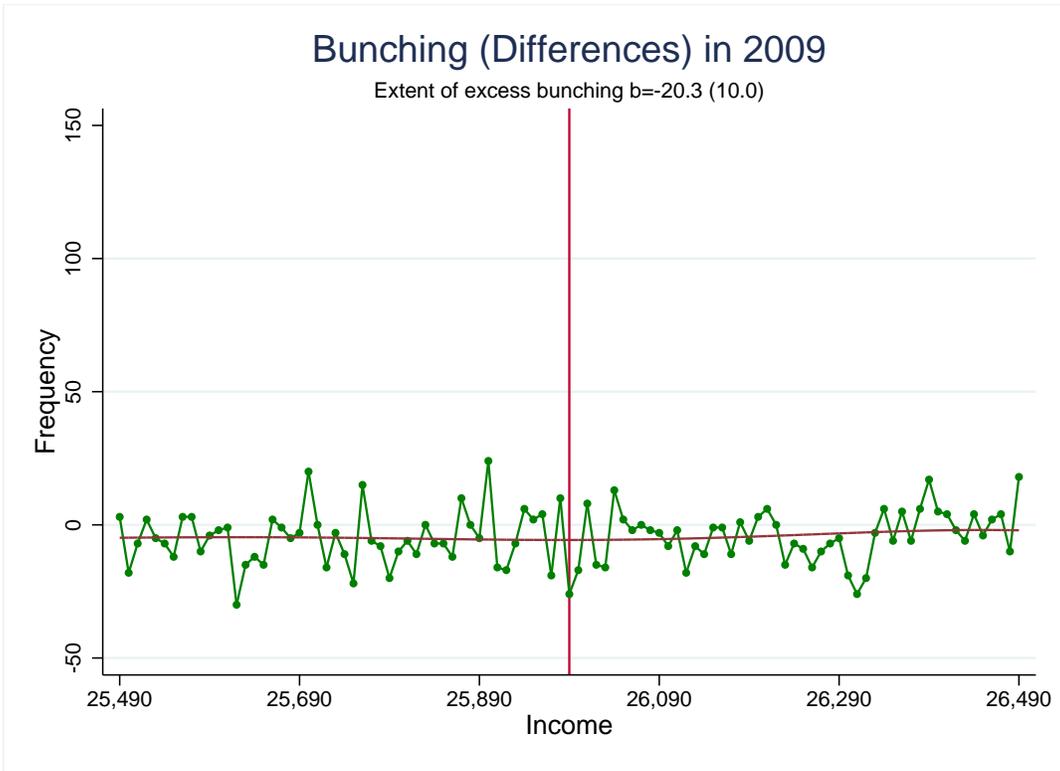
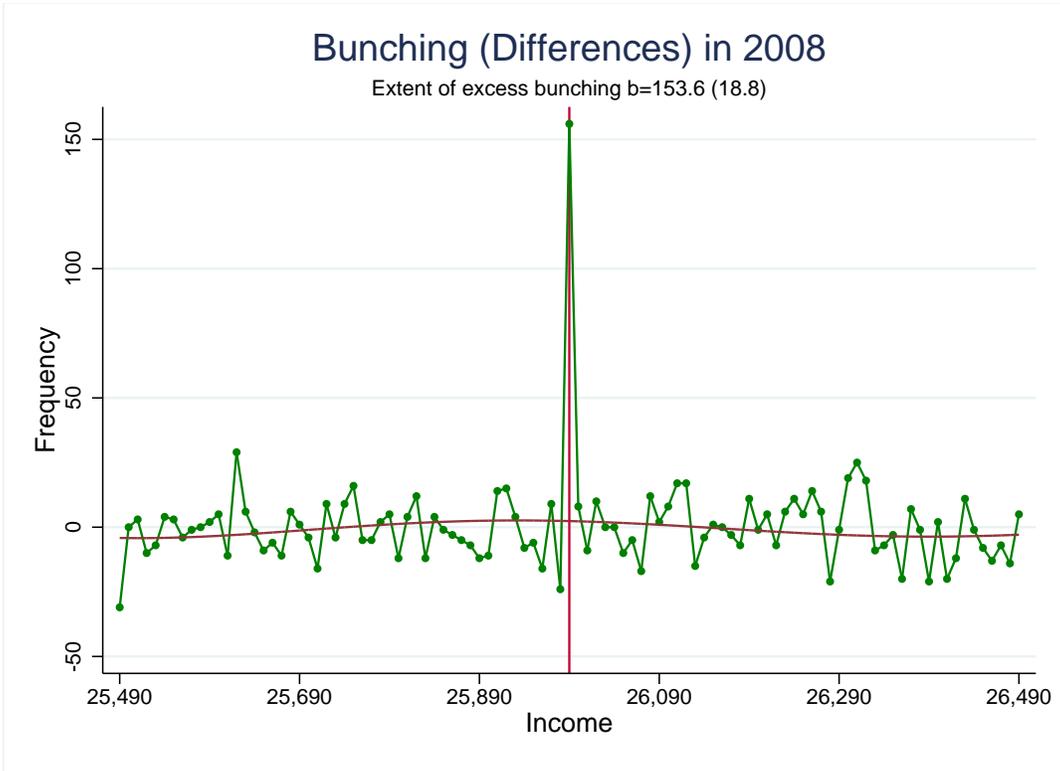


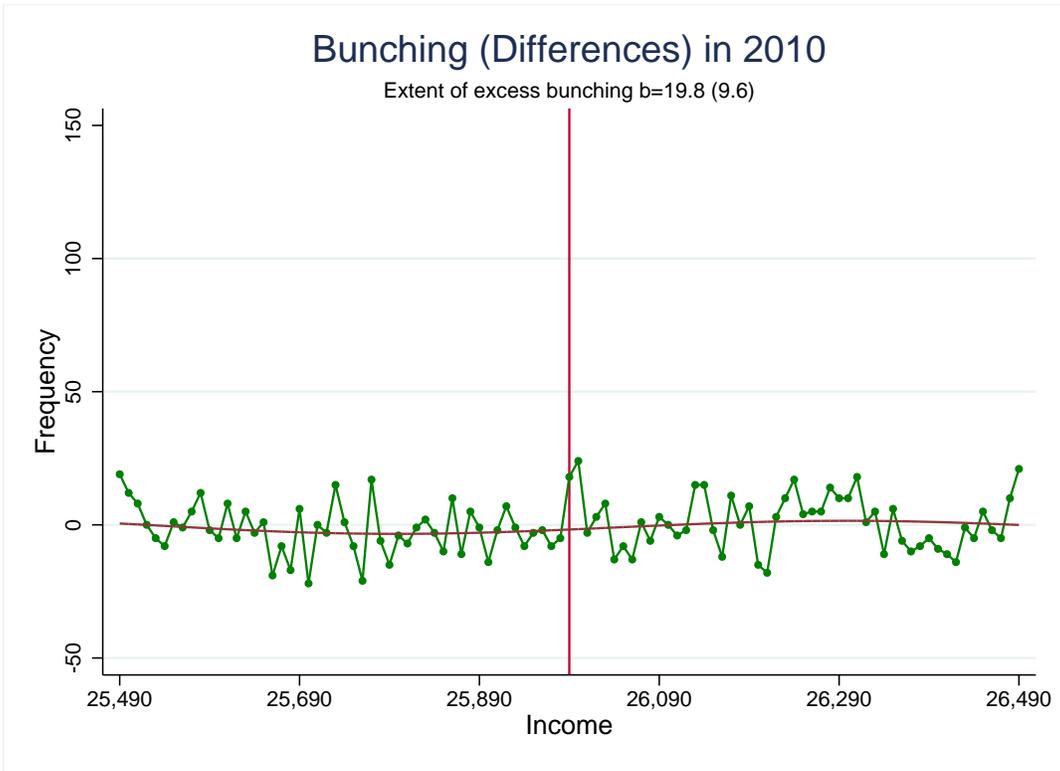
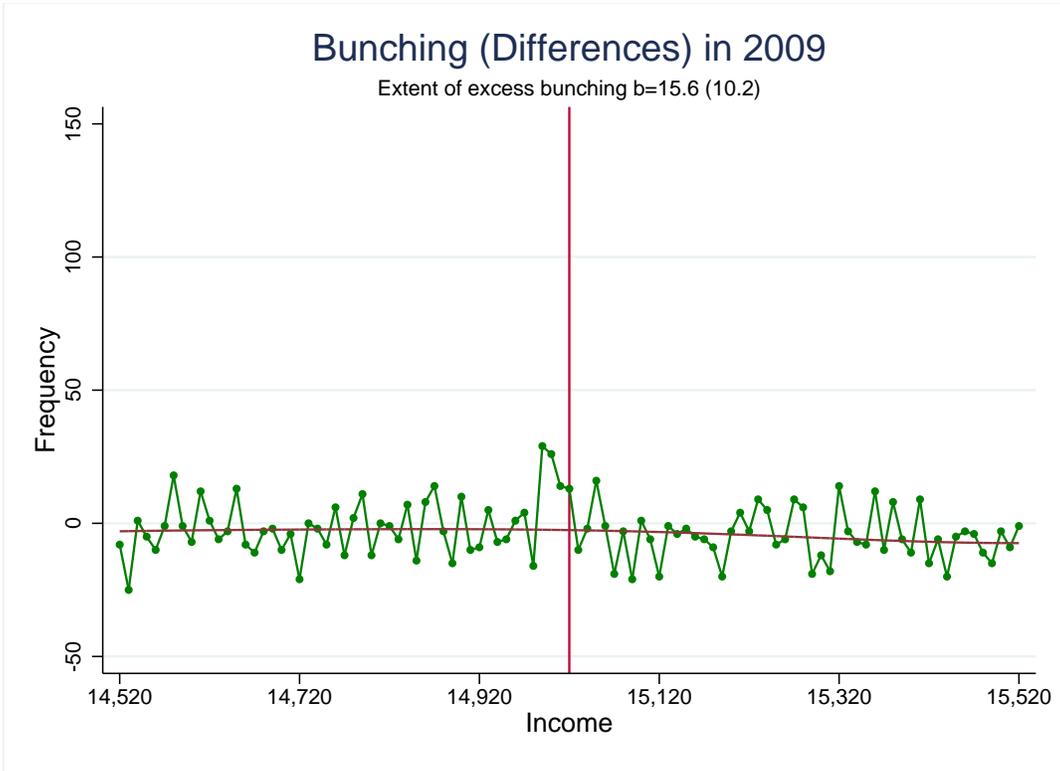


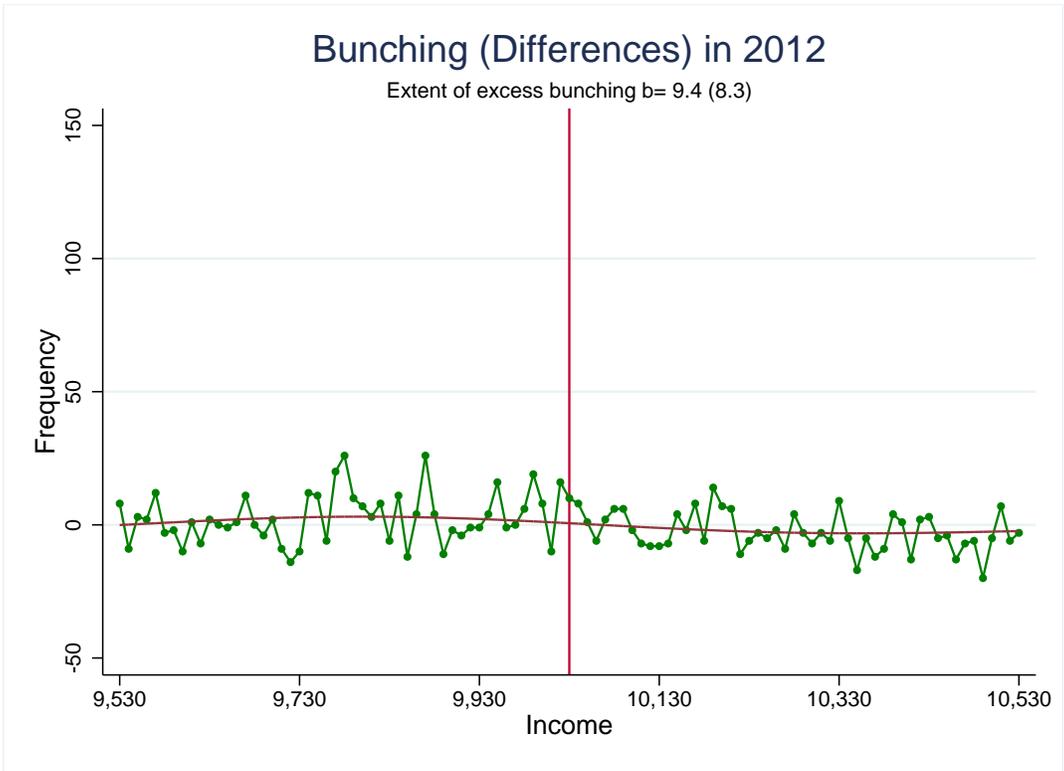
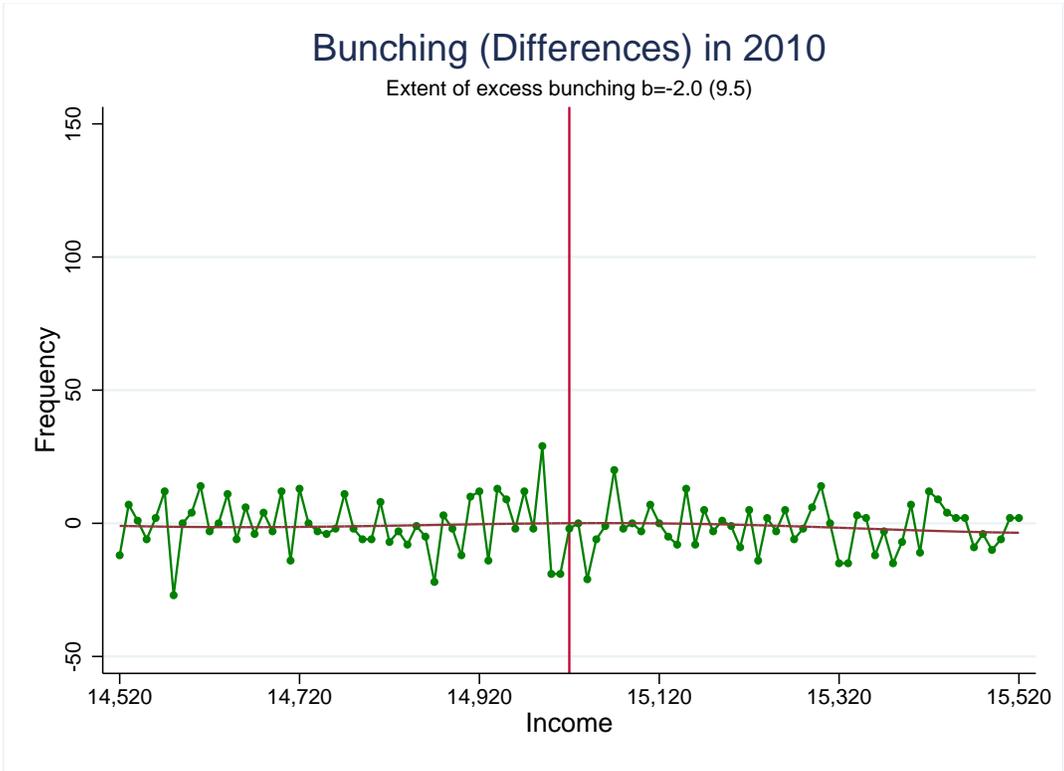


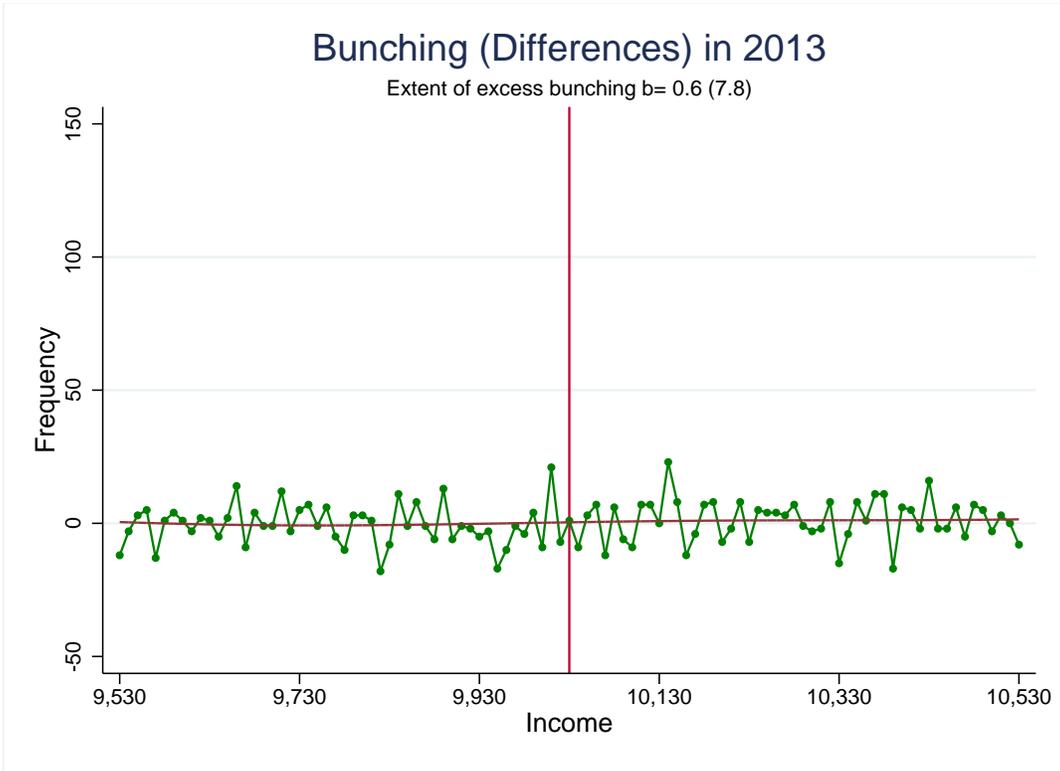
C Bunching (Differences) Figures



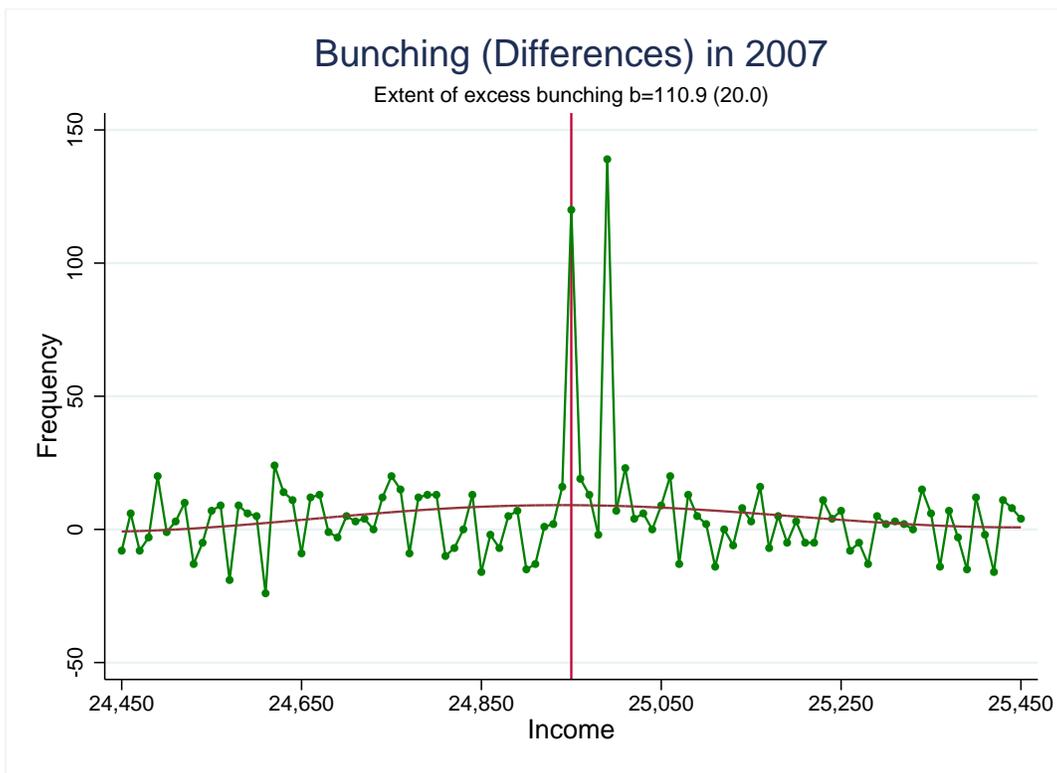
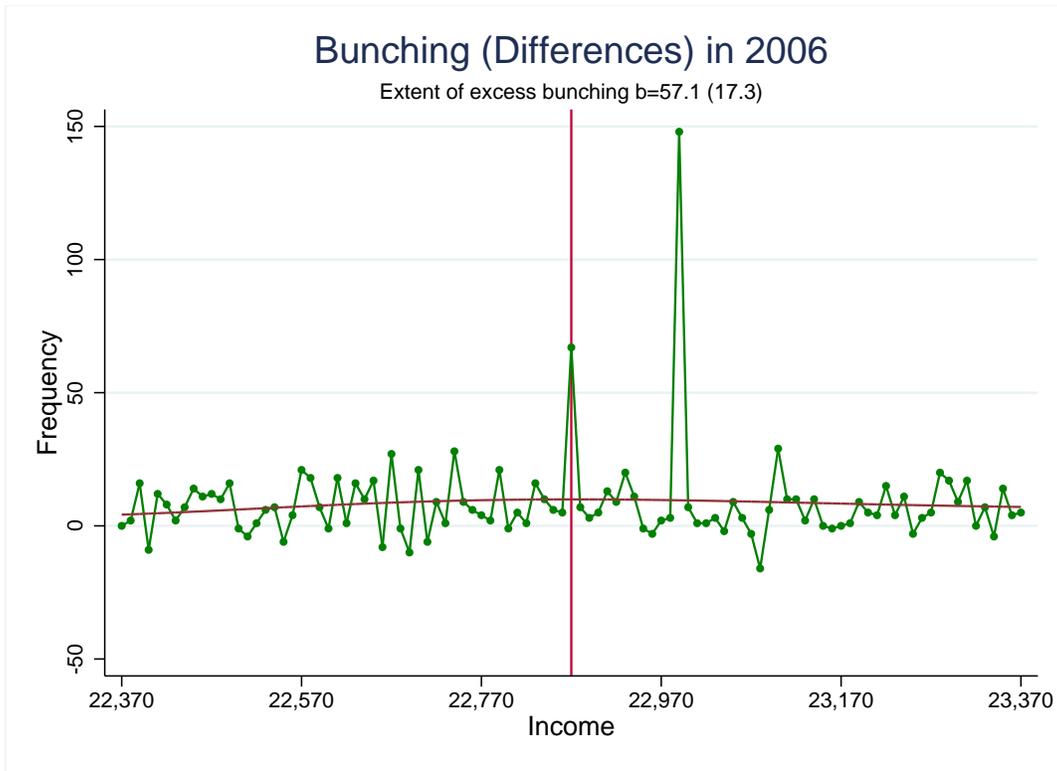


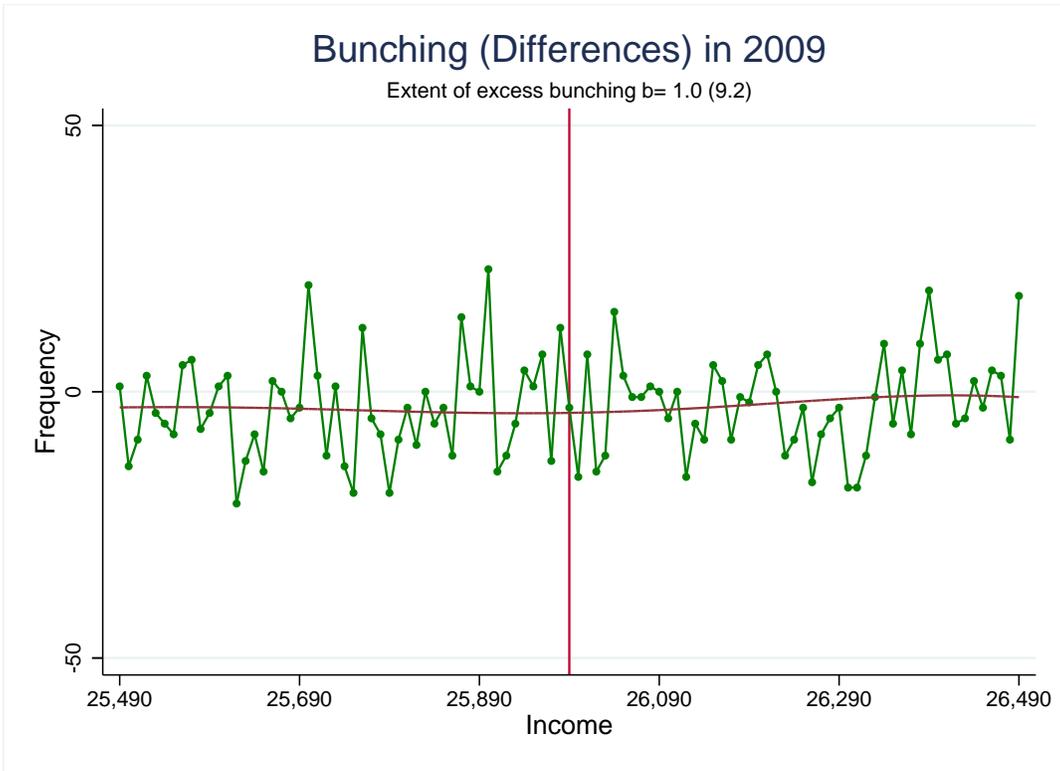
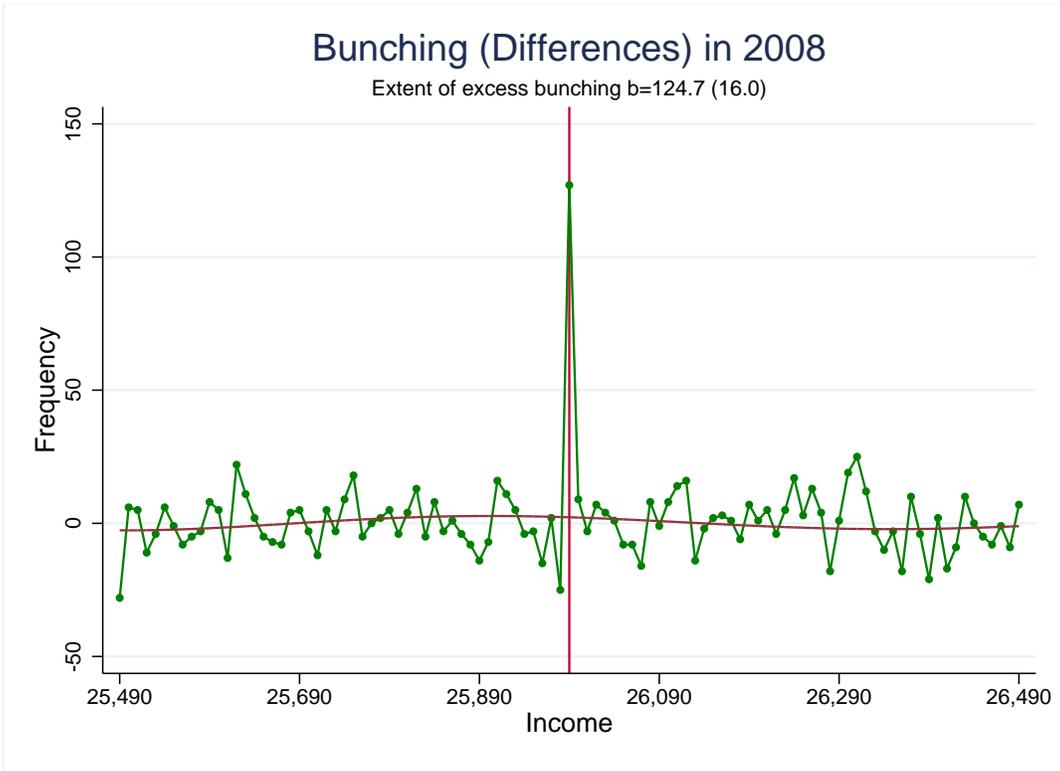


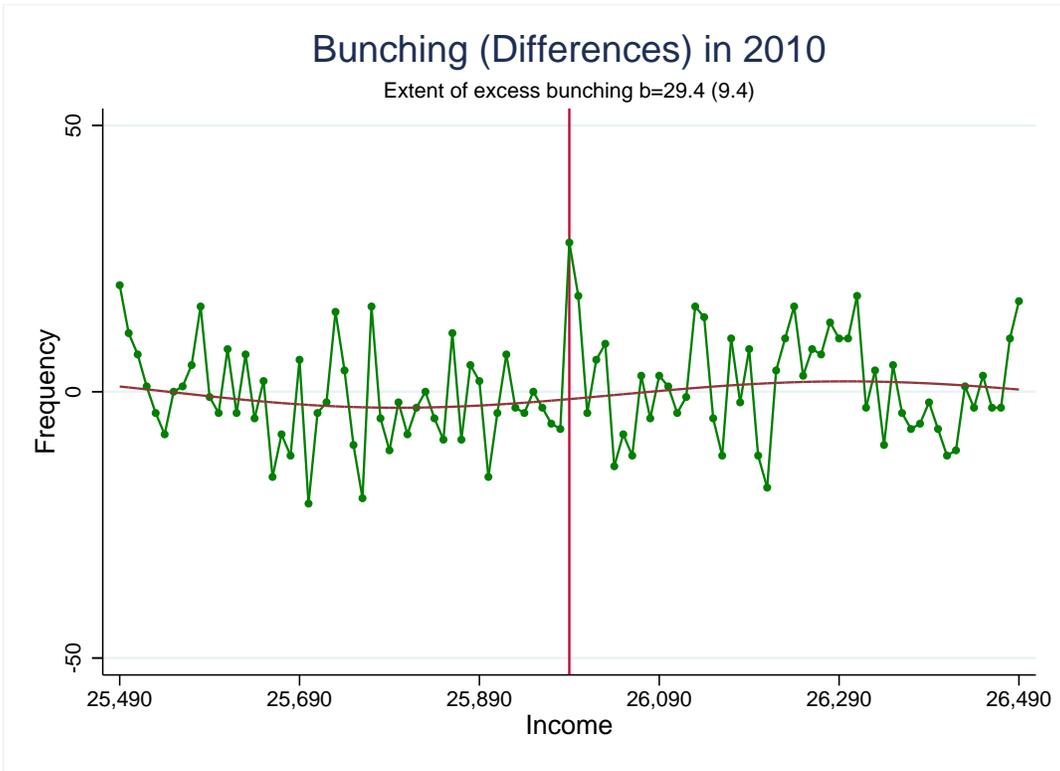
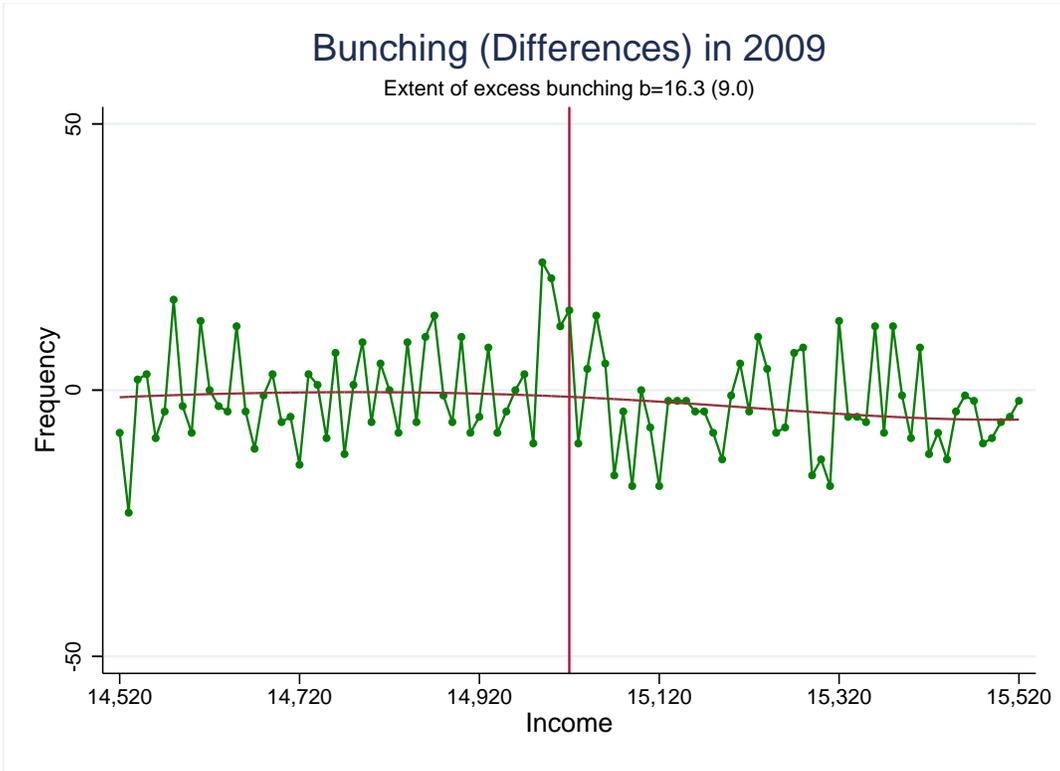


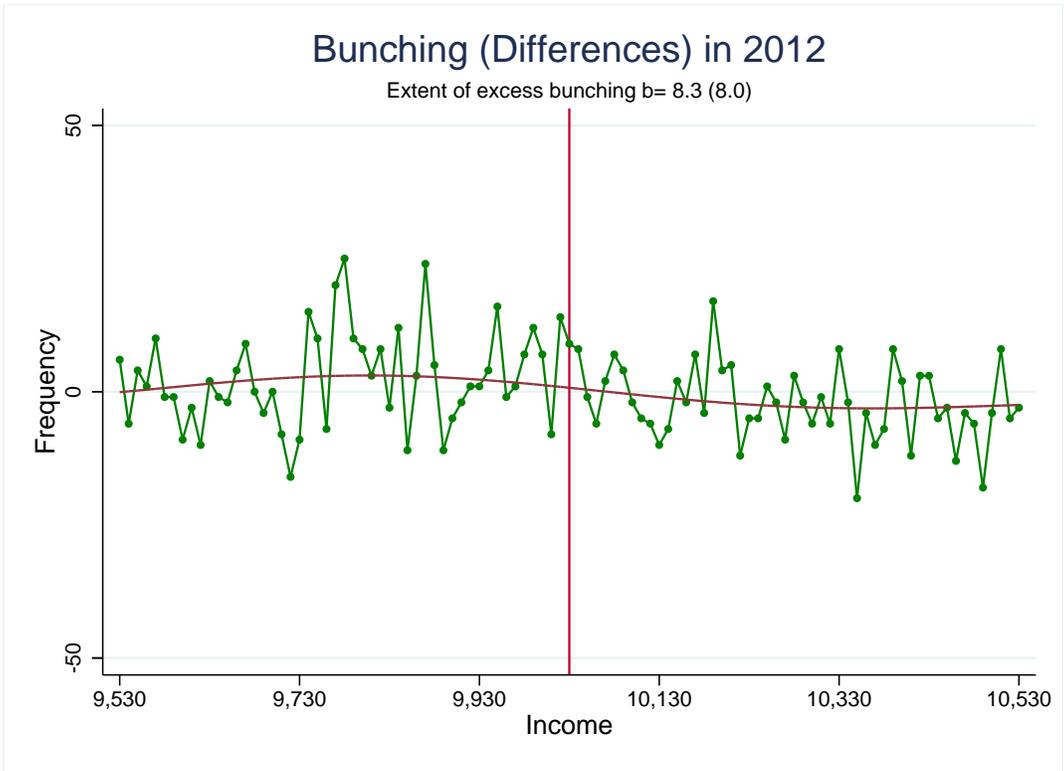
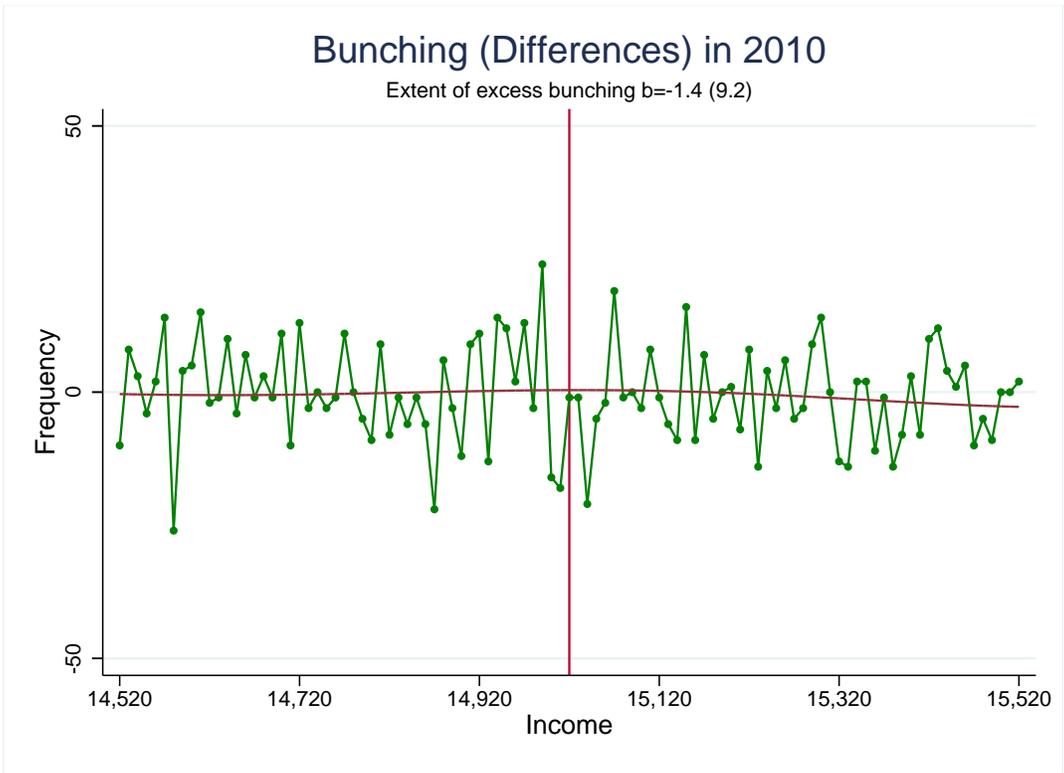


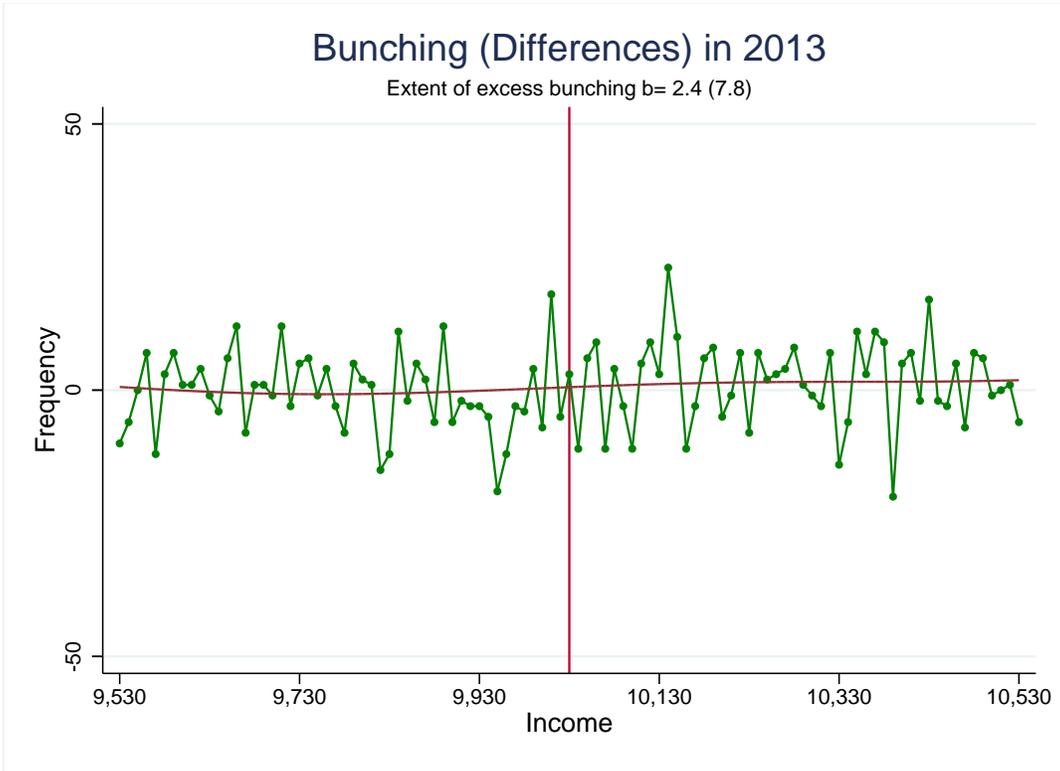
D Bunching (Differences) Figures — Excluding Construction



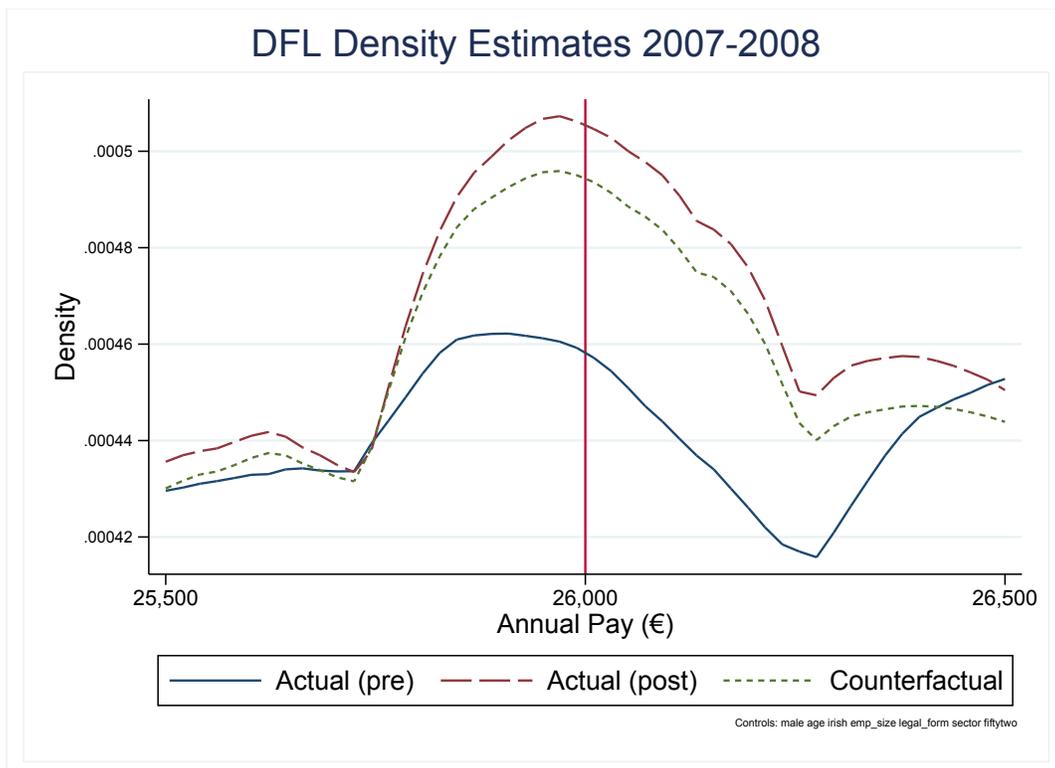
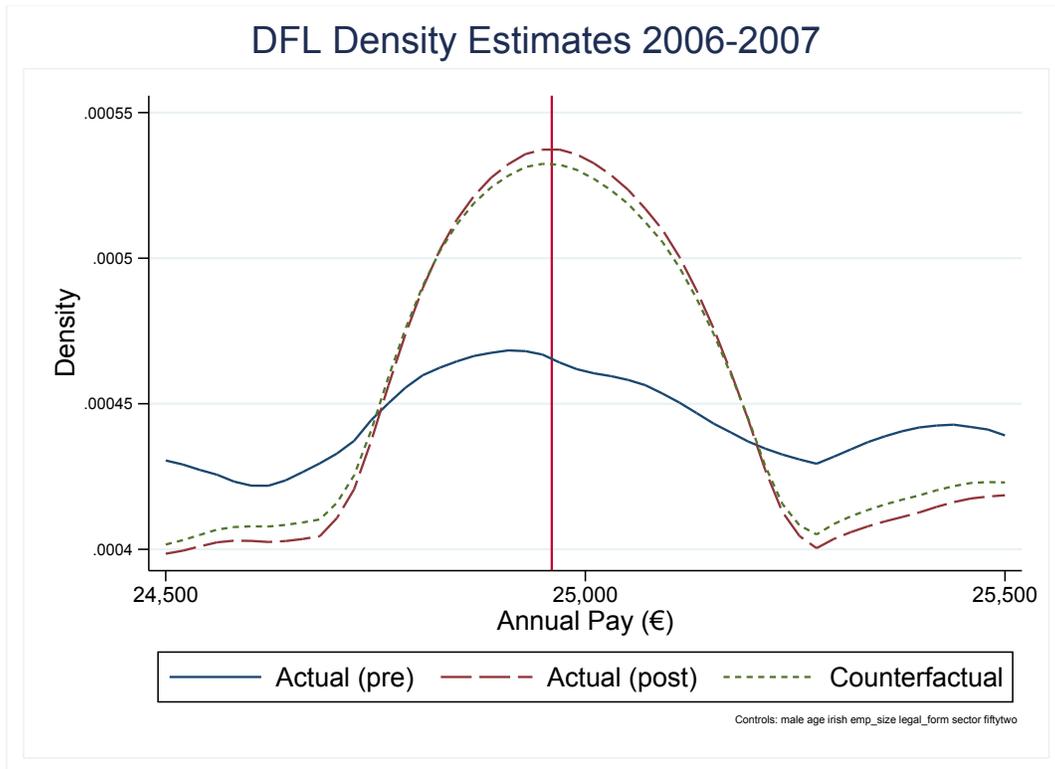




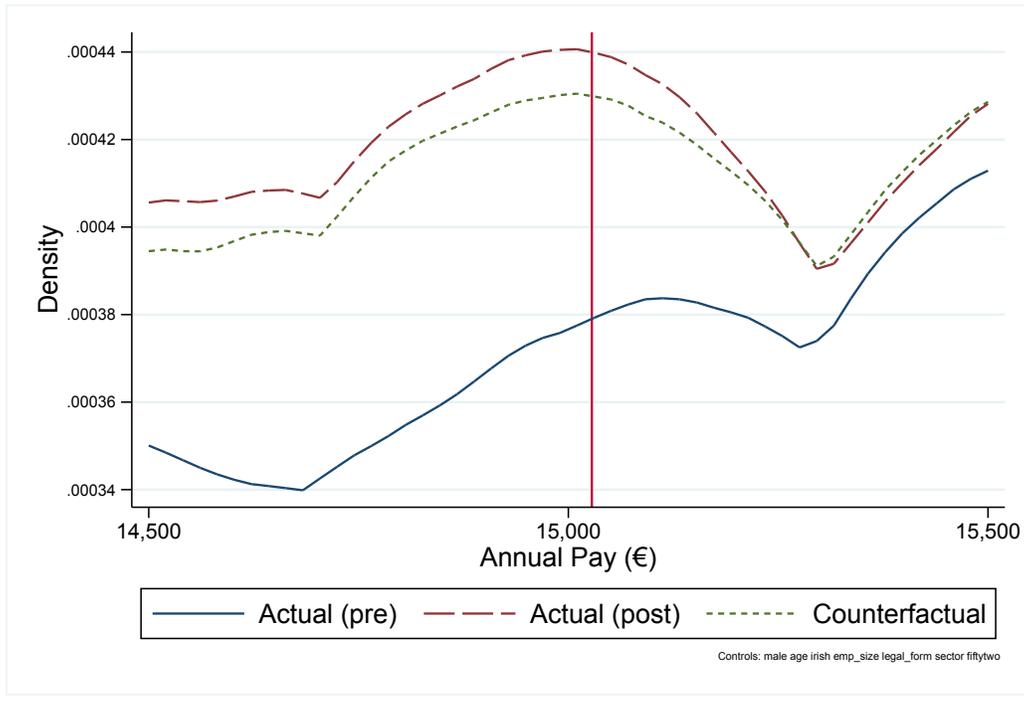




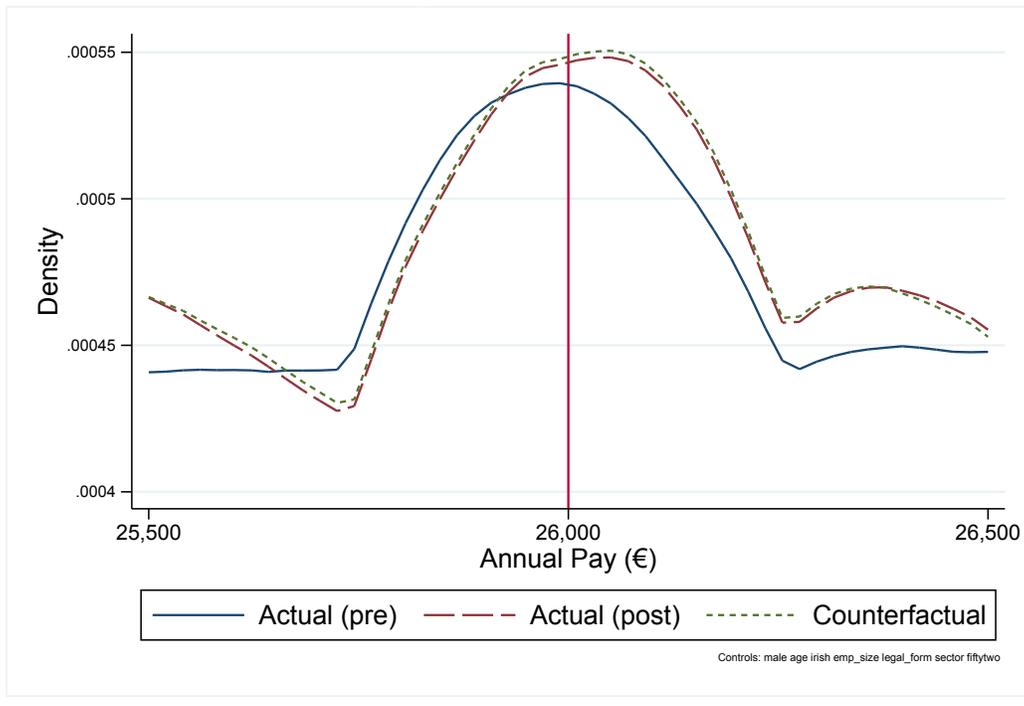
E DFL Counterfactual Figures



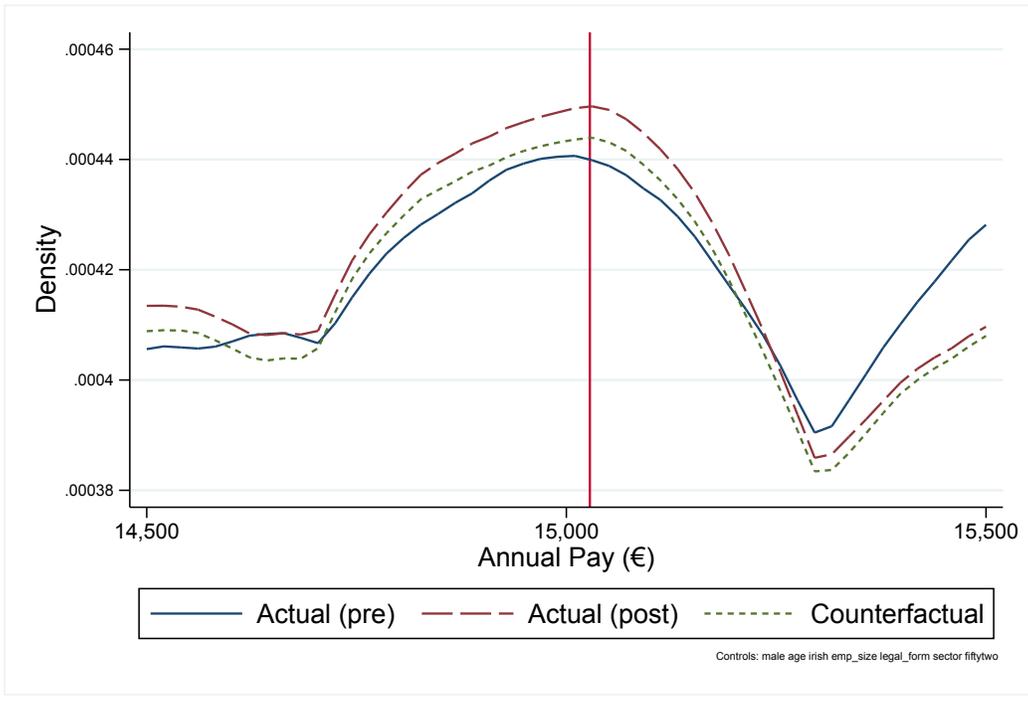
DFL Density Estimates 2008-2009



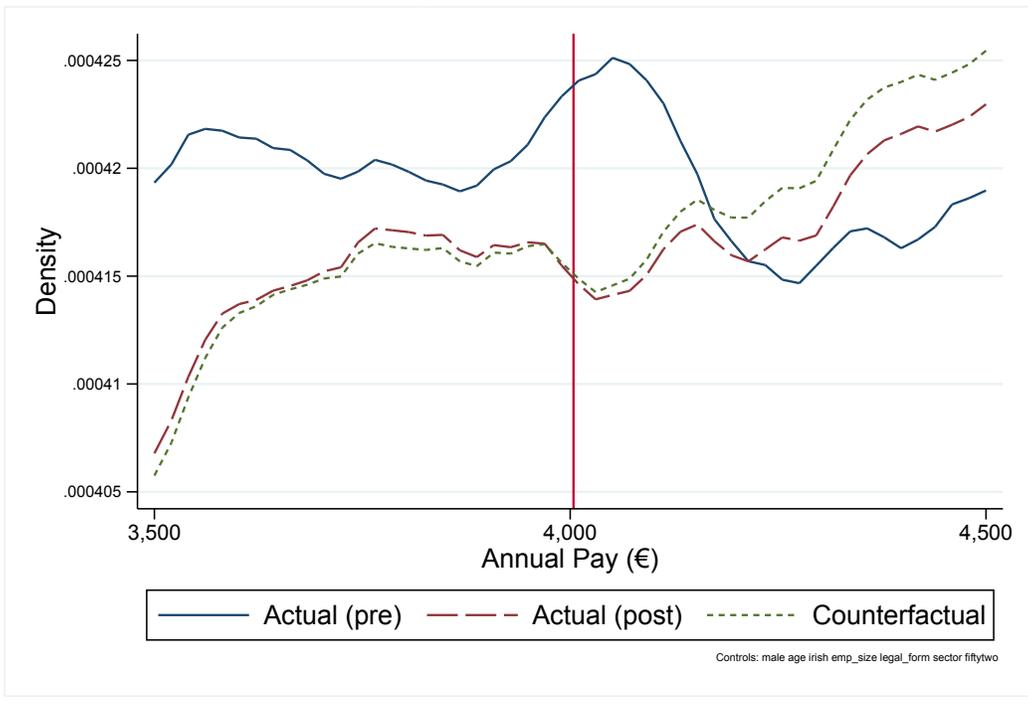
DFL Density Estimates 2009-2010



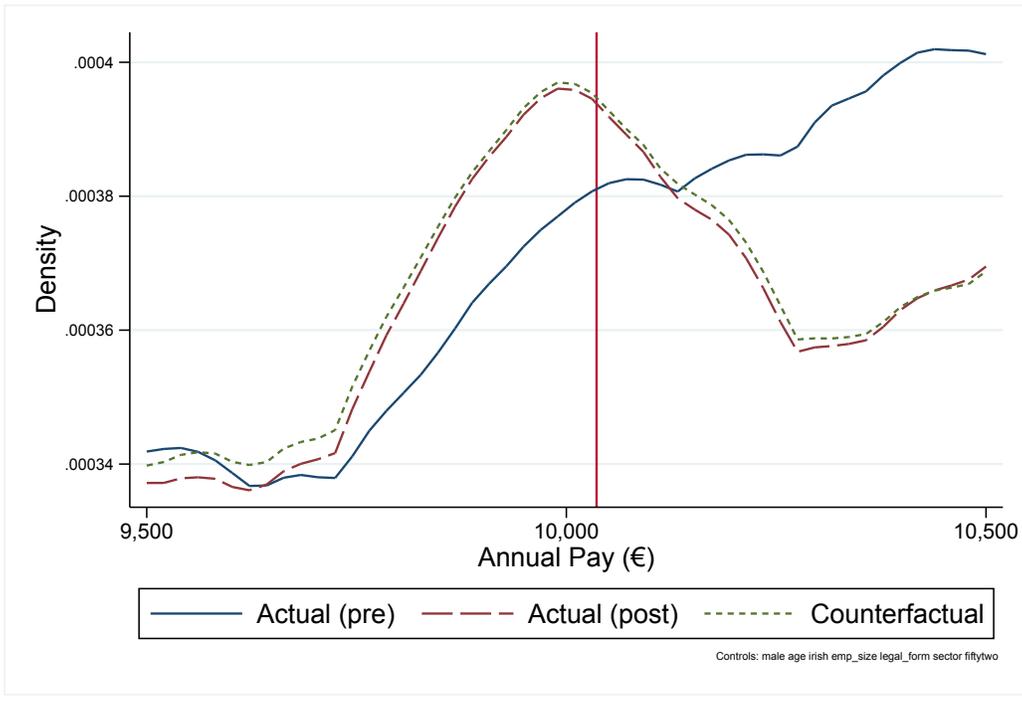
DFL Density Estimates 2009-2010



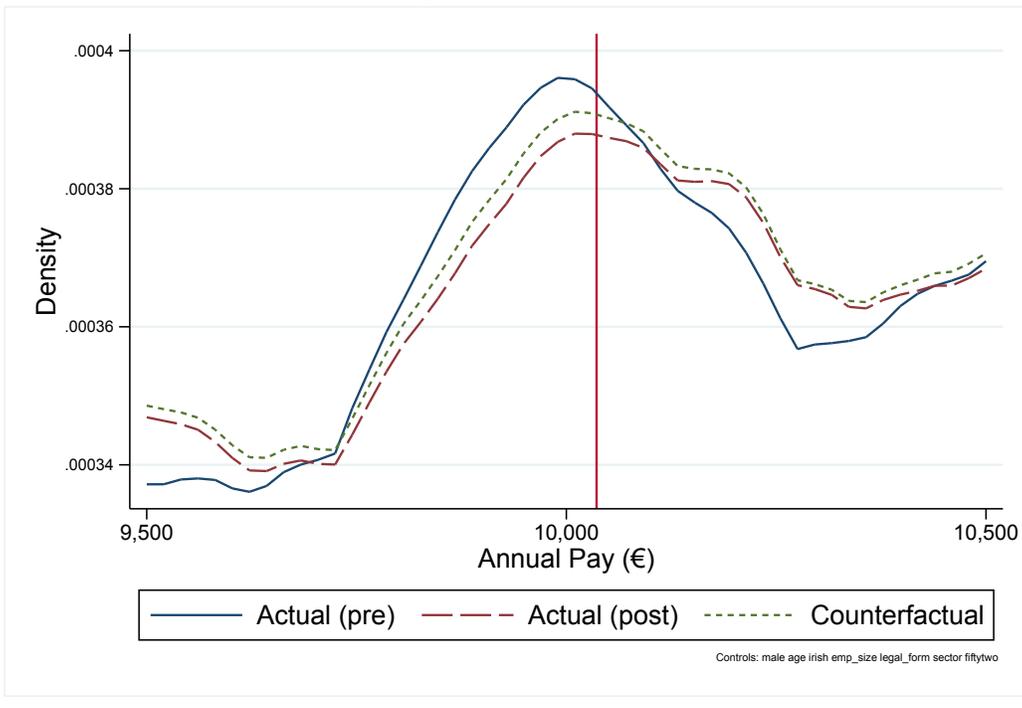
DFL Density Estimates 2010-2011



DFL Density Estimates 2011-2012



DFL Density Estimates 2012-2013



F Round-Number Bunching

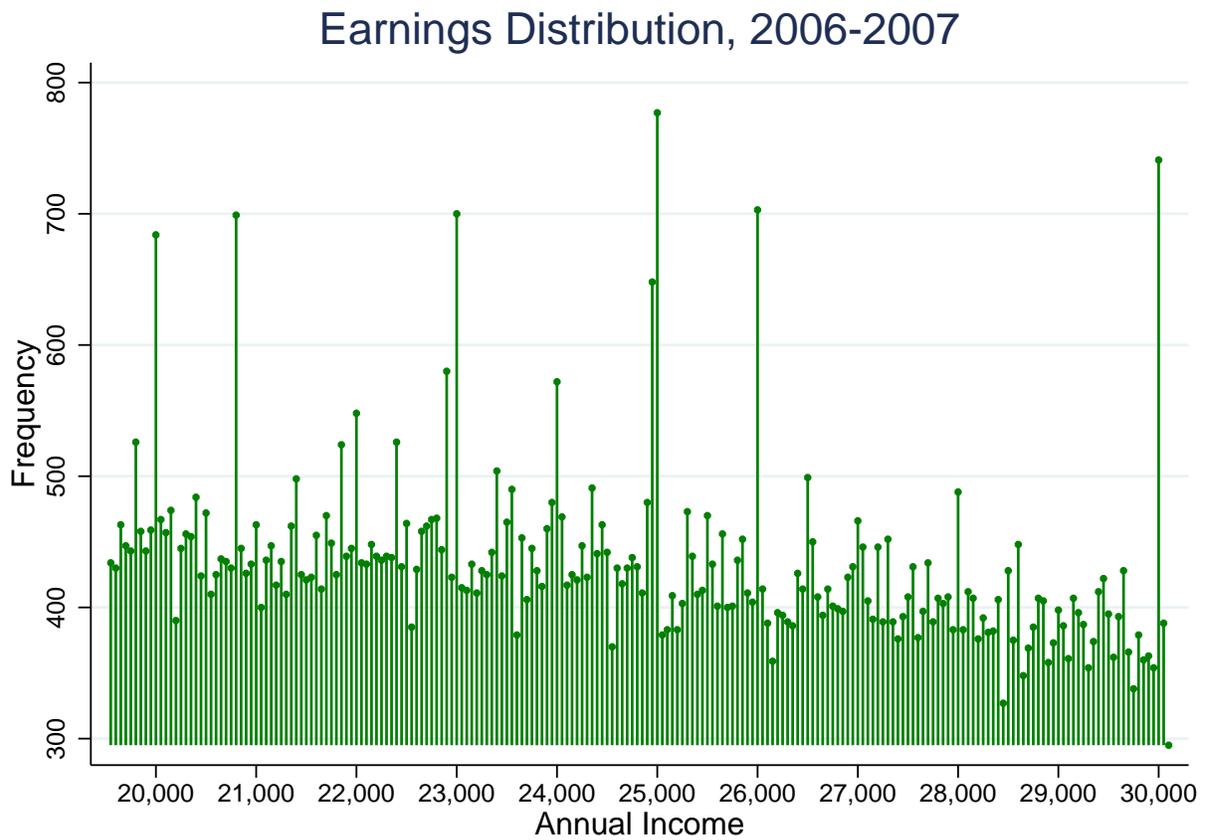


Figure 8: Histogram of earnings in €50 bins, pooled data from 2006 and 2007. There is clear round-number bunching for multiples of €1,000.

G Miscellaneous

Figure 9: Spatial Distribution of Nonfilers by County: 2007

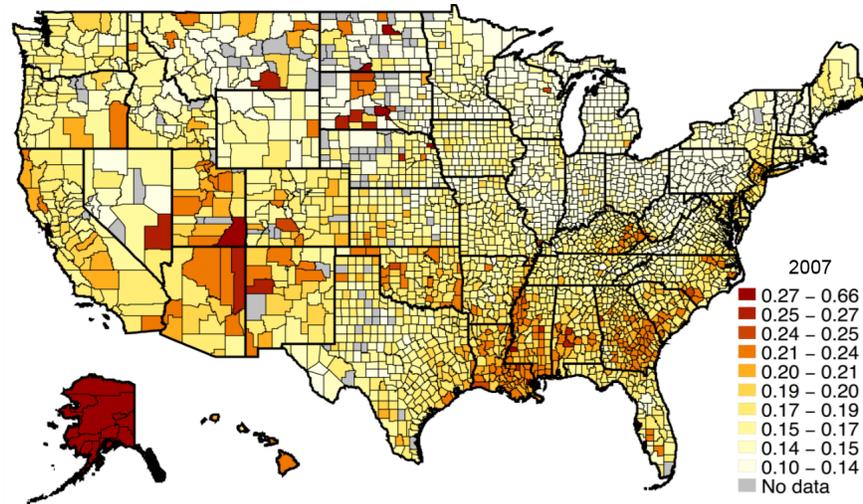
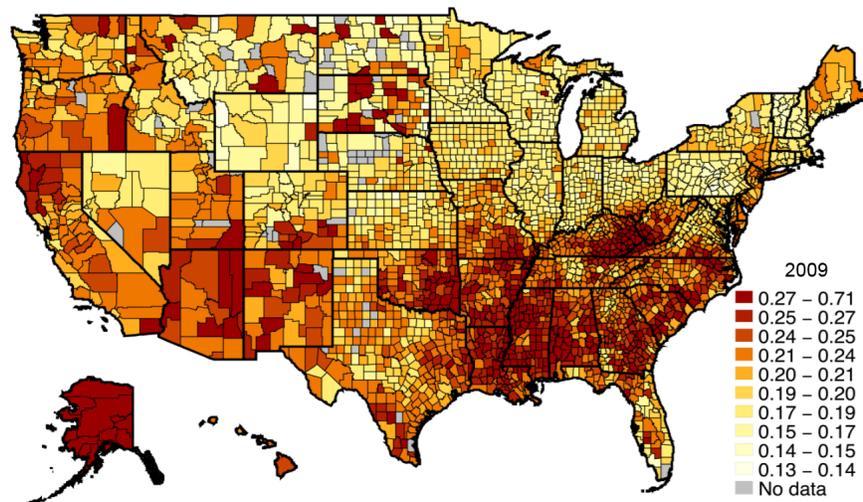


Figure 10: Spatial Distribution of Nonfilers by County: 2009



These figures, adapted from Guyton, Manoli, Schafer and Sebastiani (2016) show the fraction of income tax/EITC nonfilers by county pre- and post-2008. A substantial increase in nonfiling is apparent across the country, indicating the recession affected tax planning behaviour in the United States.

Figure 11: Frequency of Google searches for relevant taxes in Ireland

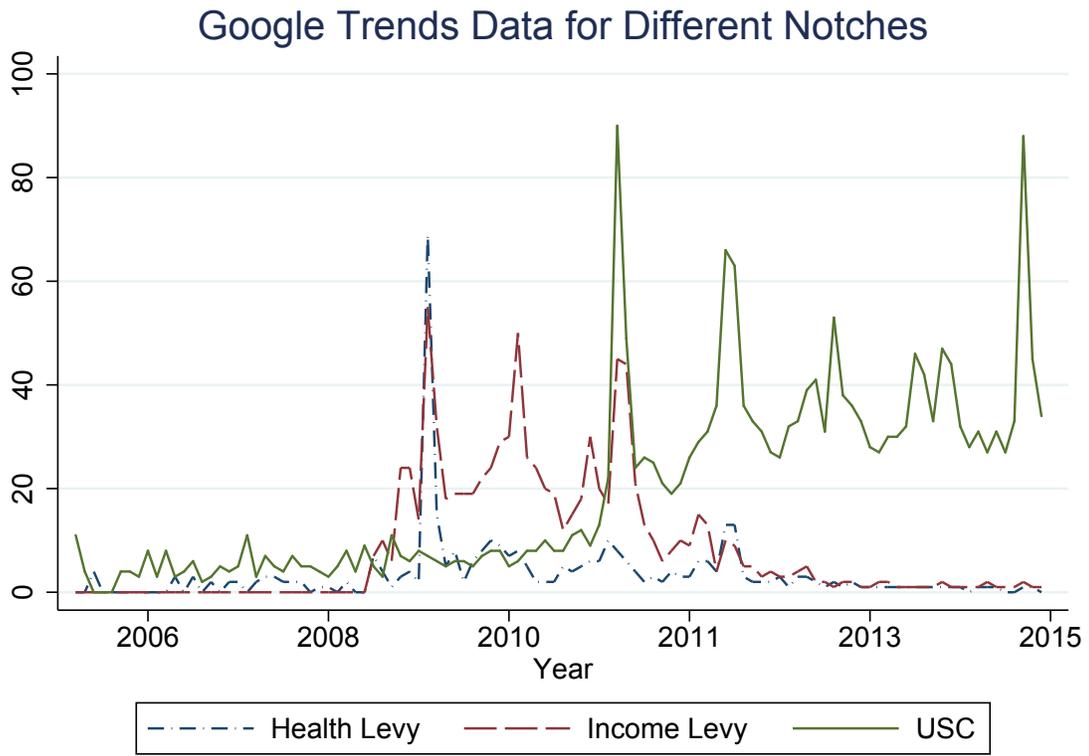


Figure 12: Declines in bunching subsequent to notch removal

Disappearance in Differences

