

The London School of Economics and Political Science

**Essays on the Economics of Gender  
Identity and Behavioural Responses  
to Tax Policy**

Panos Mavrokonstantis

A thesis submitted to the Department of Economics of  
the London School of Economics for the degree of  
Doctor of Philosophy

London, July 2018

*To my parents*

# Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent.

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 30,000 words.

## **Statement of Conjoint Work**

I confirm that Chapter 3 was jointly co-authored with Sarah Clifford and I contributed 50% of this work.

## **Disclaimers**

Chapter one of this thesis uses data provided by the UK Data Service. The views expressed in this chapter are solely those of the author and do not necessarily reflect the views of the original data creators, depositors or copyright holders, the funders of the data collections, the UK Data Archive, or the UK Data Service.

Chapters two and three use administrative data provided by the Tax Department of the Republic of Cyprus. The views expressed in these chapters are solely those of the authors and do not necessarily reflect the views of the data providers.

# Acknowledgements

I would like to express my gratitude to my supervisors Frank Cowell and Johannes Spinnewijn for all their guidance and support throughout these years. Thank you for always being willing to listen to my research ideas and provide me with the encouragement needed to complete this PhD. Special thanks also to Frank for always accepting to support and sign my numerous data access applications without hesitation.

I'm grateful to Nick, Joe and the rest of IT staff for their technical help and for never refusing to install whatever programme I would next request. I also thank Mark Wilbor for all his admin support and for providing me with the endless letters I would request to prove my student status throughout these years.

I am grateful to all my PhD peers, and especially Pedro Alves, Michel Azulai, Sebastian Camarero-Garcia, Sarah Clifford, Miguel Espinosa, Kilian Huber, Dana Kassem, Milad Khatib-Shahidi, Davide Porcellacchia, Arthur Seibold, Xuezhu Shi, Eddy Tam, Kostas Tokis and Maria Uribe for their friendship and support. Special thanks also to my high-school friends for always supporting my academic pursuits, and especially Marcos Charalambides for his invaluable computational assistance and Alexandra Patsalides for her proofreading efforts.

This thesis would not have been written without the contribution of various data providers. I want to thank the UK Data Service for granting me access to the confidential version of the Next Steps survey and to all the staff that handled my output requests. I am also grateful to the Tax Department of the Republic of Cyprus for providing access to their administrative datasets, and especially to Elina Symeonidou for preparing and explaining the data. I am also indebted to the Ministry of Labour, Welfare and Social Insurance for giving me access to their archives of trade union collective bargaining agreements, and



to numerous trade union representatives for insightful discussions on the Cypriot labour market.

My work has benefited from discussions with faculty in the LSE Public Economics Programme, as well as seminar and conference participants at the LSE, ECINEQ conference Luxembourg, EPRU Copenhagen, LAGV Marseille, EALE Ghent, ECINEQ Winter School Meeting Canazei, TADC London, HO PhD Symposium London, ZEW MaTaX Mannheim and NTA Philadelphia.

I also gratefully acknowledge financial assistance from the Onassis Foundation of which I am a proud scholar.

I am eternally grateful to my parents for their unconditional love and support throughout my life. Without their diachronic encouragement to always pursue my dreams, I would have never come this far. I also cannot thank them enough for acting as my research assistants in the last few years, always being willing to find whatever random information about the Cypriot tax system I would discover that I needed to know for my research. I also thank my sister for her endless care and support in my endeavours.

Finally, I want to thank my beloved wife Vicky for all her love and support, as well as the endless hours she spent discussing my random research ideas, reading my drafts, and sitting down for my practice presentations. Above all, I want to thank her for inspiring me, through her drive and ambition since our Warwick days, to achieve the best, and for giving me the extra push when times were tough. I would not have survived this without her.

# Abstract

This thesis presents three essays. The first examines the effect of having a breadwinner mother on children's gender norms in England. It shows that, while boys with breadwinning mothers are less likely to develop traditional views, girls are more likely to become traditional, in opposition to their family's norm but in line with society's. It then develops a model of gender identity, showing that the results can be explained by girls' weaker preference for conformity to their family. Finally, it employs conformity measures to show that this prediction holds empirically, and implements a regression discontinuity design to estimate the effect of living in such norm-minority families on conformity.

The second essay studies costs in adjusting taxable income, using tax reforms and data from the Republic of Cyprus. Combining reduced-form evidence with a structural model of labour supply, it estimates these costs at CYP 79 for salary earners, and CYP 5 for the self-employed (the equivalent of 135 and 8.5 Euros respectively). It then examines the mechanisms driving the observed responses, and proposes a method that estimates asymmetries in adjustment costs to infer the underlying source of frictions.

Using the same data, the third chapter examines the importance of tax policy design for tax enforcement and behavioural responses in charitable contributions. It shows three sets of results. First, exploiting salary-dependent thresholds for third-party reporting of charitable contributions, it estimates that reported donations increase by about 0.7 pounds when taxpayers can claim 1 pound more without providing documentation. Second, it finds that at least 40 percent of these observed responses are purely due to changes in reporting, using a reform that retroactively shifted the reporting threshold. Last, it estimates a tax price elasticity of reported charitable donations at -0.5, and shows that this policy parameter is highly sensitive to the reporting environment.

# Contents

<b>1</b>	<b>Breadwinning Mothers and Children's Gender Norms</b>	<b>16</b>
1.1	Introduction . . . . .	16
1.2	Literature Review . . . . .	19
1.3	The Theoretical Framework . . . . .	22
1.3.1	The Socialisation Process . . . . .	22
1.3.2	Formalising the Socialisation Process . . . . .	23
1.3.3	Comparative Statics . . . . .	24
1.4	The Empirical Framework . . . . .	25
1.5	Data . . . . .	27
1.5.1	The Social Norm . . . . .	27
1.5.2	The Next Steps Survey . . . . .	29
1.5.3	Choice of Variables . . . . .	30
1.6	Results . . . . .	31
1.6.1	Main Results . . . . .	31
1.6.2	Other Outcomes Related to Gender Norms . . . . .	33
1.6.3	Are Nontraditional Families Truly Transmitting Nontraditional Views?	34
1.6.4	How do These Results Square With the Literature on Maternal Labour Supply? . . . . .	36
1.7	Examining the Mechanism: A Behavioural Explanation . . . . .	37
1.7.1	Preference for Conformity to the Family's Norm . . . . .	38
1.7.2	Why Are Girls in Nontraditional Families Less Conformist? . . . .	38
1.7.3	Identifying the Mechanism Using a Regression Discontinuity Design	39

1.7.4	The RDD Identification Assumption . . . . .	40
1.7.5	Regression Discontinuity Estimates . . . . .	43
1.7.6	Placebo Tests . . . . .	45
1.8	Conclusion . . . . .	45
	References . . . . .	47
	Main Tables . . . . .	56
	Main Figures . . . . .	62
	Appendix 1A: Further Predictions of the Model . . . . .	69
	Appendix 1B: Discussion of Covariates . . . . .	71
	Appendix 1C: Extra Tables . . . . .	74
<b>2</b>	<b>Bunching and Adjustment Costs: Evidence from Cypriot Tax Reforms</b>	<b>93</b>
2.1	Introduction . . . . .	93
2.2	Institutional Context . . . . .	97
2.3	Data . . . . .	99
2.4	Conceptual Framework and Reduced-Form Evidence . . . . .	99
2.4.1	Introduction to the Bunching Framework . . . . .	100
2.4.2	Estimating the Bunching Mass . . . . .	101
2.4.3	Graphical Evidence . . . . .	104
2.4.4	Introducing Adjustment Costs . . . . .	106
2.5	Estimation of Structural Parameters . . . . .	107
2.5.1	Parametrising the Utility Function . . . . .	107
2.5.2	Results . . . . .	109
2.5.3	Heterogeneity Analysis . . . . .	110
2.6	Mechanisms . . . . .	112
2.7	Extensions . . . . .	115
2.7.1	Alternative Non-Parametric Structural Elasticity Bounds . . . . .	115
2.7.2	Asymmetric Adjustment Costs . . . . .	116
2.7.3	Implications for Policy Evaluation . . . . .	119
2.8	Conclusion . . . . .	122
	References . . . . .	124

Main Tables . . . . .	128
Main Figures . . . . .	133
Appendix 2A: Further Results . . . . .	142
Appendix 2B: A Potential Extension to the Model . . . . .	150
Appendix 2C: Collective Bargaining Agreements . . . . .	153

### **3 Enforcement and the Importance of Tax Policy Design for Elasticities:**

<b>Evidence from the Republic of Cyprus</b>	<b>155</b>
3.1 Introduction . . . . .	155
3.2 Institutional Context and Data . . . . .	159
3.2.1 Institutional Details . . . . .	159
3.2.2 Data . . . . .	160
3.3 Behavioural Responses to the Enforcement Environment . . . . .	162
3.3.1 Regression Discontinuity Design Estimates . . . . .	163
3.4 Separating Real and Reporting Responses . . . . .	167
3.4.1 Bunching Results . . . . .	168
3.4.2 Robustness Analysis . . . . .	170
3.4.3 Heterogeneity Analysis . . . . .	172
3.5 The Elasticity of Giving w.r.t. Price and Tax Policy Design . . . . .	172
3.5.1 Elasticity Results for the Whole Period 1999-2010 . . . . .	175
3.5.2 Elasticities and Reporting Thresholds . . . . .	177
3.5.3 Elasticities and Enforcement . . . . .	179
3.6 Conclusion . . . . .	180
References . . . . .	182
Appendix 3A: Extra Tables . . . . .	197
Appendix 3B: Extra Figures . . . . .	217

# List of Tables

1.1	Main Results. Dependent Variable: Traditional Norm . . . . .	56
1.2	Other Outcomes . . . . .	57
1.3	Robustness. Dependent Variable: Traditional Norm . . . . .	58
1.4	Second Shift. Dependent Variable: Mother Does Second Shift . . . . .	59
1.5	Incorporating Hours Worked. Dependent variable: Traditional Norm . .	60
1.6	Conformity to Family . . . . .	61
1.7	RD Results: Nontraditional Family Treatment Effect . . . . .	62
1C.1	Summary Statistics: Original Vs Final Sample ( <i>Next Steps</i> ) . . . . .	74
1C.2	Summary Statistics of FRS and SPI data . . . . .	75
1C.3	Summary Statistics for Next Steps Data . . . . .	76
1C.4	Full Estimates of Main Results. Dependent variable: Traditional Norm .	79
1C.5	Covariate Balance Tests . . . . .	84
1C.6	Main Results, Probit and Logit Specifications. Dependent Variable: Traditional Norm . . . . .	87
1C.7	Other Outcomes, Probit and Logit Specifications. Dependent Variable: Traditional Norm . . . . .	88
1C.8	Robustness Checks, Probit and Logit Specifications. Dependent Variable: Traditional Norm . . . . .	89
1C.9	Probit and Logit Specifications. Dependent Variable: Mother Does Second Shift . . . . .	90
1C.10	Incorporating Mother Hours Worked, Probit and Logit Specifications. Dependent Variable: Traditional Norm . . . . .	91
1C.11	Conformity to Family, Probit and Logit Specifications. . . . .	92

2.1	Summary Statistics . . . . .	128
2.2	Results for Structural Parameters . . . . .	129
2.3	Heterogeneity Analysis of Structural Parameters for Salary Earners . . . .	130
2.4	Bounds on Structural Elasticity Using a Non-Parametric Approach . . . .	131
2.5	Policy Simulations . . . . .	132
2A.1	Robustness Checks for Bunching Estimates . . . . .	143
2C.1	Collective Bargaining Agreement, Banking Sector . . . . .	153
2C.2	Collective Bargaining Agreement, Construction Sector . . . . .	154
3.1	Summary Statistics . . . . .	185
3.2	RD Estimates at Notches in Donation Schedule . . . . .	186
3.3	Elasticity of Donations w.r.t. Price (1999-2010) . . . . .	187
3.4	Elasticity of Donations w.r.t. Price - Bunchers vs. Non-bunchers (2003-2007)	188
3.5	Buncher Stickiness - Linear Probability Model . . . . .	189
3.6	Elasticity of Donations w.r.t. Price by Period . . . . .	190
3A.1	RD Estimates - Including Rounders . . . . .	197
3A.2	RD Estimates - Years 1999-2002 . . . . .	198
3A.3	RD Estimates - Including Individuals with Some Non-salary Income . . . .	199
3A.4	RD Estimates - Professional Taxes and Union Fees Residualised . . . . .	200
3A.5	RD Estimates - Only Highly Unionised Sectors Using Residualised Out- come Variable . . . . .	201
3A.6	RD Estimates - Only Highly Unionised Sectors . . . . .	202
3A.7	RD Estimates - Only Sectors with Low Unionisation . . . . .	203
3A.8	RD Estimates - Only Males . . . . .	204
3A.9	RD Estimates - Only Females . . . . .	205
3A.10	RD Estimates - Ages 25-39 . . . . .	206
3A.11	RD Estimates - Ages 40-54 . . . . .	207
3A.12	Elasticity of Donations w.r.t. Price (Sample: All Highly Unionised Sectors, 1999-2010) . . . . .	208
3A.13	Elasticity of Donations w.r.t. Price (Sample: All, 1999-2010) . . . . .	209
3A.14	Poisson Regression (1999-2010) . . . . .	210

3A.15 Elasticity of Donations w.r.t. Price - Bunchers Vs. Non-bunchers (Sample: All Highly Unionised Sectors, 2003-2007) . . . . .	211
3A.16 Elasticity of Donations w.r.t. Price - Bunchers vs. Non-bunchers (Sample: All, 2003-2007) . . . . .	212
3A.17 Buncher Stickiness - Linear Probability Model (Sample: All Highly Unionised Sectors) . . . . .	213
3A.18 Buncher Stickiness - Linear Probability Model (Sample: All) . . . . .	214
3A.19 Elasticity of Donations w.r.t. Price by Period (Sample: All Highly Unionised Sectors) . . . . .	215
3A.20 Elasticity of Donations w.r.t. Price by Period (Sample: All) . . . . .	216



# List of Figures

1.1	Social Norms and Gender Inequality . . . . .	63
1.2	The Social Norm . . . . .	64
1.3	The Distribution of the Mother's Income Share: Next Steps . . . . .	65
1.4	The Distribution of the Wife's Income Share: Further Datasets . . . . .	66
1.5	Regression Discontinuity Graphs: Conformity to Family . . . . .	67
1.6	Placebo Tests . . . . .	68
2.1	Marginal Tax Rates, 2003-2007 . . . . .	133
2.2	Optimisation and Earnings Distributions Under Linear and Nonlinear Tax Schedules . . . . .	134
2.3	How Frictions Attenuate the Bunching Region . . . . .	135
2.4	Bunching of Salary Earners at the CYP 10,000 Kink . . . . .	136
2.5	Bunching of Salary Earners at the CYP 20,000 Kink . . . . .	137
2.6	Bunching of Self-employed at the CYP 10,000 Kink . . . . .	138
2.7	Bunching of Self-employed at the CYP 20,000 Kink . . . . .	139
2.8	Bunching Patterns Among Salary Earning Potential Bunchers at CYP 10,000 Pooled Over 2004-06 . . . . .	140
2.9	De-Bunching Patterns in 2007 Among Salary Earners Bunching at CYP 10,000 by 2006 . . . . .	141
2.10	Bunching Patterns Among Self-Employed Potential Bunchers at CYP 20,000 Pooled Over 2004-06 . . . . .	142
2A.1	Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Ages 25-39 . . . . .	144

2A.2	Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Ages 40-54 . . . . .	145
2A.3	Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Males . . . . .	146
2A.4	Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Females . . . . .	147
2A.5	Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Workers Not in Highly Unionised Sectors . . . . .	148
2A.6	Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Workers in Highly Unionised Sectors . . . . .	149
3.1	Reporting Thresholds Over the Sample Period . . . . .	191
3.2	Schedule of Marginal Tax Rates (Years 1999-2010) . . . . .	192
3.3	Yearly Average Donations Among Donors . . . . .	193
3.4	Average Donations by Income 1999-2001 . . . . .	194
3.5	Bunching Around Reporting Thresholds . . . . .	195
3.6	Bunching Estimates over Time . . . . .	196
3B.1	Information on Tax Returns Regarding Thresholds . . . . .	217
3B.2	Schedule of Professional Taxes (Before 2003) . . . . .	218
3B.3	Density of Salary Income Between 1999-2001 . . . . .	219
3B.4	Robustness Check: Smoothness of Covariates . . . . .	220
3B.5	Bunching at CYP 150 with Estimated Counterfactual, Main Sample . . .	221
3B.6	Bunching at CYP 175 with Estimated Counterfactual, Main Sample . . .	222
3B.7	Donations Between 2003-2010 of Those Bunching at CYP 150 in 2007 . .	223
3B.8	Donations Between 2003-2010 of Those Bunching at CYP 175 in 2008 . .	224
3B.9	Fraction Reporting Specific Amounts of Donations Over Time . . . . .	225
3B.10	When Did People File Their Taxes for 2008? . . . . .	226
3B.11	Salary Growth Rates of 2007 Bunchers . . . . .	227
3B.12	Bunching Around Reporting Thresholds, Excluding Highly Unionised Sectors	228
3B.13	Bunching at CYP 150 with Counterfactual, Excluding Highly Unionised Sectors . . . . .	229

3B.14 Bunching at CYP 175 with Counterfactual, Excluding Highly Unionised Sectors . . . . .	230
3B.15 Bunching Estimates over Time, Excluding Highly Unionised Sectors . . .	231
3B.16 Bunching of Donations, Removing Union Fees . . . . .	232
3B.17 Bunching at CYP 150 with Counterfactual, Removing Union Fees . . . .	233
3B.18 Bunching at CYP 175 with Counterfactual, Removing Union Fees . . . .	234
3B.19 Bunching Estimates over Time, Removing Union Fees . . . . .	235
3B.20 Bunching Estimates over Time, Main Sample, Robustness Checks . . . .	236
3B.21 Bunching Estimates over Time, Heterogeneity Analysis by Sex . . . . .	237
3B.22 Bunching Estimates over Time, Heterogeneity Analysis by Age . . . . .	238

# Chapter 1

## Breadwinning Mothers and Children's Gender Norms

### 1.1 Introduction

Significant attention has been given to the recent rise of female breadwinners in the UK (IPPR 2015). The proportion of mothers earning more than 50% of the family's income has doubled between 1997-2014, increasing from 10% to 20%.<sup>1</sup> The increase in this nontraditional family structure is challenging the age-old traditional model of the male breadwinner and female caregiver. An unexplored question is whether this recent phenomenon, currently driven by a minority of the population, can change the norms about the appropriate roles of men and women in society.

In this paper, I examine the extent to which gender norms are passed on from one generation to the next when a family's norms oppose those of the society in which the family is embedded. Do children growing up in nontraditional families (i.e. where the mother is the breadwinner) adopt their family's values or those of society? Using data from England, I find unusual results for girls. While boys raised in nontraditional families are less likely to develop traditional norms, girls raised in nontraditional families are actually more likely to do so, in opposition to their family's norm but in line with society's. Examining further outcomes associated with gender norms, I also find that girls raised in nontraditional

---

<sup>1</sup>Author's calculations based on the Family Resources Survey.

families are less likely to state that being able to earn high wages is important for them and are less likely to want to study science at university. Employing a model of gender identity to examine the gender-socialisation process, I argue that my results can be explained by a weaker preference among girls for conforming to the family. I then use empirical conformity measures to show that the theoretical predictions do indeed hold in the data. Specifically, I find that girls in nontraditional families are less likely to have chosen their studies based on what their parents wished, and are more likely to argue with them. Finally, I examine the mechanism driving this heterogeneity in conformity. Using quasi-experimental variation in having a mother breadwinner, I employ a regression discontinuity design to show that this weaker preference is caused by the treatment of living in a nontraditional family, making girls react to their family's norm-minority status and adopt more traditional views.

These findings have strong policy implications given that the individual gender attitudes developed by the end of the teenage years are significant determinants of labour supply in adulthood (Johnston et al. 2014). Moreover, the prevailing gender norm of society is a strong predictor of gender inequality in the labour market. Consider, for instance, Figure 1.1a, which shows the relation between the gender pay gap and the prevalence of traditional gender norms across OECD countries. There is a strong positive relation between the proportion of a country's population that believes women should not work full-time, and the percentage difference in the median wage between men and women. Similar conclusions can be drawn from Figures 1.1b and 1.1c. The more traditional a country is, the lower female labour supply is, both along the extensive (Figure 1.1b) and intensive margins (Figure 1.1c). These stylised facts elucidate how important gender norms are to the analysis of gender inequality. Understanding their intergenerational transmission process can help identify policies that promote more egalitarian social norms, the benefits of which go beyond reducing labour market inequalities. More equal gender norms can aid in removing barriers to human development (UNDP 2014) and mitigate the adverse consequences of gender inequality traps (in particular gender gaps in education) on economic development (Dollar and Gatti 1999; Klasen 1999; Knowles et al. 2002; World Bank 2006). The modernisation of gender norms and the associated increase in female labour-market participation may also act as a key mechanism to fight poverty and promote rural development,

measures that the Food and Agriculture Organization (2009) has strongly advocated. Indeed, recent work by Kleven and Landaïs (2017) shows that egalitarian gender norms are strong predictors of economic development.

This study makes several contributions. First, it adds to a growing body of work exploring psychological attributes and preferences as drivers of gender inequality. These include risk preferences, attitudes towards competition or negotiation, and altruism.<sup>2</sup> This paper adds to these by exploring a new mechanism: gender socialisation and preferences for conformity to existing gender norms.

Second and most important, this study shows that explicitly defining parents' gender values, by distinguishing between nontraditional and traditional families, is crucial for assessing whether intergenerational transmission of gender norms is successful. The main approach in the literature is to examine the relation between mothers' and daughters' labour supply (Del Boca et al. 2000; Fernandez et al. 2004; Morrill and Morrill 2013; Olivetti et al. 2016; etc.). The finding of a positive correlation has been interpreted as evidence that nontraditional gender attitudes are successfully transmitted. I argue that this interpretation is problematic. Daughters of mothers who worked may well have a higher likelihood to also work, compared to daughters of mothers who did not work. However, the daughters of working mothers may still be working and earning *less than their husbands*. Hence, what the literature describes as evidence of the transmission of nontraditional values may still be capturing the propagation of traditional, male-breadwinner family norms. I find evidence consistent with this reasoning. I replicate the finding that girls are *less* likely to develop traditional views if their mother works. However, I also find that they are *more* likely to become traditional if their mother works or earns *more than their father*. To my knowledge, this is the first study to highlight this dichotomy by examining the effect of living in a nontraditional versus a traditional family on gender norms. The only related study is by Bertrand et al. (2015), which does not have an intergenerational approach but examines between-spouse outcomes. My study tries to fill this gap.

Lastly, in testing the assumptions of my identification strategy I uncover a caveat related to the findings of the Bertrand et al. (2015) influential study on the distribution of wives'

---

<sup>2</sup>For a review, see Bertrand (2010).

income shares. In their study on the US, they find a sharp drop in the density of wives' income shares at the 0.5 threshold, which is interpreted as an aversion to the wife earning more than the husband. I show that at least in the UK, this phenomenon is purely driven by families with earnings from self-employment, who exhibit very significant bunching at the threshold beyond which the wife starts to earn more. In families with only wage income however, no such sorting is found. This non-finding is consistent with significant earnings adjustment frictions highlighted in recent public finance literature (Chetty et al. 2011; Kleven and Waseem 2013; Gelber et al. 2016).

The paper proceeds as follows. Section 1.2 reviews the literature. Section 1.3 describes the theoretical framework using a model of gender identity, and Section 1.4 discusses the first empirical approach. Section 2.3 describes the data and presents evidence on the social norm regarding gender roles in England. Section 1.6 discusses the results, and Section 1.7 employs a regression discontinuity design to examine the mechanism driving the main findings. Section 1.8 concludes.

## 1.2 Literature Review

A vast literature in sociology and social psychology exists on theories of socialisation and social identity (Tajfel 1978; Lytton and Romney 1991; Lorber 1994; Epstein and Ward 2011). The aim of this literature is to explain how individuals develop their understanding about what behaviours and opinions are considered appropriate by society. The main principle is that socialisation, i.e. the process of learning about such values through human interaction, is the main way cultural norms are developed, adopted and transmitted.

Although studies on the role of socialisation in the intergenerational transmission of cultural values exist in the literature of other social sciences, economists have only recently begun to explore the topic. The seminal work bridging this gap between economics and other disciplines is that of Akerlof and Kranton (2000; 2002; 2010), who translated theories of social identity into an economics framework, giving birth to what is now known as Identity Economics. Various approaches to modelling identity have ensued.<sup>3</sup> Among others,

---

<sup>3</sup>For a comprehensive survey of identity models, see Costa-Font and Cowell (2015).

Benabou and Tirole (2007) propose a model where individuals hold a range of individual beliefs that they both value and can invest in. Klor and Shayo (2010) model identity as status, while Bisin et al. (2011) introduce the concept of oppositional identities. Although the models differ, the common element in this literature is the introduction of identity considerations into a neoclassic framework in which a person's self-image is valued and becomes a crucial element of her utility function.

A burgeoning empirical literature has developed that examines the effect of culture and its transmission. However, little work exists on the intergenerational transmission of explicit gender norms. Most studies that do consider gender norms focus rather on the effect of norms on some other outcomes (predominantly labour supply), without examining how gender norms are formed in the first place. For instance, Bertrand et al. (2015) examine the effect of traditional gender norms (defined as aversion to the wife earning more than the husband) on marriage and labour-market outcomes. While they argue that this aversion is induced by gender-identity norms, they do not assess how these norms are formed or transmitted intergenerationally, but simply take them as given. They show that women are less likely to work if their potential income exceeds their husbands'; if they do work, women are more likely to earn less than their potential income. Further, such female breadwinner families face a higher likelihood of divorce and lower marriage satisfaction.

Using a different approach, Alesina and Giuliano (2010) assess family culture by studying the importance of family ties for economic behaviour. They find that stronger family ties are associated with lower female labour force participation. Their explanation is that strong family ties require an adult family member to stay home and 'manage' the family institution, and this burden falls on women, who are consequently excluded from the formal labour market. In a related paper on culture, Alesina et al. (2013) look at the historical origins of gender norms. They provide evidence for intergenerational cultural persistence, showing that attitudes towards women are more traditional among descendants of societies that practised plough agriculture. Plough agriculture, in contrast to shifting cultivation, was more capital intensive and therefore required brawn-intensive labour, favouring men and leading to gender-based division of labour. As a result, the authors find historical plough use to be negatively related to modern-day attitudes towards gender inequality,



female labour force participation, and female participation in politics and firm ownership. Fernandez and Fogli (2009) examine how ancestral culture is related to fertility and labour market outcomes of second-generation American women. They find that women work more (have more children) in cases in which their country of ancestry had higher historic female labour force participation (total fertility rate). In a different setting, Kleven et al. (2018) focus on the child penalty in Denmark and find strong intergenerational transmission from parents to daughters, but not to sons.

Another strand of the literature looks at the relation between the labour supply of individuals and that of their parents. The overall findings indicate a positive correlation, implying that gender norms are transmitted from one generation to the next. For instance, Fernandez et al. (2004) find that women are more likely to work if their mother-in-law also worked, while Morrill and Morrill (2013) show that there is a positive relationship between mothers' and daughters' labour supply. In line with these findings, Del Boca et al. (2000) also find that women's labour supply is related to that of both their mothers and mothers-in-law, while Olivetti et al. (2016) find a positive relation between daughters' number of hours worked and the hours worked by both their own mothers and the mothers of their childhood peers.

Other studies move away from labour market outcomes and analyse gender norms by exploring marital satisfaction among males. Butikofer (2013) finds marital satisfaction to be lower when a man's wife works and contributes to household income, but only for men raised in a traditional family where the mother did not work. Similar findings are reported by Bonke (2008) and Bonke and Browning (2009). This result provides evidence that gender socialisation at the family level, and at a young age, is crucial in determining lifelong gender norms.

The relation between gender norms and educational outcomes has also been studied recently. For instance, Gonzales de San Roman and de la Rica Goiricelaya (2012) examine the cross-country gender gap in test scores revealed by PISA data. They find gender norms to be an important determinant; in particular, mothers' labour force participation is positively related to daughters' test performance. In a different setting, Blunch and Das (2014) show that increased access to education for girls explains much of the rise in

egalitarian views towards female access to education in Bangladesh. Similarly, employing data from 157 countries, Cooray and Potrafke (2010) find conservative culture and religion to be the primary obstacle towards gender equality in education.

Last but not least, some papers investigate the relationship between parental and child gender attitudes. Both Farre and Vella (2013) and Johnston et al. (2014) find strong correlations in these attitudes in cases in the USA and UK, respectively. In line with previous findings, these two studies also find a positive association between boys' attitudes during childhood and their future wives' labour supply in adulthood.

## **1.3 The Theoretical Framework**

This section describes the process through which children develop their beliefs. Socialisation is first explained, followed by a simple model that formalises this process through a utility-maximising framework. Some theoretical predictions follow, which will be useful in interpreting the empirical results in subsequent sections.

### **1.3.1 The Socialisation Process**

Children develop gender norms through socialisation (Epstein and Ward 2011; etc.). Social psychology identifies two main sources of socialisation: the family (vertical socialisation) and society at large (horizontal socialisation). Children acquire norms by interacting with and observing the particular behaviours of these two social institutions.

The family is, of course, the first point of contact with the outside world. Therefore, parents play a key role in socialisation. Parents are assumed to be altruistic towards their children, and have preferences regarding the norms children develop. In particular, they aim to socialise their children to their own values. Parents choose how much effort to exert to increase the probability that vertical socialisation is successful. Effort in this setting can manifest in two main ways. The first is that parents can express their beliefs through direct discussion with their children, which requires investment in 'family time'. Effort can also take the form of parents' actions, particularly the role adopted by each parent. Children thereby receive signals about appropriate gender roles by observing their parents'

household responsibilities, in particular who the breadwinner is and who is responsible for household production and child care.

As children grow older and start to interact with a social circle beyond their families, they are also exposed to what society at large deems to be the appropriate role of women. The primary sources of horizontal socialisation are the child's school, peer group, and exposure to social norms through mass media.

Each child has some preference for conforming to the family and society. This preference will depend not only on individual characteristics, but also on the probability that family socialisation was successful. This probability in turn, will depend on family and social norms; families following the social norm will have a higher chance of transmitting their values because they will not have to overcome an opposing social norm.

Having learnt what the family and society believe to be the appropriate gender roles, the child chooses her own gender values. In doing so, the child takes into account the cost of deviating from these norms. This can be thought of as a psychological cost of interacting with others who do not share the same beliefs, and arises from self-image concerns, that is, concerns about how personal views will be judged by others. The stronger the preference for conformity is, the higher is the cost of deviating from the prescribed norms.

### 1.3.2 Formalising the Socialisation Process

To formalise this process, I build on the Georgiadis and Manning (2013) model of national identity. In my setting, each child has the following utility function:

$$U = -\frac{1}{2} \left[ c_F(x_F, x_S, Z)(x - x_F)^2 + c_S(x_F, x_S, Z)(x - x_S)^2 \right] \quad (1.1)$$

where:

$x \in [0, 1]$  represents the child's choice; higher  $x$  indicates more traditional beliefs  
 $x_F \in [0, 1]$  is the child's family norm; higher  $x_F$  indicates more traditional beliefs  
 $x_S \in [0, 1]$  is society's norm; higher  $x_S$  indicates more traditional beliefs  
 $c_F \in (0, 1]$  is how strongly the child wants to conform to the family  
 $c_S \in (0, 1]$  is how strongly the child wants to conform to the society  
 $Z$  is a vector of characteristics affecting preferences for conformity

The intuition of the model is simple. The larger the distance between the child's belief ( $x$ ) and that of the family's ( $x_F$ ) and society's ( $x_S$ ), the larger the psychological cost is. This cost is increasing in the preference for conformity to the family ( $c_F$ ) and society ( $c_S$ ). The stronger this preference is, the more costly it is to deviate from a norm to which the child wants to conform. Taking this into account, the child chooses her optimal belief according to the following rule:

$$x^* = \underset{x}{argmax} U = x(c_F, c_S, x_F, x_S, Z) = \frac{c_F(x_F, x_S, Z)x_F + c_S(x_F, x_S, Z)x_S}{c_F(x_F, x_S, Z) + c_S(x_F, x_S, Z)} \quad (1.2)$$

Hence, the child chooses how traditional her gender view will be by weighting the family and social norms by her preferences for conforming to each.

### 1.3.3 Comparative Statics

I now use the model to predict how a change in conformity preferences will affect  $x^*$ , as this will be useful for interpreting my empirical results. The comparative statics are as follows:

$$\frac{\partial x^*}{\partial c_F} = \frac{c_S(x_F, x_S, Z)(x_F - x_S)}{(c_F(x_F, x_S, Z) + c_S(x_F, x_S, Z))^2} \quad (1.3)$$

(1.3) shows the response of  $x^*$  to a stronger preference for conformity to the family.<sup>4</sup> Its sign depends on which institution is more traditional, society or the family. As I will show in Section 1.5.1, the social norm in England regarding the role of mothers with young

---

<sup>4</sup>An analogous result holds for the case of preferences for conformity to society.

children is very traditional, while in the empirical framework (Section 1.4), the family norm I will examine will represent very nontraditional views. Thus, without loss of generality, these norms can be modelled by setting  $x_F = 0$  and  $x_S = 1$ . Given these restrictions, the prediction is that a stronger preference for conforming to the family's norm leads to a decrease in  $x^*$ , that is to a *less traditional* norm. The intuition is that the more strongly a child wants to conform to her family's relatively less traditional norm, the more she has to differentiate herself from the traditional social norm and adopt a more nontraditional belief. Otherwise, the psychological cost of deviating from the less traditional norm, to which she now wants to conform more strongly, will increase. I revisit this result in Section 1.7 where I rationalise my finding that girls growing up in nontraditional families develop more traditional views.<sup>5</sup>

## 1.4 The Empirical Framework

My aim is to estimate  $x^* = x(x_F, x_S, Z)$ . The Next Steps survey (described in the next section) provides data on children's gender norms through responses to the following question: “*Women should never work full-time when they have young children. Do you agree with this statement?*” Stating that a woman should *never* work full-time if she has young children reflects a very traditional view of gender roles; within the language of the model, this implies  $x^* = 1$ . Based on this, I define the outcome variable capturing children's norms as follows:

$$\text{Traditional Norm} = \begin{cases} 1 & \text{if Agree } (x^* = 1) \\ 0 & \text{if Disagree } (x^* < 1) \end{cases}$$

As the aim is to examine the intergenerational transmission of gender norms, I have to assess how successful vertical and horizontal gender socialisation are. Since the children in the survey are all from England, they are all exposed to the same social norm.<sup>6</sup> The variation of interest will therefore come from differences in family norms. Because I am

---

<sup>5</sup>In Appendix 1A, I also discuss the prediction  $\frac{\partial x^*}{\partial x_F}$  and show how it matches my empirical results.

<sup>6</sup>The social norm may of course still vary by region. To account for this, regional fixed effects are included in all empirical specifications.

particularly interested in examining the development of children's norms when the family and social norm are *oppositional*, I define the family norm in a way that translates into the model as  $x_F = 0$ . I do so using the following definition:

$$\text{Nontraditional Family} = \begin{cases} 1 & \text{if Mother Earns More than Father} \\ 0 & \text{otherwise} \end{cases}$$

I thus want to examine how children's gender norms develop when living in a nontraditional family but a traditional society (evidence on the English society's gender norm is shown in Section 1.5.1). Moreover, I am interested in between-sex heterogeneity in gender socialisation. I will therefore estimate the following (linear) probability model:<sup>7</sup>

$$\begin{aligned} \text{Pr}(\text{Traditional Norm})_i = & \beta_0 + \beta_1 \text{Nontraditional Family}_i + \beta_2 \text{Female}_i + \\ & \beta_3 \text{Nontraditional Family}_i \times \text{Female}_i + Z_i' \zeta + u_i \end{aligned} \quad (1.4)$$

where  $Z$  is a vector of covariates to be controlled for. The identifying assumption here is selection on observables, i.e.  $\mathbb{E}(u_i | \text{Nontraditional Family}_i, \text{Female}_i, Z_i) = 0$ . While this is not a weak assumption, it is not implausible given the vast array of characteristics that  $Z$  will include. Nevertheless, the primary aim of this empirical strategy is not causal inference but rather to examine the relationship between having a breadwinner mother and children's gender values. Section 1.7 will present another identification strategy - a regression discontinuity design - to provide evidence that the proposed mechanism driving these primary results does indeed have a causal interpretation.

---

<sup>7</sup>For robustness, I also estimate every linear probability model presented in this paper using probit and logit specifications, with results shown in Appendix 1C.

## 1.5 Data

### 1.5.1 The Social Norm

What is the social norm to which the children in my sample are exposed? I answer this using information from the International Social Survey Programme (ISSP), which collects data internationally on a broad range of social issues. Questions related to gender norms are included in the 1988, 1994, 2002, and 2012 surveys. Responses from the 2002 survey were chosen, since this is the period that coincides with the Next Steps survey, with children in the sample aged between 15 and 16 at the time. The ISSP 2002 contains data from a representative sample of 1960 observations from the UK. Because I study views on the appropriate roles of men and women in childrearing and in providing income, I focus on heterosexual couples with children and analyse the labour supply of both the father and the mother. In particular, I look at their responses to the following four questions:

*“Did you work full-time, part-time, or not at all when... ”*

1. *“...you had no children”*
2. *“...you had a child under school age”*
3. *“...your youngest child was still in school”*
4. *“...your children left the home”*

These four cases track the parents’ labour market activity from the time before the child is born (case 1), up to adulthood (case 4). Cases 1 and 4 focus on parents with no dependent children. Case 2 concerns parents with children under four years of age, and case 3 focuses on parents of school-age children. Figure 1.2 shows the responses to each question for both the fathers and mothers. Figure 1.2a reflects the first case, showing that there is no prescription against women working full-time when they have no children. When they have no child care responsibilities, over 80% of women work full-time, and only a small minority stay at home (less than 9%). Compared to the fathers’ labour supply however, the changes in the mothers’ labour supply when they have children are striking. Figure 1.2b reveals that when women have children below the age of four, the majority

(53%) do not work at all, while the percentage of those working full-time drops from 82% to 14%. In the same circumstance, the proportion of fathers working full-time remains steadily above 90%. These figures reveal how traditional the social norm in the UK is. There is a strong prescription that the role of mothers with young children is not at the workplace.

This is evident despite the large opportunity cost of doing so. Note that maternity leave in the UK is not very generous. Mothers receive a replacement rate of 90% for just the first six weeks, and then a fixed amount of £100 per week for up to 18 weeks (in the period studied).<sup>8</sup> Despite the large amount of foregone earnings of mothers staying at home beyond the first six weeks,<sup>9</sup> most women do not work for up to four years after child birth. This further highlights that the social norm is for the father to work full-time and provide for the family, and for the mother to reduce her labour supply in order to provide child care.

This prescription holds for all women who have dependent children. Figure 1.2c shows that the majority of women (nearly 80%) still do not work full-time, even though their children are older but still of school age. Working full-time only becomes the mode again when the children have grown up and left the home (Figure 1.2d), although the proportion of women in this category is still low (46%) compared to the first case. It is also interesting to note how stable men's labour supply is. While the age of a child strongly predicts the mother's labour supply, fathers always work full-time regardless of the presence and age of children. Overall, these figures reveal that the children in my sample are growing up in a society with very traditional gender norms.

---

<sup>8</sup>In 2002, maternity (job-protected) leave could be taken for up to 39 weeks. Maternity *pay* could only be received for 18 weeks, at a rate of 90% of previous earnings for the first six weeks, and a fixed amount capped at £100 per week for the remaining period. A 2003 reform increased maternity leave to 52 weeks and pay to 26 weeks, but the 90% replacement rate was still paid for only the first six weeks, with the capped amount for the remainder. Statutory paternity leave did not exist in 2002.

<sup>9</sup>Of course, some women may also receive more generous maternity leave and pay directly from their employer. However, evidence shows that this is also quite limited in duration. The most recent survey for UK employers (Mapper 2017) shows that most do not offer anything beyond what is mandated. Among those who do, the vast majority do not offer anything beyond the first 26-39 weeks. In the previous decade, such private provisions would likely had been even less generous.



### 1.5.2 The Next Steps Survey

I will estimate (1.4) using data from the Next Steps survey.<sup>10</sup> Next Steps follows a birth cohort of 8682 individuals between 2004-2010 across seven annual waves. All individuals were born in England between 1989-1990. The first wave of data collection took place in 2004, when cohort members were attending year nine in school and were aged 13-14. Due to the nature of my question, I focus on dual- and heterosexual-parent families (in which parents are married or cohabiting). I also drop families in which both parents earn zero income. Otherwise, these would be classified as ones without a female breadwinner; the problem is that in these families, there is neither a male nor a female breadwinner. This would make gender socialisation opaque as there is no breadwinner for the child to observe.

These restrictions drop approximately 40% of the sample (the vast majority being single-parent families or dual families with both parents long-term unemployed). As is the case with all surveys, Next Steps suffers from unit and item non-response. As a consequence, the estimation sample is restricted to the 1487 cohort-members for whom there is no missing information on the variables of interest. While this may raise concerns about sample selection, I check and confirm that my final sample is very similar to the original sample on all key outcome variables of interest. The comparisons are shown in Table 1C.1 of Appendix 1C. Importantly, the proportion of nontraditional families is approximately the same before and after the sample restrictions (14.3% Vs 14.5%). What is further reassuring is that this proportion is also nearly identical to the estimate of 14.7% that I find from a separate, much larger UK dataset (the Family Resources Survey), which I discuss further in Section 1.7.4. This provides even more confidence that I am not picking up a specific sub-sample of nontraditional families.

---

<sup>10</sup>Survey previously known as the Longitudinal Study of Young People in England. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

### 1.5.3 Choice of Variables

#### Main Variables of Interest

The outcome variable of interest is ‘Traditional Norm.’ It is a binary variable derived from responses to the question “*Women should never work full-time when they have young children. Do you agree with this statement?*” as explained in Section 1.4. Responses to this question are taken from the 2010 wave. The timing is ideal as it is asked at an age by which children’s socialisation has been completed; I thus avoid ascribing norms to possibly ‘transitory’ beliefs. The timing is also supported by UK evidence showing that beliefs at this age strongly predict individuals’ future labour supply (Johnston et al. 2014). The main independent variable will be a dummy named ‘Nontraditional Family,’ as also defined in Section 1.4. In the sample, 14.5% of the children live in such families (i.e., the mothers are the breadwinners), while 31.1% of them hold traditional gender views. The ‘Nontraditional Family’ status is derived from data on gross earnings of the father and mother of each child for the year 2004. While data on incomes exists for years 2004-2007, the 2004 wave was chosen to minimise the amount of dropped observations due to item non-response. I explain how I account for possible transitory income affecting my results in Section 1.6.1.

#### Control Variables

A wide range of control variables is included in the regressions. The aim is to control, as much as possible, for variables that are correlated with both the child expressing traditional norms and living in a nontraditional family. The variables can be grouped into four categories: family, child and geographical characteristics, and proxies for parental-socialisation effort. Family characteristics include family structure and income, household size, parental marital status, age, religion, education and social class. Child characteristics include country of birth, religion, ethnicity, birth weight, disability status, and whether a child has ever attended a religious or single-sex school. Geographic characteristics comprise of region and area-type controls, while parental-socialisation effort proxies measure frequency of parent-child interactions and parental strictness. A detailed discussion of the intuition behind the

inclusion of each variable is presented in Appendix 1B. Variable definitions and associated summary statistics are shown in Table 1C.3 of Appendix 1C.

## 1.6 Results

### 1.6.1 Main Results

Table 1.1 presents the main results. The first column essentially provides summary statistics for the unconditional difference in the mean probability of developing traditional norms between boys and girls in nontraditional and traditional families. These results uncover significant heterogeneity; I find evidence of intergenerational persistence of gender norms, but only for boys. Boys in nontraditional families are 3.2% points less likely than boys in traditional families to develop traditional norms. For girls however, results are somewhat unexpected. Girls in nontraditional families are 4.4% points *more* likely to develop traditional norms, compared to girls in traditional families. Hence, it seems that girls react to the family norm when it opposes the social norm, and are more likely to adopt the latter instead, rendering vertical socialisation unsuccessful.

The immediate concern is that these results may be driven by omitted variable bias. A wide range of controls is next introduced, with results shown in column (2). The full estimates of all covariates' coefficients are shown in Table 1C.4 of Appendix 1C. Despite the vast number of control variables, the previous findings cannot be explained away. If anything, the between-group comparisons are now even larger and statistically significant at lower percentage levels. For robustness, I have also estimated a version of (2) that includes interaction terms between each control variable and the 'Female' dummy, and have also run (2) separately for boys and girls (after dropping interaction terms). The results are virtually unchanged, and are available upon request.<sup>11</sup>

Which background characteristics are associated with developing more traditional views? Starting with family structure and parental characteristics, I find household size to be associated with children expressing more traditional beliefs, although the marital status of

---

<sup>11</sup>I have also repeated these checks for all subsequent columns and tables where controls are included, with the findings being highly robust.

parents appears to be insignificant. While the age of each parent is also statistically significant, its economic significance is too negligible to have any meaningful interpretation, due to the tiny coefficient sizes.

An important characteristic that proxies for family culture is religion. As expected, I find this to be a highly statistically significant determinant of children’s beliefs. Not only is the specific religion important, but the relation between child beliefs and parental religion seems to differ by parent sex. For instance, compared to Christian fathers and mothers, the likelihood for traditional norms is lower with a Muslim father but higher with a Muslim mother. Further, it also matters whether parents state that religion is important to their way of life. This result is in line with previous findings (Guiso et al. 2003), as prescriptions on appropriate behaviours and social roles characterise religious beliefs.

Besides parents’ religion, I also find a significant relationship between parental socio-economic characteristics and norms, which also confirms previous research (Dee 2004; Stankov 2009; Kanazawa 2010). Children of parents with lower education, lower socio-economic class, and lower income are all more likely to develop traditional norms.

With regard to child attributes, country of birth and ethnicity seem to matter. The likelihood of developing traditional norms is lower for UK-born children, while it is the highest for whites, consistent with earlier research findings (Dugger 1988). The only exception is the black African category, while the most traditional ethnic group appears to be the Bangladeshi. Consistent with the effect of parental religiosity, children who state that religion is important to their way of life are more likely to develop traditional norms. Furthermore, having special education needs (SEN) is positively related to expressing traditional views. This result may be driven by the fact that the majority of SEN teachers and caregivers are female (Department for Education 2014), which may affect an SEN-status child’s view on appropriate gender roles. Last, I consider the proxies for parental-socialisation effort. Conditional on all the other covariates, parental effort does not seem to affect the likelihood of expressing traditional views.

In the remaining columns, I show results from robustness checks. First, I address a possible concern regarding the information used to define a nontraditional family. As explained previously, these are taken from just one wave (the first), and therefore, may be

affected by transitory income, hence not reflecting the true earnings trajectory of a family. To account for this, the specification in column (3) introduces controls for whether the employment status of the mother and father has remained constant over all the years for which employment data exist. This spans five years before the survey and up to wave 4. The results are very stable, providing confidence that such confounders cannot explain the results.

Next, what if children’s gender norms are actually affected by the children’s observation of parental differences in either education or job status, which can be correlated with who the breadwinner is, driving the results I am finding? To rule this out, I control for both of these factors in turn in columns (4) and (5). Neither the relative education nor the job status of the mother matters. The mother’s role as breadwinner is still what drives the results.

Further, what if the results are driven by the fact that in some families, one of the parents does not work? This would imply that what matters for socialisation may not be that the mother earns more than the father (and vice versa), but that (s)he is simply the only parent working. To rule this out, I run the same regression on the sub-sample of children whose parents both work. As column (6) shows, results are again not affected. If anything, the between-group differences are even larger for only dual-earner families. Finally, I re-estimate all regressions using a probit and logit specification and find nearly identical estimates.<sup>12</sup>

## 1.6.2 Other Outcomes Related to Gender Norms

I now consider, as an extension, two further outcomes associated with gender norms. They are derived from the following survey questions: “*Do you agree that having a job that pays well is important?*” (asked in wave 1) and “*Would you like to study for a science degree at university?*” (asked in wave 3). Both are coded as binary variables (Yes versus No). The first question is related to gender norms because it captures how children envision their financial independence in adulthood. For example, if children have traditional norms, I expect to find boys more likely to agree that high earnings are important, due to their

---

<sup>12</sup>Results from these robustness checks are shown in Appendix 1C.

future role as breadwinners. Similarly, girls with traditional norms, expecting their spouses to be the breadwinner, should place less importance on high future earnings.

The second question is related to the educational gender gap in science (and STEM subjects more generally) (OECD 2012). In fact, a significant proportion of the gender pay gap among university graduates has been attributed to gender gaps in entry into science degree programmes (Brown and Cororan 1997; Hunt et al. 2012; Weinberger 1999). The sciences have diachronically been considered as ‘masculine’ disciplines and have led to stereotypes about the appropriate qualifications and thus professions for each sex. It therefore becomes important to examine whether family socialisation exacerbates this phenomenon by also propagating traditional gender norms in this dimension. To examine the effect of family socialisation on these outcomes, I re-estimate (1.4) using each of these further outcomes as the dependent variable.

The results are shown in Table 1.2. I again find that girls in nontraditional families are more likely to express traditional views, in opposition to their family’s norm but in line with society’s. Specifically, they are 4.2% points less likely to believe that high wages are important, and 10% points less likely to want to study science, compared to girls in traditional families.

### **1.6.3 Are Nontraditional Families Truly Transmitting Nontraditional Views?**

The results suggest that girls growing up in nontraditional families develop views opposite to what we may expect. Instead of adopting their parents’ nontraditional values, they become more traditional, along several dimensions. One explanation is that parental actions do not coincide with beliefs, and so the families I am categorising as nontraditional may still be transmitting traditional values. Consider, for instance, families that hold traditional gender views but are forced by their economic circumstances to have a female breadwinner. Not only will their actions not match their beliefs, but their dissatisfaction with this situation may make it even clearer to children that having breadwinning mothers is not the accepted norm. If this is true, it may explain why girls in nontraditional families develop more traditional views.

While this explanation is possible, a plethora of evidence suggests that beliefs and actions are indeed aligned. First, Figures 1.1b and 1.1c imply this alignment: there is strong cross-country correlation between gender views and female labour supply. Fortin (2015), who focuses on the US, reaches the same conclusion. She shows that changes in gender-role attitudes are the strongest predictor of changes in female labour force participation (particularly the recent ‘opting-out’ phenomenon). Further evidence comes from Johnston et al. (2014), who use data from the 1970 British Cohort Study. They find a strong correlation between the mother’s gender views and her hours worked. Interestingly, they also find children’s beliefs to be strongly related to maternal labour supply, over and above the mother’s gender views. This evidence emphasises that views and actions are aligned, but also that parents’ actions do matter for how children are socialised, irrespective of parental beliefs.

Nevertheless, I also examine my data for the possibility that female breadwinners are not actually transmitting nontraditional views. Because the survey does not ask parents about their own gender views, I approach this differently. I look at two situations in which families may be ‘forced’ to have a female breadwinner: (1) the father cannot work due to a disability, and (2) the family is very poor. To test whether my results are driven by such factors, I split my sample into groups that differ by how constrained families are based on these factors, and repeat my analysis for each sub-sample.

Table 1.3 shows the results. Each column corresponds to a different sub-sample. The results show that constrained families do not drive the results. If anything, the opposite holds. Columns (1) and (2), for instance, show that results are driven purely by families where the father is *not* disabled. Columns (3) and (4) split the sample into those above and below the median family income. Again, I find that the results are driven by families that are not financially constrained. In the last column, I restrict the sample to families in which both the father is not disabled, and total family income is above the median level. Again, the main finding holds. These results show that my findings cannot be explained by female breadwinner families being ‘forced’ into nontraditional-family status.

Another explanation may be that female breadwinners do second shifts, i.e. try to ‘compensate’ for their ‘violation’ of traditional norms by taking on a more traditional

gender role in the household through increased domestic work. I therefore examine whether second shifts are more prevalent among nontraditional families. I create a dummy for whether the mother provides any positive amount of domestic work and use this measure as the outcome in specification (1.4). Table 1.4 shows evidence against second shifts. The difference in the probability of second shifts between nontraditional and traditional is approximately zero and not statistically significant. This finding is consistent with other UK evidence. Female relative earnings and domestic work are not positively related in either the 2000 Time Use Survey (Washbrook 2007) or the British Household Panel Survey (Kan 2008).

#### **1.6.4 How do These Results Square With the Literature on Maternal Labour Supply?**

The literature has focused on the relation between mothers' labour force participation and either daughters' gender attitudes or daughters' own labour supply (Del Boca et al. 2000; Fernandez et al. 2004; Morrill and Morrill 2013; Olivetti et al. 2016). These studies find a positive relation, suggesting the transmission of nontraditional gender norms. Here, I show that focusing only on labour supply, without taking into account which parent works or earns more, can be misleading. Table 1.5 shows the results. Column (1) replicates previous approaches. It shows a regression of the probability that a child develops traditional norms on a dummy for whether the mother works, a dummy for the child being female, and their interaction. In line with previous work, daughters of working mothers are (by 9% points) less likely to develop traditional norms, compared to daughters of non-working mothers, with this difference being highly statistically significant. Based on these results, we would conclude that maternal labour force participation is enough to successfully transmit more egalitarian gender values. Now consider column (2), which replaces the 'mother works' with a 'mother works more (hours) than father' dummy (and its associated interaction term). We now reach the opposite conclusion. The first thing to note is that there is now no statistically significant difference in the probability to develop traditional views between daughters of mothers working less than the fathers, and of mothers not working at all. Secondly, compared to both of these groups, daughters of mothers who work more



than the father are actually 9.2% points *more* likely to develop traditional norms. Hence, what matters is not whether mothers work, but whether they work more than the fathers.

To examine whether this depends on breadwinner status, I introduce the dummy for ‘Nontraditional family’ and its interaction with the child being female. As column (3) shows, each of the associated variables is individually statistically significant, though there is no statistically significant difference between nontraditional and traditional families (p-value is 0.23). Finally, I introduce in column (4) the covariates used in previous tables to control for the usual omitted variable bias concerns. Conditional on covariates for parental employment, family, child, geographic characteristics and parental socialisation effort, I find that mothers relative hours have no statistically significant relation with children’s beliefs. What actually matters is not whether the mother works more, but whether she earns more than the father. Specifically, the probability that girls adopt traditional views is 4.5% points higher among nontraditional families, with the difference being statistically significant at the 1% level. Overall, the findings of this section call into question previous results relying on extensive margin labour supply correlations. Failing to account for whether the mother earns more than the father can lead to misleading conclusions regarding the intergenerational transmission of gender norms.

## **1.7 Examining the Mechanism: A Behavioural Explanation**

The results reveal that family socialisation is not successful for girls in nontraditional families. The natural next question is then: why are girls brought up in nontraditional families more likely to reject their family norm and to develop traditional norms instead? This section explores preferences for conformity as a possible explanation, and develops an identification strategy using a regression discontinuity design to estimate the effect of living in a nontraditional family on these preferences.

### 1.7.1 Preference for Conformity to the Family's Norm

One key parameter in the model presented in Section 1.3 is the strength of the preference for conformity to the family. The model showed that, when the family norm is less traditional than the social norm, a stronger preference for conformity to the family implies the development of relatively less traditional norms. In the same way, a weaker preference for conformity to the family then implies more traditional norms, which is what I observe for girls. Could a weaker preference for conformity to the family's norm therefore explain my results? To test this prediction, I run the specification from (1.4) but replace the dependent variable with conformity measures.

Two such measures of conformity to the family are explored. The first comes from the first wave, at a point in time when children have already chosen which subjects to study at the General Certificate of Secondary Education (GCSE) level (the last stage of compulsory education in the UK). This is an important decision as it affects what one can then study at the General Certificate of Education (GCE) Advanced Level, which directly affects university entrance. The particular measure I will exploit is agreement with the statement "*I chose what to study at GCSE level based on what my parents wanted.*" The second is a measure of whether children argue with their parents, and is available from the second wave. Both conformity indicators are constructed as binary (Yes versus No) variables.

Table 1.6 shows the results from re-estimating the main specification but using each of these alternative measures as the dependent variable. The findings support the model predictions, confirming that girls in modern families are less conformist to their family. Compared to girls in traditional families, girls in nontraditional families are less likely to have had their GCSE subjects chosen by their parents by 7.4% points, and are more likely to argue with them by 6.5% points.

### 1.7.2 Why Are Girls in Nontraditional Families Less Conformist?

If girls in nontraditional families are more traditional because they are less conformist to their family, *why* are they less conformist? Is there an underlying cause making them less conformist? To answer this, I draw on research in social psychology. Beginning with the

seminal work and experiments of Asch (1951; 1952; etc.) the literature has established that people have a strong preference to conform to the view of the majority. With regard to between-gender differences in conformity preferences, it has also been shown that girls are much more susceptible to the majority’s views than boys are (Eagly and Carli 1981; Santee and Jackson 1982). I combine these findings with the fact that girls in nontraditional families are growing up in a norm-minority family, in the sense that only about 14.5% of families with dependent children in the UK have female breadwinners.<sup>13</sup> Based on these facts, I formulate the following hypothesis: living in a nontraditional family causes girls to reject their family’s norm-minority status, because it violates the oppositional social norm. Hence, growing up in a nontraditional family makes girls less conformist to their family, which in turn makes them adopt more traditional gender views. I next test this hypothesis.

### 1.7.3 Identifying the Mechanism Using a Regression Discontinuity Design

I estimate the effect of having a breadwinner mother on children’s preferences for conformity by exploiting quasi-experimental variation in nontraditional family status and applying a regression discontinuity design. This is possible because growing up in a nontraditional family is a deterministic function of the relative family income earned by the mother, which thereby creates a discontinuity in the probability of treatment.

Within the RD design language, the treatment is living in a nontraditional family, and the assignment variable is the mother’s share ( $s_i$ ) of family income, defined as:

$$s_i = \frac{Mother's\ Income_i}{Mother's\ Income_i + Father's\ Income_i} \quad (1.5)$$

I exploit the jump in treatment at the 0.5 threshold of the assignment variable:

---

<sup>13</sup>This statistic is based on the Next Steps Survey. Using the UK Family Resources Survey, I reassuringly find a nearly identical estimate (14.7%).

$$Pr(Treatment)_i = \begin{cases} 1 & \text{if } s_i > 0.5 \\ 0 & \text{if } s_i \leq 0.5 \end{cases} \quad (1.6)$$

and use this to estimate different versions of:

$$Pr(Conform to Family)_i = \tau_0 + \tau_1 Treatment_i + f(s_i, \chi) + Treatment_i * f(s_i, \psi) + \varpi_i \quad (1.7)$$

$f(s_i, ;)$  is a polynomial function with parameter vector  $\chi$  that controls for the assignment variable and  $\psi$  that controls for the interaction between the assignment variable and treatment status.  $\tau_1$  is the causal effect of living in a nontraditional family on the probability of conforming. For robustness, some specifications will also include the same vector of controls  $Z$  used in the previous analysis.

#### 1.7.4 The RDD Identification Assumption

Identification of  $\tau_1$  requires local random assignment of the assignment variable. This means that families cannot precisely choose where they locate around the 0.5 threshold of the mother's income share (Lee and Lemieux 2010). In a recent paper using US data, Bertrand et al. (2015) (henceforth BKP) show that the density of wives' income shares exhibits a drop around 0.5. While the BKP finding would invalidate my empirical strategy, there is no reason why it should always necessarily hold in other countries. Eriksson and Stenberg (2015), for instance, repeat the BKP exercise for Sweden and find no evidence of sorting at the 0.5 threshold.

In this section I show evidence against sorting and in support of my identifying assumption. I begin by examining the probability density function of the mother's income share to graphically test for potential manipulation at the cutoff. If parents can precisely choose the mother's income share, I should observe sorting around the cutoff. Figure 1.3 plots the number of families in each bin of  $s_i$  (with size 0.02) and shows that the assignment variable varies smoothly across the 0.5 threshold. I also perform a McCrary test and fail to reject the null of no discontinuity at the cutoff. Both the graphical and statistical evidence support the identifying assumption.

A concern is that I may not be detecting any sorting either because of small sample size, or because of measurement error. I address this by drawing on two other UK datasets to examine the distribution of wives' income shares. The first is the much larger, Family Resources Survey (FRS).<sup>14</sup> The FRS includes cross-sections of more than 20 thousand households per year and contains detailed information on respondents' (and importantly spouses') incomes. Due to its high quality and sample size, it serves as the primary source of information for the Department for Work and Pensions (DWP) to guide UK welfare policy (DWP 2015). The second source is administrative data from the Survey of Personal Incomes, a dataset generated from tax records of the former UK Inland Revenue (known today as Her Majesty's Revenue and Customs).<sup>15</sup> I use the fact that prior to the introduction of independent taxation in 1990, couples had to file taxes jointly, enabling me to identify husbands and wives in administrative datasets. I use the only available dataset from this period - the SPI for the 1985-86 tax year - to examine the distribution of wives' incomes. Table 1C.2 in Appendix 1C provides summary statistics for each dataset.

Following BKP, I calculate individual total income as the sum of wage and self-employment income, and estimate the wife's income share. To make the FRS sample comparable to the families in my Next Steps survey, I restrict my analysis to dual-parent, heterosexual families, in which there is at least one dependent child and at least one parent working, and consider the years 1997-2004. For the SPI, I restrict my sample to couples where both are below pensioner age.<sup>16</sup>

Figure 1.4 plots the distribution of the wife's income share for different samples in the FRS and SPI. I start by following BKP and do not differentiate between wage and self-employment income. Figures 1.4a and 1.4b show the distribution for the full sample in the FRS and SPI respectively. Similar to BKP's findings, I find evidence for significant bunching at (or just below) the 0.5 threshold, implying that the assignment variable is manipulated. The sorting can be observed visually and is also confirmed by the McCrary test. Sorting is much more evident in the Survey of Personal Incomes, consistent with

<sup>14</sup>Department of Work and Pensions, Office for National Statistics. Social and Vital Statistics Division, NatCen Social Research. (2016). *Family Resources Survey*. UK Data Service.

<sup>15</sup>Inland Revenue. Statistics Division. (1989). *Survey of Personal Incomes, 1985-1986: Public Use Tape*. UK Data Service.

<sup>16</sup>As is usual with administrative tax records, the SPI does not contain any further information on demographics, number of children, etc.

administrative data being more accurate and hence more suited to uncover manipulation, as expected.

However, a concern with the BKP method is that no distinction is made between wage and self-employment income; these are aggregated into a single measure of income. Of course, self-employed individuals have much more control over their exact incomes, but those with only wage income may only be able to sort at shares of 0 (or 1) by having only the husband (wife) work (this is also evident in the Next Steps distribution with sorting at 0). It is therefore unclear whether the BKP results are driven by the self-employed.

To clarify this, I next split each sample into two groups: families with at least one self-employed spouse, and families with no self-employed spouses. Figures 1.4c and 1.4e (1.4d and 1.4f) show the corresponding densities for the FRS (SPI). I now find a large discrepancy that speaks against BKP's results. While there is very significant bunching among the self-employed at the 0.5 cutoff (Figures 1.4c and 1.4d), there is none among wage-earning families (Figures 1.4e and 1.4f). The McCrary test again confirms the visual evidence. BKP's findings are therefore not universal; in the UK, at least, they only hold for the self-employed.

This finding is in line with the recent literature in public finance that looks at 'bunching' at kinks and notches in income tax schedules.<sup>17</sup> While economic theory predicts that workers should bunch at notches and convex kinks, the evidence shows that such behavioural responses are mostly driven by the self-employed. The lack of bunching among wage earners is attributed to earnings adjustment frictions that constrain workers from freely choosing their exact earnings (Chetty et al. 2011; Kleven and Waseem 2013; Gelber et al. 2016). This also follows from a related literature highlighting significant limitations in the choice workers have over their hours worked. While the self-employed can choose their own work hours, wage earners cannot freely choose from a variety of job hours packages; hours requirements are usually fixed, dictated by employment contracts, and vary little across jobs (Dickens and Lundberg 1993; Blundell et al. 2008). To expect, then, that couples can precisely choose where to locate around the 0.5 threshold would require not only that workers can choose their work hours and earnings precisely, but that they can do so also

---

<sup>17</sup>For a review, see Kleven (2016).

in response to their spouse's earnings.<sup>18</sup>

The evidence presented in this section supports the RDD identifying assumption regarding the assignment variable, at least among wage earners. While the self-employed are not a large group in the Next Steps data (they are approximately 10%), I drop them from my analysis to ensure my results are not affected from any potential sorting. Before proceeding with estimation, I also test whether baseline covariates exhibit discontinuities across the threshold. If treatment is randomised, these should be locally balanced on each side of the 0.5 cutoff. I test this by estimating different versions of (3.1), using each element of  $Z$  as the outcome variable. Results are shown in Table 1C.5 of Appendix 1C and confirm that baseline covariates are balanced overall. I find only a very small proportion of covariates to exhibit any discontinuity that is statistically significant. Out of 61 total covariates, just one is significant at the 1% and four at the 5% significance level. This very small number of significant discontinuities is consistent with the rate of false positives expected given the large number of covariates considered (Lee and Lemieux 2010). Applying a Bonferroni-type correction would render the number of significant discontinuities even smaller. Furthermore, in the results that follow, I show that estimates of the treatment effect are not significantly affected by the inclusion of controls, providing further evidence that baseline covariates are balanced around the threshold.

### 1.7.5 Regression Discontinuity Estimates

I first examine whether a discontinuity in the probability of conforming to the family can be identified visually at the 0.5 cutoff. Consider Figure 1.5 which shows the relationship between the share of mother's income and the two conformity measures. In both cases, there is a sharp, visible discontinuity in the probability of each statement being true around the cutoff. Compared to those just below the cutoff, girls whose mothers have a share of family income just above 0.5 are much less likely to have chosen their GCSE subjects based on what their parents wanted (Figure 1.5a). They are also also much more likely to argue

---

<sup>18</sup>Moreover, to do so requires individuals to know exactly what their spouse earns. A recent survey in the UK (Noodle 2016) finds that half of married people do not know their spouse's earnings. Hence besides search costs and hours constraints, another reason why couples may not sort is due to information frictions.

with them (Figure 1.5b).<sup>19</sup>

I next estimate the size of these discontinuities. Table 1.7 shows the results of the estimated treatment effect, split in two panels: panel A for the GCSE choice outcome, and panel B for the arguing with parents outcome. To assess the robustness of the estimates, results are shown for a range of specifications: with and without controls, with a first and second order polynomial of the assignment variable, and using an unrestricted as well as optimal bandwidth.<sup>20</sup>

Panel A shows that the treatment effect of living in a nontraditional family always has the expected sign, is large and highly statistically significant for the choice of GCSE choice outcome, regardless of the specification. The estimates are robust to the inclusion of controls and the choice of polynomial order of the assignment variable. Further, estimates from the optimal bandwidth specifications are always larger than the unrestricted versions. This is expected, since zooming in around the cutoff eliminates the influence of observations further away. As can be seen in Figure 1.5, observations further away are on average closer to the mean probability for the full sample, compared to the size of the jump at the cutoff.

Similar patterns are observed in panel B. Compared to the choice of GCSE outcome, the treatment effect on arguing with parents is smaller in size, and in some specifications not as highly statistically significant. With no controls, a 2nd order polynomial and optimal bandwidth (column (2)), the effect is not significant at the 10% level, though it still is at the 15% level (p-value 0.13).

Overall, the results support my hypothesis of a causal effect of living in a nontraditional family. The nontraditional family treatment causes girls to have a weaker preference for conformity to the family. This, in turn, makes them adopt more traditional views regarding gender roles.

---

<sup>19</sup>An unsurprising feature of the graphical evidence worth noting is that, as we move towards the top end of mother's income share levels, the bin means become very noisy. As depicted previously in Figure 1.3, the number of observations per bin falls as the share increases; families where mothers earn more than 80% of family income are extremely rare. This results in single observations having a disproportionate influence on bin means at these extreme income shares.

<sup>20</sup>The optimal bandwidth is chosen as the one that minimises the mean square error of the estimator.



### 1.7.6 Placebo Tests

I next conduct some further robustness checks. If my identification strategy is valid, each of the outcome variables should not exhibit discontinuous jumps at income shares where no jumps should exist. I test this by taking in turn each multiple of 0.05 in the interval  $s_i \in [0.1, 0.6]$  and define treatment as having an income share above that hypothetical value. I then repeat the analysis by estimating the discontinuity at each of these placebo cutoffs. Figure 1.6 plots the treatment effect estimates and the associated 95% confidence intervals for each outcome and placebo cutoff. The true cutoff at 0.5 is marked with a dashed grey line. Both Figures 1.6a and 1.6b support the validity of my identification strategy. There is no statistically significant discontinuity at any of these 11 placebo cutoffs, except for the true cutoff at 0.5, for each conformity measure.

## 1.8 Conclusion

In this paper, I examine whether gender norms are passed on from parents to children when the norms of the family oppose those of the society in which the family is embedded. While boys raised in nontraditional families (i.e. in which the mother is the breadwinner) are less likely to develop traditional norms, girls raised in nontraditional families are actually more likely to do so, in opposition to their family's norm but in line with society's. I argue, based on the predictions of a gender-identity model, that these results can be explained by a weaker preference among girls for conforming to the family, and show empirical evidence confirming this prediction. Drawing on research from social psychology and using a regression discontinuity design, I then show that this weaker preference is, in fact, driven by a causal effect of living in a nontraditional family.

My results question the literature on the relation between mothers' and daughters' labour supply or gender attitudes (Del Boca et al. 2000; Fernandez et al. 2004; Morrill and Morrill 2013; Olivetti et al. 2016). Previous studies' finding of a positive relation is interpreted as transmission of nontraditional gender norms. I replicate their finding when only considering mothers' labour supply, but show that results are the opposite when I consider whether the mother works or earns more than the father does. Hence, I show that

a positive relation between mothers' and daughters' labour supply is not sufficient for the transmission of nontraditional norms. What matters is not just whether mothers work, but whether they work and earn *more than their husbands*.

My results also highlight a caveat related to the Bertrand et al. (2015) finding of sorting at the 0.5 threshold of wives' income shares in the US. I show that at least in the UK, this is purely driven by families with earnings from self-employment. No such sorting is found for families with only wage income, consistent with significant constraints in adjusting earnings levels or hours worked highlighted in recent public finance literature.

My findings reveal that horizontal socialisation is very important for the development of girls' gender norms. In fact, it is so strong that it leads to 'reactionary' behaviour by girls when their families violate the traditional social norm. If we want to reduce gender inequalities by promoting more egalitarian gender norms, we must therefore focus on changing the social norm. Moreover, a critical assessment of the current UK family policy seems pertinent. At the time of writing, the UK's Statutory Paternity Pay entitles fathers to just two weeks of paid paternity leave, payable at the minimum of 90% of previous weekly earnings and £139.58. While this can be extended up to 26 weeks, it is conditional on the mother returning to work before the end of her maternity leave period, payable again at just £139.58 per week. As a result, less than 1% of fathers take leave beyond two weeks (TUC 2015). The current policy does not provide adequate incentives to fathers, especially in cases where their foregone earnings are higher than their wives'. If the social norm defining child care as an exclusively maternal responsibility is to change, a more generous paternity leave policy may be required.

# References

- [1] Akerlof, G.A. and Kranton, R.E. (2000) ‘Economics and Identity’, *Quarterly Journal of Economics*, vol. 115(3), pp. 715-753.
- [2] Akerlof, G.A. and Kranton, R.E. (2002) ‘Identity and Schooling: Some Lessons for the Economics of Education’, *Journal of Economic Literature*, vol. 40(3), pp. 1167-1201.
- [3] Akerlof, G.A. and Kranton, R.E. (2010) *Identity Economics*, New Jersey: Princeton University Press.
- [4] Alesina, A. and Giuliano, P. (2010) ‘The Power of the Family’, *Journal of Economic Growth*, vol. 15, pp. 93-125.
- [5] Alesina, A., Giuliano P. and Nunn, N. (2013) ‘On the Origins of Gender Roles: Women and the Plough’, *Quarterly Journal of Economics*, vol. 128(2), pp. 469-530.
- [6] Asch, S.E. (1951) ‘Effects of Group Pressure on Modification and Distortion of Judgments’, in (Guetzkow, H., eds.), *Groups, Leadership and Men*, pp. 177-190, Oxford: Carnegie Press.
- [7] Asch, S.E. (1952) *Social Psychology*, New Jersey: Prentice Hall.
- [8] Asgari, S., Dasgupta, N. and Cote, N.G. (2010) ‘When Does Contact with Successful Ingroup Members Change Self-Stereotypes?’ *Social Psychology*, vol. 41(3), pp. 203-211.
- [9] Becker, G. (1964) *Human Capital*, New York: Columbia University Press.
- [10] Bem, S L. (1985) ‘Androgyny and Gender Schema Theory: A Conceptual and Empirical Integration’, *Psychology and Gender*, vol. 32, pp. 179-226.

- [11] Benabou, R. and Tirole, J. (2007) ‘Identity, Dignity and Taboos: Beliefs as Assets’, CEPR Discussion Paper 6123.
- [12] Bertrand, M. (2010) ‘New Perspectives on Gender’, *Handbook of Labor Economics*, vol. 4b, pp. 1545-1592.
- [13] Bertrand, M., Kamenica, E. and Pan, J. (2015) ‘Gender Identity and Relative Income Within Households’, *Quarterly Journal of Economics*, vol. 130(2), pp. 571-614.
- [14] Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2011) ‘Formation and Persistence of Oppositional Identities’, *European Economic Review*, vol. 55, pp. 1046-1071.
- [15] Blau, F.D. and Kahn, L.M. (2000) ‘Gender Differences in Pay’, *Journal of Economic Perspectives*, vol. 14(4), pp. 75-99.
- [16] Blunch, N.H. and Das, M.B. (2014) ‘Changing Norms About Gender Inequality in Education: Evidence from Bangladesh’, IZA Discussion Paper No. 8365.
- [17] Blundell, R., Brewer, M. and Francesconi, M. (2008) ‘Job Changes and Hours Changes: Understanding the Path of Labor Supply Adjustment’, *Journal of Labor Economics*, vol. 26(3), pp. 421-453.
- [18] Bonke, J. (2008) ‘Income Distribution and Financial Satisfaction Between Spouses in Europe’, *The Journal of Socio-Economics*, vol. 37, pp. 2291-2303.
- [19] Bonke, J. and Browning, M. (2009) ‘The Distribution of Well-Being and Income Within the Household’, *Review of the Economics of the Household*, vol. 7, pp. 31-42.
- [20] Brown, C. and Cororan, M. (1997) ‘Sex-Based Differences in School Content and the Male-Female Wage Gap’, *Journal of Labor Economics*, vol. 15(3), pp. 431-465.
- [21] Butikofer, A. (2013) ‘Revisiting ‘Mothers and Sons’ Preference Formation and the Female Labor Force in Switzerland’, *Labour Economics*, vol. 20, pp. 82-91.

- [22] Carrell, S.E., Page, M.E. and West, J.E. (2010) ‘Sex and Science: How Professor Gender Perpetuates the Gender Gap’, *Quarterly Journal of Economics*, vol. 125(3), pp. 1101-1144.
- [23] Chetty, R., Friedman, J., Olsen, T. and Pistaferri, L. (2011) ‘Adjustment Costs, Firm Responses and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records’, *Quarterly Journal of Economics*, vol. 126(2), pp. 749-804.
- [24] Cooksey, E.C. and Fondell, M.M. (1996) ‘Spending Time With his Kids: Effects of Family Structure on Fathers’ and Children’s Lives’, *Journal of Marriage and the Family*, vol. 58, pp. 693-707.
- [25] Cooray, A. and Potrafke, N. (2010) ‘Gender Inequality in Education: Political Institutions or Culture and Religion?’ Konstanz Working Paper 2010-01.
- [26] Costa-Font, J. and Cowell, F.A. (2015) ‘Social Identity and Redistributive Preferences’, *Journal of Economic Surveys*, vol. 29, pp. 357-374.
- [27] Deary, I.J., Batty, G.D. and Gale, C.R. (2008) ‘Bright Children Become Enlightened Adults’, *Perspectives on Psychological Science*, vol. 19(1), pp. 1-6.
- [28] Dee, T.S. (2004) ‘Are There Civic Returns to Education?’, *Journal of Public Economics*, vol. 88, pp. 1697- 1720.
- [29] Del Boca, D., Locatelli, M. and Pasqua, S. (2000) ‘Employment Decisions of Married Women: Evidence and Explanations’, *Labour*, vol. 14(1), pp. 35-52.
- [30] Department for Education (2014) ‘School Workforce in England: November 2013’, Statistical First Release 11/2004.
- [31] Department for Work and Pensions (2015) ‘Family Resources Survey – United Kingdom 2013/14’, National Statistics Release 06/2015.
- [32] Dickens, W.T. and Lundberg, S.J. (1993) ‘Hours Restrictions and Labor Supply’, *International Economic Review*, vol. 34(1), pp. 169-192.

- [33] Dollar, D. and Gatti, R. (1999) ‘Gender Inequality, Income, and Growth: Are Good Times Good For Women’, World Bank Working Paper Series No. 1.
- [34] Dugger, K. (1988) ‘Social Location and Gender-Role Attitudes: A Comparison of Black and White Women’, *Gender and Sociology*, vol. 2(4), pp. 425-448.
- [35] Eagly, A.H. and Carli, L.L. (1981) ‘Sex of Researchers and Sex-Typed Communications as Determinants of Sex Differences in Influenceability: A Meta-Analysis of Social Influence Studies’, *Psychological Bulletin*, vol. 90(1), pp. 1-20.
- [36] Eckstein, Z. and Lifshitz, O. (2011) ‘Dynamic Female Labor Supply’, *Econometrica*, vol. 79(6), pp. 1675-1726.
- [37] Eriksson, K.H and Stenberg, A. (2015) ‘Gender Identity and Relative Income Within Households: Evidence from Sweden’, IZA Discussion Paper No. 9533.
- [38] Epstein, M. and Ward, M.L. (2011) ‘Exploring Parent-Adolescent Communication About Gender. Results from Adolescent and Emerging Adult Samples’, *Sex Roles*, vol. 65, pp. 108-118.
- [39] FAO (2009) ‘Women and Rural Employment: Fighting Poverty by Redefining Gender Roles’, *Economic and Social Perspectives*, Policy Brief 5.
- [40] Farre, L. and Vella, F. (2013) ‘The Intergenerational Transmission of Gender Role Attitudes and its Implications for Female Labor Force Participation’, *Economica*, vol. 80, pp. 219-247.
- [41] Fernandez, R. and Fogli, A. (2009) ‘Culture: An Empirical Investigation of Beliefs, Work and Fertility’, *American Economic Journal: Macroeconomics*, vol. 1(1), pp. 146-177.
- [42] Fernandez, R., Fogli, A. and Olivetti, C. (2004) ‘Mothers and Sons: Preference Formation and Female Labor Force Dynamics’, *Quarterly Journal of Economics*, vol. 119(4), pp. 1249-1299.

- [43] Fortin, N.M. (2015) ‘Gender Role Attitudes and Women’s Labor Market Participation: Opting-Out, AIDS, and the Persistent Appeal of Housewifery’, *Annals of Economics and Statistics*, vol. 117-118, pp. 379-401.
- [44] Georgiadis, A. and Manning, A. (2013) ‘One Nation Under a Groove? Understanding National Identity’, *Journal of Economic Behavior and Organization*, vol. 93, pp. 166-185.
- [45] Goldin, C. (1995) ‘The U-Shaped Female Labor Force Function in Economic Development and Economic History’, NBER Working Paper No. 4707.
- [46] Goldin, C. and Katz, L.F. (2002) ‘The Power of the Pill: Oral Contraceptives and Women’s Career and Marriage Decisions’, *Journal of Political Economy*, vol. 110(4), pp. 730-770.
- [47] Gonzalez de San Roman, A. and de la Rica Goiricelaya, S. (2012) ‘Gender Gaps in PISA Test Scores: The Impact of Social Norms and the Mother’s Transmission of Role Attitudes’, IZA Discussion Paper No. 6338.
- [48] Gelber, A.M., Jones, D. and Sacks, D.W. (2016) ‘Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test’, Working Paper.
- [49] Guiso, L., Sapienza, P. and Zingales, L. (2003) ‘People’s Opium? Religion and Economic Attitudes’, *Journal of Monetary Economics*, vol. 50, pp. 225-282.
- [50] Heineck, G. (2004) ‘Does Religion Influence the Labor Supply of Married Women in Germany?’ *Journal of Socio-Economics*, vol. 33, pp. 307-328.
- [51] Hognas, R.S. and Carlson, M.J. (2012) ‘Like Parent, Like Child? The Intergenerational Transmission of Nonmarital Childbearing’, *Social Science Research*, vol. 41(6), pp. 1480-1494.
- [52] Hunt, J., Garant, J.P., Herman, H. and Munroe, D.J. (2012) ‘Why Don’t Women Patent?’, NBER Working Paper No. 17888.

- [53] IPPR (2015) ‘Who’s Breadwinning in Europe? A Comparative Analysis of Maternal Breadwinning in Great Britain and Germany’, IPPR Report.
- [54] Jacob, R., Zhu, P., Somers, M.A. and Bloom, H. (2012) ‘A Practical Guide to Regression Discontinuity’, MDRC.
- [55] Johnston, D.W., Schurer, S. and Schields, M.A. (2014) ‘Maternal Gender Role Attitudes, Human Capital Investment, and Labour Supply of Sons and Daughters’, *Oxford Economics Papers*, vol. 66, pp. 631-659.
- [56] Kan, M.Y. (2008) ‘Does Gender Trump Money? Housework Hours of Husbands and Wives in Britain’, *Work, Employment and Society*, vol. 22(1), pp. 45-66.
- [57] Kanazawa, S. (2010) ‘Why Liberals and Atheists are More Intelligent’, *Social Psychology Quarterly*, vol. 73, pp. 33-57.
- [58] Kane, E.W. (2000) ‘Racial and Ethnic Variations in Gender-Related Attitudes’, *Annual Review of Sociology*, vol. 26, pp. 419-439.
- [59] Keshet, J.K. (2001) ‘Uncounted Stepfamilies: Remarried Fathers and their Wives’, in (Cath, S.H. and Shopper, M., eds.), *Stepparenting: Creating and Recreating Families in America Today*. New Jersey: The Analytic Press.
- [60] Klasen, S. (1999) ‘Does Gender Inequality Reduce Growth and Development? Evidence from Cross-Country Regressions’, World Bank Working Paper Series No. 7.
- [61] Kleven, H. and Waseem, M. (2013) ‘Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan’, *Quarterly Journal of Economics*, vol. 128, pp. 669-723.
- [62] Kleven, H. (2016) ‘Bunching’, *Annual Review of Economics*, vol. 8, pp. 435-464.
- [63] Kleven, H. and Landais, C. (2017) ‘Gender Inequality and Economic Development: Fertility, Education and Norms’, *Economica*, vol. 84(334), pp. 180-209.
- [64] Kleven, H., Landais, C. and Sogaard, J.E. (2018) ‘Children and Gender Inequality: Evidence from Denmark’, NBER Working Paper 24219.



- [65] Klor, E.F. and Shayo, M. (2010) ‘Social Identity and Preferences Over Redistribution’, *Journal of Public Economics*, vol. 92, pp. 944-961.
- [66] Knowles, S., Lorgelly, P. and Owen, D. (2002) ‘Are Educational Gender Gaps a Brake on Economic Development? Some Cross-Country Empirical Evidence’, *Oxford Economic Papers*, vol. 54, pp. 118–149.
- [67] Lee, D.S. and Lemieux, T. (2010) ‘Regression Discontinuity Designs in Economics’, *Journal of Economic Literature*, vol. 48, pp. 281-355.
- [68] Lehrer, E.L. (1995) ‘The Effects of Religion on the Labor Supply of Married Women’, *Social Science Research*, vol. 24, pp. 281-301.
- [69] Lorber, J. (1994) *Paradoxes of Gender*, New Haven: Yale University Press.
- [70] Lytton, H. and Romney, D.M. (1991) ‘Parents’ Differential Socialization of Boys and Girls: A Meta-Analysis’, *Psychological Bulletin*, vol. 109(2), pp. 267-296.
- [71] Mammen, K. and Paxson, C. (2000) ‘Women’s Work and Economic Development’, *Journal of Economic Perspectives*, vol. 14(4), pp. 141-164.
- [72] Maneschild, P. and Haraldsson, B. (2007) ‘Religious Norms and Labour Supply of Married Women in Sweden’, *Finnish Economic Papers*, vol. 20, pp. 41-56.
- [73] Mapper (2017) ‘Maternity and Paternity Pay and Enhanced Rights’, <http://www.mapper.uk.com/index.php/employers/> (last accessed: 31 May 2017).
- [74] McCrary, J. (2008) ‘Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test’, *Journal of Econometrics*, vol. 142, pp. 698–714.
- [75] Mincer, J. and Polachek, S. (1974) ‘Family Investments in Human Capital: Earnings of Women’, *Journal of Political Economy*, vol. 82(2), pp. S76-S108.
- [76] Morrill, M.S. and Morrill, T. (2013) ‘Intergenerational Links in Female Labor Force Participation’, *Labour Economics*, vol. 20, pp. 38-47.

- [77] Neighbourhood Renewal Unit (2004) *The English Indices of Deprivation 2004*, London: The Office of the Deputy Prime Minister.
- [78] Noodle (2016) ‘Brits in the dark about their partner’s debts’, <http://www.blog.noodle.co.uk/brits-in-the-dark-about-their-partners-debts/> (last accessed: 31 May 2017).
- [79] OECD (2004) ‘Female Labour Force Participation: Past Trends and Main Determinants in OECD Countries’, OECD Working Papers Series No. 376.
- [80] OECD (2012) *Closing the Gender Gap: Act Now*, OECD Publishing.
- [81] Olivetti, C., Patacchini, E. and Zenou, Y. (2016) ‘Mothers, Peers and Gender Identity’, Working Paper, Boston College.
- [82] Parker, K. and Wang, W. (2013) ‘Modern Parenthood-Roles of Moms and Dads Converge as they Balance Work and Family’, PEW Research Center Report.
- [83] Polachek, S.W. (2008) ‘Earnings Over the Life Cycle: The Mincer Earnings Function and Its Applications’, *Foundations and Trends in Microeconomics*, vol. 4(3), pp. 165-272.
- [84] Ramey, C.T., Campbell, F.A. and Ramey, S.L. (1999) ‘Early Intervention: Successful Pathways to Improving Intellectual Development’, *Developmental Neuropsychology*, vol. 16, pp. 385-392.
- [85] Santee, R.T. and Jackson, S.E. (1982) ‘Identity Implications of Conformity: Sex Differences in Normative and Attributional Judgments’, *Social Psychology Quarterly*, vol. 45(2), pp. 121-125.
- [86] Stankov, L. (2009) ‘Conservatism and Cognitive Ability’, *Intelligence*, vol. 37, pp. 294-304.
- [87] Tajfel, H. (1978) *Differentiation Between Social Groups: Studies in the Psychology of Intergroup Relations*, Academic Press.

- [88] TUC (2015) ‘Statistical Briefing on Fathers’ Leave and Pay’, Trade Union Congress Report.
- [89] UNDP (2014) *Human Development Report 2014. Sustaining Human Progress: Reducing Vulnerabilities and Building Resilience*, New York: UNDP.
- [90] Vella, F. (1994) ‘Gender Roles and Human Capital Investment: The Relationship Between Traditional Attitudes and Female Labour Market Performance’, *Economica*, vol. 61, pp. 191-211.
- [91] Washbrook, E. (2007) ‘Explaining the Gender Division of Labour: The Role of the Gender Wage Gap’, CMPO Working Paper Series No. 07/174.
- [92] Weinberger, C.J. (1999) ‘Mathematical College Majors and the Gender Gap in Wages’, *Industrial Relations*, vol. 38(4), pp. 407-413.
- [93] World Bank (2006) *World Development Report 2006: Equity and Development*, Washington: The World Bank.

# Main Tables

Table 1.1: Main Results. Dependent Variable: Traditional Norm

Variables	(1)	(2)	(3)	(4)	(5)	(6)
Nontraditional family	-0.032* (0.014)	-0.024** (0.009)	-0.022** (0.008)	-0.022** (0.007)	-0.016** (0.006)	-0.091*** (0.020)
Female	-0.153*** (0.012)	-0.138*** (0.009)	-0.139*** (0.010)	-0.139*** (0.010)	-0.139*** (0.010)	-0.135*** (0.018)
Nontraditional family $\times$ Female	0.076** (0.029)	0.081*** (0.020)	0.082*** (0.023)	0.082*** (0.022)	0.083** (0.024)	0.181*** (0.046)
Father employment: stable			0.012 (0.023)	0.012 (0.023)	0.012 (0.023)	0.009 (0.023)
Mother employment: stable			0.006 (0.028)	0.006 (0.028)	0.006 (0.028)	-0.000 (0.019)
Mother more educated				-0.003 (0.009)		
Mother higher job status					-0.027 (0.018)	
Controls	No	Yes	Yes	Yes	Yes	Yes
$\beta_1 + \beta_3$	0.044	0.056	0.060	0.060	0.066	0.090
p-value( $\beta_1 + \beta_3 = 0$ )	0.11	0.0022	0.0076	0.011	0.020	0.061
$R^2$	0.024	0.11	0.11	0.11	0.12	0.13
Observations	1487	1485	1485	1485	1485	1125

*Notes:* The table shows the coefficient estimates for various versions of the main specification, where different controls are included. Besides the main controls listed in the table, baseline controls include family and child characteristics, parental socialisation effort, and region and area fixed effects (as defined in Section 1.5.3). Column (6) runs the same specification as shown in (3) but on the subsample of families where both parents work.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1.2: Other Outcomes

Variables	(1) Believe High Wage Important	(2) Want to Study Science
Nontraditional family	0.021 (0.013)	0.106** (0.034)
Female	-0.050* (0.023)	-0.062*** (0.008)
Nontraditional family $\times$ Female	-0.063*** (0.011)	-0.210*** (0.050)
Controls	Yes	Yes
$\beta_1 + \beta_3$	-0.042	-0.10
p-value( $\beta_1 + \beta_3 = 0$ )	0.092	0.024
$R^2$	0.13	0.22
Observations	1484	490

*Notes:* This table shows estimates for variants of the main specification where two alternative measures of gender norms are considered as dependent variables. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1.3: Robustness. Dependent Variable: Traditional Norm

Variables	(1) Not Disabled	(2) Disabled	(3) Above Median	(4) Below Median	(5) Not Disabled & Above Median
Nontraditional family	-0.047*** (0.011)	0.083 (0.115)	-0.084*** (0.013)	0.043 (0.033)	-0.072*** (0.012)
Female	-0.145*** (0.010)	-0.169 (0.111)	-0.158*** (0.025)	-0.129* (0.055)	-0.164*** (0.036)
Nontraditional family $\times$ Female	0.118*** (0.022)	-0.116 (0.130)	0.157*** (0.008)	0.033 (0.077)	0.142*** (0.010)
Controls	Yes	Yes	Yes	Yes	Yes
$\beta_1 + \beta_3$	0.070	-0.033	0.073	0.076	0.071
p-value( $\beta_1 + \beta_3 = 0$ )	0.0025	0.48	0.0008	0.17	0.00
$R^2$	0.12	0.63	0.18	0.19	0.18
Observations	1353	132	820	665	767

*Notes:* This table shows results from estimating the main specification on different sub-samples. Column (1) restricts the sample to those where the father is not disabled, (2) to where the father is disabled, (3) to those with family income above the median level, (4) to those with family income below the median, and (5) to families with a non-disabled father and family income above the median. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 1.4: Second Shift. Dependent Variable: Mother Does Second Shift

Variables	(1)
Nontraditional family	-0.021 (0.027)
Female	-0.007 (0.012)
Nontraditional family $\times$ Female	0.030 (0.018)
Controls	Yes
$\beta_1 + \beta_3$	0.0093
p-value( $\beta_1 + \beta_3 = 0$ )	0.83
$R^2$	0.099
Observations	1485

*Notes:* This table shows estimates of the probability that a mother does a second shift (i.e. performs domestic duties). Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1.5: Incorporating Hours Worked. Dependent variable: Traditional Norm

Variables	(1)	(2)	(3)	(4)
Mother works	0.028 (0.031)	-0.026 (0.042)	-0.021 (0.040)	0.125 (0.069)
Female	-0.037 (0.042)	-0.167 (0.089)	-0.162 (0.086)	-0.060 (0.114)
Mother works $\times$ Female	-0.117** (0.043)	-0.005 (0.087)	-0.017 (0.082)	-0.110 (0.109)
Mother works more than father		-0.081*** (0.023)	-0.076*** (0.021)	-0.046 (0.046)
Mother works more than father $\times$ Female		0.173** (0.066)	0.161** (0.063)	0.106 (0.089)
Nontraditional family			-0.023* (0.011)	-0.031*** (0.006)
Nontraditional family $\times$ Female			0.054** (0.020)	0.075*** (0.013)
Controls	No	No	No	Yes
$\beta_{Works} + \beta_{Works \times Female}$	-0.090	-0.031	-0.038	0.015
p-value( $\beta_{Works} + \beta_{Works \times Female} = 0$ )	0.002	0.550	0.460	0.750
$\beta_{WorksMore} + \beta_{WorksMore \times Female}$		0.092	0.085	0.060
p-value( $\beta_{WorksMore} + \beta_{WorksMore \times Female} = 0$ )		0.074	0.090	0.210
$\beta_{NontraditionalFamily} + \beta_{NontraditionalFamily \times Female}$			0.032	0.045
p-value( $\beta_{NontraditionalFamily} + \beta_{NontraditionalFamily \times Female} = 0$ )			0.230	0.005
$R^2$	0.025	0.029	0.029	0.029
Observations	1487	1487	1487	1487

*Notes:* This table shows estimates from specifications where the dependent variable is the probability that a child develops traditional gender norms, and the main regressors are various parental labour market measures. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.



Table 1.6: Conformity to Family

Variables	(1) Parents Chose GCSE Subjects	(2) Argue with Parents
Nontraditional family	0.084** (0.028)	0.028 (0.030)
Female	0.004 (0.021)	0.021 (0.013)
Nontraditional family $\times$ Female	-0.159*** (0.045)	0.038*** (0.007)
Controls	Yes	Yes
$\beta_1 + \beta_3$	-0.074	0.065
p-value( $\beta_1 + \beta_3 = 0$ )	0.023	0.076
$R^2$	0.095	0.078
Observations	1485	1434

*Notes:* This table shows estimates from specifications where the dependent variable is an indicator for a child's conformity to her family. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of 'Nontraditional family' and  $\beta_3$  of 'Nontraditional family  $\times$  Female'. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1.7: RD Results: Nontraditional Family Treatment Effect

	(1)	(2)	(3)	(4)
<b>Panel A: Parents Chose GCSE Subjects</b>				
Unrestricted BW	-0.349*** (0.075)	-0.398*** (0.103)	-0.349*** (0.085)	-0.467*** (0.119)
Optimal BW	-0.495*** (0.170)	-0.653*** (0.238)	-0.408** (0.162)	-0.588** (0.235)
1st-order Polynomial	Yes	No	Yes	No
2nd-order Polynomial	No	Yes	No	Yes
Controls	No	No	Yes	Yes
Observations	560	560	559	559
<b>Panel B: Argue with Parents</b>				
Unrestricted BW	0.157* (0.090)	0.275** (0.128)	0.183* (0.101)	0.236* (0.140)
Optimal BW	0.420*** (0.142)	0.311 (0.205)	0.295** (0.142)	0.374** (0.173)
Observations	552	552	551	551
Controls	No	No	Yes	Yes
1st-order Polynomial	Yes	No	Yes	No
2nd-order Polynomial	No	Yes	No	Yes

*Notes:* This table shows results from estimating various versions of the RD specification: with and without controls, with a first and second order polynomial of the assignment variable, and using an unrestricted as well as optimal bandwidth. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, region and area fixed effects. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

# Main Figures

Figure 1.1: Social Norms and Gender Inequality

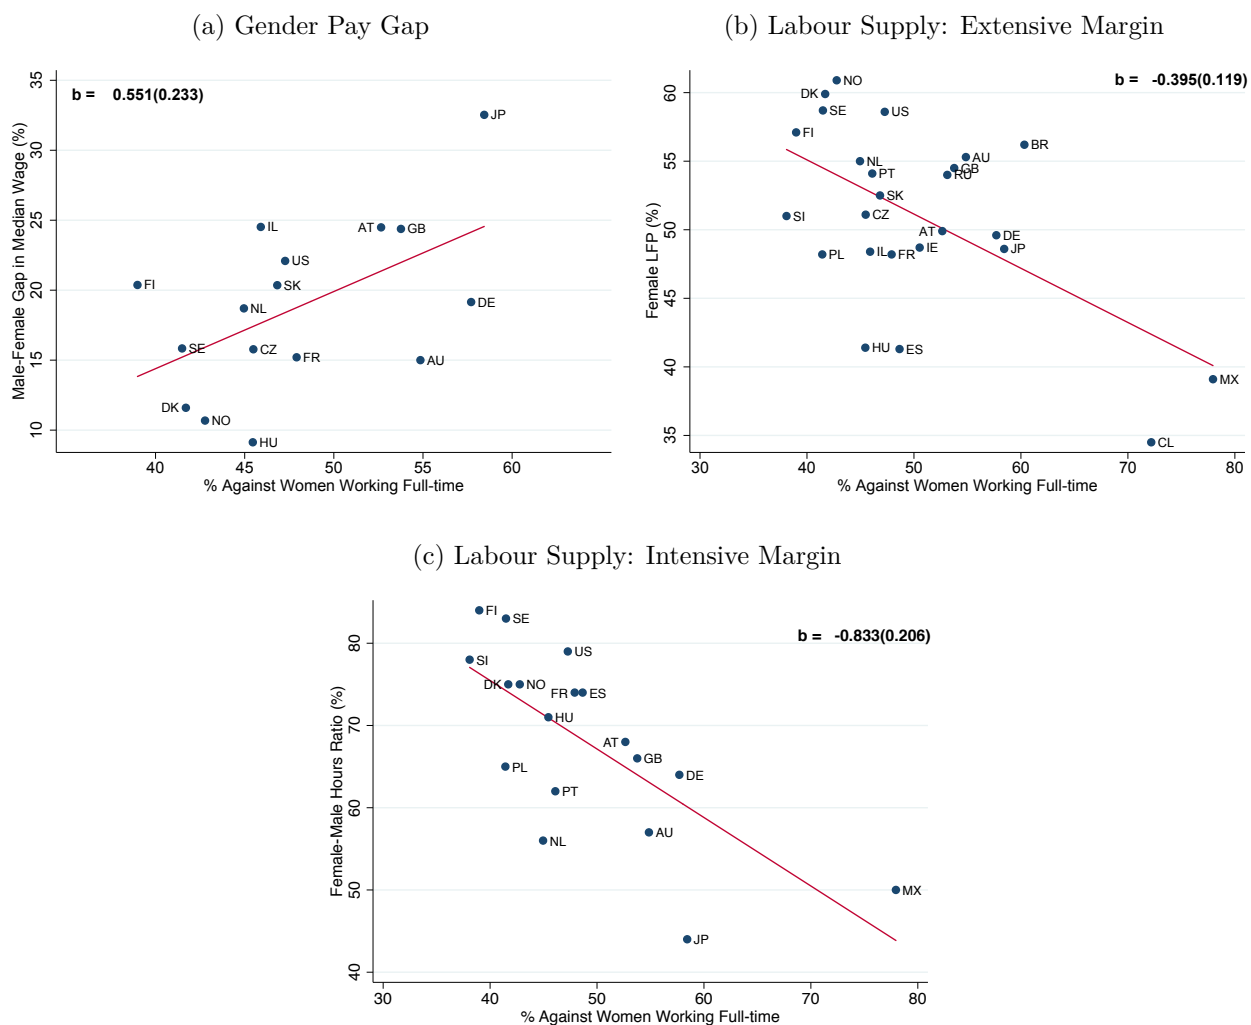
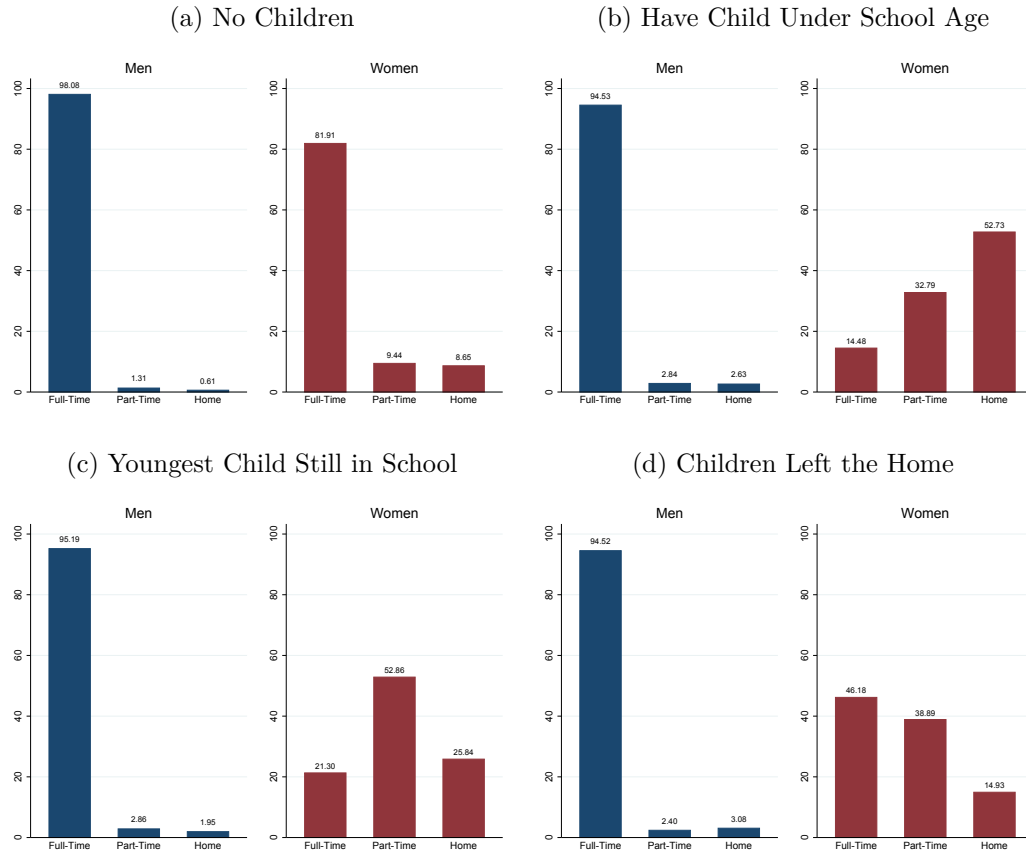
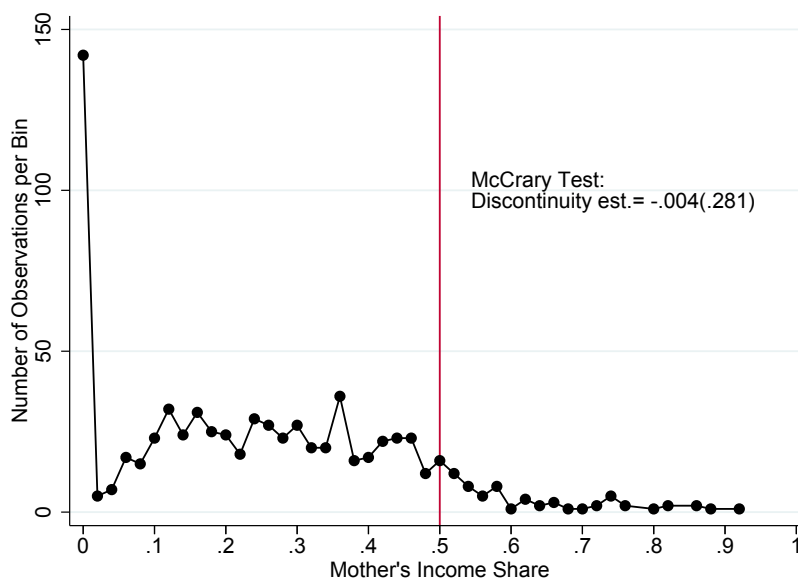


Figure 1.2: The Social Norm



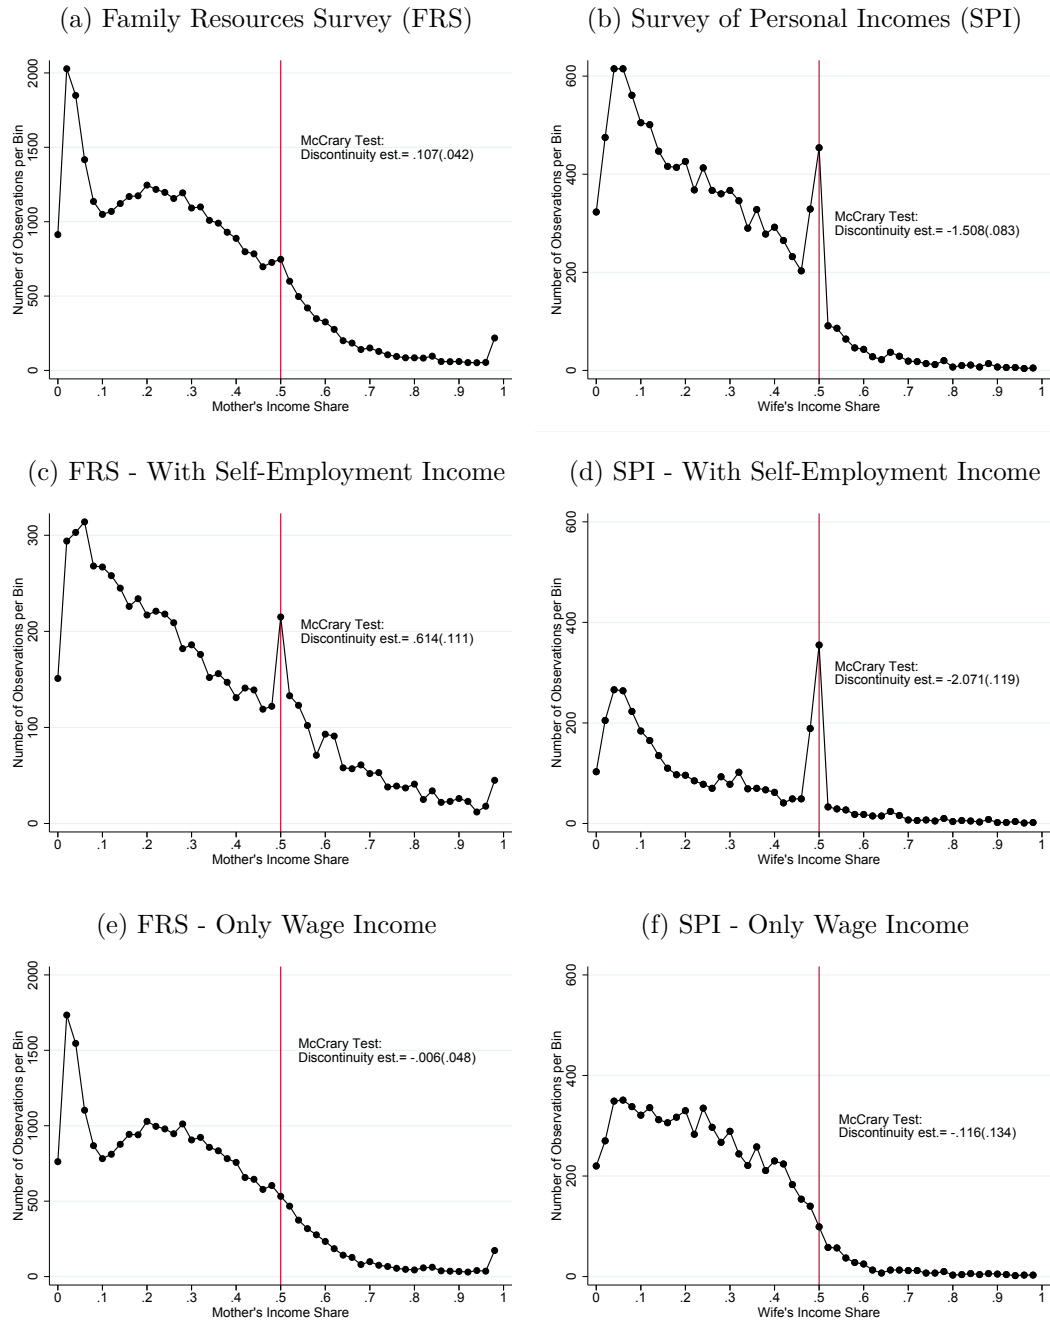
*Notes:* This figure shows the labour supply of men and women in heterosexual couples, using data from the 2002 International Social Survey Programme. It plots the proportion of men and women working full-time, part-time or not at all in each of the following cases: (a) when they have no children, (b) when they have a child under school age (i.e. below 4), (c) when their youngest child is above 4 but still within schooling age, and (d) when their children have grown up and left the home. Source: ISSP Research Group (2013): International Social Survey Programme: Family and Changing Gender Roles III - ISSP 2002. GESIS Data Archive, Cologne. ZA3880 Data file Version 1.1.0, doi:10.4232/1.11564.

Figure 1.3: The Distribution of the Mother's Income Share: Next Steps



*Notes:* This figure plots the distribution of the assignment variable (mother's income share) in the Next Steps data, using a bin size of 0.02. It also shows the results from a McCrary test, which tests the null of no discontinuity at the 0.5 threshold (vertical line) of the assignment variable. The null of no discontinuity cannot be rejected at the 10% level. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.


Figure 1.4: The Distribution of the Wife's Income Share: Further Datasets




*Notes:* These figures plot the distribution of the assignment variable across different samples. Figures (a), (c) and (e) are from the Family Resources Survey, and (b), (d) and (f) from the Survey of Personal Incomes. Figures (a) and (b) include all families, (c) and (d) only those with self-employment income, and (e) and (f) those with only wage income. The bin size is 0.02. Each figure reports the result of the McCrary test, which tests the null of no discontinuity at the 0.5 threshold (vertical line) of the assignment variable. Sources: Department of Work and Pensions, Office for National Statistics. Social and Vital Statistics Division, NatCen Social Research. (2016). *Family Resources Survey*. UK Data Service, and Inland Revenue. Statistics Division. (1989). *Survey of Personal Incomes, 1985-1986: Public Use Tape*. UK Data Service.

Figure 1.5: Regression Discontinuity Graphs: Conformity to Family

(a) Parents Chose GCSE Subjects




(b) Argue with Parents



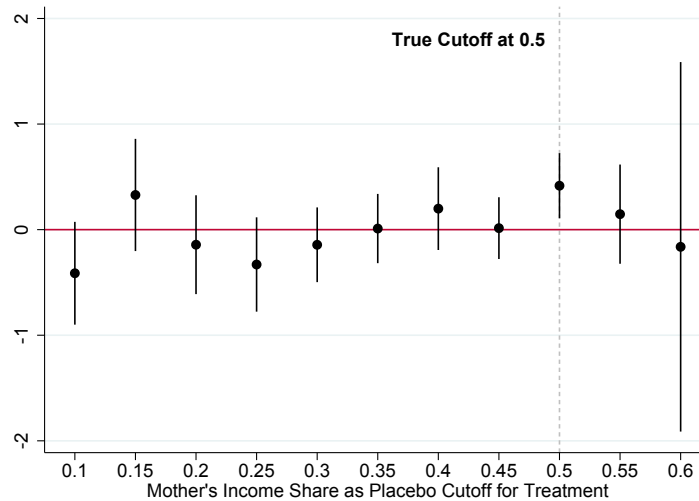
*Notes:* The figure shows, for females, the probability that they agree with each of the two conformity measures, above and below the 0.5 cutoff. The first is whether children chose what GCSE subjects to study based on what their parents wanted, and the second is whether children argue with their parents. Observations are divided into bins of width 0.05. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Figure 1.6: Placebo Tests

(a) Parents Chose GCE Subjects



(b) Argue with Parents



*Notes:* These figures show estimates of jumps in the outcome variable, and their associated 95% confidence intervals, for a wide range of placebo treatment cutoffs between 0.1 and 0.6 of the mother's income share. The dashed grey line marks the true cutoff at 0.5. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.



## Appendix 1A: Further Predictions of the Model

Section 1.3 presented a simple model of gender identity. It was shown that children choose their optimal belief about gender roles using the following rule:

$$x^* = \underset{x}{argmax} U = x(x_F, x_S, Z) = \frac{c_F(x_F, x_S, Z)x_F + c_S(x_F, x_S, Z)x_S}{c_F(x_F, x_S, Z) + c_S(x_F, x_S, Z)} \quad (1A.1)$$

The main text focused on  $\frac{\partial x^*}{\partial c_F}$  and showed that this prediction holds in the data. In this appendix I briefly discuss another comparative static and show that it also confirmed by the empirical results. Here I look at  $\frac{\partial x^*}{\partial x_F}$ , i.e. how living in a more traditional family changes the child's optimal belief. Recall that higher values of  $x^*$ ,  $x_F$  and  $x_S$  indicate more traditional values, and that in the empirical specification,  $x_S > x_F$ . Given this setup, the prediction is:<sup>21</sup>

$$\begin{aligned} \frac{\partial x^*}{\partial x_F} = & \left[ (c_F(x_F, x_S, Z) + c_S(x_F, x_S, Z))^2 \right] \times \\ & \left[ c_F^2(x_F, x_S, Z) + c_F(x_F, x_S, Z)c_S(x_F, x_S, Z) + \right. \\ & \left. (x_S - x_F) \left( c_F(x_F, x_S, Z) \frac{\partial c_S}{\partial x_F} - c_S(x_F, x_S, Z) \frac{\partial c_F}{\partial x_F} \right) \right] \quad (1A.2) \end{aligned}$$

The sign of (1A.2) depends on the signs of both  $\frac{\partial c_S}{\partial x_F}$  and  $\frac{\partial c_F}{\partial x_F}$ . While there is no proxy in the data to estimate  $\frac{\partial c_S}{\partial x_F}$ , the evidence from the social psychology literature (Asch 1951; 1952; etc.) implies  $\frac{\partial c_S}{\partial x_F} \geq 0$ . The reasoning is as follows. The literature has established that individuals have strong preferences for conforming to society. Furthermore, the social norm regarding gender roles in the UK is very traditional. Hence, the more traditional a family's gender norm is, the more the family norm aligns with the social norm that individuals have a strong preference to conform to. Hence, the strength of their preference for conforming to the social norm must be non-decreasing in  $x_F$ .

---

<sup>21</sup>Here I have for simplicity assumed that the family norm does not depend on the social norm. Introducing this dependence does not change the prediction.

The prediction therefore ultimately depends on the sign of  $\frac{\partial c_F}{\partial x_F}$ , which can be estimated from the data. In Section 1.6, I found that  $\frac{\partial x^*}{\partial x_F} > 0$  for boys, and  $\frac{\partial x^*}{\partial x_F} < 0$  for girls. To show that these empirical results are consistent with the model, first consider boys. A sufficient condition for  $\frac{\partial x^*}{\partial x_F} > 0$  is  $\frac{\partial c_F}{\partial x_F} < 0$ , i.e. that boys in nontraditional families have stronger preferences for conforming to their family than boys in traditional families. The results of Table 1.6 confirm this. The coefficient on the ‘Nontraditional family’ dummy is positive, confirming that boys in such families have stronger preferences for conformity.

Now consider the prediction for girls. A necessary (though not sufficient condition) for  $\frac{\partial x^*}{\partial x_F} < 0$  is  $\frac{\partial c_F}{\partial x_F} > 0$ , i.e. that girls in nontraditional families have weaker preferences for conforming to their family than girls in traditional families. This was the focus of all of Section 1.7, which confirmed this weaker preference.

## Appendix 1B: Discussion of Covariates

As discussed in Section 1.5.3, the main regressions include controls for family, child and geographical characteristics, and proxies for parental-socialisation effort. The aim is to control, as much as possible, for variables that are correlated with both the child expressing traditional views and living in a nontraditional family. In what follows I explain in more detail which variables I control for and why I consider it important to do so in the context of my research question.

The family characteristics of immediate concern are cultural and socio-economic attributes. A key variable controlled for is parental religion. As religions offer various prescriptions on the appropriate behaviour of individuals in society, including gender roles, parents' religious views can affect not only how much mothers work and earn, but also the way parents socialise their children (Guiso et al. 2003; Lehrer 1995). Moreover, this effect may depend on how important religion is to the parents' way of life; religious parents are more likely than non-practising individuals to conform to the gender roles prescribed by their religion (Guiso et al. 2003; Heineck 2004). I account for this by controlling for the strength of parental religiosity. I also control for whether a child has ever attended a religious or single-sex school, as both reflect parents' related social preferences.

To account for family structure, I control for whether the child lives with her biological parents and whether the parents are married. Evidence (Parker and Wang 2013) shows that marital status affects parents' beliefs about the ideal relative labour supply, which may determine who the breadwinner is. Growing up in a 'non-conventional' family of unmarried parents may, however, also affect how children view social norms, thereby affecting their preference for conformity to them. For instance, living in a 'non-conventional' family can affect the extent to which children consider social norms (which the parents are not following) desirable. Indeed, a positive relation between living in 'non-conventional' families and adopting 'non-conventional' norms is found by Hognas and Carlson (2012), who show that children raised by unmarried parents have a higher likelihood of non-marital childbearing in adulthood.

To further account for family structure, I also control for household size. Larger households imply a higher financial burden and, therefore, a greater likelihood that both parents

are working. Nonetheless, larger households (and, hence, more siblings, chores, and so forth) to take care of may increase the need for the child's contribution to household tasks, thereby affecting the child's belief about her gender's role within the family.

Among parent-specific characteristics, parental age is a crucial factor that needs to be accounted for. As age increases, so does the probability that a parent is working. Moreover, due to the stylised fact of a concave earnings profile over the life cycle (Polachek 2008), earnings also increase with age for working adults. I therefore control for parental age because it directly affects the relative family labour supply and income.

I also control for parental education, social class and total family income. Education not only increases earnings through higher human capital (Becker 1964) but can also directly affect social attitudes; higher education is associated with more liberal social views (Dee 2004; Stankov 2009; Kanazawa 2010). The same holds for social class. Although higher class status has historically been associated with more conservative beliefs, social class in modern times (as measured by the Goldthorpe scale using occupation types) captures job prestige and is highly correlated with education. For instance, Guiso et al. (2003) find that higher income (which is associated with social class, as this is defined by occupation) leads to more liberal views on gender equality. Thus, I expect income to have similar effects as education.

To better capture between-family income heterogeneity, I include an additional variable: the Index of Multiple Deprivation (IMD). This is defined at the Super Output Area (SOA) level. SOAs divide England into 32,482 local geographical districts - each the size of a neighbourhood - and has been devised to improve the reporting of small-area statistics by the UK Office of National Statistics. The IMD accounts for deprivation in the following dimensions: income, employment, health and disability, education skills and training, barriers to housing and services, crime, and living environment.<sup>22</sup> The most deprived SOA is given an IMD rank of 1, and the least deprived a rank of 32,482.

Geographic characteristics may also confound my results and are therefore controlled for. The region and type of area (urban versus rural) may affect the employment opportunities and thus relative earnings of mothers. In addition, areas may vary in how traditional

---

<sup>22</sup>For a more detailed explanation of this index, see Neighbourhood Renewal Unit (2004).

or modern the prevailing local norms are. For instance, larger and more urban areas are usually associated with more secular and liberal views (Maneschiold and Haraldsson 2007).

Besides family and geographic characteristics, I also account for between-child heterogeneity. First, I control for ethnicity. Research shows that ethnicity has an important effect on attitudes towards gender roles (Kane 2000). Compared to non-whites, whites are more critical of maternal employment and consider it more harmful for young children (Dugger 1988). Beyond ethnicity, I also control for whether the child was born in the UK, the child's religion and the importance of religion to her life.

Furthermore, child characteristics related to cognitive abilities may also have a confounding effect. Research relates higher intelligence to more socially liberal views (Deary et al. 2008); evidence also shows that low birth weight is associated with development delays and thereby lower intellectual capacity (Ramey et al. 1999). Since the data do not contain objective measures of innate ability, the best way to control for this possible confounder is to indirectly account for it by controlling for birth weight. Further proxies are also available in the data, such as whether the child has special education needs and disabilities. All of these variables are included to control for cognitive abilities as accurately as possible.

Last, as I am interested in examining how successfully parents socialise their children, I account for an important determinant of success: parental effort in socialisation. Since no perfect measure exists, I use proxies capturing the quality and quantity of contact time between parents and children, which has been emphasised in the literature as an important factor in gender stereotyping (Bem 1985; Cooksey and Fondell 1996; Asgari et al. 2010; Carrell et al. 2010). For more details of these proxies, see Table 1C.3 in Appendix 1C.

## Appendix 1C: Extra Tables

Table 1C.1: Summary Statistics: Original Vs Final Sample (*Next Steps*)

Variables	Initial Sample		Final Sample	
	Mean	s.d.	Mean	s.d.
Nontraditional family	0.143	0.35	0.145	0.352
Mother's income share	0.279	0.254	0.287	0.251
Traditional norm	0.349	0.477	0.311	0.463
High wage important	0.656	0.475	0.641	0.48
Want to study science	0.399	0.49	0.409	0.492
Parents chose GCSE subjects	0.359	0.48	0.374	0.484
Argue with parents	0.578	0.494	0.605	0.489
Female	0.516	0.500	0.524	0.500
Observations	8682		1487	

*Notes:* This table shows summary statistics for the main variables of interest in the Next Steps Survey, comparing those from the original sample to those used in the final sample after imposing sample restrictions. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.2: Summary Statistics of FRS and SPI data

Variables	Family Resources Survey		Survey of Personal Incomes	
	Mean	s.d.	Mean	s.d.
Nontraditional Family	0.161	0.367	0.028	0.165
Wife's Share	0.294	0.215	0.078	0.166
Have Self-Employment Income	0.223	0.473	0.284	0.451
Husband's Earnings	27688.063	36409.407	12794.532	12446.757
Wife's Earnings	10503.258	12134.177	1512.643	3814.021
Years	1997-2004		1985-1986	
Observations	33309		39021	

*Notes:* This table shows summary statistics for the Family Resources Survey and the Survey of Personal Incomes datasets. Sources: Department of Work and Pensions, Office for National Statistics. Social and Vital Statistics Division, NatCen Social Research. (2016). *Family Resources Survey*. UK Data Service, and Inland Revenue. Statistics Division. (1989). *Survey of Personal Incomes, 1985-1986: Public Use Tape*. UK Data Service.

Table 1C.3: Summary Statistics for Next Steps Data

Variable	Description	Mean	s.d.
Nontraditional family	1 if mother earns more than father	0.145	0.352
Mother's income share	Mother's income divided by family income	0.287	0.251
Traditional norm	1 if child agrees that women should never work full-time when they have young children	0.311	0.463
High wage important	1 if child believes that having a job that pays well is important	0.641	0.480
Want to study science	1 if child wants to study science at university	0.409	0.492
Parents chose GCSE subjects	1 if child studied what her parents wanted at GCSE	0.374	0.484
Argue with parents	1 if child argues with parents	0.605	0.489
Parents not married	1 if child's parents are not married	0.059	0.236
Both parents biological	1 if both are child's biological parents	0.894	0.307
Household size	No. of people in child's household	4.593	1.119
Father age	Father's age in years	42.843	5.252
Mother age	Mother's age in years	44.379	6.032
Religion important to parents	1 if parents state religion is important to their way of life	0.582	0.493
Father religion: muslim	1 if father is Muslim	0.076	0.265
Father religion: minority	1 if father is Sikh, Hindu, Buddhist, Jewish, Atheist	0.075	0.263
Mother religion: muslim	1 if mother is Muslim	0.078	0.268
Mother religion: minority	1 if mother is Sikh, Hindu, Buddhist, Jewish, Atheist	0.179	0.383
Family education: below degree	1 if highest education between parents is higher education below degree level	0.202	0.401
Family education: GCE A Level	1 if highest education between parents is GCE A Level or equivalent (post-compulsory)	0.210	0.407
Family education: GCSE	1 if highest education between parents is GCSE level grades A-C, or equivalent	0.238	0.426
Family education: level 1	1 if highest education between parents is at vocation level/GCSE grades D- or equivalent	0.036	0.187
Family education: other	1 if highest education between parents is informal	0.009	0.097
Family education: none	1 if neither parent has any education qualifications	0.059	0.236
Father SC: lower manag.	1 if father's occupation is Lower Managerial or Professional	0.297	0.457
Father SC: intermediate	1 if father's occupation is Intermediate	0.177	0.382
Father SC: small employer	1 if father's occupation is Small Employer or Own Accounts Worker	0.038	0.192
Father SC: lower supervisor	1 if father's occupation is Lower Supervisory or Technical	0.085	0.279
Father SC: semi-routine	1 if father's occupation is Semi-Routine	0.206	0.405



Continuation of Table 1C.3

Variable	Description	Mean	s.d.
Father SC: routine	1 if father's occupation is Routine	0.087	0.283
Father SC: long-term u/e	1 if father is long-term unemployed/never worked	0.031	0.173
Mother SC: lower manag.	1 if mother's occupation is Lower Managerial or Professional	0.258	0.438
Mother SC: intermediate	1 if mother's occupation is Intermediate	0.077	0.266
Mother SC: small employer	1 if mother's occupation is Small Employer or Own Accounts Worker	0.079	0.270
Mother SC: lower supervisor	1 if mother's occupation is Lower Supervisory or Technical	0.144	0.351
Mother SC: semi-routine	1 if mother's occupation is Semi-Routine	0.102	0.302
Mother SC: routine	1 if mother's occupation is Routine	0.118	0.322
Mother SC: long-term u/e	1 if mother is long-term unemployed/never worked	0.048	0.213
Father employment: stable	1 if father did not change employment 5 years before wave 1 up to wave 4	0.751	0.433
Mother employment: stable	1 if mother did not change employment 5 years before wave 1 up to wave 4	0.480	0.500
Mother more educated	1 if mother has higher level of education than father	0.293	0.455
Mother higher job status	1 if mother has higher job status level than father (defined by occupation type)	0.497	0.500
Family income	Total annual family income (sum of father's and mother's income)	39.964	31.855
Rank of IMD	Rank of Index of Multiple Deprivation	17577.731	9427.898
Female	1 if child is female	0.524	0.500
Child born in UK	1 if child was born in the UK	0.962	0.190
Child ethnicity: mixed	1 if child's ethnicity is mixed	0.024	0.152
Child ethnicity: indian	1 if child's ethnicity is Indian	0.074	0.262
Child ethnicity: pakistani	1 if child's ethnicity is Pakistani	0.043	0.203
Child ethnicity: bangladeshi	1 if child's ethnicity is Bangladeshi	0.015	0.121
Child ethnicity: black caribbean	1 if child's ethnicity is Black Caribbean	0.011	0.103
Child ethnicity: black african	1 if child's ethnicity is Black African	0.014	0.118
Child ethnicity: other	1 if child's ethnicity is none of above	0.022	0.147
Religion important to child	1 if child states that religion is important to her way of life	0.165	0.372
Child religion: muslim	1 if child is Muslim	0.077	0.267
Child religion: minority	1 if child is Sikh, Hindu, Buddhist, Jewish, Atheist	0.124	0.330
Child birth weight	Child birth weight in kg	3.317	0.594

Continuation of Table 1C.3

Variable	Description	Mean	s.d.
Child has SEN statement	1 if child has Special Education Needs	0.139	0.346
Child has disability	1 if child has a disability	0.028	0.166
Ever attended single sex school	1 if child ever attended a single sex school	0.137	0.343
Ever attended religious school	1 if child ever attended a religious school	0.160	0.367
Family evenings: often	1 if parents spend evenings with child at least 3 times per month	0.941	0.236
Family evenings: sometimes	1 if parents spend evenings with child once per month or less	0.028	0.166
Curfew: always	1 if parents always set curfew for child	0.968	0.175
Family activities: often	1 if go out together as family at least 3 times per month	0.905	0.294
Family activities: sometimes	1 if go out together as family once per month or less	0.079	0.270
Talk about school: often	1 if parents frequently talk to child about school day	0.485	0.500
North East	1 if child grew up in the North East	0.062	0.241
North West	1 if child grew up in the North West	0.150	0.357
Yorkshire and the Humber	1 if child grew up in Yorkshire and the Humber	0.102	0.303
East Midlands	1 if child grew up in East Midlands	0.093	0.290
West Midlands	1 if child grew up in West Midlands	0.135	0.342
East of England	1 if child grew up in the East of England	0.104	0.305
London	1 if child grew up in London	0.106	0.307
South East	1 if child grew up in the South East	0.159	0.366
South West	1 if child grew up in the South West	0.090	0.286
Urban	1 if child grew up in an urban area	0.787	0.410
Town	1 if child grew up in a town	0.089	0.285
Village	1 if child grew up in a village	0.081	0.272
Hamlet	1 if child grew up in a hamlet (smaller than a village)	0.026	0.158

*Notes:* This table shows summary statistics for the main variables of interest in the Next Steps Survey. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.4: Full Estimates of Main Results. Dependent variable: Traditional Norm

Variables	(1)	(2)	(3)	(4)	(5)	(6)
Nontraditional family	-0.032*	-0.024**	-0.022**	-0.022**	-0.016**	-0.091***
	(0.014)	(0.009)	(0.008)	(0.007)	(0.006)	(0.020)
Female	-0.153***	-0.138***	-0.139***	-0.139***	-0.139***	-0.135***
	(0.012)	(0.009)	(0.010)	(0.010)	(0.010)	(0.018)
Nontraditional family $\times$ Female	0.076**	0.081***	0.082***	0.082***	0.083**	0.181***
	(0.029)	(0.020)	(0.023)	(0.022)	(0.024)	(0.046)
Father employment: stable			0.012	0.012	0.012	0.009
			(0.023)	(0.023)	(0.023)	(0.023)
Mother employment: stable			0.006	0.006	0.006	-0.000
			(0.028)	(0.028)	(0.028)	(0.019)
Mother more educated				-0.003		
				(0.009)		
Mother higher job status					-0.027	
					(0.018)	
Both parents biological		0.035	0.034	0.034	0.032	0.039
		(0.038)	(0.039)	(0.039)	(0.040)	(0.049)
Parents not married		0.029	0.029	0.029	0.030	0.079*
		(0.027)	(0.027)	(0.027)	(0.027)	(0.034)
Household size		0.019**	0.020**	0.020**	0.020**	0.010
		(0.007)	(0.008)	(0.008)	(0.008)	(0.013)
Father age		0.002***	0.002***	0.002**	0.002***	0.005***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Mother age		-0.002	-0.002	-0.002	-0.002	-0.001
		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Religion important to parents		0.053***	0.052***	0.053***	0.052***	0.054***
		(0.008)	(0.008)	(0.008)	(0.008)	(0.014)
Father religion: muslim		-0.178***	-0.177***	-0.178***	-0.175***	-0.416**
		(0.031)	(0.033)	(0.032)	(0.035)	(0.126)
Father religion: minority		0.339***	0.338***	0.338***	0.337***	0.429***
		(0.050)	(0.049)	(0.049)	(0.049)	(0.025)
Mother religion: muslim		0.072	0.077	0.077	0.072	0.497***

Continuation of Table 1C.4

Variables	(1)	(2)	(3)	(4)	(5)	(6)
		(0.040)	(0.047)	(0.047)	(0.045)	(0.063)
Mother religion: minority	-0.151***	-0.152***	-0.152***	-0.151***	-0.132***	
	(0.004)	(0.002)	(0.002)	(0.003)	(0.014)	
Family education: below degree	-0.011	-0.011	-0.010	-0.012	-0.031	
	(0.016)	(0.016)	(0.016)	(0.016)	(0.029)	
Family education: GCE A Level	0.005	0.005	0.005	0.004	0.015	
	(0.018)	(0.018)	(0.018)	(0.018)	(0.022)	
Family education: GCSE	0.022	0.022	0.022	0.021	0.041*	
	(0.027)	(0.027)	(0.027)	(0.027)	(0.021)	
Family education: level 1	0.051	0.050	0.050	0.049	0.075	
	(0.046)	(0.046)	(0.046)	(0.047)	(0.058)	
Family education: other	0.171***	0.170***	0.170***	0.165***	0.437***	
	(0.020)	(0.017)	(0.018)	(0.016)	(0.071)	
Family education: none	-0.024	-0.023	-0.024	-0.024	0.012	
	(0.026)	(0.025)	(0.026)	(0.025)	(0.023)	
Father SC: lower manag.	0.014	0.014	0.015	0.014	-0.047**	
	(0.022)	(0.022)	(0.022)	(0.024)	(0.018)	
Father SC: intermediate	0.033*	0.034*	0.034*	0.028	-0.038*	
	(0.017)	(0.018)	(0.018)	(0.018)	(0.020)	
Father SC: small employer	0.103***	0.104***	0.104***	0.098***	0.053	
	(0.027)	(0.027)	(0.027)	(0.026)	(0.040)	
Father SC: lower supervisor	0.051	0.052	0.052	0.046	-0.035	
	(0.029)	(0.031)	(0.031)	(0.031)	(0.045)	
Father SC: semi-routine	0.127***	0.128***	0.128***	0.115***	0.093***	
	(0.013)	(0.011)	(0.011)	(0.016)	(0.018)	
Father SC: routine	0.149***	0.150***	0.149***	0.136***	0.121*	
	(0.024)	(0.025)	(0.025)	(0.025)	(0.051)	
Father SC: long-term u/e	0.118***	0.118***	0.118***	0.100***		
	(0.017)	(0.016)	(0.016)	(0.013)		

Continuation of Table 1C.4

Variables	(1)	(2)	(3)	(4)	(5)	(6)
Mother SC: lower manag.		0.020 (0.011)	0.020 (0.011)	0.021 (0.011)	0.028* (0.014)	0.034** (0.010)
Mother SC: intermediate		0.050** (0.014)	0.050** (0.014)	0.051** (0.015)	0.060** (0.017)	0.109*** (0.020)
Mother SC: small employer		0.114*** (0.027)	0.116*** (0.026)	0.117*** (0.024)	0.127*** (0.021)	0.123*** (0.032)
Mother SC: lower supervisor		0.012 (0.017)	0.014 (0.017)	0.015 (0.017)	0.026 (0.021)	0.017 (0.012)
Mother SC: semi-routine		0.055 (0.031)	0.057* (0.029)	0.058* (0.027)	0.071** (0.023)	0.067** (0.025)
Mother SC: routine		0.059* (0.027)	0.061* (0.027)	0.062** (0.026)	0.080** (0.025)	0.112** (0.039)
Mother SC: long-term u/e		-0.003 (0.041)	-0.002 (0.040)	-0.002 (0.039)	0.001 (0.039)	
Family income		-0.000* (0.000)	-0.000* (0.000)	-0.000* (0.000)	-0.000* (0.000)	-0.000 (0.000)
Rank of IMD		-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Child born in UK		-0.064* (0.031)	-0.065** (0.026)	-0.065** (0.026)	-0.066** (0.025)	-0.097** (0.038)
Child ethnicity: mixed		-0.049*** (0.010)	-0.049*** (0.007)	-0.049*** (0.007)	-0.045*** (0.007)	-0.101*** (0.005)
Child ethnicity: indian		-0.217*** (0.024)	-0.216*** (0.025)	-0.216*** (0.024)	-0.216*** (0.025)	-0.297*** (0.044)
Child ethnicity: bangladeshi		-0.163*** (0.016)	-0.163*** (0.015)	-0.163*** (0.015)	-0.161*** (0.015)	-0.430*** (0.049)
Child ethnicity: black caribbean		-0.196*** (0.007)	-0.200*** (0.013)	-0.199*** (0.013)	-0.203*** (0.015)	-0.191*** (0.014)
Child ethnicity: black african		0.101***	0.104***	0.103***	0.101***	-0.037***

Continuation of Table 1C.4

Variables	(1)	(2)	(3)	(4)	(5)	(6)
		(0.021)	(0.028)	(0.026)	(0.026)	(0.010)
Child ethnicity: other		0.023	0.024	0.024	0.019	-0.126***
		(0.032)	(0.033)	(0.032)	(0.031)	(0.035)
Religion important to child		0.080***	0.080***	0.080***	0.080***	0.055
		(0.017)	(0.016)	(0.016)	(0.016)	(0.033)
Child religion: muslim		0.068	0.063	0.063	0.064	
		(0.050)	(0.057)	(0.058)	(0.058)	
Child religion: minority		-0.038	-0.039	-0.038	-0.039	-0.063***
		(0.049)	(0.051)	(0.050)	(0.050)	(0.013)
Child birth weight		-0.009	-0.009	-0.009	-0.009	-0.018
		(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Child has SEN statement		0.076***	0.077***	0.077***	0.078***	0.062**
		(0.017)	(0.019)	(0.019)	(0.020)	(0.020)
Child has disability		0.005	0.007	0.007	0.008	0.039
		(0.041)	(0.038)	(0.038)	(0.038)	(0.043)
Ever attended single sex school		-0.002	-0.001	-0.001	-0.001	-0.008
		(0.011)	(0.011)	(0.011)	(0.011)	(0.020)
Ever attended religious school		-0.011	-0.012	-0.012	-0.011	-0.007
		(0.013)	(0.014)	(0.014)	(0.013)	(0.012)
Family evenings: often		-0.014	-0.015	-0.015	-0.015	-0.005
		(0.057)	(0.058)	(0.058)	(0.059)	(0.071)
Family evenings: sometimes		-0.126	-0.128	-0.128	-0.128	-0.096
		(0.083)	(0.085)	(0.085)	(0.086)	(0.088)
Curfew: always		0.107	0.108	0.108	0.107	0.115
		(0.058)	(0.059)	(0.059)	(0.060)	(0.107)
Family activities: often		0.015	0.014	0.014	0.008	0.106
		(0.074)	(0.074)	(0.074)	(0.070)	(0.099)
Family activities: sometimes		0.046	0.044	0.044	0.040	0.109
		(0.083)	(0.081)	(0.081)	(0.079)	(0.099)

Continuation of Table 1C.4

Variables	(1)	(2)	(3)	(4)	(5)	(6)
Talk about school: often		-0.013 (0.011)	-0.013 (0.011)	-0.013 (0.011)	-0.014 (0.011)	-0.004 (0.018)
$\beta_1 + \beta_3$	0.044	0.056	0.060	0.060	0.066	0.090
p-value( $\beta_1 + \beta_3 = 0$ )	0.11	0.0022	0.0076	0.011	0.02	0.061
$R^2$	0.024	0.11	0.11	0.11	0.12	0.13
Observations	1487	1485	1485	1485	1485	1125

*Notes:* This table shows the coefficient estimates for various versions of the main specification, where different controls are included. Column (6) runs the same specification as shown in (3) but on the subsample of families where both parents work.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.5: Covariate Balance Tests

Variables	(1) Point Estimate	(2) s.e.
Both parents biological	-0.084	0.103
Parents not married	0.104	0.089
Household size	-0.81**	0.396
Father age	-2.118	1.48
Mother age	-1.62	2.215
Religion important to parents	0.152	0.238
Father religion: muslim	0.066	0.106
Father religion: minority	-0.048	0.105
Mother religion: muslim	0.066	0.106
Mother religion: minority	-0.017	0.15
Family education: below degree	0.260**	0.128
Family education: GCE A Level	-0.214	0.139
Family education: GCSE	-0.141	0.095
Family education: level 1	-0.013	0.014
Family education: other	-0.019	0.02
Family education: none	-0.132	0.118
Father SC: lower manag.	0.432***	0.133
Father SC: intermediate	0.111**	0.048
Father SC: lower supervisor	-0.151**	0.076
Father SC: semi-routine	-0.028	0.063
Father SC: routine	-0.162	0.124
Mother SC: lower manag.	0.211	0.196
Mother SC: intermediate	0.096	0.111
Mother SC: lower supervisor	-0.239*	0.127
Mother SC: semi-routine	-0.084	0.086
Mother SC: routine	-0.145	0.157
Family income	13.595	9.311
Rank of IMD	997.992	3534.826
Child born in UK	-0.263*	0.145
Child ethnicity: mixed	0.125	0.11



Continuation of Table 1C.5

Variables	(1) Point Estimate	(2) s.e.
Child ethnicity: indian	-0.055	0.093
Child ethnicity: bangladeshi	0.096	0.096
Child ethnicity: black caribbean	0.027	0.095
Child ethnicity: black african	0.001	0.011
Child ethnicity: other	-0.031	0.054
Religion important to child	0 .027	0.133
Child religion: muslim	0.066	0.106
Child religion: minority	-0.116	0.109
Child birth weight	-0.068	0.177
Child has SEN statement	0.033	0.123
Child has disability	-0.029	0.074
Ever attended single sex school	0.091	0.165
Ever attended religious school	0.127	0.101
Family evenings: often	0.189*	0.103
Family evenings: sometimes	-0.115	0.071
Curfew: always	0.004	0.013
Family activities: often	0.131	0.122
Family activities: sometimes	-0.225	0.147
Talk about school: often	0.314	0.199
North East	0.002	0.117
North West	-0.125	0.139
Yorkshire and the Humber	-0.154	0.151
East Midlands	0.156	0.105
West Midlands	0.244	0.158
East of England	-0.136	0.136
London	-0.024	0.133
South East	-0.132*	0.073
Urban	0.107	0.175
Town	-0.206	0.136

Continuation of Table 1C.5

Variables	(1)	(2)
	Point Estimate	s.e.
Village	0.032	0.116
Hamlet	0.001	0.002

*Notes:* This table shows results from balance tests on all covariates used in the main section. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.6: Main Results, Probit and Logit Specifications. Dependent Variable: Traditional Norm

Variables	(1)	(2)	(3)	(4)	(5)	(6)
<b>Panel A: Probit Specification</b>						
Nontraditional family	-0.029** (0.013)	-0.020** (0.008)	-0.017*** (0.006)	-0.017*** (0.005)	-0.011*** (0.003)	-0.079*** (0.021)
Female	-0.151*** (0.012)	-0.138*** (0.010)	-0.139*** (0.009)	-0.139*** (0.009)	-0.139*** (0.009)	-0.136*** (0.016)
Nontraditional family $\times$ Female	0.076*** (0.029)	0.072*** (0.020)	0.075*** (0.022)	0.075*** (0.021)	0.075*** (0.023)	0.175*** (0.045)
$\beta_1 + \beta_3$	0.047	0.053	0.058	0.058	0.065	0.096
p-value( $\beta_1 + \beta_3 = 0$ )	0.063	0.000	0.000	0.000	0.002	0.0120
<b>Panel B: Logit Specification</b>						
Nontraditional family	-0.028** (0.013)	-0.020** (0.009)	-0.017** (0.007)	-0.017*** (0.006)	-0.011*** (0.004)	-0.080*** (0.020)
Female	-0.151*** (0.012)	-0.137*** (0.010)	-0.138*** (0.010)	-0.138*** (0.010)	-0.138*** (0.010)	-0.135*** (0.017)
Nontraditional family $\times$ Female	0.077*** (0.029)	0.076*** (0.021)	0.078*** (0.023)	0.078*** (0.022)	0.079*** (0.024)	0.179*** (0.046)
$\beta_1 + \beta_3$	0.048	0.056	0.061	0.061	0.068	0.099
p-value( $\beta_1 + \beta_3 = 0$ )	0.061	0.000	0.000	0.000	0.001	0.010
Family characteristics	No	Yes	Yes	Yes	Yes	Yes
Child characteristics	No	Yes	Yes	Yes	Yes	Yes
Socialisation effort	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	Yes	Yes	Yes	Yes	Yes
Area Type FE	No	Yes	Yes	Yes	Yes	Yes
Parental employment	No	No	Yes	Yes	Yes	Yes
Mother higher education	No	No	No	Yes	No	No
Mother higher job status	No	No	No	No	Yes	No
Observations	1487	1463	1463	1463	1463	1100

*Notes:* This table shows average marginal effects from probit and logit estimates of the main specification.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.7: Other Outcomes, Probit and Logit Specifications. Dependent Variable: Traditional Norm

Variables	(1) Believe High Wage Important	(2) Want to Study Science
<b>Panel A: Probit Specification</b>		
Nontraditional family	0.021 (0.015)	0.113*** (0.019)
Female	-0.050** (0.020)	-0.052*** (0.009)
Nontraditional family $\times$ Female	-0.066*** (0.009)	-0.237*** (0.040)
$\beta_1 + \beta_3$	-0.044	-0.12
p-value( $\beta_1 + \beta_3 = 0$ )	0.042	0.002
<b>Panel B: Logit Specification</b>		
Nontraditional family	0.025* (0.014)	0.109*** (0.017)
Female	-0.051** (0.020)	-0.054*** (0.008)
Nontraditional family $\times$ Female	-0.068*** (0.009)	-0.231*** (0.041)
$\beta_1 + \beta_3$	-0.043	-0.12
p-value( $\beta_1 + \beta_3 = 0$ )	0.045	0.005
Controls	Yes	Yes
Observations	1466	471

*Notes:* This table shows average marginal effects from probit and logit estimates, for variants of the main specification where two alternative measures of gender norms are considered as dependent variables. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.8: Robustness Checks, Probit and Logit Specifications. Dependent Variable: Traditional Norm

Variables	(1) Not Disabled	(2) Disabled	(3) Above Median	(4) Below Median	(5) Not Disabled & Above Median
<b>Panel A: Probit Specification</b>					
Nontraditional family	-0.038*** (0.008)	-0.010 (0.067)	-0.065*** (0.010)	0.038 (0.023)	-0.056*** (0.011)
Female	-0.144*** (0.010)	-0.214** (0.096)	-0.161*** (0.028)	-0.133*** (0.050)	-0.165*** (0.036)
Nontraditional family $\times$ Female	0.106*** (0.022)	-0.251*** (0.057)	0.148*** (0.006)	0.032 (0.063)	0.132*** (0.011)
$\beta_1 + \beta_3$	0.067	-0.26	0.083	0.070	0.077
p-value( $\beta_1 + \beta_3 = 0$ )	0.000	0.000	0.000	0.130	0.000
<b>Panel B: Logit Specification</b>					
Nontraditional family	-0.039*** (0.007)	0.192*** (0.065)	-0.068*** (0.014)	0.036 (0.024)	-0.060*** (0.014)
Female	-0.142*** (0.011)	-0.050 (0.066)	-0.158*** (0.030)	-0.134*** (0.048)	-0.164*** (0.038)
Nontraditional family $\times$ Female	0.109*** (0.022)	-0.132 (0.113)	0.153*** (0.006)	0.045 (0.063)	0.137*** (0.011)
$\beta_1 + \beta_3$	0.070	0.060	0.084	0.081	0.077
p-value( $\beta_1 + \beta_3 = 0$ )	0.000	0.240	0.000	0.073	0.000
Controls	Yes	Yes	Yes	Yes	Yes
Observations	1334	124	789	642	741

*Notes:* This table shows average marginal effects from probit and logit estimates using different subsamples. Column (1) restricts the sample to those where the father is not disabled, (2) to where the father is disabled, (3) to those with family income above the median level, (4) to those with family income below the median, and (5) to families with a non-disabled father and family income above the median. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.9: Probit and Logit Specifications. Dependent Variable: Mother Does Second Shift

Variables	(1) Probit Specification	(2) Logit Specification
Nontraditional family	-0.032 (0.029)	-0.033 (0.032)
Female	-0.011 (0.015)	-0.011 (0.015)
Nontraditional family $\times$ Female	0.041** (0.017)	0.042*** (0.014)
$\beta_1 + \beta_3$	0.009	0.008
p-value( $\beta_1 + \beta_3 = 0$ )	0.830	0.850
Controls	Yes	Yes
Observations	1452	1452

*Notes:* This table shows average marginal effects from probit and logit estimates of the probability that a mother does a second shift (i.e. performs domestic duties). Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.10: Incorporating Mother Hours Worked, Probit and Logit Specifications. Dependent Variable: Traditional Norm

Variables	(1)	(2)	(3)	(4)
<b>Panel A: Probit Specification</b>				
Mother works	0.025 (0.028)	-0.024 (0.038)	-0.019 (0.037)	0.124** (0.059)
Female	-0.035 (0.039)	-0.162* (0.085)	-0.157* (0.082)	-0.065 (0.103)
Mother works $\times$ Female	-0.118*** (0.040)	-0.008 (0.082)	-0.021 (0.077)	-0.107 (0.101)
Mother works more		-0.074*** (0.022)	-0.070*** (0.020)	-0.047 (0.044)
Mother works more $\times$ Female		0.170*** (0.064)	0.158** (0.061)	0.114 (0.083)
Nontraditional family			-0.021** (0.010)	-0.028*** (0.005)
Nontraditional family $\times$ Female			0.054*** (0.021)	0.067*** (0.015)
$\beta_{Works} + \beta_{Works \times Female}$	-0.093	-0.033	-0.040	0.016
p-value( $\beta_{Works} + \beta_{Works \times Female} = 0$ )	0.000	0.497	0.393	0.794
$\beta_{WorksMore} + \beta_{WorksMore \times Female}$		0.095	0.087	0.068
p-value( $\beta_{WorksMore} + \beta_{WorksMore \times Female} = 0$ )		0.026	0.039	0.084
$\beta_{NontraditionalFamily} + \beta_{NontraditionalFamily \times Female}$			0.033	0.039
p-value( $\beta_{NontraditionalFamily} + \beta_{NontraditionalFamily \times Female} = 0$ )			0.198	0.001
<b>Panel B: Logit Specification</b>				
Mother works	0.025 (0.028)	-0.024 (0.038)	-0.019 (0.036)	0.121* (0.062)
Female	-0.034 (0.038)	-0.163* (0.086)	-0.157* (0.083)	-0.060 (0.106)
Mother works $\times$ Female	-0.118*** (0.040)	-0.008 (0.083)	-0.021 (0.078)	-0.109 (0.104)
Mother works more		-0.073*** (0.021)	-0.069*** (0.019)	-0.040 (0.045)
Mother works more $\times$ Female		0.170*** (0.064)	0.157** (0.061)	0.108 (0.085)
Nontraditional family			-0.021** (0.010)	-0.028*** (0.004)
Nontraditional family $\times$ Female			0.055*** (0.021)	0.071*** (0.016)
$\beta_{Works} + \beta_{Works \times Female}$	-0.094	-0.031	-0.040	0.011
p-value( $\beta_{Works} + \beta_{Works \times Female} = 0$ )	0.000	0.530	0.410	0.810
$\beta_{WorksMore} + \beta_{WorksMore \times Female}$		0.097	0.089	0.068
p-value( $\beta_{WorksMore} + \beta_{WorksMore \times Female} = 0$ )		0.027	0.039	0.089
$\beta_{NontraditionalFamily} + \beta_{NontraditionalFamily \times Female}$			0.035	0.043
p-value( $\beta_{NontraditionalFamily} + \beta_{NontraditionalFamily \times Female} = 0$ )			0.170	0.002
Controls	No	No	No	Yes
Observations	1487	1487	1487	1464

*Notes:* This table shows average marginal effects from probit and logit estimates of specifications where the dependent variable is the probability that a child develops traditional gender norms, and the main regressors are various parental labour market measures. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of ‘Nontraditional family’ and  $\beta_3$  of ‘Nontraditional family  $\times$  Female’. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.

Table 1C.11: Conformity to Family, Probit and Logit Specifications.

Variables	(1) Parents Chose GCSE Subjects	(2) Argue with Parents
<b>Panel A: Probit Specification</b>		
Nontraditional family	0.079*** (0.026)	0.032 (0.028)
Female	0.005 (0.023)	0.021* (0.013)
Nontraditional family $\times$ Female	-0.156*** (0.051)	0.031*** (0.005)
$\beta_1 + \beta_3$	-0.077	0.063
p-value( $\beta_1 + \beta_3 = 0$ )	0.013	0.036
<b>Panel B: Logit Specification</b>		
Nontraditional family	0.079*** (0.026)	0.033 (0.029)
Female	0.004 (0.023)	0.020 (0.012)
Nontraditional family $\times$ Female	-0.158*** (0.050)	0.033*** (0.006)
$\beta_1 + \beta_3$	-0.079	0.066
p-value( $\beta_1 + \beta_3 = 0$ )	0.012	0.033
Controls	Yes	Yes
Observations	1449	1419

*Notes:* This table shows average marginal effects from probit and logit estimates of specifications where the dependent variable is an indicator for a child's conformity to her family. Controls include all variables for parental employment stability, family and child characteristics, parental socialisation effort, and region and area fixed effects.  $\beta_1$  refers to the coefficient estimate of 'Nontraditional family' and  $\beta_3$  of 'Nontraditional family  $\times$  Female'. Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Source: Department for Education, NatCen Social Research. (2013). *First Longitudinal Study of Young People in England: Waves One to Seven, 2004-2010: Secure Access*. [data collection]. 2nd Edition. UK Data Service. SN: 7104, <http://doi.org/10.5255/UKDA-SN-7104-2>.



# Chapter 2

## Bunching and Adjustment Costs: Evidence from Cypriot Tax Reforms

### 2.1 Introduction

The canonical model of labour supply assumes that workers can frictionlessly adjust their earnings. Recent studies challenge this view, providing evidence that workers face significant adjustment costs in responding to tax policy changes (Chetty et al. 2011; Gelber et al. 2017; Kleven and Waseem 2013). Such frictions reflect a variety of factors attenuating short-term adjustments to tax changes, and may include factors such as the cost of renegotiating contracts or of searching for new jobs, inattention to reforms, or a lack of knowledge of tax incentives. Incorporating such adjustment costs into the canonical model can explain several empirical puzzles in public finance, such as the large difference between estimates of micro and macro labour supply elasticities (Chetty 2012; Chetty et al. 2013), the difference between elasticities estimated using bunching techniques at kinks versus notches, or the lack of taxable income responses to jumps in marginal tax rates around income tax bracket thresholds (Chetty et al. 2011; Saez 2010). Moreover, the existence of frictions has implications for optimal tax policy and welfare evaluation (Farhi and Gabaix 2015; Werquin 2015). Yet, little is known about their size or source, with the literature providing direct estimates of such adjustment costs still in its infancy.

In this paper, I study adjustment costs using administrative tax data from the Republic

of Cyprus and quasi-experimental variation in the location of kinks in the personal income tax schedule, both cross-sectionally and over time. The discontinuities in marginal tax rates at bracket thresholds induce workers to locate around these cutoffs, allowing me to identify behavioural responses and adjustment costs using the recently developed bunching approach (Chetty et al. 2011; Kleven and Waseem 2013; Saez et al. 2010). To track the full trajectory of responses to changes in marginal tax rates, I focus on two kinks that were created at taxable income thresholds that did not initially feature any discontinuities. These new kinks were kept in place for three years and were then subsequently removed.

I begin by presenting several reduced-form findings. I show compelling graphical evidence of behavioural responses through bunching around kinks, i.e. spikes in the distribution of taxable income at thresholds where taxpayers cross from one bracket to another. Further, I provide clear evidence of adjustment costs by examining the distribution of taxable income when a kink is removed (flattened out) by a tax reform. While theory predicts that in a frictionless world, any former excess mass should immediately disappear, workers continue to bunch at former bracket thresholds that no longer exhibit jumps in the marginal tax rate. This holds true not only for salary earners, but also for the self-employed. Moreover, I document the dynamic adjustment patterns in bunching and de-bunching. Bunching to a new kink starts immediately but keeps growing even three years after its introduction, while it takes two years for residual bunching to be completely eliminated through de-bunching from a former kink.

Next, I specify a model that allows me to translate this reduced-form evidence into an estimate of the structural elasticity of taxable income and an adjustment cost that can be used for welfare analysis. I find that both salary earners and self-employed face significant adjustment costs, but that these are much larger for salary earners. Heterogeneity analysis further reveals that adjustment costs are particularly high for females and for those working in sectors characterised by low unionisation rates. To understand the mechanisms driving the adjustments in taxable incomes, I examine whether the observed responses are due to changes in earnings or deductions. The primary mechanism through which salary earners adjust appears to be salary growth, with deductions playing only a minor role. The response of the self-employed however is completely different. Bunching is driven by

earnings decreases rather than increases, and deductions adjustments play a much more important role in targeting kinks. This suggests that bunching patterns are capturing real responses for salary earners, but mainly avoidance responses among the self-employed.

The remaining part of the paper discusses how current approaches to estimating adjustment costs can be extended to the study of asymmetries. The asymmetry I consider is between responses to the creation of a kink, and responses to the elimination of a former kink. I outline how studying asymmetric responses by the direction of the tax change can be informative on the underlying source of the friction, which other studies have largely treated as a black box. I consider the two main types of frictions typically mentioned in this literature: information (or salience) and real search costs, and propose a third: reference-dependence. I finally show, through policy simulations, how distinguishing between settings with symmetric and asymmetric adjustment costs can have significant policy implications.

This paper is closely related to a very scarce literature that combines reduced-form bunching techniques<sup>1</sup> with structural estimation of adjustment costs (Gelber et al. 2017; Zaresani 2017), to which it makes several contributions. Firstly, it leverages a richer source of identifying variation and estimates adjustment costs in a much cleaner setting. Previous studies all share the limitation of relying on bunching moments at already existing kinks. In such cases, it is unclear how the bunching developed in the first place, what the dynamic adjustments were that resulted in the bunching they initially observe, and hence what the bunching mass would be before the creation of such kinks. In contrast, I am able to observe the full trajectory of bunching dynamics, from *before* a kink is created, to *after* it has been removed, which allows for a more credible link between the reduced-form evidence and the structural model. Secondly, I examine the mechanisms driving the observed responses, which other studies have largely treated as a black box. Thirdly, my analysis is conducted on a sample that is much more representative of the average worker than that used in previous work. Gelber et al. (2017) for instance look at US workers just before retirement (aged 65 and above), while Zaresani (2017) only analyses disability insurance beneficiaries and among one specific geographic region in Canada. While interesting in their own right,

---

<sup>1</sup>See Kleven (2016) for a review.

the selectivity of these samples makes it hard to generalise the results from these studies to the wider working population. Fourthly, besides examining salary earners, my study also considers the self-employed and provides the first estimates of adjustment costs for this type of workers in this literature. Finally, I propose an extension of the current approach to one that considers asymmetries in adjustment costs, and outline how such asymmetries can also be informative on the type of friction faced.

Besides the work on bunching and adjustment costs, my paper is also related to several other literatures, such as studies of hours' constraints as a determinant of labour supply choices (Altonji and Paxson 1988; Dickens and Lunberg 1993; Steward and Swaffield 1997). The direct estimation of such adjustment costs has however remained relatively unexplored due to significant identification challenges. Beyond compelling quasi-experimental variation arising from appropriate institutional settings, credible identification also requires high-quality data. This has only recently been made possible via the increasing availability of administrative data. Without such accurate data, it is difficult to separate optimisation frictions from basic measurement error in survey responses.

Further, it adds to the voluminous research on the elasticity of taxable income (for a survey, see Saez et al. 2012). This literature is predominantly focused on the United States or Scandinavian countries, and has mainly considered tax reforms affecting the top of the income distribution (Aarbu and Thoresen 2001; Blomquist and Selin 2010; Feldstein 1995; Gelber 2012; Goolsbee 2000; Gruber and Saez 2002; Kleven and Schultz 2014; etc.). The Cypriot setting provides insights about behavioural responses to taxation in a new context: that of a small open economy, with an intermediate level of GDP/capita compared to other developed economies.

Moreover, it is related to studies examining optimisation frictions in health insurance plan choices (Abaluck and Gruber 2011; Ketcham et al. 2012; Kling et al. 2012; Handel 2013; Handel and Kolstad 2015), pension plan choices (Illanes 2016), home purchases (Bradley 2017), and grocery store shopping (Chetty et al. 2009; Clerides and Courty 2017), among others. It is also related to a theoretical literature highlighting the implications of such optimisation frictions for welfare and optimal policy (Chetty et al. 2009; Farhi and Gabaix 2015; Spinnewijn 2017; Werquin 2015), and to a literature studying asymmetric

responses to price changes in other settings (Benzarti et al. 2017; Doerrenberg et al. 2016; Kube et al. 2013).

Lastly, it contributes to the very limited work on the Cypriot labour market. Previous work has focused on estimating labour supply functions (Pashardes and Polycarpou 2010) or examining the determinants of labour supply among welfare recipients (Pashardes and Polycarpou 2012). The current study uses first-time access to high-quality administrative tax data to provide the first estimates of both micro and macro estimates of the elasticity of taxable income, as well as labour adjustment costs, for Cyprus. This opens the possibility for incorporating these structural parameters in welfare analysis of various policies, especially in light of the recent Cypriot bail-in programme.

The paper proceeds as follows. Section 2.2 outlines the institutional context and Section 2.3 describes the data. Section 2.4 presents the theoretical framework and reduced-form evidence of behavioural responses and adjustment costs, while Section 2.5 discusses the structural estimation method and results. The underlying mechanism driving the observed responses is examined in Section 2.6, while Section 2.7 discusses asymmetries in adjustment costs and presents policy simulations. Finally, Section 2.8 concludes.

## 2.2 Institutional Context

Cyprus has a simple personal income tax system with a progressive rate structure and a single measure of taxable income. Taxable income is defined as the sum of individual income from employment, business income, pension income, foreign income and capital income, net of deductions. The main deductions include most types of capital income, trade union subscriptions, charitable donations, social insurance contributions, pension fund contributions, and life and medical insurance premia. Both salary earners and the self-employed are subject to the same tax schedule, which is not indexed to inflation. The local currency until 2007 was the Cypriot Pound (CYP), after which it was replaced by the Euro.<sup>2</sup>

All salary earners with gross income above the personal allowance must file a tax

---

<sup>2</sup>The permanent exchange rate was locked on 7th December 2007 at €1 = CYP 0.585274.

return.<sup>3</sup> In practice, many individuals earning below the exemption threshold also file because it is a requirement for access to government welfare programmes. Individuals earning any income from self-employment must always file, irrespective of the amount earned. The enforcement system involves third-party reporting for employees, and tax withholding for both employees and self-employed. Tax returns are filed individually.

Due to its progressive structure, the Cypriot income tax system creates convex kinks in workers' budget sets.<sup>4</sup> This provides an opportunity to examine bunching responses by considering reforms to such kinks. To implement my strategy in a credible way, I require a series of reforms that fulfil certain requirements. First, I need to be able to observe a case where a kink is introduced at an earnings level which did not previously feature any discontinuity. Second, this kink needs to stay in place for several years, so that bunching dynamics can be traced out. With only a single year, it is difficult to know what the counterfactual bunching mass would have been the year after. In particular, it is not possible to know whether the excess mass would have grown even more, thereby signifying adjustment costs in the first year, or whether the excess mass in the first year captures the full adjustment. Third, the same kink needs to be eliminated through a later reform. Otherwise, one would have to rely on comparing bunching to one kink and debunching from a different kink, which would require imposing strong assumptions on the distribution of structural elasticities across different income levels. Fourth, there must not be any significant changes to the tax base during these years, otherwise the year-by-year bunching estimates will not be comparable, as they depend on the institutional context in the same way that the elasticity of taxable income does.

Given all these requirements, I exploit a set of reforms in 2004 and 2007 which created and subsequently eliminated kinks at taxable income levels of CYP 10,000 and CYP 20,000. The variation in tax rates during these years is shown in Figure 2.1. This period is particularly appropriate, not only because it fulfils the exact requirements, but also because it does so for two kinks, one at a low and another at a high income level. This enables me to extend my analysis to more than one kink while at the same time holding constant the

---

<sup>3</sup>The personal allowance is equivalent to the first bracket of the tax schedule which always features a zero marginal tax rate.

<sup>4</sup>Historically, there are no more than four brackets and therefore three kinks.

macroeconomic and institutional conditions affecting each of them. In this way, I avoid any confounding issues that would arise, had I had to rely on comparing different kinks from different reforms.

## 2.3 Data

I have obtained first-time access to a matched employer-employee dataset covering the universe of tax filers in the Republic of Cyprus. The dataset contains information from the main tax return form (the IR1A), as well as basic firm and demographics characteristics. I use data covering the years 2003-2008, and restrict my estimation sample to working individuals aged 25-54 to avoid confounding effects from part-time work or early retirement. I focus on two types of workers: salary earners, who only have earnings from employment, and the self-employed, who only have earnings from self-employment. I do not consider a tiny proportion (1% of the sample) that has income from both salaried employment and self-employment. These restrictions produce a sample of 760,277 observations.

Summary statistics of the sample are shown in Table 2.1. On average, tax filers earned CYP 13,780 of gross income and received CYP 1,560 of deductions, leading to a taxable income of CYP 12,220 and a tax bill of CYP 830 per year. Females comprise about 38% of the sample, the average age is 41, and about 91% are salaried employees. Finally, the largest represented sector is the public sector, followed by services, trade and finance.

## 2.4 Conceptual Framework and Reduced-Form Evidence

Before estimating adjustment costs, I first examine whether there is any reduced-form evidence of frictions in the first place. To motivate this analysis, I set ideas by presenting the predictions of a simple model of earnings supply without frictions.<sup>5</sup>

---

<sup>5</sup>This follows basic nonlinear budget set analysis (for a survey, see Moffitt 1990).

### 2.4.1 Introduction to the Bunching Framework

Consider the following model for earnings choices. Each worker has preferences over consumption  $c$  and taxable income  $z$ , and differs by her ability  $n$ . Utility  $u(c, z; n)$  is increasing and concave in consumption ( $u_c > 0$ ,  $u_{cc} < 0$ ), and decreasing and convex in earnings ( $u_z < 0$ ,  $u_{zz} > 0$ ), capturing the idea that labour supply  $z/n$  is costly. Moreover,  $u_{zn} < 0$ , i.e. the marginal disutility of earnings falls with ability. Individual ability is the only source of heterogeneity and is distributed according to a smooth density function  $f(n)$ . Individuals pay tax on income  $z$  captured by  $T(z)$ , which may be a nonlinear function. Denoting virtual income by  $R$ ,<sup>6</sup> the worker's optimisation problem can be stated as:

$$\begin{aligned} \max_{c, z} \quad & U(c, z; n) \\ \text{s.t.} \quad & c = (1 - \tau)z + R \end{aligned}$$

The usual first-order condition  $(1 - \tau)u_c + u_z = 0$  implicitly defines the optimal earnings supply function  $z = z((1 - \tau); n)$ . The trade-off between more consumption and earnings is given by  $n$ , the opportunity cost of leisure. Under a smooth ability distribution, and a linear tax schedule with constant marginal tax rate  $T(z) = \tau_0 \forall z$ , this gives rise to a smooth earnings distribution with cdf  $H_0(z)$  and pdf  $h_0(z)$ . This scenario is depicted in Figure 2.2. Figure 2.2a shows that under a linear tax schedule, workers simply locate smoothly along the budget constraint according to their heterogeneous ability, with higher ability individuals optimising at higher earnings levels.<sup>7</sup> Figure 2.2b shows the associated density, which as a result is also smooth.

Now consider the creation of a second bracket, increasing the marginal tax rate from  $t_0$  to  $t_1$  for earnings above  $z^*$ . This nonlinear schedule features a convex kink, shown in Figure 2.2c, and affects the FOC for those initially earning above  $z^*$ , who adjust (decrease) their earnings. This discontinuous jump in the slope of the budget constraint at  $z^*$  now causes a discontinuous jump in the earnings distribution. As Figure 2.2d shows, individuals in a

<sup>6</sup>Following the public finance literature, I use virtual income  $R = z - T(z) - [1 - T'(z)]z$  to re-write the nonlinear budget constraint in linearised form.

<sup>7</sup>This always holds under the weak assumption of the Spence-Mirrlees (single-crossing) condition on preferences, where the marginal rate of substitution between consumption and pre-tax earnings is decreasing in ability.



range  $z \in (z^*, z^* + \Delta z^*]$  will optimise under the new tax schedule by bunching (from above) at the kink level  $z^*$ . The bunching range can be identified by a marginal buncher, i.e. the individual with the highest earnings under the linear tax schedule who chooses to bunch under the nonlinear schedule. As drawn in Figure 2.2c, this is the individual tangent both at  $z^* + \Delta z^*$  under the linear schedule, and the upper part of the nonlinear budget set at  $z^*$ .

To summarise, this simple nonlinear budget set analysis predicts that under a linear (smooth) tax schedule, the earnings density should be smooth, and under a nonlinear (kinked) income tax schedule, there should be bunching (excess mass) at bracket cutoffs, given by  $B = \int_{z^*}^{z^* + \Delta z^*} h_0(z) dz$ .

## 2.4.2 Estimating the Bunching Mass

A key component of examining responses to the introduction and elimination of kinks is measuring the generated bunching mass. To estimate  $\hat{B}$ , I follow the methodology introduced by Saez (2010) and Chetty et al. (2011). The aim is to estimate what the counterfactual distribution  $h_0(z)$  would have been in the absence of the kink at  $z^*$ , and compare this to the observed distribution in the presence of the kink.  $\hat{B}$  is then the difference between the empirical and counterfactual density. To proceed, I group individuals into earnings bins  $j$  of width  $\delta$ , and fit a flexible polynomial to the empirical distribution. Bunching can then be estimated by running the following regression:

$$c_j = \sum_{i=0}^p \beta_i (z_j)^i + \sum_{i=z_L}^{z_U} \gamma_i \mathbb{1}[z_j = i] + \sum_{r \in R} \rho_r \mathbb{1}\left[\frac{z_j}{r} \in \mathbb{N}\right] + \sum_{k \in K} \theta_k \mathbb{1}\left[z_j \in K \wedge z_j \notin [z_L, z_U]\right] + v_j \quad (2.1)$$

where  $c_j$  is the count of individuals in bin  $j$ ,  $z_j$  is the earnings level in bin  $j$ , and  $p$  is the order of the fitted polynomial. In my main specification I use  $\delta = 50$  and  $p = 9$ . To account for potential diffuse bunching, I define the bunching region to be  $z \in [z_L, z_U]$  around  $z^*$ , where  $z_L$  and  $z_U$  are determined visually. Because kinks are located at round numbers, and round numbers may act as reference points, I also control for round-number bunching. Otherwise, my estimates will be biased if workers (also) bunch at a round-numbered kink

for reasons beyond the jump in marginal tax rate at that kink. I control for this flexibly, including dummies for bins outside the bunching region that are multiples of each element of the set  $R = \{500, 1000\}$ . This allows for the fact that some round numbers may be “rounder” than others and will therefore induce more bunching. Further, I also include fixed effects for bins in the set of bracket thresholds  $K = \{z_1^*, z_2^*, z_3^*\}$  whenever they lie outside the excluded region but are included in the estimation of the counterfactual. Such “nearby” kinks would otherwise bias the counterfactual estimate, especially if they are also at round numbers, producing spikes in the (otherwise) smooth counterfactual.

The estimated counterfactual distribution is the predicted bin count from (2.1), including the contributions of the round number and nearby kink dummies, but excluding the contribution of the excluded range, i.e.:

$$\hat{c}_j = \sum_{i=0}^p \hat{\beta}_i (z_j)^i + \sum_{r \in R} \hat{\rho}_r \mathbb{1} \left[ \frac{z_j}{r} \in \mathbb{N} \right] + \sum_{k \in K} \hat{\theta}_k \mathbb{1} \left[ z_j \in K \wedge z_j \notin [z_L, z_U] \right] \quad (2.2)$$

An initial estimate of the excess mass is then the difference between the actual and counterfactual bin counts, given by:

$$\hat{B}^0 = \sum_{j=z_L}^{z_U} (c_j - \hat{c}_j) \quad (2.3)$$

This initial estimate overestimates  $\hat{B}$ , because it does not account for the fact that bunchers at  $z^*$  are coming from counterfactual earnings levels above this threshold. Hence, the estimated counterfactual will underestimate the count of workers in bins above  $z^*$ , had there been no kink. I address this by shifting upwards the counterfactual distribution to the right of the kink, until the count of workers under the empirical and counterfactual distribution are equal. I estimate the following specification:

$$\begin{aligned}
c_j \left( 1 + \mathbb{1}[j > z_U] \frac{\hat{B}^0}{\sum_{j=z_U+1}^{\infty} c_j} \right) = \\
\sum_{i=0}^p \beta_i(z_j)^i + \sum_{i=z_L}^{z_U} \gamma_i \mathbb{1}[z_j = i] + \sum_{r \in R} \rho_r \mathbb{1} \left[ \frac{z_j}{r} \in \mathbb{N} \right] + \sum_{k \in K} \theta_k \mathbb{1} \left[ z_j \in K \wedge z_j \notin [z_L, z_U] \right] + v_j
\end{aligned} \tag{2.4}$$

This is estimated by iteration, recomputing  $\hat{B}$  in each run until a fixed point is reached.  $\hat{B} = \sum_{j=z_L}^{z_U} (c_j - \hat{c}_j)$  is then based on this corrected counterfactual. To make this comparable across different kinks, I scale it by the average counterfactual frequency in the excluded range,  $\hat{c}_0$ , to get the *normalised* excess mass: :

$$b = \frac{\hat{B}}{\hat{c}_0} \tag{2.5}$$

where  $\hat{c}_0 = \left[ \frac{z_U - z_L}{\delta} \right]^{-1} \sum_{j=z_L}^{z_U} \hat{\beta}_i(z_j)$ . I calculate bootstrapped standard errors using 200 bootstrap samples with replacement. The standard deviation of the distribution of these estimates gives the standard error of  $b$ .

Using  $\hat{c}_0$  I can estimate  $h_0(z^*) = \hat{c}_0/\delta$  and  $\Delta z^* = \hat{B}/\hat{h}_0(z^*)$  which identifies the marginal buncher. The earnings response of the marginal buncher can then be used to estimate the (observed) elasticity of taxable income w.r.t. the net-of-tax rate:<sup>8</sup>

$$\varepsilon_{observed} = \frac{\Delta z^*/z^*}{\Delta t/(1-t)} \tag{2.6}$$

---

<sup>8</sup>As long as the kink  $\Delta t$  is small, this will represent a compensated elasticity. With large kinks, the estimate will be a weighted average of compensated and uncompensated elasticities. For a formal treatment see Kleven (2016).

### 2.4.3 Graphical Evidence

I next present reduced-form evidence for the existence of both behavioural responses and adjustment costs.<sup>9</sup> In all analyses that follow, I consider salary earners and the self-employed separately. I start with salary earners and examine responses to the CYP 10,000 kink. This was introduced in 2004, remained in place until 2006 and was removed in 2007. In this case, the frictionless model predicts no bunching in 2003, bunching during 2004-2006, and immediate de-bunching in 2007. Figure 2.4 shows the distribution of taxable income and the estimated counterfactual around CYP 10,000 for years 2003-2008.<sup>10</sup> Each sub-figure reports the estimate of the normalised excess mass, and provided that a kink is in place in that year and bunching is statistically significant, the estimate of the observed elasticity. The location of the kink is marked with a solid vertical line. Starting with the 2003 distribution, there is no evidence of bunching at CYP 10,000. Compare this to 2004, the first year when the tax schedule features a kink at CYP 10,000. There is now a clear spike at the kink's location, with an estimated normalised excess mass  $b$  of 0.87, implying that there are 87% more workers locating at the kink than there would have been in the counterfactual scenario of no jump in the marginal tax rate at that income level. Bunching appears to grow in the following years, with  $b$  increasing to 0.99 in 2005 and 1.64 in 2006. The most striking evidence of adjustment costs comes from the 2007 distribution, which still exhibits significant bunching at CYP 10,000 although there is no kink at that income level. While the bunching drops from 2006, implying that some workers do adjust, a significant amount do not. Residual bunching at an income level where the tax schedule is smooth goes against the model, which predicts complete and immediate de-bunching. Residual bunching is eliminated by 2008, showing that it takes longer than what a frictionless model would predict for the distribution to become smooth around the former kink.

The responses of salary earners around the CYP 20,000 kink are considered next. Figure 2.5 shows that in this case, there are no bunching responses in any year. While the lack of

---

<sup>9</sup>Appendix Table 2A.1 shows that the patterns in the bunching estimates presented in this section are robust to alternative specifications of bandwidth, bin size and polynomial order.

<sup>10</sup>The vertical dashed line below  $z^*$  demarcates the lower bound of the excluded region,  $z_L$ . In all cases,  $z_U = z^*$  because there is never any visible diffuse bunching mass to the right of the kink under consideration.

bunching in 2003 and 2007-2008 is consistent with the lack of a kink in those years, it is clearly at odds with the model for 2004-2006. This implies that the adjustment cost is so large that it outweighs any gain from bunching for any individual in the bunching region around this kink. An alternative explanation that does not require frictions could be that workers' elasticity of earnings at such high income levels is zero, leading to no behavioural responses, with or without adjustment costs. This explanation is however highly unlikely, given the extensive evidence that high income earners tend to be the most responsive to tax changes (Saez et al. 2012). Interestingly, this heterogeneity in bunching across the income distribution is the opposite to that found by the Chetty et al. (2011) study on Denmark, who found salary earners to bunch only at the top kink.

I next repeat the analysis for the self-employed. Figure 2.6 shows the patterns around the CYP 10,000 kink, and Figure 2.7 around the CYP 20,000 kink. In contrast to salary earners, the self-employed do not respond by bunching at the CYP 10,000 kink. While there are some small spikes during some of the years between 2004-2006, they are not much larger than the spike in 2003 and can purely be explained by round-number bunching. Further, the estimate of the normalised excess mass is not statistically significantly different from zero in any year. On the other hand, the self-employed exhibit very strong responses to the CYP 20,000 kink. Bunching first appears in the first year the kink is introduced, grows steadily until 2006, exhibits residual bunching in 2007 and finally disappears in 2008. This pattern resembles that of salary earners' bunching at CYP 10,000 and again goes against a frictionless model. The patterns however differ in two important ways. Firstly, the spike is now more sharp, with no visible diffuse bunching around the kink, implying that the self-employed are more able to precisely target the kink. Secondly, the size of the normalised excess mass is now also much larger, growing from 5.11 to 10.5, an order of magnitude larger than the largest estimate for salary earners. Such differences in the magnitude of the responses between salary earners and self-employed have also been documented in previous studies (Chetty et al. 2011; Kleven and Waseem 2013), and can be attributed to the fact that the self-employed have more flexible ways of responding to tax reforms than salary earners.

## 2.4.4 Introducing Adjustment Costs

The graphical evidence speaks against the frictionless model, which assumes that everyone in the range  $z \in (z^*, z^* + \Delta z^*]$  can costlessly bunch when a kink is introduced at  $z^*$  and de-bunch when it is removed. In reality, this response is attenuated by various frictions, such as the cost of adjusting labour supply (e.g. renegotiating contracts with employers or searching for new jobs offering the desired hours-wage package), acquiring information about the tax code, etc. Following Gelber et al. (2017), I adopt a simple model to show how introducing a fixed cost can reconcile the patterns in the data.<sup>11</sup> Specifically, adjusting earnings now entails a fixed cost  $\phi$ , with utility given by:

$$U(c_t, z_t; n) - \phi \mathbb{1}(z_t \neq z_{t-1}) \quad (2.7)$$

This attenuates bunching because for some individuals, the gain in utility from re-optimising from an income level  $z \in (z^*, z^* + \Delta z^*]$  to  $z^*$  is smaller than the cost required to do so. This is explained in Figure 2.3a. As the gain from re-optimising is increasing in the size of the adjustment,<sup>12</sup> there is a marginal individual with earnings level  $\underline{z}$  who is indifferent between not adjusting and staying at  $\underline{z}$  (but dropping from consumption point A to B) or incurring the cost and moving to  $z^*$  (point C). The marginal buncher with the highest ability level still bunches from  $z^* + \Delta z^*$ , but individuals with incomes  $z \in (z^*, z^* + \underline{z}]$  now do not adjust. Earnings  $\underline{z}$  of the marginal non-buncher are implicitly defined by the following condition:

$$U(\underline{z}, \tau_1) = U(z^*, \tau_0) - \phi \quad (2.8)$$

Compared to the frictionless case, bunching with frictions will therefore be:<sup>13</sup>

---

<sup>11</sup>A further extension of the model is presented in Appendix 2B.

<sup>12</sup>As Gelber et al. (2017) show, this requires the size of the optimal adjustment to increase in ability at a rate faster than the decrease in the marginal utility of consumption. This holds for instance in the case of a quasi-linear utility function, which is typically assumed in the public finance literature and which will also be employed here for structural estimation.

<sup>13</sup>Note that the approximation assumes a constant initial density  $h_0$ . This is only used here to simplify the exposition, but is not imposed in the empirical implementation that follows, which allows for curvature.

$$B_+ = \int_{\underline{z}}^{z^* + \Delta z^*} h_0(z) dz \approx h_0(z^*)(z^* + \Delta z^* - \underline{z}) \quad (2.9)$$

where  $B_+$  represents the bunching mass at an income level  $z^*$  that features a kink. Following the same logic, frictions also attenuate de-bunching from a kink that is removed. De-bunching requires increasing earnings. For some, the gain from doing so will not be enough to offset the adjustment cost. As shown in Figure 2.3b, there will be a marginal “non de-buncher” indifferent between incurring the cost to increase earnings to  $\bar{z}$  (point B), or staying at  $z^*$  (point A). Everyone with counterfactual earnings below  $\bar{z}$  stays at  $z^*$ , and everyone above adjusts, where  $\bar{z}$  is implicitly defined by:

$$U(z^*, \tau_0) = U(\bar{z}, \tau_0) - \phi \quad (2.10)$$

Compared to the frictionless case, the remaining bunching is now:

$$B_- = \int_{\underline{z}}^{\bar{z}} h_0(z) dz \approx h_0(z^*)(\bar{z} - \underline{z}) \quad (2.11)$$

where  $B_-$  represents the bunching mass left over at an income level  $z^*$  that no longer features a kink.

## 2.5 Estimation of Structural Parameters

### 2.5.1 Parametrising the Utility Function

The previous analysis showed that the amount of bunching depends not only on the (structural) elasticity, but also on the size of the adjustment cost. Given the available variation in tax rates and bunching moments, it is possible to estimate both by specifying a functional form for  $U(c, z; n)$ . Following the typical approach in the public finance literature, I adopt the following quasi-linear parametrisation:

$$U(c, z; n) = c - \frac{n}{1 + \frac{1}{\varepsilon}} \left( \frac{z}{n} \right)^{1 + \frac{1}{\varepsilon}} \quad (2.12)$$

Optimal earnings supply is  $z = n(1 - \tau)^\varepsilon$ . Denote by  $n^*$  the ability of the marginal buncher earning  $z^* + \Delta z^* = n^*(1 - \tau_0)^\varepsilon$  under  $\tau_0$  and  $z^* = n^*(1 - \tau_1)^\varepsilon$  under  $\tau_1$ . The two tangency conditions that must hold for this buncher can be combined to define the elasticity as:

$$\varepsilon = \frac{\ln\left(1 + \frac{\Delta z^*}{z^*}\right)}{\ln\left(\frac{1-\tau_0}{1-\tau_1}\right)} \quad (2.13)$$

which implies:

$$\Delta z^* = z^* \left( \left( \frac{1 - \tau_0}{1 - \tau_1} \right)^\varepsilon - 1 \right) \quad (2.14)$$

Using (2.9), (2.11) and (2.14), the earnings levels of the marginal non-buncher and non de-buncher can be expressed as:

$$\underline{z} = z^* \left( \frac{1 - \tau_0}{1 - \tau_1} \right)^\varepsilon - \frac{B_+}{h_0(z^*)} \quad (2.15)$$

$$\bar{z} = \frac{B_-}{h_0(z^*)} + \underline{z} \quad (2.16)$$

Finally, the two indifference conditions can be written as:

$$\phi + (1 - t_1)(\underline{z} - z^*) + \left( \frac{1 - t_0}{1 + \frac{1}{\varepsilon}} \right) \left( z^{*1+\frac{1}{\varepsilon}} \underline{z}^{-\frac{1}{\varepsilon}} - \underline{z} \right) = 0 \quad (2.17)$$

$$\phi + (1 - t_0) \left( z^* - \left( \frac{1}{1 + \frac{1}{\varepsilon}} \right) z^{*1+\frac{1}{\varepsilon}} \bar{z}^{-\frac{1}{\varepsilon}} - \frac{\bar{z}}{1 + \varepsilon} \right) = 0 \quad (2.18)$$

where (2.17) is the indifference condition for the marginal non-buncher when the kink is in place, and (2.18) is that of the marginal non de-buncher when the kink is removed. These two equations describe a system of two equations for a given kink in two unknown structural parameters,  $\varepsilon$  and  $\phi$ , which can be solved for numerically by drawing on two



empirical moments associated with bunching to and de-bunching from that kink,  $\hat{B}_+$  and  $\hat{B}_-$ .

## 2.5.2 Results

As the model requires empirical bunching moments, I focus on the CYP 10,000 kink for salary earners, and the 20,000 kink for the self-employed. In each case, I draw on moments from 2004 and 2007. These are the first years each kink was introduced and then eliminated, allowing me to estimate the *immediate* adjustment costs in each case. Results are shown in Table 2.2. It reports the observed and structural elasticities, as well as the estimated fixed cost, at CYP 10,000 for salary earners (panel A) and CYP 20,000 for the self-employed (panel B). Given the bunching variation used, the estimated adjustment cost can be interpreted as the cost of adjusting earnings in the first year after a tax reform.

Salary earners are found to face an average cost of CYP 79 when adjusting earnings. This is a very significant cost, especially when compared to their tax liability, which is approximately zero at their level of taxable income. Moreover, accounting for this adjustment cost uncovers a structural elasticity of 0.35, which is orders of magnitude larger than their observed elasticity. This elucidates how a small observed elasticity is consistent with a large structural elasticity and a large adjustment cost, and shows how important it is to incorporate adjustment costs when estimating behavioural responses.

Similar effects are found when considering the self-employed. The presence of adjustment costs again creates a “wedge” between the observed and structural elasticity, which are estimated at 0.064 and 0.271 respectively. The main difference is that the self-employed face a very low adjustment cost, estimated at about CYP 5. This low cost was also implied by the reduced-form evidence. The much stronger bunching responses to a kink of a much smaller size, together with no diffuseness in bunching around the kink, suggest that any adjustment costs at this income level play a much weaker role. This difference is particularly important, especially when considering the difference in disposable incomes between workers at CYP 10,000 and CYP 20,000. This discrepancy in costs is likely due to self-employment being a more flexible form of work, especially in terms of adjusting hours, and more access to tax avoidance behaviours.

The heterogeneity in these estimates shows that it is possible for some group of workers (salary earners) to exhibit a lower observed elasticity than another (self-employed), but have larger structural elasticities, as long as they face large adjustment costs. This is especially important for welfare analysis. The results show that bunching has not reached its steady state by 2006, since in the long-run, the short-run attenuated elasticities should converge to their long-run structural equivalents. Ignoring adjustment costs would imply that salaried workers simply respond less to reforms, where in fact they are simply further away from their steady state than the self-employed.

### 2.5.3 Heterogeneity Analysis

The previous analysis implicitly assumed a homogeneous elasticity and adjustment cost (for each of the two types of worker). In the presence of underlying heterogeneity, the results can be interpreted as the average elasticity and adjustment cost among the set of bunchers (Kleven 2016). To see this, consider a joint distribution of abilities  $n$ , elasticities  $\varepsilon$  and adjustment costs  $\phi$  given by  $\hat{f}(n, \varepsilon, \phi)$  and a corresponding joint distribution of earnings, elasticities and adjustment costs given by  $\hat{h}_0(z, \varepsilon, \phi)$ , implying  $h_0(z) = \int_{\phi} \int_{\varepsilon} \hat{h}(z, \varepsilon, \phi) d\varepsilon d\phi$ . Then, denoting the response of the marginal buncher and non-buncher at each elasticity level  $\varepsilon$  by  $\Delta z_{\varepsilon}^*$  and  $\underline{z}_{\varepsilon}$ , the expression for bunching from (2.9) can be re-written as:

$$B_+ = \int_{\phi} \int_{\varepsilon} \int_{\underline{z}}^{z^* + \Delta z^*} h_0(z) dz d\varepsilon d\phi \approx h_0(z^*) (z^* + \mathbb{E}[\Delta z_{\varepsilon}^*] - \mathbb{E}[\underline{z}]) \quad (2.19)$$

Similarly, denoting the response of the marginal non de-buncher at each elasticity level by  $\bar{z}_{\varepsilon}$ , the expression from (2.11) for the residual bunching after a kink is removed can be re-written as:

$$B_- = \int_{\phi} \int_{\varepsilon} \int_{\underline{z}}^{\bar{z}} h_0(z) dz d\varepsilon d\phi \approx h_0(z^*) (\mathbb{E}[\bar{z}_{\varepsilon}] - \mathbb{E}[\underline{z}_{\varepsilon}]) \quad (2.20)$$

The bunching moments used to estimate the model parameters are then based on the responses of the average buncher, non-buncher and non de-buncher.

To check for any underlying heterogeneity, I repeat the structural estimation on subsamples by sex, age, and by whether workers are in highly unionised sectors. I focus only

on salary earners, as the sample size of the self-employed around CYP 20,000 is too small to conduct bunching estimation on sub-samples. Table 2.3 shows the results.<sup>14</sup>

The estimates reveal significant heterogeneity depending on the dimension considered. Age for instance does not appear to play an important role. Younger workers (aged 25-39) do seem to have slightly higher structural elasticities and adjustment costs than older workers (aged 40-54), but the estimates are rather similar. What is more striking is the heterogeneity by sex. Despite having very similar observed elasticities, females have a structural elasticity of 0.558, which is much larger than that of males, estimated at 0.293. This is rationalised by a much higher adjustment cost, estimated at CYP 130 compared to CYP 65 for men. While a large difference in elasticities is consistent with existing evidence (Evers et al. 2008; Saez et al. 2013), this is the first study to document such a large difference in adjustment costs.

The most significant dimension of heterogeneity however appears to be whether workers are in highly unionised sectors. In Cyprus, unionisation is heavily concentrated in just a few industries: the public sector, commercial banking, hotel services and construction. Unionisation rates in the public sector and commercial banking are nearly 100% due to automatic enrolment, while they are over 75% in hotel services and construction (Ioannou and Sonan 2014). As I do not observe unionisation status directly, I instead group individuals by whether they are in a highly unionised sector or not,<sup>15</sup> which I define as one of the four aforementioned industries. The results show that workers not in highly unionised sectors have both, much larger structural elasticities (0.57), and higher adjustment costs (CYP 133), compared to those in highly unionised sectors (0.344 and CYP 78 respectively). Examining their bunching patterns (Figures 2A.5 and 2A.6), we also find very little bunching among highly unionised sector workers between 2004-2006, and *no* residual bunching in 2007, which explains the lower adjustment cost.

The bunching patterns are the opposite to the Chetty et al. (2011) study on Denmark, which found that bunching was higher among highly-unionised sectors because of unions' responses to tax reforms. Their argument was that the higher the proportion of workers

---

<sup>14</sup>The yearly empirical and counterfactual distributions for each sub-sample are shown in Appendix Figures 2A.1-2A.6.

<sup>15</sup>For the public sector and commercial banking, this will be equivalent since nearly 100% are unionised.

with taxable incomes potentially affected by tax reforms, the larger the propensity for unions to cater to these workers' preferences due to aggregate bunching. One reason why the Cypriot results may be different is that the majority of unionised workers are simply not affected by the kink under consideration. This would imply that unionised workers simply bunched elsewhere, which is not however the case. In fact, the explanation seems to be that Cypriot trade unions prioritise pay raises over any responses to marginal tax reforms, which I have confirmed through conversations with representatives of the unions covering the highly unionised sectors.<sup>16</sup> I have also obtained access to the collective bargaining agreements in the relevant sectors covering the period studied in this paper, which again confirm this reasoning. As an illustration, Appendix Tables 2C.1 and 2C.2 show the main terms of the agreements in the commercial banking and construction sector. Rather than focusing on particular salary ranges, both agreements contain equal salary increases across the board, from the lowest skilled employees to the highest. Hours remain fixed and so do not contribute to any adjustment. More importantly, salary raises do not only occur in years with tax changes, but also in years without, even though there is no automatic indexation of thresholds to adjust for inflation. Unions therefore do not attempt to avoid bracket creep, which is a very strong indication that they are not taking members' tax incentives into account but strive to achieve consistent salary raises.

## 2.6 Mechanisms

Are adjustments driven primarily by income or deductions responses? Distinguishing between the two can provide useful insights into the mechanisms behind the observed responses. For instance, it can show whether bunching is capturing real or avoidance responses, with the former being driven by mainly salary adjustments and the latter driven mainly by deductions adjustments. Further, understanding which is more important can also tell us about the type of adjustment cost driving the inertia in bunching.

I first examine the bunching response to a kink in place. To do so, I focus on each year's new potential bunchers. I define new potential bunchers for a given year as those

---

<sup>16</sup>These include ETYK, PASYDY, PEO and SEK. ETYK covers the banking sector, PASYDY the public sector, and PEO and SEK the remaining private sector.

who in the previous year appear in the bandwidth but not yet in the bunching region, and pool these observations across years 2004-2006. I use this sample to investigate bunching patterns among these potential bunchers according to their salary and deductions growth.<sup>17</sup> Second, I investigate the de-bunching response by conditioning my sample on those who have bunched by 2006, and check how earnings and deductions adjustments affect the observed de-bunching in 2007. The analysis is conducted for both salary earners and self-employed (separately).

Figure 2.8 looks at the bunching patterns for salary earners. Results show that salary growth is the main driver of bunching. For instance, in Figure 2.8a we see a large difference in the bunching response between potential bunchers who experience salary growth and those who do not; salary growth leads to more pronounced targeting of the kink. On the other hand, the difference between those with and without growth in deductions is much smaller (Figure 2.8b). Importantly, the patterns are not in line with what we would expect, had the adjustment cost required to bunch be driven by the ability to change deductions. Intuitively, engaging in avoidance behaviour by filing more deductions should be costlier than filing less. Hence, if bunching captures individuals who incurred such an adjustment cost in order to target the kink, then bunching should be stronger among those increasing deductions, which is not the case. This is also confirmed if we examine deduction responses after conditioning on the type of salary growth (Figures 2.8c and 2.8d).

The first-order effect of salary changes is even starker when we analyse the 2007 de-bunching responses of the 2006 bunchers (Figure 2.9). As Figure 2.9a shows, salary growth completely determines whether previous bunchers are able to de-bunch. Residual bunching is now purely driven by workers experiencing no growth in salaries, compared to the bunching case (Figure 2.8a) where salary growth determined the intensive but not the extensive margin of bunching. The second-order importance of deductions is again exemplified by the sub-cases shown in Figures 2.9b, 2.9c, and 2.9d.<sup>18</sup>

---

<sup>17</sup>In the results that follow, I compare those with positive to those with zero growth, with the latter including those with negative growth. This is done for reasons of parsimony, as the patterns between salary earners exhibiting zero and negative growth are identical. Among the self-employed moreover, the “no growth” category is essentially equivalent to the “negative growth” case, since less than 1% of self-employed stay at the same value of earnings or deductions year by year.

<sup>18</sup>Figure 2.9d only shows the case of no salary nor deductions growth, because there is no individual in the bandwidth who has experienced no salary growth but positive deductions growth.

When we consider the self-employed,<sup>19</sup> we find the opposite results. As shown in Figure 2.10a, bunching is now more pronounced among those with *no* growth in earnings. Further, deductions now play a much bigger role. In contrast to the case of salary earners, bunching among the self-employed who do not exhibit growth in deductions is now larger than those with deductions growth by an order of magnitude (Figure 2.10b). Furthermore, the type of deductions adjustment now matters far more, even when conditioning on positive earnings growth (Figure 2.10c). The most striking difference with salary earners is that the strongest bunching response among self-employed comes from those who exhibit a decrease in both earnings and deductions.

These discrepancies imply that salary earners and self-employed use different means to bunch, and hence face different types of adjustment costs. Overall, the ability of salary earners to respond seems to depend primarily on salary adjustments, and specifically, *increasing* their salary. This shows that the adjustment cost precluding bunching (and de-bunching) is related to the ability to increase earnings, rather than changing deductions in order to target the kink. Moreover, while salary adjustments are important in both bunching and de-bunching responses, they appear to play a much stronger role in de-bunching, as they completely determine whether previous bunchers stay at the former kink or not. The primarily role of earnings adjustments for salary earners implies that their bunching patterns represent real responses and that the main friction they face is a real search/re-negotiation cost related to increasing earnings. The self-employed on the other hand seem to respond by adjusting both earnings and deductions. Moreover, self-employed bunch by *decreasing*, rather than increasing, their earnings, and by decreasing deductions. With deductions now playing a much larger role, a large proportion of the bunching responses among the self-employed likely represents avoidance behaviour. To the extent that earnings are also self-reported, the reduction in earnings also likely represents a reporting response, rather than real changes in earnings.

---

<sup>19</sup>Note that the de-bunching analysis cannot be implemented for the self-employed due to an insufficiently low number of observations.

## 2.7 Extensions

### 2.7.1 Alternative Non-Parametric Structural Elasticity Bounds

The model presented incorporates frictions by imposing some structure on the adjustment cost. While this is modelled using the simplest possible specification (that of a fixed cost parameter), it of course comes at the cost of making certain parametric assumptions.

Another approach, suggested by Chetty (2012), is to remain agnostic about the structure of adjustment costs and to instead estimate bounds for the true structural elasticity. Chetty (2012) shows how to derive closed-form solutions for these bounds as functions of the observed elasticity  $\hat{\varepsilon}$ , the size of the marginal tax rate change  $\Delta \ln(1-t)$ , and the degree of optimisation frictions  $\delta$ , defined as the utility cost of not responding to a tax change as a percentage of earnings. For the functional form of preferences that I have adopted in this paper, the lower and upper bounds,  $\varepsilon_L$  and  $\varepsilon_U$ , are given by:

$$\varepsilon_L = \hat{\varepsilon} + \frac{4\delta}{(\Delta \ln(1-t))^2}(1-\rho) \quad (2.21)$$

$$\varepsilon_U = \hat{\varepsilon} + \frac{4\delta}{(\Delta \ln(1-t))^2}(1+\rho) \quad (2.22)$$

where:

$$\rho = \sqrt{1 + \frac{1}{2} \frac{\hat{\varepsilon}}{\delta} (\Delta \ln(1-t))^2} \quad (2.23)$$

For comparison purposes, I also implement this method for each of the two kinks analysed in this paper. Table 2.4 shows what these alternative bounds are, using common values of  $\delta$  used in the literature (Chetty 2012; Bastani and Selin 2014). Reassuringly, my structural estimates using my adopted functional form lie within the Chetty (2012) predicted bounds. These bounds however are not very informative; they are too wide to offer accurate guidance as to what a reasonable structural elasticity could be assumed to be. The lower bounds are approximately zero, and the upper bounds lie above unity, blowing up to unrealistically high values, especially in the case of the CYP 20,000 kink.

This highlights that structural assumptions regarding the adjustment cost are necessary

to make progress in a setting with tax kinks. As we can see, larger kinks lead to tighter bounds, so with large enough kinks this method could be informative. Unfortunately however, it is unusual for taxable income schedules to feature sufficiently large kinks that would make this agnostic approach reliable.

### 2.7.2 Asymmetric Adjustment Costs

This paper has estimated adjustment costs using two empirical bunching moments: one when a kink is in place, and one after a kink has been removed. The implicit assumption is that adjusting to a kink, and adjusting from a previous kink, entails similar adjustment costs. To the extent that these differ, the estimates can be interpreted as an average between the two.

While no study has considered asymmetric adjustment costs, these may be empirically relevant if tax changes in different directions may require different means of adjustment, some more costly than others. For instance, it may be that reducing taxable income is achieved by increasing deductions, but increasing taxable income is done by searching for a higher wage-hours package, which may be costlier to achieve. Even when both responses are driven by the same mechanism, say a change in salary, it may be that adjusting salaries in one direction is costlier than in the opposite case.

The possibility for asymmetries arises from the fact that, compared to the counterfactual case of no kink, creating a kink should lead to bunching from above, i.e. a movement towards a *lower* than otherwise earnings level. When a kink is removed however, this should lead to de-bunching to a *higher* than otherwise earnings level. To the extent that adjustment costs are asymmetric by the direction of the optimal response, this becomes policy relevant when the fiscal authority is deciding between increasing or decreasing tax rates, or implementing a combination of the two at different parts of the income distribution that may exhibit different elasticities and adjustment costs. Understanding what factors inhibit taxable incomes from responding to tax changes is also germane to the correct design of any policy intervention aiming to ameliorate such frictions. For instance, if workers cannot adjust their taxable incomes due to real search costs, an information intervention to increase tax salience will be ineffective in inducing the anticipated response.



In this paper, I unfortunately do not have access to enough variation in bunching moments that would make the estimation of such asymmetries feasible. Nevertheless, I propose how one could do so in future research. To proceed, one would need to modify the structural model presented in Section 2.5 to allow for adjustment costs that vary by whether they relate to a bunching or de-bunching moment, as follows:

$$\phi_{bunch} + (1 - t_{1k})(z_k - z_k^*) + \left(\frac{1 - t_{0k}}{1 + \frac{1}{\varepsilon_k}}\right) \left(z_k^{*1+\frac{1}{\varepsilon_k}} z_k^{-\frac{1}{\varepsilon_k}} - z_k\right) = 0 \quad (2.24)$$

$$\phi_{de-bunch} + (1 - t_{0k}) \left(z_k^* - \left(\frac{1}{1 + \frac{1}{\varepsilon_k}}\right) z_k^{*1+\frac{1}{\varepsilon_k}} \bar{z}_k^{-\frac{1}{\varepsilon_k}} - \frac{\bar{z}_k}{1 + \varepsilon_k}\right) = 0 \quad (2.25)$$

where  $k$  stands for a particular kink. As the number of parameters to be estimated has now increased, so has the requirement for variation in bunching moments across different kink sizes. Specifically, one would need to observe (statistically significant) bunching and de-bunching at (at least) two differentially sized kinks.<sup>20</sup> This also entails one has data spanning a long enough time period to be able to observe bunching dynamics and evidence of inertia at each kink.<sup>21</sup> The model would then contain a set of pairs of equations (2.24) and (2.25), one pair for each kink, and associated bunching and de-bunching moments for each one.

Given this setup and having estimated  $\phi_{bunch}$  and  $\phi_{de-bunch}$ , one could then exploit any uncovered asymmetries to get further insights into the sources of the underlying frictions. In what follows I review the main candidate explanations for the source of frictions typically discussed in this literature, search costs and misinformation/inattention, and discuss a third possibility, reference dependence. I focus on salary earners because the intuition assumes salary responses are real.<sup>22</sup>

The first possibility is that adjustment costs are due to real search costs. In this case,

---

<sup>20</sup>This is the main limitation faced by the current study as well as the vast majority of bunching papers, where bunching (among the same sample of workers) tends to be only observed at (at most) one kink.

<sup>21</sup>For instance, a case where one observes immediate bunching to and immediate de-bunching from a previous kink cannot be used for estimation, since this reveals no inertia, and hence no adjustment cost. On the other hand, a case of positive bunching moments but no residual bunching (immediate de-bunching from a former kink) would imply that  $\phi_{bunch} > 0$  but  $\phi_{de-bunch} = 0$ , uncovering an asymmetry.

<sup>22</sup>Self-employed individuals may change both their reported earnings and deductions without any real changes taking place, making it harder to make inferences from any asymmetries.

workers can adjust their taxable income by changing hours, wages, or both. If hours in labour contracts are constrained by market forces to match aggregate preferences (Chetty et al. 2011), workers need to re-negotiate wages, or search for a new job matching their preferences. In more flexible labour markets, workers can negotiate on both dimensions. In either case, workers will face adjustment costs because adjusting taxable income affects firms' production costs. In general, it will be costlier to re-negotiate for a contract, or search for a new job, that offers higher wage-hours packages than lower ones, due to the higher cost to firms. This will be especially more costly when contracts are constrained to match aggregate preferences.

**Prediction 1:** *If frictions are due to search/re-negotiation costs, adjustment costs are higher when de-bunching.*

Another reason why workers may not respond to the tax incentives they face is simply because they do not know about them. This can be either because the tax system is not salient, due to complexity, or because individuals are not informed, due to inattention. In this case, there should be no difference in the adjustment cost upwards or downwards. Non-responsiveness should not depend on the direction of the desired direction of response, if workers are not aware of the incentive they are not responding to.

**Prediction 2:** *If frictions are due to (lack of) information or (in)salience, adjustment costs when bunching and de-bunching are equal.*

The third possible explanation is reference dependence over individual utility from consumption (or equivalently taxable income). If the reference point is the previous year's consumption level, bunching and de-bunching will involve asymmetries. The intuition is straightforward. When a kink is introduced, the standard bunching framework predicts that workers in a bunching range above the kink respond by reducing taxable income (and hence consumption). While this maximises utility given the new budget constraint, it also introduces a loss relative to the reference point. Hence, compared to a frictionless model, responding to a kink (by moving downwards) also entails a cost from moving to a lower consumption level. In contrast, there is no such cost when a kink is removed. In

this case, workers optimise by increasing taxable income (and hence consumption) beyond their existing reference point. Another way to see this is by introducing loss aversion to individual utility:

$$U = U(c, \frac{z}{n}) + \phi \mathbb{1}(c \neq \bar{c})$$

where:

$$\phi = \begin{cases} \eta(c - \bar{c}) & \text{if } c \geq \bar{c} \\ \lambda\eta(c - \bar{c}) & \text{if } c < \bar{c} \end{cases}$$

and  $\bar{c}$  is the reference point for consumption. Utility now features a gain-loss utility component,  $c - \bar{c}$ . When consumption grows, individuals also gain utility ( $c - \bar{c}$ ), which is weighted by  $\eta$ . When consumption drops by the same amount, individuals derive loss utility ( $c - \bar{c}$ ), weighted now by  $\lambda\eta$ , with  $\lambda > 1$  capturing loss aversion, i.e. that marginal utility is higher for losses than gains of an equal size.

**Prediction 3:** *If frictions are due to reference dependence, adjustment costs are higher when bunching.*

Each type of friction therefore has a distinct prediction about the friction asymmetry by direction of desired movement. Under search costs,  $\phi_{bunch} < \phi_{de-bunch}$ , under lack of salience or inattention,  $\phi_{bunch} = \phi_{de-bunch}$ , while under reference dependence,  $\phi_{bunch} > \phi_{de-bunch}$ .

### 2.7.3 Implications for Policy Evaluation

To show the importance of incorporating both symmetric and asymmetric adjustment costs, I simulate the effect of various policy changes on taxable income, and show how predictions depend on what the adjustment cost is. I consider a simple scenario where the tax authority introduces a kink for a certain time period and then removes it. This would resemble for instance a case where the government needs to increase tax rates temporarily to raise additional tax revenue.

As before, I assume individuals are heterogeneous in their ability, preferences take the usual iso-elastic and quasi-linear form, and changing earnings entails an adjustment cost  $\phi$ . For tractability, I consider a stylised three-period setting. In the initial period 0, marginal tax rates are everywhere zero, and so individuals are distributed according to their ability. In period 1, a kink is introduced at  $z^*$ , increasing the marginal tax rate above the threshold from 0 to some value  $\tau > 0$ . In period 2, the marginal tax rate below  $z^*$  is unchanged, but is reduced to 0 above it, thereby removing the kink.

I simulate outcomes for two different values of marginal tax rates,  $\tau \in \{0.05, 0.25\}$ . In each case, I estimate the share of workers bunching at the kink in period 1, and the share of bunchers de-bunching from the removed kink in period 2. I perform these calculations under four scenarios: (1) no frictions, (2) symmetric frictions, (3) asymmetric frictions where  $\phi_{bunch} < \phi_{de-bunch}$  and (4) asymmetric frictions where  $\phi_{bunch} > \phi_{de-bunch}$ . I use a value of  $\varepsilon = 0.35$  throughout,  $\phi = 79$  in case (2),  $\phi_{bunch} = 40$  and  $\phi_{de-bunch} = 120$  in case (3), and  $\phi_{bunch} = 120$  and  $\phi_{de-bunch} = 40$  in case (4). The aim is to see how changing the assumption about adjustment costs affects the predictions, holding the elasticity constant.

I proceed as follows. First, I estimate what the counterfactual earnings distribution would be under a tax system with zero marginal tax rates. I use the observed distribution of taxable income in the symmetric range  $z \in [9100, 10900]$  around CYP 10,000 from 2003, since this does not feature changes in marginal tax rates, and back out individual ability as  $n_i = z_i(1 - T'(z_i))^{-\varepsilon}$ . Counterfactual earnings are then simply  $z_{0i} = n_i$ . Next I consider the introduction of a kink at CYP 10,000 in period 1. This increases the marginal tax rate to  $\tau$  for taxable incomes above the kink, but leaves the rate unchanged to zero for those below it. I calculate  $z_{1i}$ , the optimal frictionless earnings level under this new tax schedule (which could feature a bunching or interior response) ignoring adjustment costs, and estimate the gain in utility from adjusting compared to remaining inert. Using the calibrations of the adjustment cost, I examine whether the gain in utility outweighs the adjustment cost. Then, I consider the removal of this kink in the third period, and repeat the above exercise.

Table 2.5 shows the results for each of the two marginal tax rates under consideration. The main objects of interest are the bunching share in period 1 and the de-bunching share

in period 2. The bunching share is defined as the fraction of workers who are willing to incur the adjustment cost and bunch at the kink in period 1, among those with counterfactual earnings lying within the bunching region. The de-bunching share is the fraction of those bunching in period 1 who are willing to incur the adjustment cost and de-bunch from the kink in period 2. The table also reports the mean taxable income in both periods, and the total tax revenue in period 1. These are estimated taking into account the full sample used in this calibration exercise (i.e. not just the bunchers).

Panel A shows the results for  $\tau = 0.05$ . With no adjustment cost, 100% of the potential bunchers bunch in period 1, and 100% of them also de-bunch in period 2. This is expected because adjustment is costless and utility gains from responding to the tax change are always positive. When adjustment costs are allowed to be positive ( $\phi = 79$ ), the complete opposite is found. No one bunches in period 1, and consequently there is no one to de-bunch in period 2. In this case, the kink size and therefore utility gain from bunching is not large enough to outweigh the adjustment cost of doing so. Due to no earnings responses, the mean taxable income is now higher than in the previous case of no frictions, which leads to 30% higher tax revenue. Thus, even modest adjustment costs can lead to significantly different responses, highlighting how important it is to incorporate adjustment cost in policy simulations and predictions. The last two columns consider asymmetric adjustment costs. In this case, the asymmetry doesn't matter because the utility gain is so small that any modest level of adjustment cost precludes any bunching from taking place, regardless of how it varies by direction. For the sample considered, any adjustment cost above CYP 5.21 would lead to the same outcome.

The results are very different in Panel B, which repeats the analysis for the hypothetical marginal tax rate of  $\tau = 0.25$ . This case reveals both much stronger bunching responses, and significant consequences of asymmetries. The stronger bunching responses are due to the larger kink size, which generates a larger bunching segment and thereby adjustment. This can be seen for instance in the case of no adjustment costs, shown in column (1). Just as in panel A, the bunching and de-bunching shares are 100%. The mean taxable income and total tax revenue however are now lower, because of the larger segment over which responses are generated, leading to bunching from higher counterfactual earnings

than before. When a symmetric cost  $\phi = 79$  is introduced in column (2), 57% of potential bunchers bunch at the kink in period 1, and 39% of them de-bunch in period 2. Compared to panel A, some now find it optimal to incur the adjustment cost and adjust their taxable income due to the larger kink size. Because period 1 bunchers are now coming from higher counterfactual earnings than compared to the case of Panel A, the gains in utility from adjusting are now also larger, leading to some de-bunching in period 2. The extent of this however depends on the asymmetry. As shown in column (3), when  $\phi_{bunch}$  is small, more potential bunchers respond (81% compared to 57%), but very few de-bunch (7% compared to 39%). When the asymmetry is in the opposite direction ( $\phi_{bunch} > \phi_{de-bunch}$ ), results are also opposite. As shown in column (4), very few will bunch (19%), due to the much larger adjustment cost. Among those who bunch however, everyone de-bunches in period 2. This is because in this scenario, period 1 bunchers will be coming from much higher counterfactual earnings due to much higher adjustment cost  $\phi_{bunch}$ , leading to complete de-bunching when the same individuals next face a much lower de-bunching adjustment cost. Another interesting comparison of the asymmetry case in column (4) of Panel B is with the equivalent asymmetry but with the lower  $\tau = 0.05$  of Panel A. While the response with  $\tau = 0.25$  is larger (19% Vs 0%) and leads to lower mean taxable income, total tax revenue is now also much larger because of the higher marginal tax rate.

Overall, this stylised policy simulation reveals that incorporating adjustment costs can have significant effects on policy simulations, with even moderate costs having large effects on predicted outcomes of tax reforms. Moreover, asymmetries in adjustment costs can mask significant heterogeneity in adjustments between creating a kink, and eliminating one, or increasing and decreasing marginal tax rates more generally.

## 2.8 Conclusion

In conclusion, this paper studied adjustment costs in bunching responses by exploiting a series of tax reforms in the Republic of Cyprus that created and subsequently eliminated kinks at different pre-tax income levels. Reduced-form evidence revealed both behavioural responses and significant adjustment costs, with bunching significantly attenuated by fric-

tions. A structural model of frictional earnings supply was then employed to estimate the adjustment cost using the observed bunching moments. Salary earners were found to face an adjustment cost of CYP 79, compared to only CYP 5 for the self-employed. Significant heterogeneity was also found within the sample of salary earners; females and workers not in highly unionised sectors face particularly large adjustment costs. Examining the mechanisms driving the observed responses, it was shown that salary earners rely mainly on salary changes, but the self-employed rely on both salary and deductions adjustments. Lastly, the paper discussed how the current approach to estimating adjustment costs can be extended to examining asymmetries in frictions according to the direction of the tax change and thereby the bunching response. It also outlined how these can be used to shed further light on the underlying source of frictions, and conducted simulations to illustrate how such asymmetries can affect policy predictions.

The findings of this paper have significant implications. They reveal that even modest adjustment costs have significant effects on taxable income responses to tax reforms, and therefore that estimates of short-term earnings elasticities are highly attenuated compared to the long-run (structural) elasticity. This shows that relying only on cross-sections of bunching moments without time variation to estimate elasticities can be highly misleading, because this ignores the dynamic adjustments taking place year by year. This not only cannot provide estimates of the structural elasticity, which is required for welfare calculations, but will also produce estimates of the elasticity of taxable income that highly depend on the time since the analysed kink was created. This is because this elasticity is a linear function of the observed bunching mass,<sup>23</sup> which can grow significantly in the years following the creation of a kink, as documented in this study.

Lastly, the impact of adjustment costs also has significant implications for policy. To the extent that policymakers use fiscal policy to induce short-term, rapid responses in earnings and tax revenue, it is paramount to understand the degree to which different types of workers can respond, and whether they would respond to increases or decreases in tax rates in different ways. As confirmed by simulations, both the size and asymmetry of adjustment costs can make a large difference in the predicted outcomes of tax policy.

---

<sup>23</sup>This is shown in equation (2.6).

# References

- [1] Aarbu, K.O. and Thoresen, T.O. (2001) ‘Income Responses to Tax Changes - Evidence from the Norwegian Tax Reform’, *National Tax Journal*, vol. 54(2), pp. 319-338.
- [2] Abaluck, J.T. and Gruber, J. (2011) ‘Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program’, *American Economic Review*, vol. 101(4), pp. 1180-1210.
- [3] Altonji, J.G. and Paxson, C.H. (1988) ‘Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-Offs’, *Journal of Labor Economics*, vol. 6, pp. 254-276.
- [4] Bastani, S. and Selin, H. (2014) ‘Bunching and Non-bunching at Kink Points of the Swedish Tax Schedule’, *Journal of Public Economics*, vol. 109, pp. 36-49.
- [5] Benzarti, Y., D. Carloni, J. Harju and Kosonen, T. (2017) ‘What Goes Up May Not Come Down: Asymmetric Incidence of Value Added Taxes’, NBER Working Paper No. 23849.
- [6] Blomquist, S. and Selin, H. (2010) ‘Hourly Wage Rate and Taxable Labor Responsiveness to Changes in Marginal Tax Rates’, *Journal of Public Economics*, vol. 94(11-12), pp. 878-889.
- [7] Bradley, S. (2017) ‘Inattention to Deferred Increases in Tax Bases: How Michigan Home Buyers are Paying for Assessment Limits’, *Review of Economics and Statistics*, vol. 99(1), pp. 53-66.
- [8] Chetty, R., A. Looney and Kroft, K. (2009) ‘Salience and Taxation: Theory and Evidence’, *American Economic Review*, vol. 99, pp. 1145-1177.



- [9] Chetty, R., J.N. Friedman, T. Olsen and Pistaferri, L. (2011) ‘Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records’, *Quarterly Journal of Economics*, vol. 126(2), pp. 749-804.
- [10] Chetty, R. (2012) ‘Bounds on Elasticities with Optimisation Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply’, *Econometrica*, vol. 80(3), pp. 969-1018.
- [11] Chetty, R., A. Guren, D. Manoli and Weber, A. (2013) ‘Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities’, *NBER Macroeconomics Annual 2012*, vol. 27, pp. 1-56.
- [12] Clerides, S. and Courty, P. (2017) ‘Sales, Quantity Surcharge, and Consumer Information’, *forthcoming, Review of Economics and Statistics*.
- [13] Dickens, W.T. and Lundberg, S.J. (1993) ‘Hours Restrictions and Labor Supply’, *International Economic Review*, vol. 34(1), pp. 169-192.
- [14] Doerrenberg, P., D. Duncan and Löffler, M. (2016) ‘Asymmetric Labor-Supply Responses to Wage-Rate Changes: Evidence from a Field Experiment’, IZA Discussion Paper No. 9683.
- [15] Evers, M., R.D. Mooij and Van Vuuren, D. (2008) ‘The Wage Elasticity of Labour Supply: A Synthesis of Empirical Estimates’, *De Economist*, vol. 156, pp. 25-43.
- [16] Farhi, E. and Gabaix, X. (2015) ‘Optimal Taxation with Behavioral Agents’, Working Paper.
- [17] Feldstein, M. (1995) ‘The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act’, *Journal of Political Economy*, vol. 103(3), pp. 551-572.
- [18] Gelber, A.M. (2012) ‘Taxation and Earnings of Husbands and Wives: Evidence from Sweden’, *Review of Economics and Statistics*, vol. 96(2), pp. 287-305.
- [19] Gelber, A.M., D. Jones and Sacks, D.W. (2017) ‘Estimating Earnings Adjustment Frictions: Method & Evidence from the Earnings Test’, Goldman School of Public Policy Working Paper.

- [20] Goolsbee, A. (2000) ‘What Happens When You Tax the Rich? Evidence from Executive Compensation’, *Journal of Political Economy*, vol. 108(2), pp. 352-378.
- [21] Gruber, J. and Saez, E. (2002) ‘The Elasticity of Taxable Income: Evidence and Implications’, *Journal of Public Economics*, vol. 84, pp. 1-32.
- [22] Handel, B. (2013) ‘Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts’, *American Economic Review*, vol. 103(7), pp. 2643-2682.
- [23] Handel, B. and Kolstad, J.T. (2015) ‘Health Insurance for ‘Humans’: Information Frictions, Plan Choice, and Consumer Welfare’, *American Economic Review*, vol. 105(8), pp. 2449-2500.
- [24] Illanes, G. (2016) ‘Switching Costs in Pension Plan Choice’, Working Paper, Northwestern University.
- [25] Ioannou, G. and S. Sonan (2014) ‘Trade Unions in Cyprus: History of Division, Common Challenges Ahead’, Friedrich-Ebert-Stiftung Report.
- [26] Ketcham, J.D., C. Lucarelli, E.J. Miravete and Roebuck, M.C. (2012) ‘Sinking, Swimming, or Learning to Swim in Medicare Part D’, *American Economic Review*, vol. 102(6), pp. 2639-2673.
- [27] Kleven, H. and Waseem, M. (2013) ‘Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan’, *Quarterly Journal of Economics*, vol. 128, pp. 669-723.
- [28] Kleven, H. and Schultz, E. (2014) ‘Estimating Taxable Income Responses Using Danish Tax Reforms’, *American Economic Journal: Economic Policy*, vol. 6(4), pp. 271-301.
- [29] Kleven, H. (2016) ‘Bunching’, *Annual Review of Economics*, vol. 8, pp. 435-464.
- [30] Kling, J.R., S. Mullainathan, E. Shafir, L.C. Vermeulen and Wrobel, M.V. (2012) ‘Comparison Friction: Experimental Evidence from Medicare Drug Plans’, *Quarterly Journal of Economics*, vol. 127(1), pp. 199-235.

- [31] Kube, S., M.A. Marechal and Puppe, C. (2013) ‘Do Wage Cuts Damage Work Morale? Evidence from a Natural Field Experiment’, *Journal of the European Economic Association*, vol. 11(4), pp. 853-870.
- [32] Moffitt, R. (1990) ‘The Econometrics of Kinked Budget Constraints’, *Journal of Economic Perspectives*, vol. 4(2), pp. 119-139.
- [33] Pashardes, P. and Polycarpou, A. (2010) ‘Labour Supply in Cyprus’, Economic Analysis Papers 05-10, Economics Research Centre, University of Cyprus.
- [34] Pashardes, P. and Polycarpou, A. (2012) ‘The Labour Market Behavior of Public Assistance Recipients in Cyprus’, Economic Analysis Papers 09-12, Economics Research Centre, University of Cyprus.
- [35] Saez, E. (2010) ‘Do Taxpayers Bunch at Kink Points?’ *American Economic Journal: Economic Policy*, vol. 2(3), pp. 180-212.
- [36] Saez, E., J. Slemrod and Giertz, S.H. (2012) ‘The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review’, *Journal of Economic Literature*, vol. 50(1), pp. 3-50.
- [37] Spinnewijn, J. (2017) ‘Heterogeneity, Demand for Insurance and Adverse Selection’, *American Economic Journal: Economic Policy*, vol. 9(1), pp. 308-343.
- [38] Stewart, M.B. and Swaffield, J.K. (1997) ‘Constraints on the Desired Hours of Work of British Men’, *Economic Journal*, vol. 107(441), pp. 520-535.
- [39] Werquin, N. (2015) ‘Income Taxation with Frictional Labor Supply’, Working Paper.
- [40] Zaresani, A. (2017) ‘Adjustment Costs and Incentives to Work: Evidence from a Disability Insurance Program’, Working Paper.

## Main Tables

Table 2.1: Summary Statistics

Variable	Mean	Std. Dev.
Earnings	13,779.294	8,471.519
Deductions	1,560.582	1,542.807
Taxable Income	12,219.816	7,504.831
Tax Bill	830.145	1,645.767
Female	0.384	0.486
Age	40.632	8.065
Salary Earners	0.909	0.287
Agriculture	0.004	0.064
Mining	0.003	0.053
Manufacturing	0.082	0.275
Construction	0.074	0.262
Utilities	0.022	0.146
Trade	0.106	0.308
Services	0.218	0.413
Finance	0.105	0.307
Public	0.372	0.483
Other	0.011	0.102
<b>Observations</b>	760,277	

*Notes:* This table shows summary statistics for workers aged 25-54 between 2003-2008.

Table 2.2: Results for Structural Parameters

$\varepsilon_{observed}$	$\varepsilon_{structural}$	$\phi$
<b>A: Salary Earners at the CYP 10,000 Kink</b>		
0.022 (0.002)	0.349 (0.156)	78.708 (37.292)
<b>B: Self-Employed at the CYP 20,000 Kink</b>		
0.064 (0.020)	0.271 (0.059)	5.058 (2.577)

*Notes:* This table shows the results from solving the model described in Section 2.5, based on the empirical bunching moments at CYP 10,000 and CYP 20,000. Each panel shows parameter estimates for the structural elasticity  $\varepsilon_{structural}$  and fixed adjustment cost  $\phi$ , as well the observed elasticity  $\varepsilon_{observed}$  that would be equivalent to the structural elasticity in the absence of adjustment costs. Panel A shows the results for salary earners at the CYP 10,000 kink, and panel B for the self-employed at the CYP 20,000 kink. Standard errors are shown in parentheses.

Table 2.3: Heterogeneity Analysis of Structural Parameters for Salary Earners

	$\varepsilon_{observed}$	$\varepsilon_{structural}$	$\phi$
<b>By Age</b>			
Ages 25-39	0.017 (0.003)	0.380 (0.162)	86.018 (38.912)
Ages 40-54	0.027 (0.003)	0.347 (0.173)	77.954 (41.308)
<b>By Sex</b>			
Male	0.021 (0.002)	0.298 (0.159)	66.502 (38.071)
Female	0.024 (0.004)	0.557 (0.180)	129.540 (43.105)
<b>By High Unionisation Rate Status</b>			
Highly Unionised	0.027 (0.002)	0.570 (0.140)	132.720 (33.936)
Not Highly Unionised	0.010 (0.004)	0.344 (0.186)	77.646 (44.400)

*Notes:* This table shows the results from the structural estimation of the model detailed in Section 2.5 by sex, age and unionisation status, for salary earners around the CYP 10,000 kink. It reports the observed ( $\varepsilon_{observed}$ ) and structural elasticities ( $\varepsilon_{structural}$ ), and the adjustment cost ( $\phi$ ). Standard errors are shown in parentheses.

Table 2.4: Bounds on Structural Elasticity Using a Non-Parametric Approach

$\delta$	Kink Level	$ \Delta \ln(1 - t) $	$\hat{\varepsilon}$	$\varepsilon_L$	$\varepsilon_U$
1%	CYP 10,000	0.22	0.022	0.0003	1.650
2%	CYP 10,000	0.22	0.022	0.0001	3.257
1%	CYP 20,000	0.07	0.064	0.0002	16.934
2%	CYP 20,000	0.07	0.064	0.0001	33.741

*Notes:* This table shows estimates of the upper and lower bounds of the structural elasticity,  $\varepsilon_L$  and  $\varepsilon_U$ , following the Chetty (2012) non-parametric approach. Results are shown for different values of the degree of optimisation frictions  $\delta$ , for each of the two kink levels at CYP 10,000 and CYP 20,000.  $\hat{\varepsilon}$  is the observed elasticity estimated from Section 2.4.3. For more details, see Section 2.7.1.

Table 2.5: Policy Simulations


	(1)	(2)	(3)	(4)
	<b>Adjustment Cost</b>			
	None $\phi = 0$	Symmetric $\phi = 79$	Asymmetric $\phi_{bunch} < \phi_{de-bunch}$	Asymmetric $\phi_{bunch} > \phi_{de-bunch}$
<b>Panel A: 0.05 MTR</b>				
<b>Period 1</b>				
Bunching Share (%)	100	0	0	0
Mean Taxable Income (CYP)	10,542	10,703	10,703	10,703
Total Tax Revenue (CYP)	512,940	662,345	662,345	662,345
<b>Period 2</b>				
De-Bunching Share (%)	100	-	-	-
Mean Taxable Income (CYP)	10,703	10,703	10,703	10,703
<b>Panel B: 0.25 MTR</b>				
<b>Period 1</b>				
Bunching Share (%)	100	57	81	19
Mean Taxable Income (CYP)	10,084	10,136	10,094	10,271
Total Tax Revenue (CYP)	434,604	675,210	481,569	1,304,529
<b>Period 2</b>				
De-Bunching Share (%)	100	39	7	100
Mean Taxable Income (CYP)	10,703	10,580	10,453	10,703

*Notes:* The table shows results of various policy simulations of earnings responses in the case when a kink is created in period 1 and removed in period 2. Column (1) shows the results when it is assumed that the adjustment cost is 0, column (2) when  $\phi = 79$ , column (3) when  $\phi_{bunch} = 40$  and  $\phi_{de-bunch} = 120$  and column (4) when  $\phi_{bunch} = 120$  and  $\phi_{de-bunch} = 40$ . A value of  $\varepsilon = 0.35$  is assumed throughout. Panel A shows this for a marginal tax rate of 0.05, and panel B for a marginal tax rate of 0.25. The bunching share is the fraction of individuals bunching among those with counterfactual earnings lying within the bunching region, and the de-bunching share is the fraction of those bunching at the kink in period 1 that de-bunch in period 2. For more details, see Section 2.7.3.



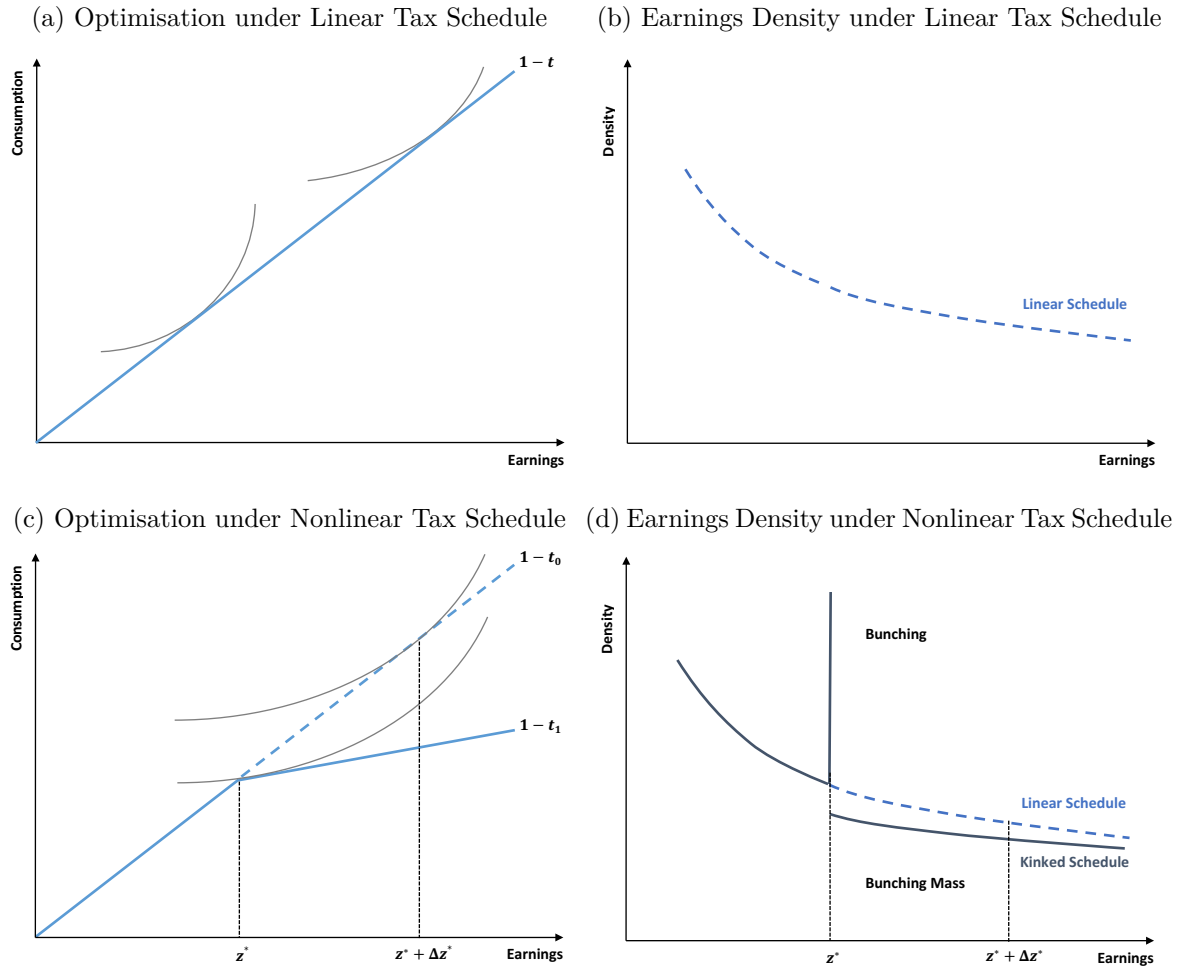
# Main Figures

Figure 2.1: Marginal Tax Rates, 2003-2007



*Notes:* This figure shows the income tax schedule over the periods 2003-2007.


Figure 2.2: Optimisation and Earnings Distributions Under Linear and Nonlinear Tax Schedules




*Notes:* These figures depict the effect of introducing a convex kink at earnings level  $z^*$ . Figures 2.2a and 2.2c present budget set diagrams for linear and nonlinear tax schedules respectively, and Figures 2.2b and 2.2d the corresponding earnings densities. For a detailed explanation, see Section 2.4.1.

Figure 2.3: How Frictions Attenuate the Bunching Region

(a) Frictions Attenuating Bunching

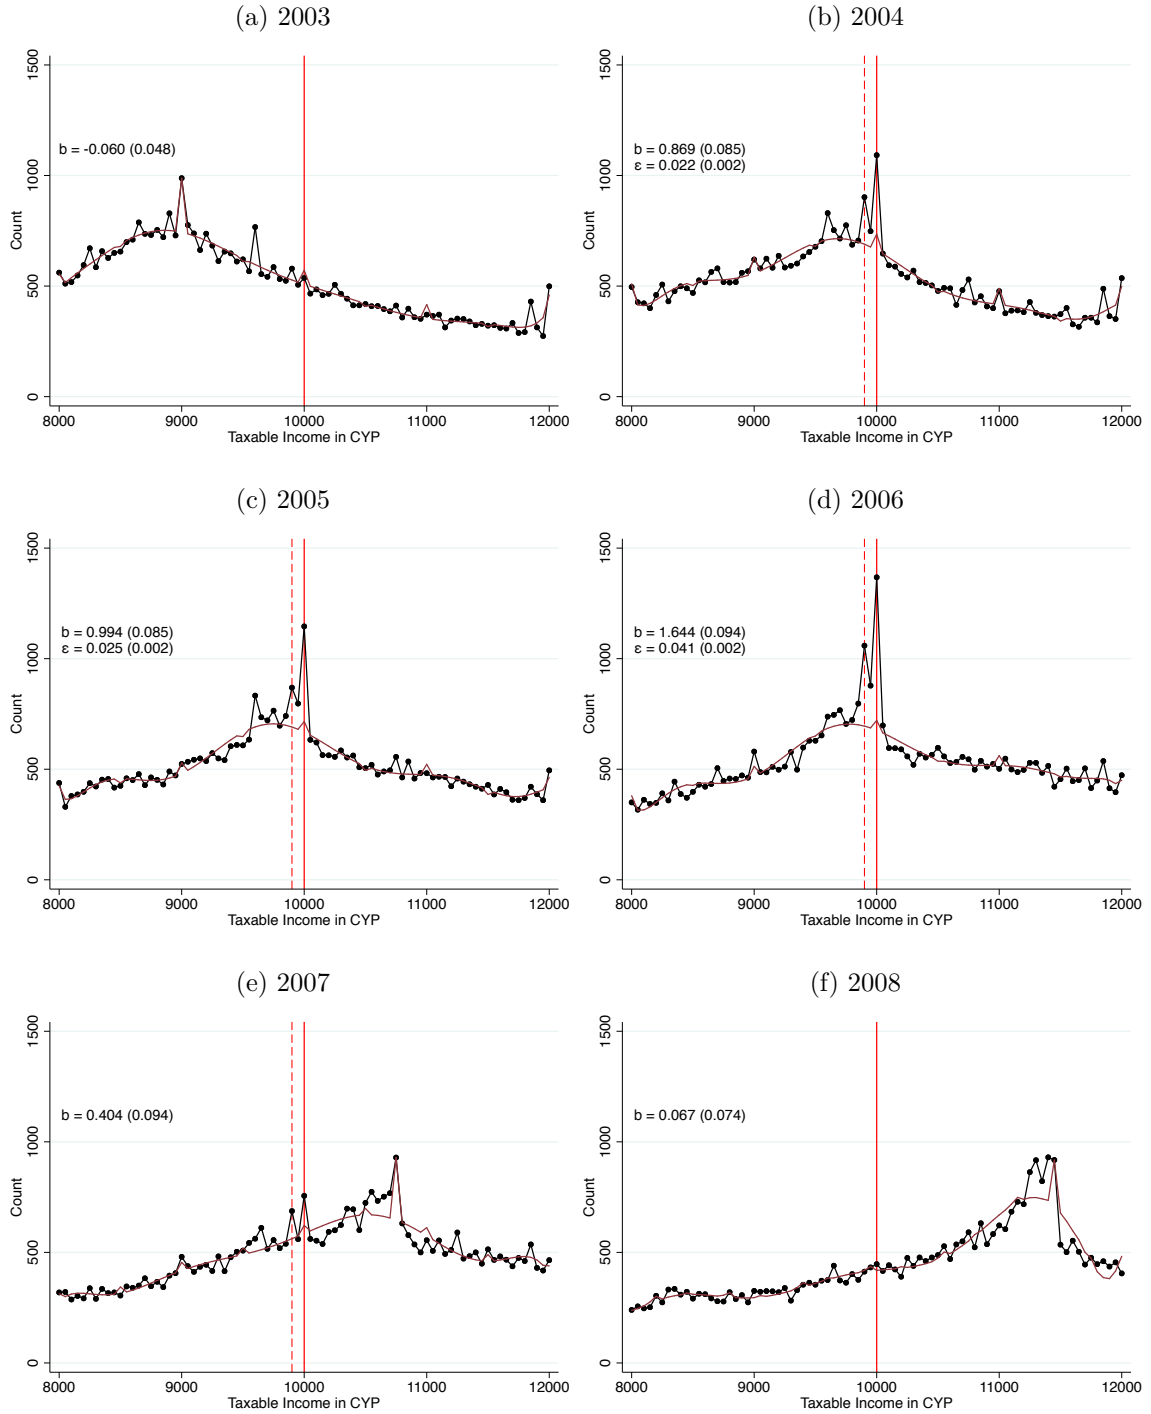


(b) Frictions Attenuating De-Bunching



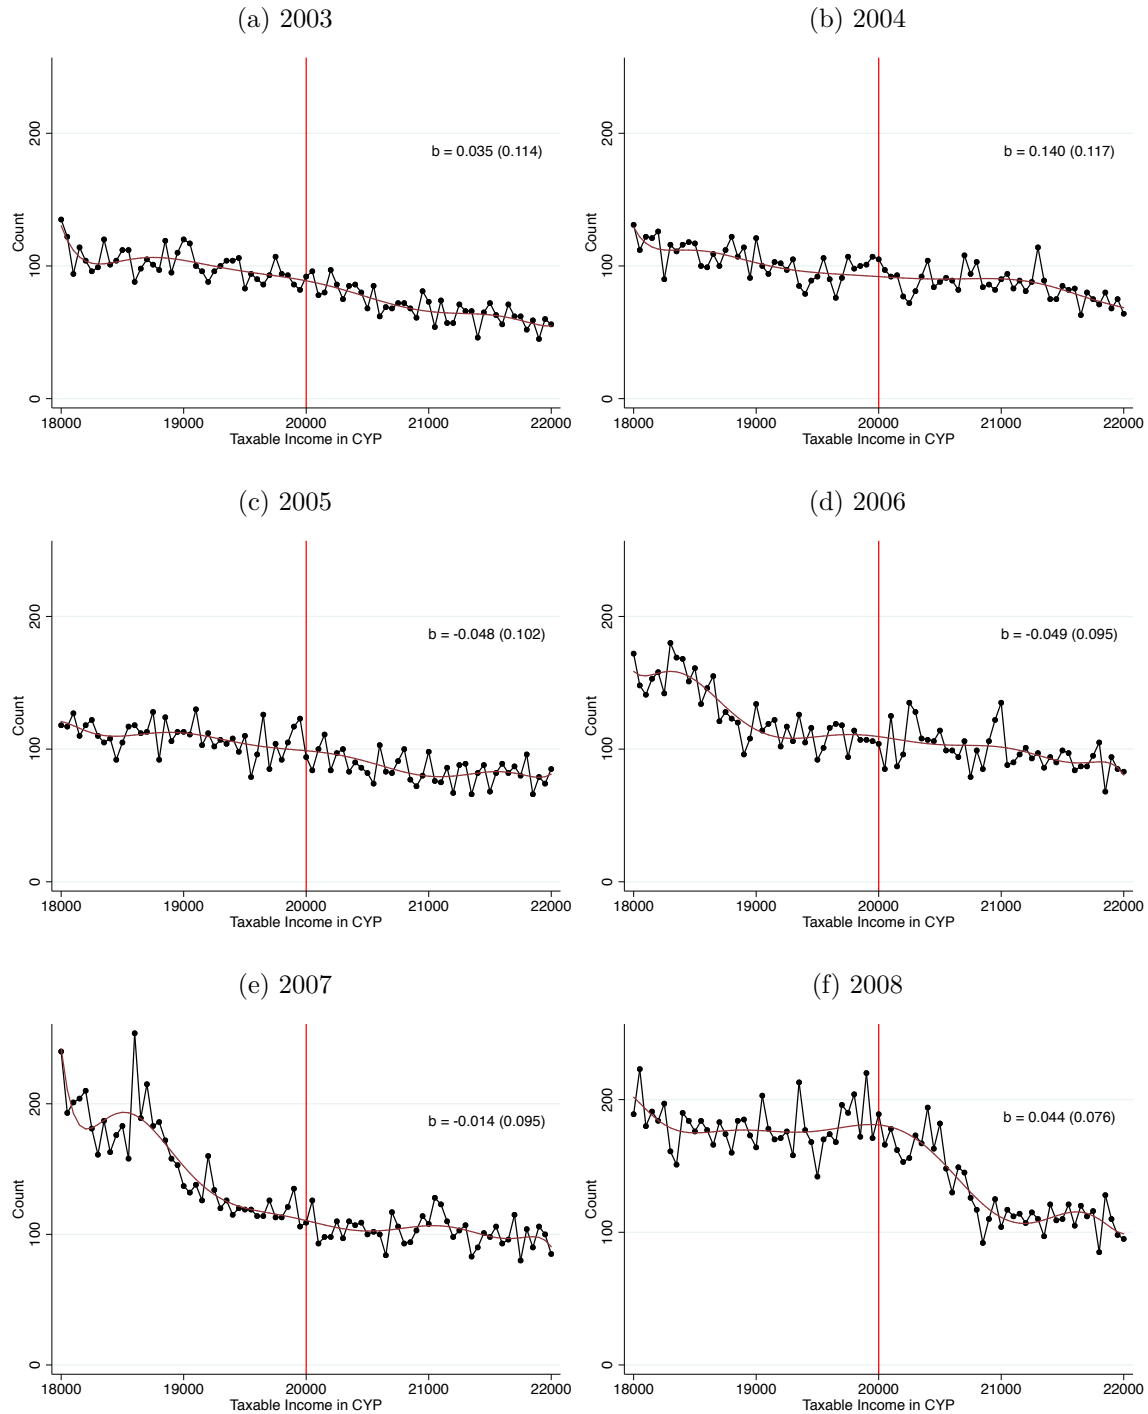
*Notes:* This figure shows how frictions attenuate the bunching region. (a) shows how bunching is attenuated when a kink is in place, and (b) how de-bunching is attenuated when a kink is removed. Each diagram is explained in detail in Section 2.4.4.

Figure 2.4: Bunching of Salary Earners at the CYP 10,000 Kink



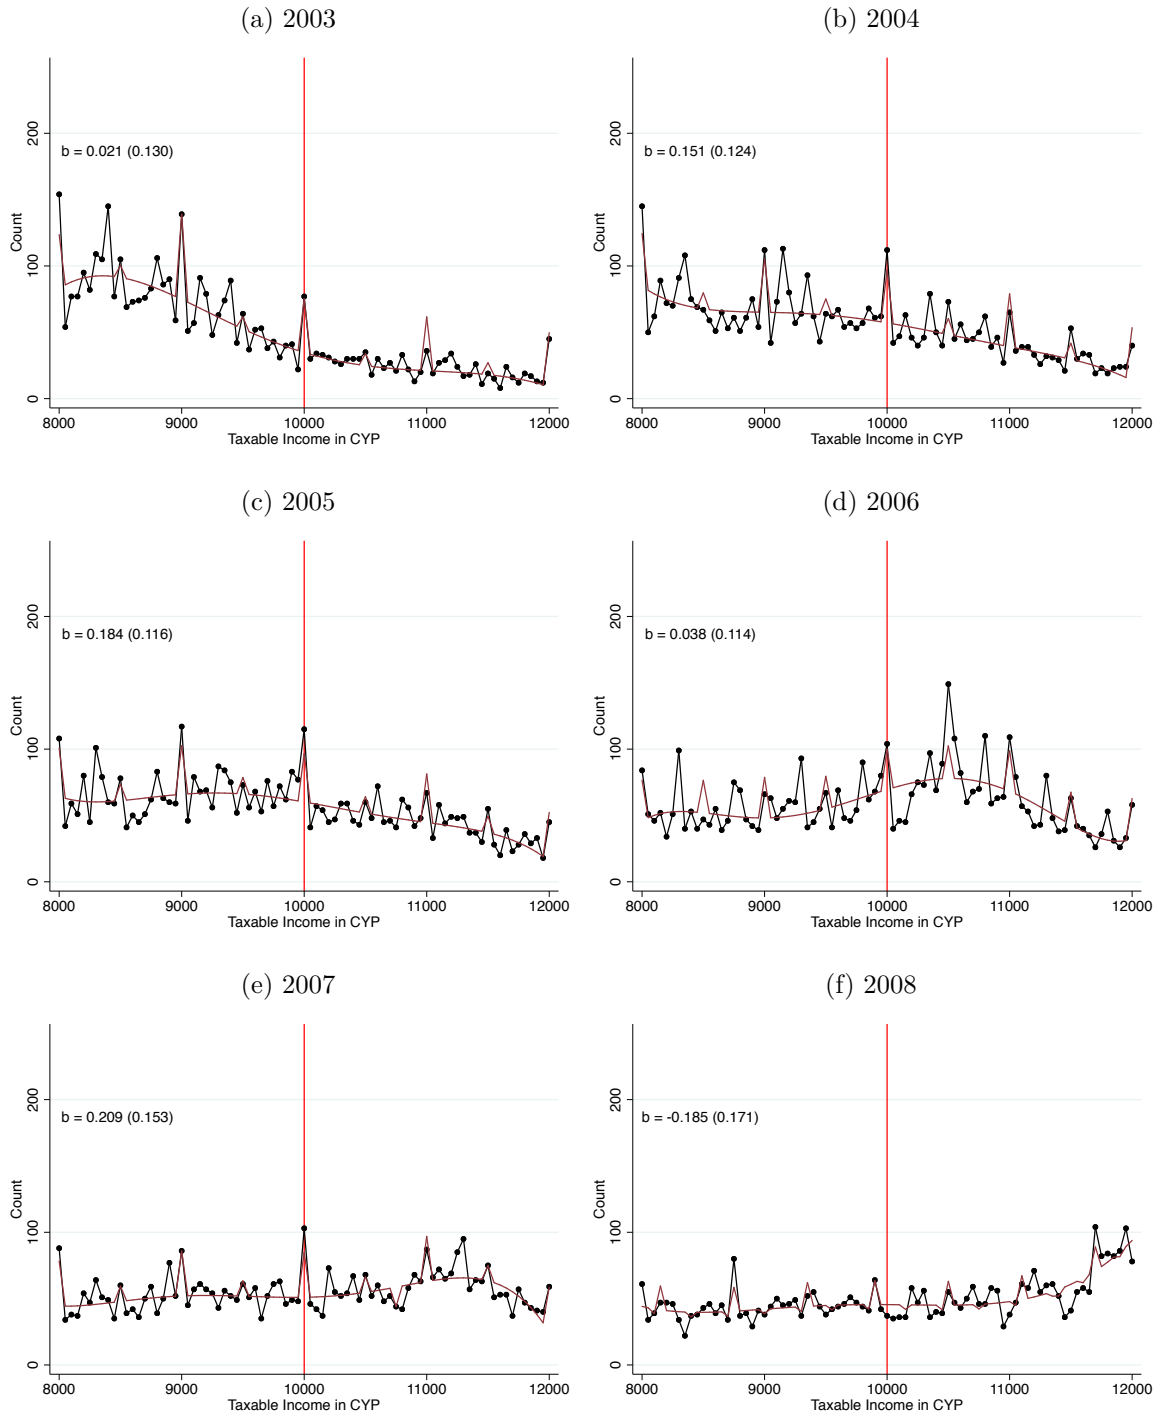
*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 10,000 level of taxable income between 2003-2008. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. In each sub-figure, solid vertical lines mark the CYP 10,000 level, while dashed lines demarcate the lower end of the bunching region. The sub-figures include the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\epsilon$ . Bootstrapped standard errors are in parentheses.

Figure 2.5: Bunching of Salary Earners at the CYP 20,000 Kink



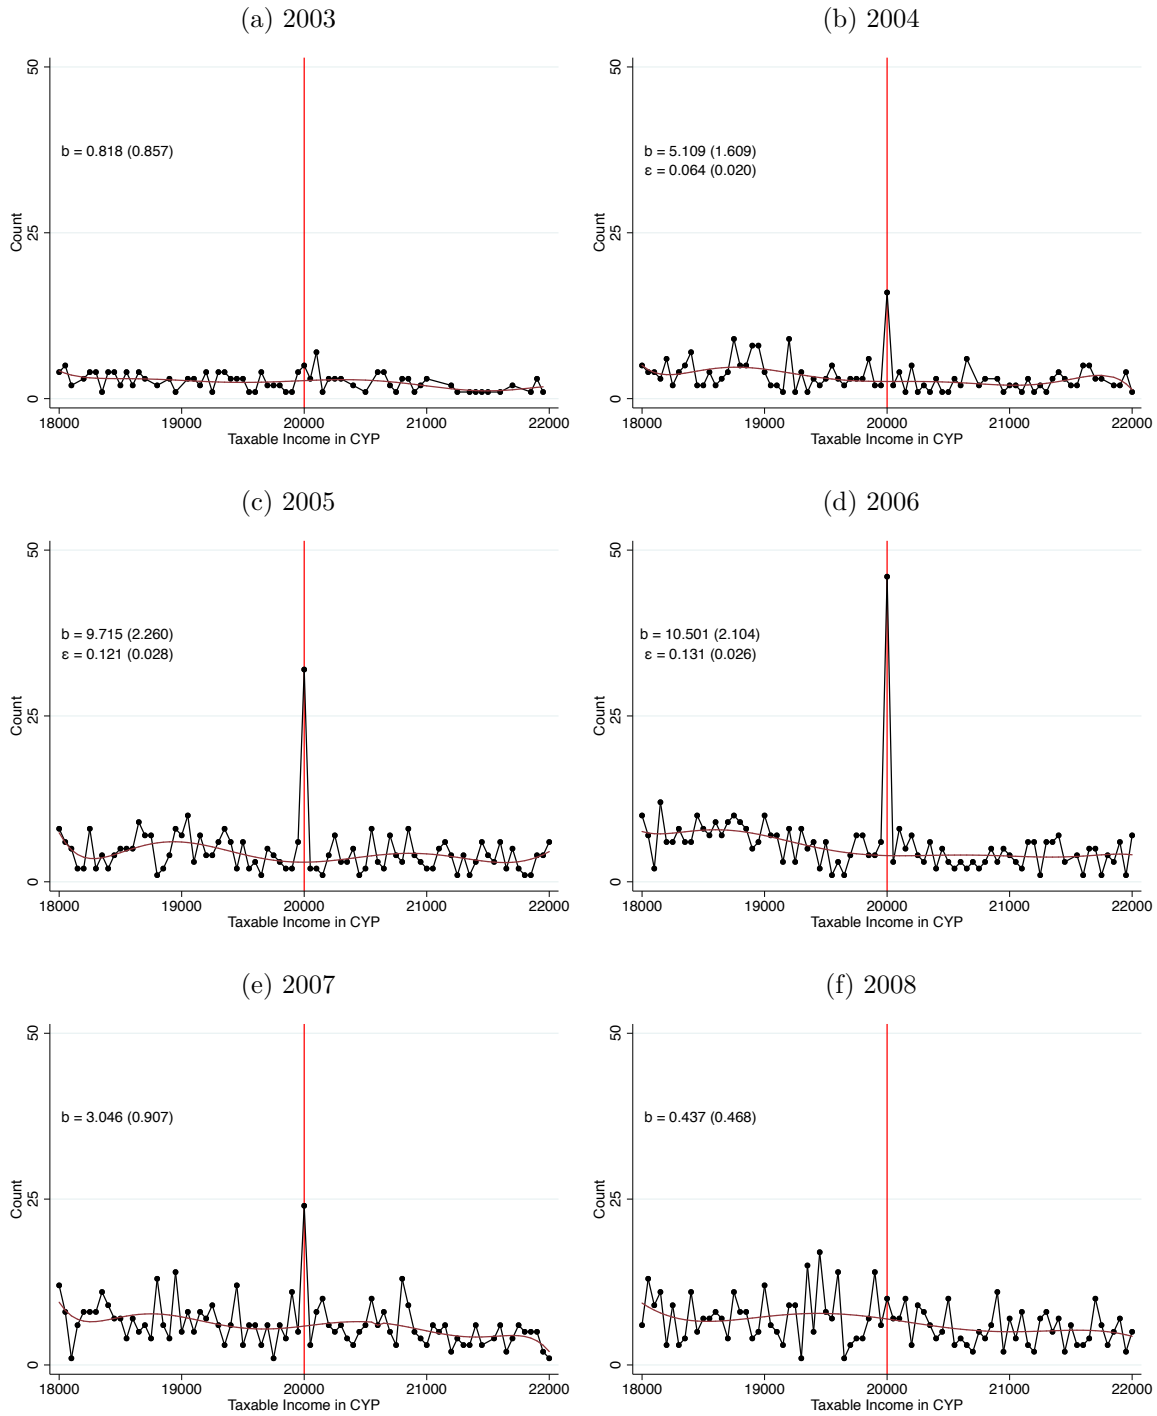
*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 20,000 level of taxable income between 2003-2008. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 20,000 in 2004, and to its removal in 2007. Each sub-figure marks the CYP 20,000 level with a solid vertical line, and includes the estimated counterfactual distribution and the normalised excess bunching mass  $b$ . Bootstrapped standard errors are in parentheses.

Figure 2.6: Bunching of Self-employed at the CYP 10,000 Kink



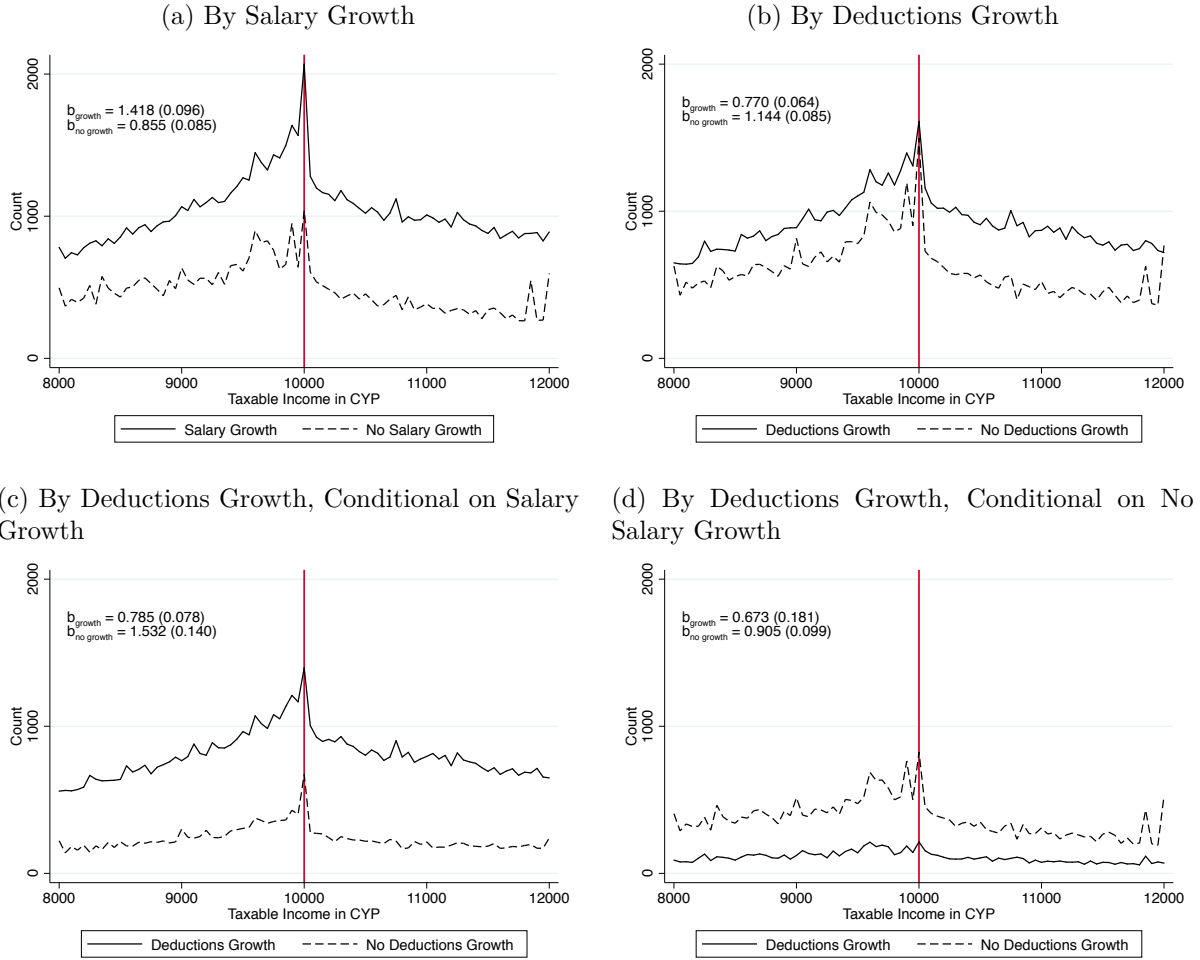
*Notes:* This figure shows the bunching dynamics of the self-employed around the CYP 10,000 level of taxable income between 2003-2008. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. Each sub-figure marks the CYP 10,000 level with a solid vertical line, and includes the estimated counterfactual distribution and the normalised excess bunching mass  $b$ . Bootstrapped standard errors are in parentheses.

Figure 2.7: Bunching of Self-employed at the CYP 20,000 Kink



*Notes:* This figure shows the bunching dynamics of the self-employed around the CYP 20,000 level of taxable income between 2003-2008. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 20,000 in 2004, and to its removal in 2007. Each sub-figure marks the CYP 20,000 level with a solid vertical line, and includes the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\varepsilon$ . Bootstrapped standard errors are in parentheses.

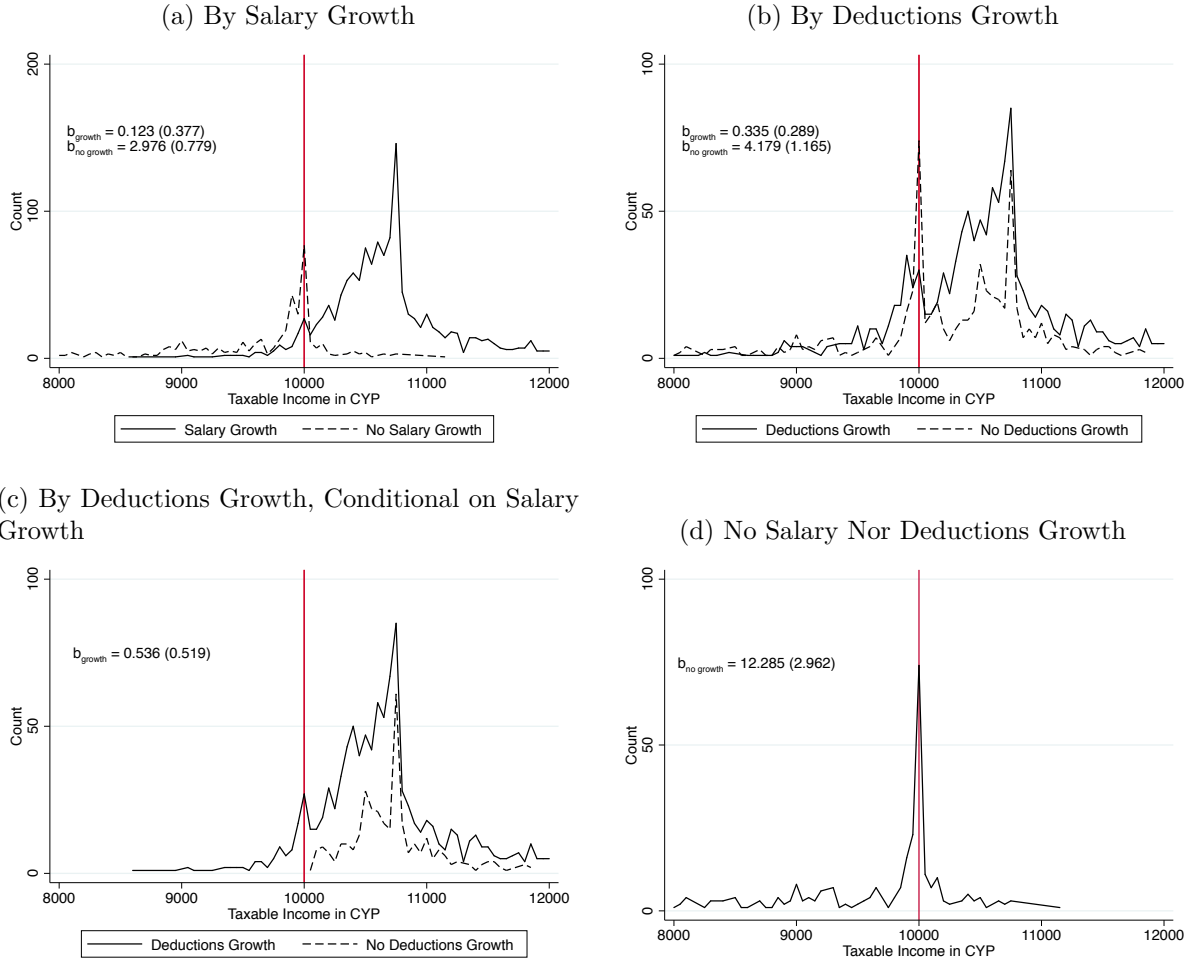
Figure 2.8: Bunching Patterns Among Salary Earning Potential Bunchers at CYP 10,000 Pooled Over 2004-06



*Notes:* This figure shows bunching patterns for the sample of potential bunchers among salary earners around the CYP 10,000 kink, pooled over 2004-2006. The sub-figures split the sample according to whether an individual exhibits (a) salary growth, (b) deductions growth, (c) deductions growth conditional on having positive salary growth, and (d) deductions growth conditional on having no salary growth. Each sub-figure also reports, for each sub-sample, the estimate for the normalised excess mass  $b$  and bootstrapped standard errors. For more details, see Section 2.6.

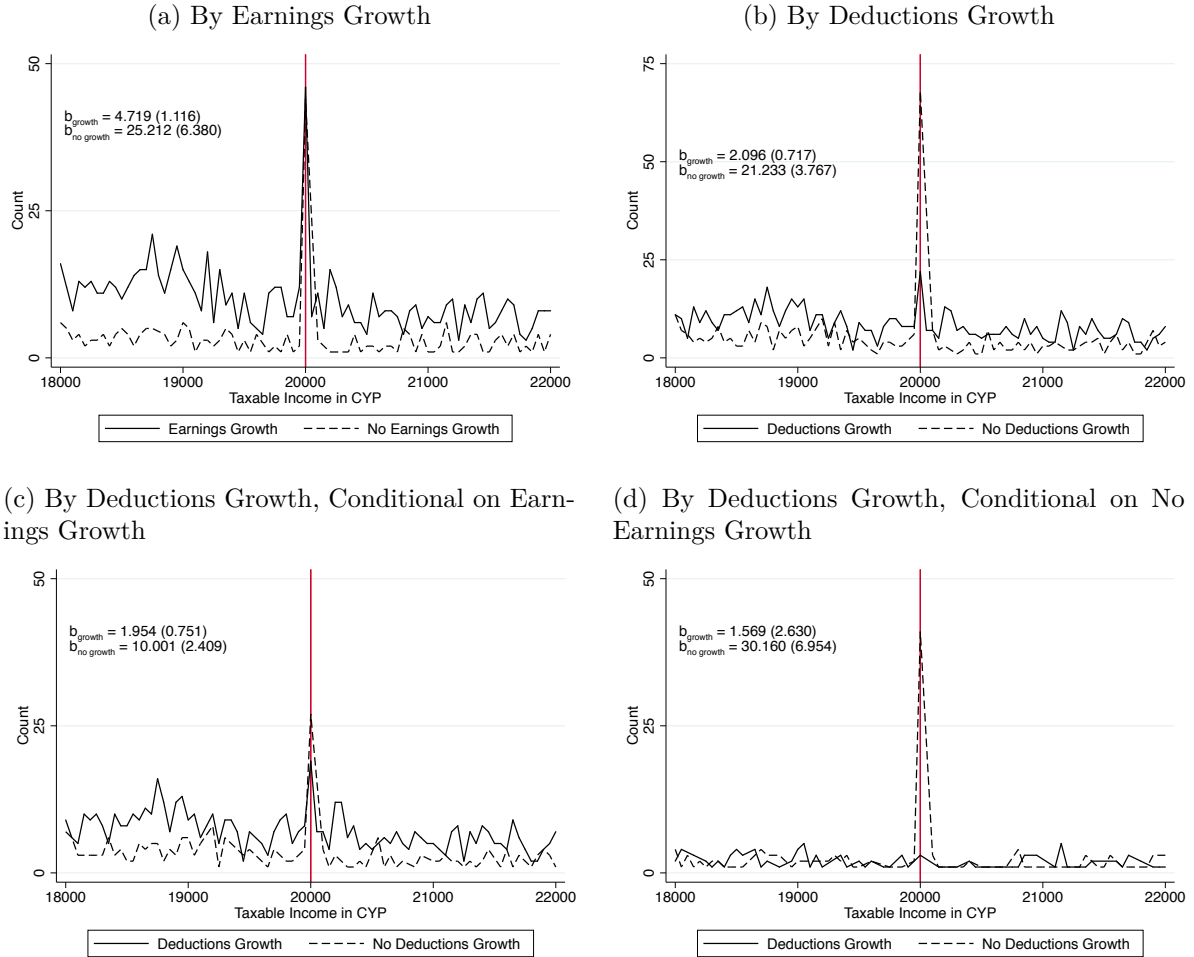


Figure 2.9: De-Bunching Patterns in 2007 Among Salary Earners Bunching at CYP 10,000 by 2006



*Notes:* This figure shows the 2007 de-bunching patterns for the sample of salary earners bunching at the CYP 10,000 kink by 2006. The sub-figures split the sample according to whether an individual exhibits (a) salary growth, (b) deductions growth, (c) deductions growth conditional on having positive salary growth, and (d) deductions growth conditional on having no salary growth. Each sub-figure also reports, for each sub-sample, the estimate for the normalised excess mass  $b$  and bootstrapped standard errors. For more details, see Section 2.6.

Figure 2.10: Bunching Patterns Among Self-Employed Potential Bunchers at CYP 20,000 Pooled Over 2004-06



*Notes:* This figure shows bunching patterns for the sample of potential bunchers among self-employed workers earners around the CYP 20,000 kink, pooled over 2004-2006. The sub-figures split the sample according to whether an individual exhibits (a) earnings growth, (b) deductions growth, (c) deductions growth conditional on having positive earnings growth, and (d) deductions growth conditional on having no earnings growth. Each sub-figure also reports, for each sub-sample, the estimate for the normalised excess mass  $b$  and bootstrapped standard errors. For more details, see Section 2.6.

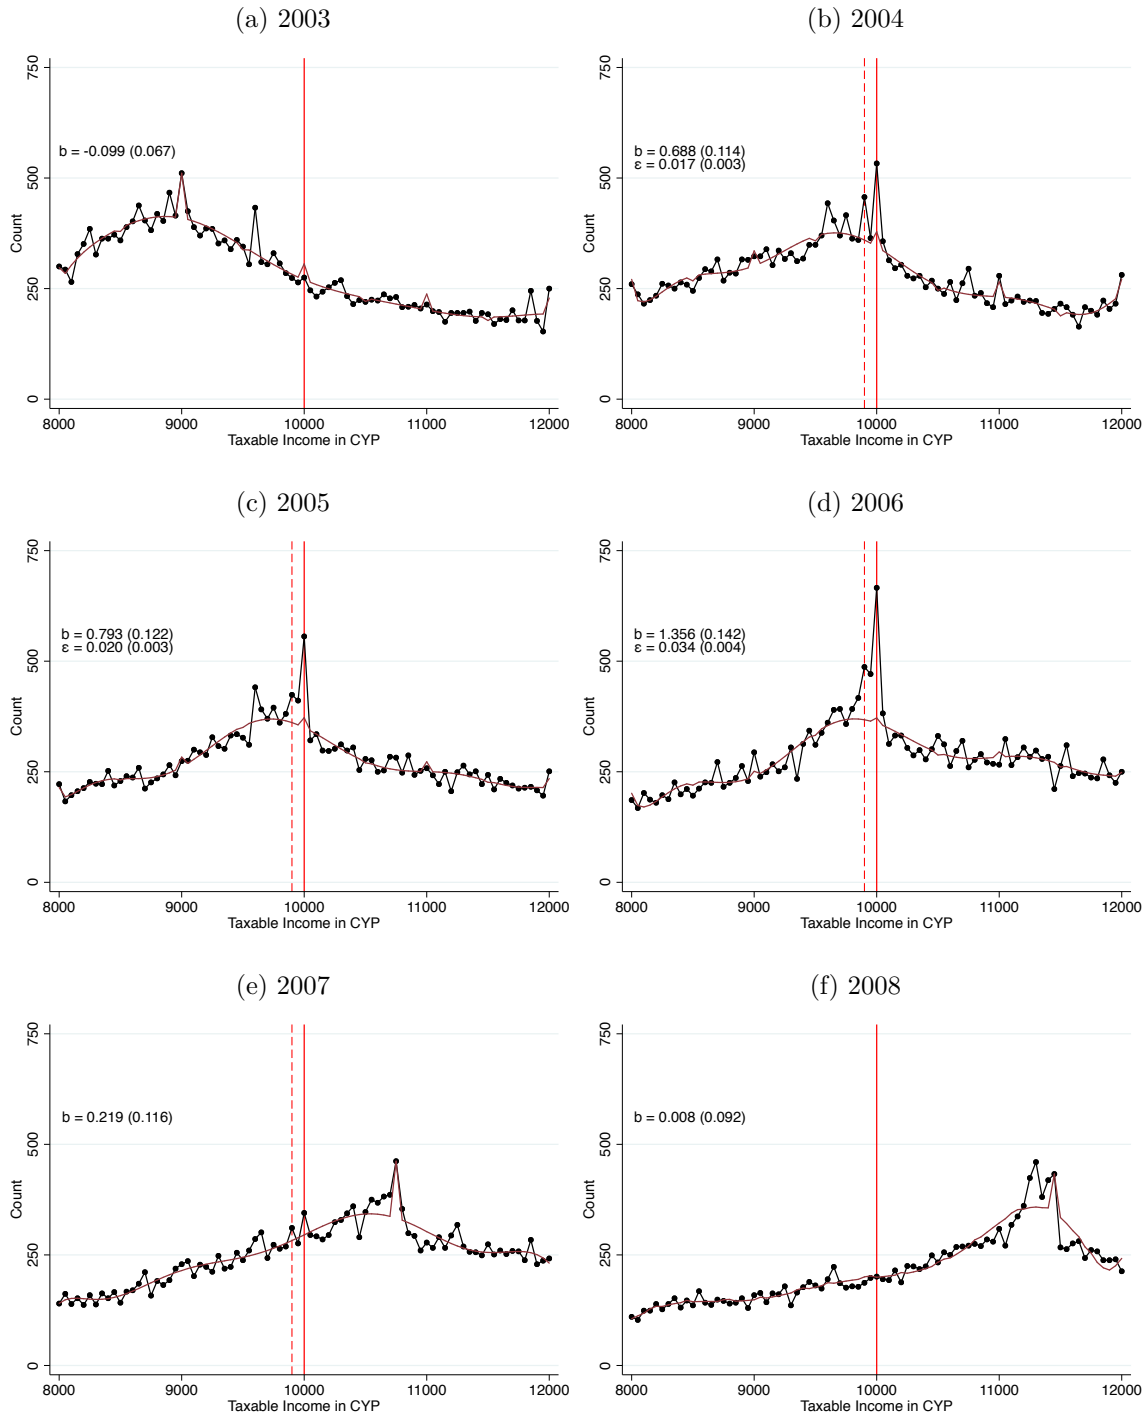
## Appendix 2A: Further Results

Table 2A.1: Robustness Checks for Bunching Estimates

	$b_{2003}$	$b_{2004}$	$b_{2005}$	$b_{2006}$	$b_{2007}$	$b_{2008}$
<b>Panel A: Salary Earners at the CYP 10,000 Kink</b>						
Bandwidth: 1000	0.072 (0.065)	0.912 (0.137)	1.076 (0.126)	1.556 (0.138)	0.630 (0.152)	0.070 (0.071)
Bin Width: 100	0.021 (0.044)	0.512 (0.073)	0.612 (0.071)	0.958 (0.072)	0.285 (0.081)	0.019 (0.051)
Poly. Order: 6	-0.068 (0.054)	0.848 (0.089)	0.993 (0.093)	1.522 (0.096)	0.555 (0.101)	0.044 (0.070)
<b>Panel B: Salary Earners at the CYP 20,000 Kink</b>						
Bandwidth: 1000	-0.073 (0.092)	0.021 (0.119)	-0.152 (0.092)	-0.222 (0.192)	0.001 (0.117)	0.036 (0.114)
Bin Size: 100	-0.151 (0.174)	0.181 (0.127)	0.050 (0.086)	-0.020 (0.069)	0.022 (0.104)	-0.009 (0.072)
Poly. Order: 6	-0.073 (0.117)	0.021 (0.119)	-0.152 (0.092)	-0.222 (0.093)	-0.001 (0.127)	0.036 (0.114)
<b>Panel C: Self-Employed at the CYP 10,000 Kink</b>						
Bandwidth: 1000	-0.091 (0.136)	0.184 (0.152)	0.159 (0.151)	0.086 (0.144)	0.288 (0.182)	-0.151 (0.235)
Bin Size: 100	0.048 (0.120)	0.111 (0.116)	0.074 (0.090)	-0.206 (0.066)	0.105 (0.115)	-0.207 (0.128)
Poly. Order: 6	-0.017 (0.130)	0.189 (0.132)	0.141 (0.127)	0.157 (0.143)	0.229 (0.154)	-0.161 (0.211)
<b>Panel D: Self-Employed at the CYP 20,000 Kink</b>						
Bandwidth: 1000	0.832 (0.845)	4.572 (1.621)	9.530 (2.993)	10.195 (1.767)	2.864 (1.102)	0.478 (0.466)
Bin Size: 100	0.758 (0.583)	3.904 (2.392)	5.646 (1.493)	8.399 (3.298)	2.606 (1.250)	0.105 (0.305)
Poly. Order: 6	0.867 (0.752)	4.798 (1.940)	6.975 (3.161)	8.279 (2.147)	3.642 (1.162)	0.360 (0.495)

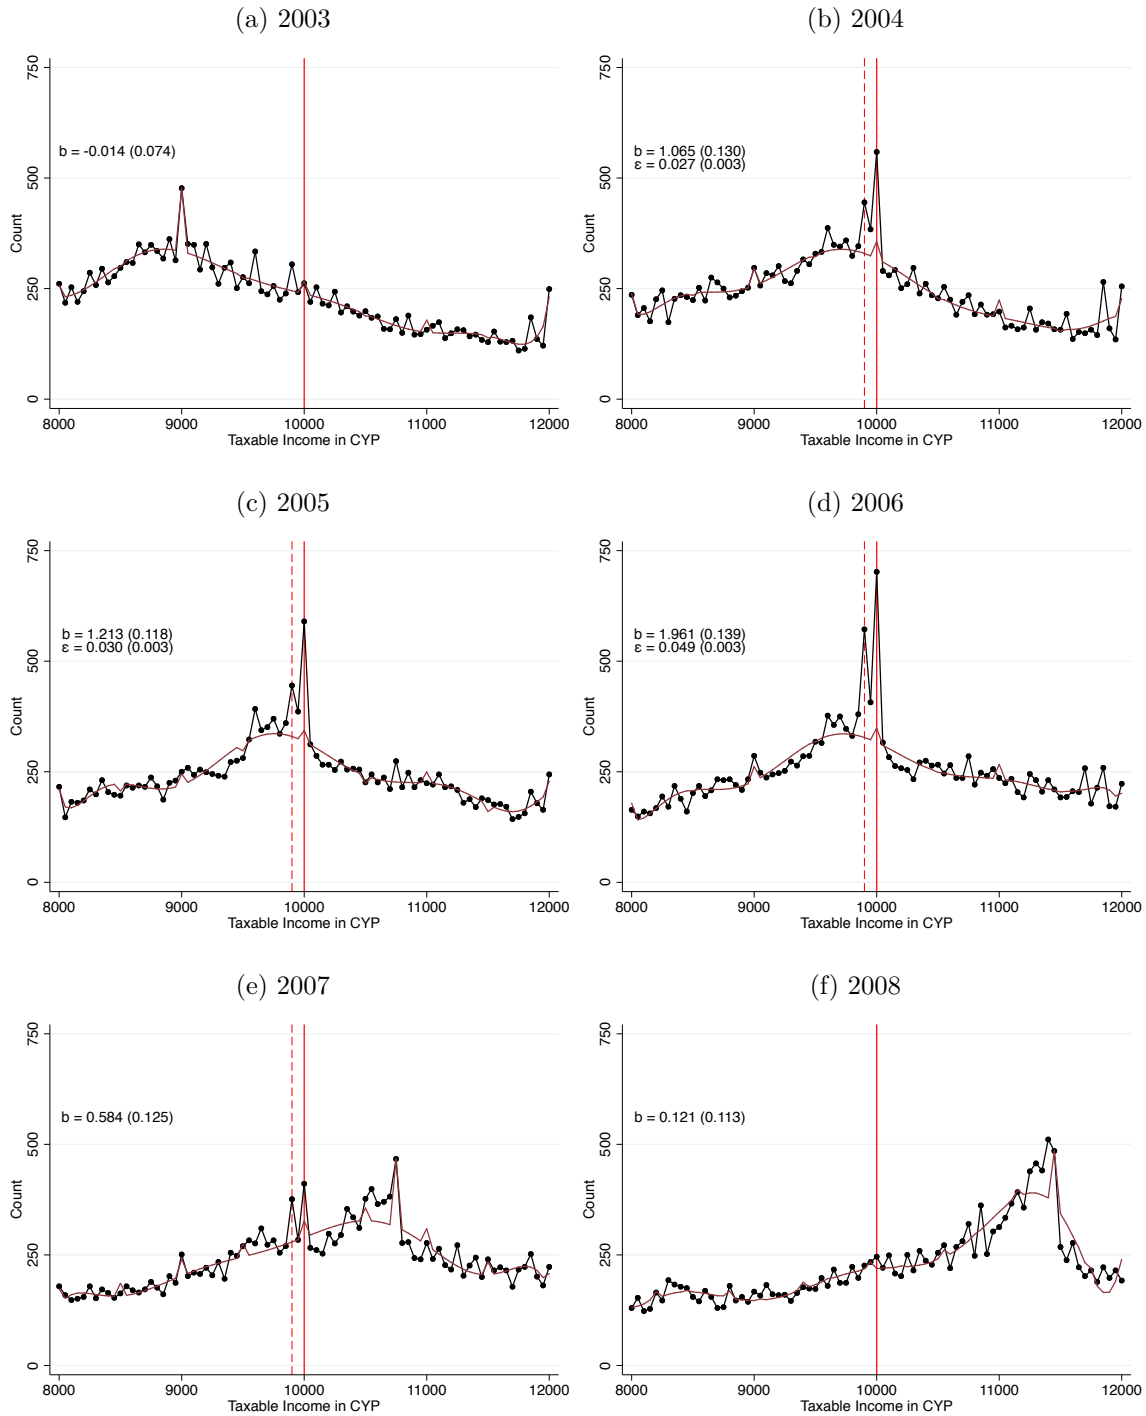
*Notes:* This table shows estimates of the yearly normalised excess mass  $b$  for a different bandwidth, bin width and polynomial order to those used in the main specification. Bootstrapped standard errors are shown in parentheses.

Figure 2A.1: Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Ages 25-39



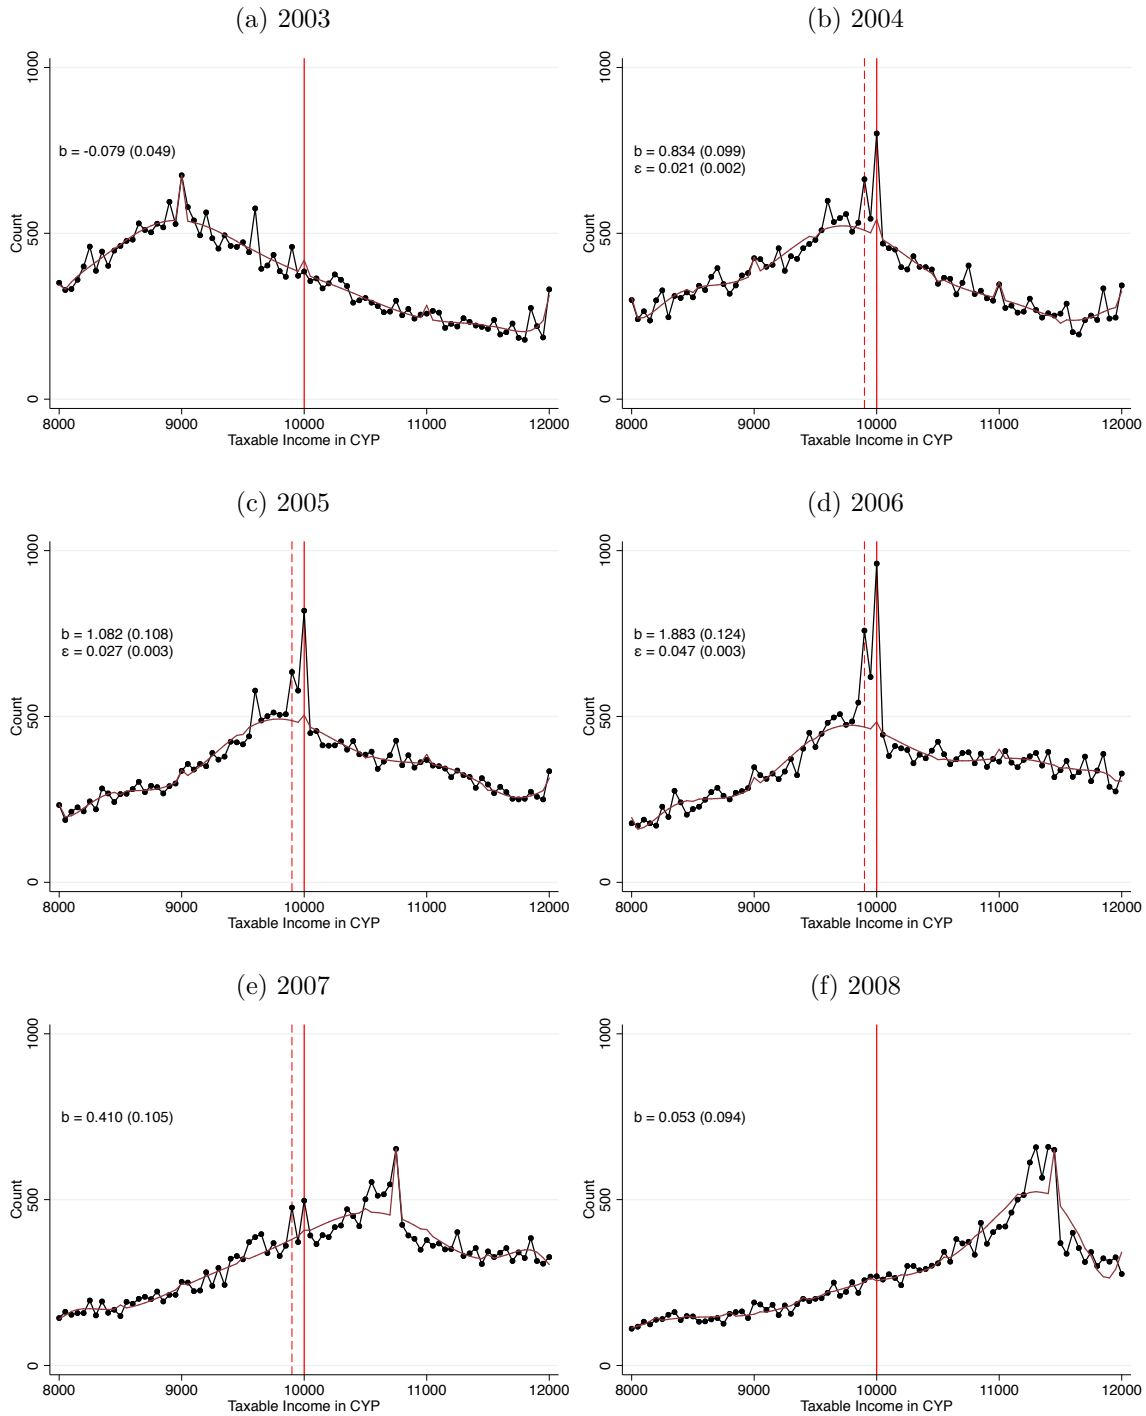
*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 10,000 level of taxable income between 2003-2008, restricting the sample to those aged 25-39. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. In each sub-figure, solid vertical lines mark the CYP 10,000 level, while dashed lines demarcate the lower end of the bunching region. The sub-figures include the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\varepsilon$ . Bootstrapped standard errors are in parentheses.

Figure 2A.2: Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Ages 40-54



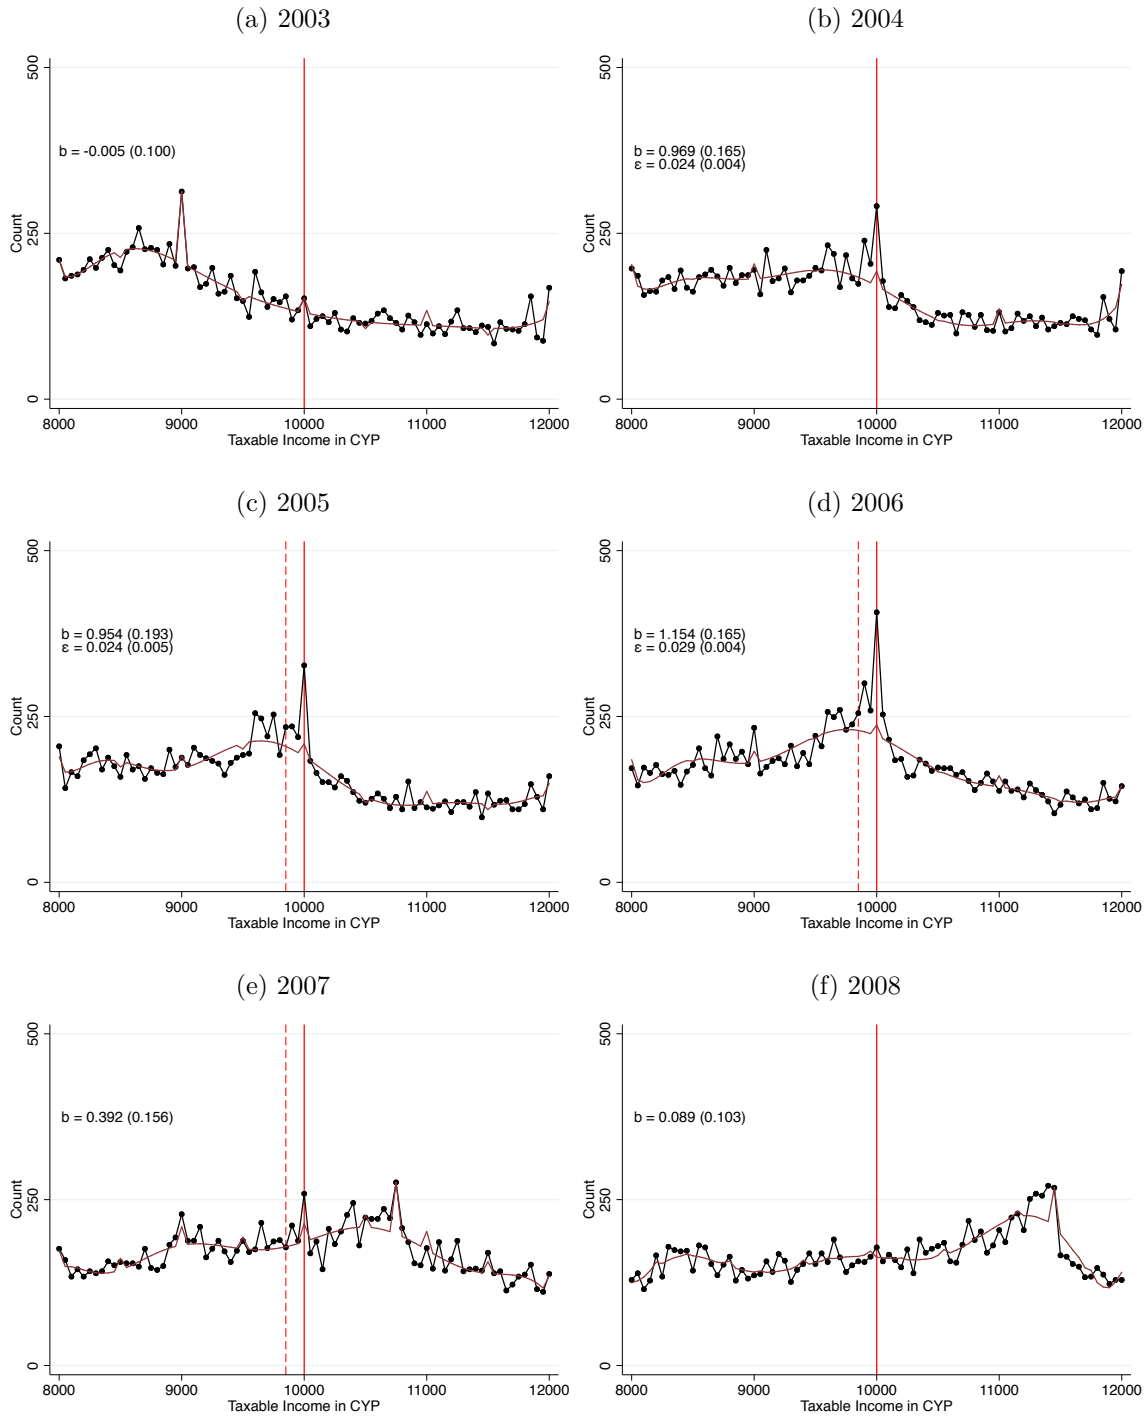
*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 10,000 level of taxable income between 2003-2008, restricting the sample to those aged 40-54. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. In each sub-figure, solid vertical lines mark the CYP 10,000 level, while dashed lines demarcate the lower end of the bunching region. The sub-figures include the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\epsilon$ . Bootstrapped standard errors are in parentheses.

Figure 2A.3: Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Males



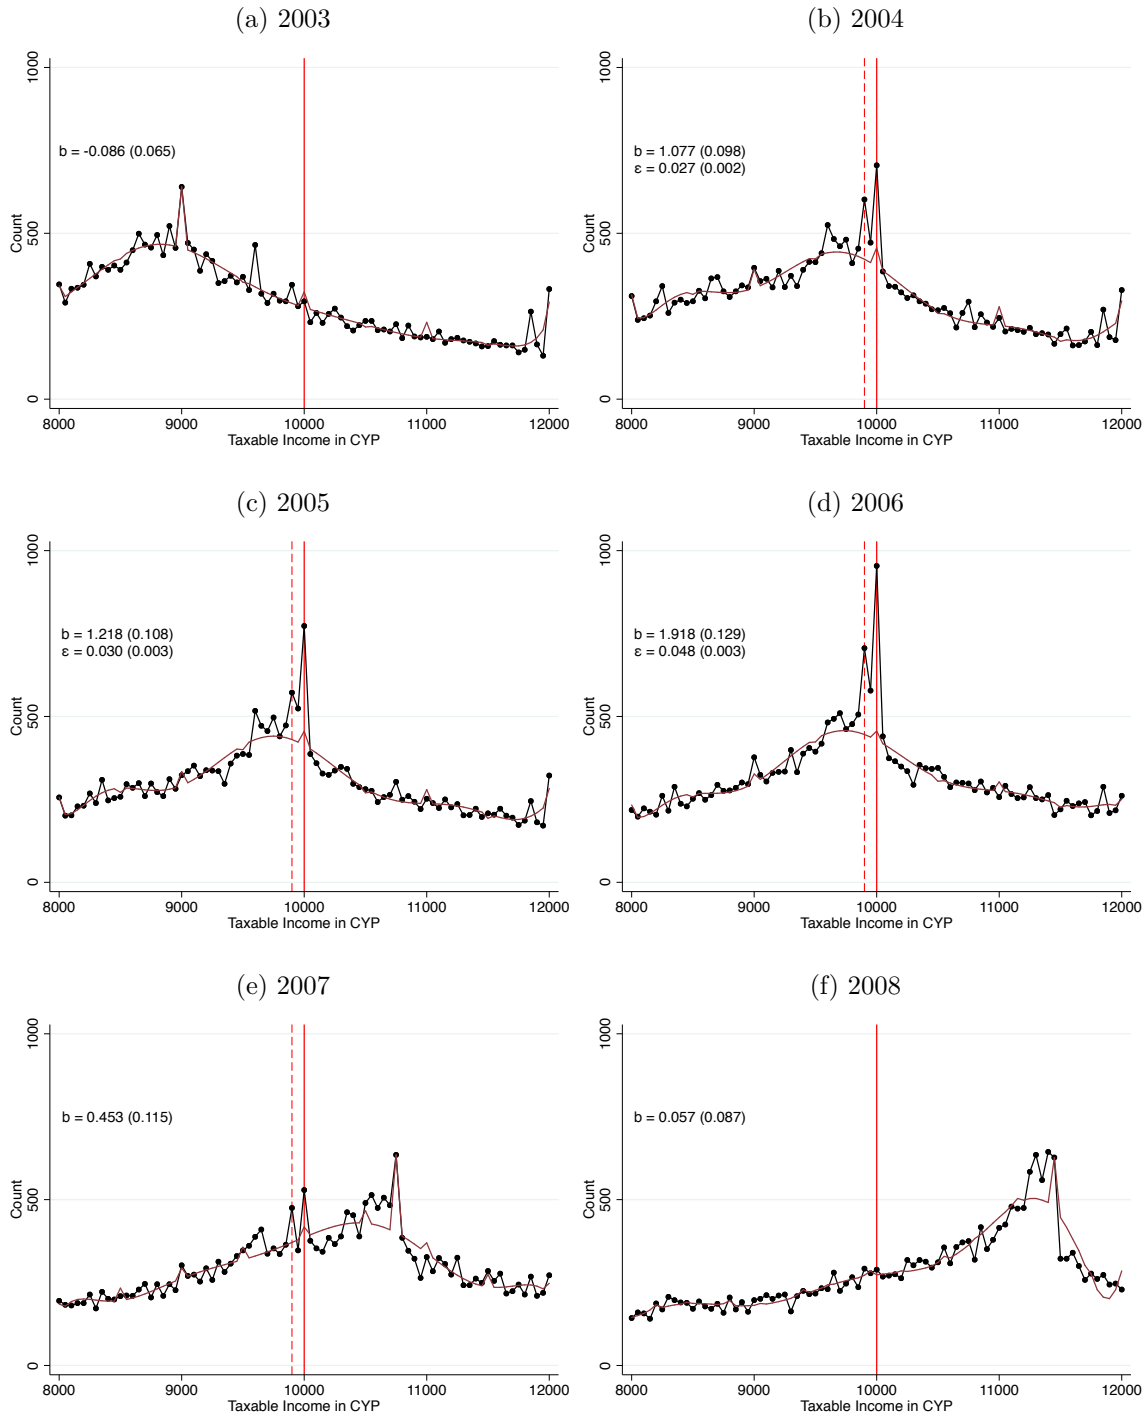
*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 10,000 level of taxable income between 2003-2008, restricting the sample to males. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. In each sub-figure, solid vertical lines mark the CYP 10,000 level, while dashed lines demarcate the lower end of the bunching region. The sub-figures include the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\epsilon$ . Bootstrapped standard errors are in parentheses.

Figure 2A.4: Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Females



*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 10,000 level of taxable income between 2003-2008, restricting the sample to females. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. In each sub-figure, solid vertical lines mark the CYP 10,000 level, while dashed lines demarcate the lower end of the bunching region. The sub-figures include the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\varepsilon$ . Bootstrapped standard errors are in parentheses.

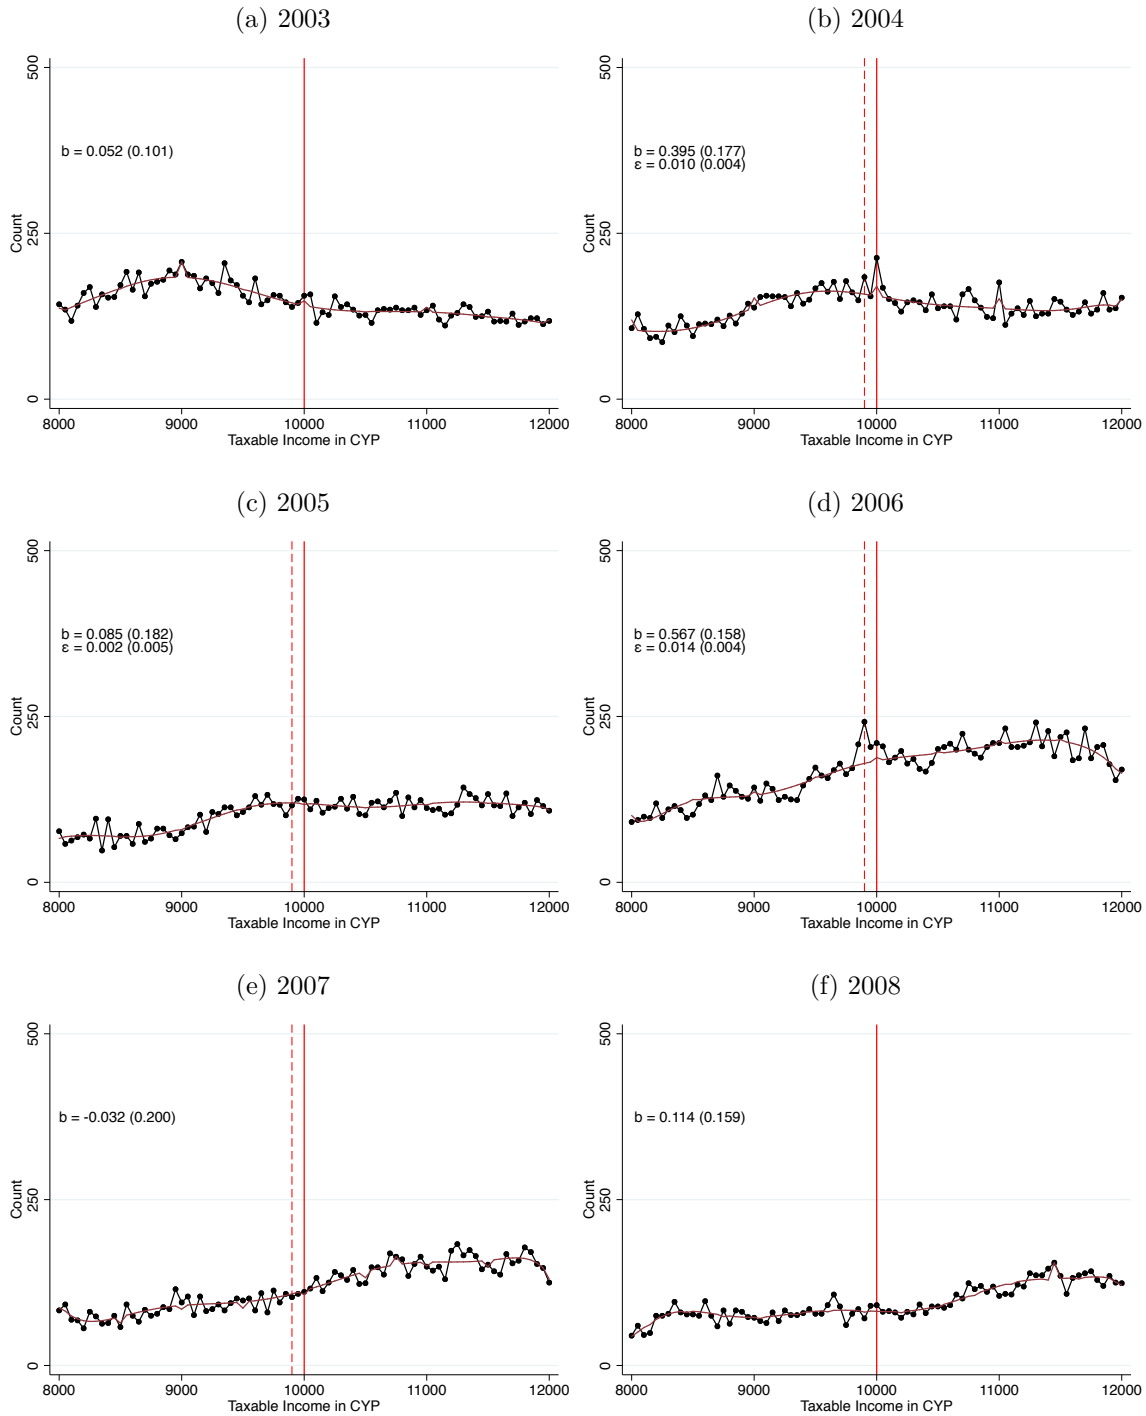
Figure 2A.5: Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink  
Workers Not in Highly Unionised Sectors



*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 10,000 level of taxable income between 2003-2008, restricting the sample to those not in highly unionised sectors. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. In each sub-figure, solid vertical lines mark the CYP 10,000 level, while dashed lines demarcate the lower end of the bunching region. The sub-figures include the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\epsilon$ . Bootstrapped standard errors are in parentheses.



Figure 2A.6: Heterogeneity Analysis: Bunching of Salary Earners at the CYP 10,000 Kink Among Workers in Highly Unionised Sectors



*Notes:* This figure shows the bunching dynamics of salary earners around the CYP 10,000 level of taxable income between 2003-2008, restricting the sample to those in highly unionised sectors. It plots the empirical distribution in bins of width CYP 50 by year, and shows how taxable incomes adjust to the introduction of a kink at CYP 10,000 in 2004, and to its removal in 2007. In each sub-figure, solid vertical lines mark the CYP 10,000 level, while dashed lines demarcate the lower end of the bunching region. The sub-figures include the estimated counterfactual distribution, the normalised excess bunching mass  $b$ , and for years 2004-2006, the observed elasticity  $\epsilon$ . Bootstrapped standard errors are in parentheses.

## Appendix 2B: A Potential Extension to the Model

In this appendix I present a theoretical extension of this paper's model. While this allows for more dynamics, it puts significantly more structure on the model, and greatly increases the bunching moments required for estimation. For this reason, it cannot be used for direct estimation in the current paper, but serves as an illustration for future research.

Following the same setup as in Section 2.5, consider the case where the adjustment cost now also includes a Calvo element. In each period  $t$ , workers draw an adjustment cost from the distribution  $\{0, \phi\}$ . With probability  $\pi_t$ , they face a finite adjustment cost of  $\phi$ , and with probability  $1 - \pi_t$ , they face no cost.

In 2004, the first year the kink is introduced, the bunching mass is:

$$B_{04} = \int_{\underline{z}}^{z^* + \Delta z^*} h_0(z) dz + (1 - \pi_{04}) \int_{z^*}^{\underline{z}} h_0(z) dz \quad (2B.1)$$

This consists of two parts. The first describes individuals with counterfactual earnings  $z \in (\underline{z}, z^* + \Delta z^*]$ , who always respond because the utility gain always outweighs the adjustment cost. The second includes those coming from  $z \in (z^*, \underline{z}]$ , who only adjust when they draw a zero adjustment cost. Otherwise, the utility gain from adjusting does not outweigh the cost. They therefore contribute to the 2004 bunching mass with probability  $1 - \pi_{04}$ , the probability of facing no cost in that year.

Similarly, the 2005 bunching mass is:

$$B_{05} = \int_{\underline{z}}^{z^* + \Delta z^*} h_0(z) dz + (1 - \pi_{04}) \int_{z^*}^{\underline{z}} h_0(z) dz + \pi_{04}(1 - \pi_{05}) \int_{z^*}^{\underline{z}} h_0(z) dz \quad (2B.2)$$

This consists of three parts. The first two are those who bunched in 2004, and the third represents new bunchers in 2005. These are workers coming from  $z \in (z^*, \underline{z}]$  who faced an adjustment cost in 2004 and did not bunch (with probability  $\pi_{04}$ ) but drew a zero cost in 2005 (with probability  $1 - \pi_{05}$ ).

Using the same reasoning, 2006 bunching is:

$$B_{06} = \int_{\underline{z}}^{z^* + \Delta z^*} h_0(z) dz + (1 - \pi_{04}) \int_{z^*}^{\bar{z}} h_0(z) dz + \pi_{04}(1 - \pi_{05}) \int_{z^*}^{\bar{z}} h_0(z) dz + \pi_{04}\pi_{05}(1 - \pi_{06}) \int_{z^*}^{\bar{z}} h_0(z) dz \quad (2B.3)$$

Bunching for year  $t \in \{2004, 2005, 2006\}$  can therefore be summarised as:

$$B_t = \int_{\underline{z}}^{z^* + \Delta z^*} h_0(z) dz + (1 - \prod_{j=1}^t \pi_j) \int_{z^*}^{\bar{z}} h_0(z) dz \quad (2B.4)$$

Finally, bunching in 2007 can be express as:

$$B_{07} = \pi_{07} \int_{\underline{z}}^{\bar{z}} h_0(z) dz + \pi_{07}(1 - \pi_{04}\pi_{05}\pi_{06}) \int_{z^*}^{\bar{z}} h_0(z) dz \quad (2B.5)$$

This again consists of two parts. The first are those coming from  $z \in (\underline{z}, \bar{z}]$ . While they always bunched at the kink in previous years regardless of the adjustment cost, they only de-bunch from the kink if they draw a zero adjustment cost in 2007. Hence, with probability  $\pi_{07}$  they draw a positive cost in 2007 and do not de-bunch. The second component relates to those from  $z \in (z^*, \underline{z}]$ . These would only contribute to residual bunching if they were able to bunch by 2006, with probability  $1 - \pi_{04}\pi_{05}\pi_{06}$ , but not de-bunch in 2007, with probability  $\pi_{07}$ . Lastly, note that those with earnings above  $\bar{z}$  always de-bunch because the gain from doing so always outweighs the adjustment cost.

Estimating this model greatly increases the necessary number of bunching moments.<sup>24</sup> In particular, this model involves estimating six parameters:  $\varepsilon, \phi, \pi_{04}, \pi_{04}\pi_{05}, \pi_{04}\pi_{05}\pi_{06}$  and  $\pi_{04}\pi_{05}\pi_{06}\pi_{07}$ . This is not feasible in this paper and so is not done, as there are only four available moments between 2004-2007. The only way to proceed would be to calibrate the values for two parameters. This would be feasible if the available bunching variation came from a kink that was in place for long enough, say  $\mathcal{T} = 10$  years. In this case, and provided that the graphical evidence revealed that the normalised excess mass by that year had stabilised to some steady-state value, then the observed elasticity in the final year could be interpreted as the long-run, structural elasticity. Further, this would also imply that all potential bunchers would have bunched by then, and so  $\prod_{j=1}^{\mathcal{T}} \pi_j \approx 0$ , providing us

---

<sup>24</sup>Having as many moments as parameters to estimate is not sufficient; there also needs to be enough variation between the moments.

with our second calibrated value. In this paper, this is clearly not satisfied as the observed elasticity in 2006 is much smaller than the estimated structural elasticity, showing that the 2006 bunching is far from its steady-state level.

This extended model has both strengths and weaknesses. Its main strength is that it can capture more nuances in bunching dynamics. Its main weakness is the requirement for more bunching moments. Another major weakness is that the implicit assumption of fixed ability is harder to justify when consider a time horizon long enough to make the calibration of the elasticity and  $\prod_{j=1}^T \pi_j$  feasible.

## Appendix 2C: Collective Bargaining Agreements

Table 2C.1: Collective Bargaining Agreement, Banking Sector

Year	Weekly Hours Winter	Weekly Hours Summer	Salary Increase (%) By Employee Category			
			Clerks	General Staff	Special Staff	Senior Management
2000	38	35	2	2	2	2
2001	38	35	2	2	2	2
2002	38	35	2	2	2	2
2003	38	35	2	2	2	2
2004	38	35	2	2	2	2
2005	38	35	0.6	0.6	0.6	0.6
2006	38	35	0.7	0.7	0.7	0.7
2007	38	35	0.7	0.7	0.7	0.7
2008	38	35	2	2	2	2
2009	38	35	2	2	2	2
2010	38	35	2	2	2	2

*Notes:* The table shows the main terms of the sectoral-level collective bargaining agreement in the banking sector, between 2000-2010. It shows the agreed working weekly hours during winter (September-April) and summer (May-August), as well as the agreed % increase in salaries for each of the following four categories of workers: clerks, general staff, special staff and senior management. Source: Ministry of Labour, Welfare and Social Insurance Archives.

Table 2C.2: Collective Bargaining Agreement, Construction Sector

Year	Weekly Hours	Salary Increase (CYP) By Employee Category			
		Skilled	Semi-Skilled	Unskilled	Apprentice
2000	38	2.80	2.80	2.80	2.80
2001	38	2.80	2.80	2.80	2.80
2002	38	4.00	4.00	4.00	4.00
2003	38	4.00	4.00	4.00	4.00
2004	38	4.00	4.00	4.00	4.00
2005	38	4.00	4.00	4.00	4.00
2006	38	4.75	4.75	4.75	4.75
2007	38	5.25	5.25	5.25	5.25
2008	38	5.00	5.00	5.00	5.00
2009	38	5.55	5.55	5.55	5.55
2010	38	5.70	5.70	5.70	5.70

*Notes:* The table shows the main terms of the sectoral-level collective bargaining agreement in the construction sector, between 2000-2010. It shows the agreed working weekly hours and the increase in salaries for each of the following four categories of workers: skilled, semi-skilled, unskilled (general workers), and apprentices. Amounts are set at the weekly level in CYP. Source: Ministry of Labour, Welfare and Social Insurance Archives.

## Chapter 3

# Enforcement and the Importance of Tax Policy Design for Elasticities: Evidence from the Republic of Cyprus

### 3.1 Introduction

Tax enforcement is an essential dimension of any tax system and has therefore been the subject of an extensive literature. The enforcement environment consists of several components such as auditing, self- versus third-party reporting, penalties and documentation requirements. Each element is arguably important, but in recent decades many countries have moved specifically towards a heavy reliance on stricter documentation requirements and third-party reporting to assure compliance. Given the significant role of this policy dimension, assessing the behavioural responses to the reporting environment is important both to understand the magnitude of misreporting and to be able to invest efficiently in enforcement initiatives.

Moreover, the enforcement environment can have a significant effect on how other traditional policy instruments, such as tax subsidies, affect the behaviour of taxpayers. Understanding the size and composition of such behavioural responses is essential for de-

signing optimal policies but also for evaluating the efficiency of a tax system. If elasticities are not structural parameters determined by preferences, but instead parameters controlled by policy choices of the government, such as the level of enforcement, this affects how we should interpret and utilise such parameters. Further, it affects the external validity of estimates across the literature that are dependent on a specific form of tax policy design.

In this paper, we address these important dimensions of a tax system. We focus on charitable contributions in the Republic of Cyprus,<sup>1</sup> which are subsidised through a tax deduction, and have been subject to varying levels of enforcement. The enforcement environment in this context is characterised by a set of thresholds determining a level of claimed deductions for which no documentation is necessary. For claimed deductions above this level, documentation from a third party must be provided.<sup>2</sup> Within this setting we present three sets of results. We start by identifying substantial reactions to changes in the enforcement environment. Exploiting salary-dependent thresholds governing the documentation requirements for claiming deductions, we show clear discontinuities at exactly these threshold values. Using a regression discontinuity approach, we find that individuals increase reporting by 0.7 pounds when 1 pound more can be claimed without providing documentation from a third party.

Next, we use a reform that retroactively shifted the location of a reporting threshold to separate real from reporting responses. Exploiting the time-profile of responses using bunching techniques, we find that at least 36 percent of responses to such threshold changes are purely changes in reporting behaviour.<sup>3</sup> This separation of reporting and real responses is crucial in a setting where we expect positive externalities from real behaviour, such as expenditures on charitable contributions, investments in education, professional training, retirement savings, etc. In such cases, the goal of the fiscal authority is not only limited to raising tax revenue, but also to encourage real responses through tax incentives.

In the final section, we turn to estimating the elasticity of charitable contributions with

---

<sup>1</sup>Henceforth, we simply use the term Cyprus to refer to the Republic of Cyprus.

<sup>2</sup>Note that this type of enforcement initiative is very closely related to the standard concept of *third party reporting*. While charities are not required to report directly to the tax authority on behalf of the tax-payer, the authority here still requires documentation from a third party, but collected and filed by the tax-payer himself.

<sup>3</sup>This wedge can be driven by over-reporting due to the discontinuity in detection probabilities, or under-reporting due to discontinuity in the filing hassle cost such thresholds generate.



respect to the price of giving. Using quasi-experimental variation in tax prices generated by reforms to the income tax schedule, we find a price elasticity of about  $-0.5$ . We then show that the presence of reporting thresholds has a very significant effect on our estimates. Price elasticities double if we disregard individuals who bunch around these thresholds. Further, we estimate elasticities across different sub-periods with varying enforcement levels, and show suggestive evidence that this parameter is sensitive to even minor changes in the strictness of the enforcement environment. Our findings show that elasticities are sensitive to a design feature which is relatively common to many tax systems, reporting thresholds, because people who bunch around these thresholds are very insensitive to price changes.

Our paper contributes to several strands of literature. First, we contribute to a growing literature measuring behavioural responses to enforcement initiatives. Using our unique setting, we can show clear discontinuities around thresholds in the reporting environment. While responses to enforcement initiatives have been documented earlier (Alm et al. 2009; Kleven et al. 2011; Phillips 2014; Skov & Gillitzer 2018)<sup>4</sup> we do not have to rely on comparisons between different types of income or potentially selected audit data. Further, we are able to exploit a unique reform which retroactively changed the reporting environment on earlier years' contributions, to credibly show that a large fraction of responses is pure reporting. Carrillo et al. (2017) also disentangle pure reporting responses by analysing a policy intervention concerning the reporting environment of firms in Ecuador. However, the case of firm profit is very different from our setting, because firms can offset any change in income with changes in costs. We consider deductions in the personal income tax schedule, and hence study a very different tax environment.

Second, we contribute to a small and emerging empirical literature on the importance of tax system design features for estimated elasticities. The concept that features of the tax system, such as the size of the tax base, matters for the size of behavioural elasticities has been addressed theoretically by Slemrod (1994) and Slemrod & Kopczuk (2002). Kopczuk (2005) shows empirically that the elasticity of reported income with respect to tax rates depends on the level of deductions in the tax system. Mishra et al. (2008) look

---

<sup>4</sup>Skov & Gillitzer (2018) consider a similar tax design to us, but examine the effect of moving from no third-party reporting (unlimited reporting) to automatic third-party reporting. Our variation does not change the self-reporting feature of the system, but affects the requirement for receipts and hence examines pure enforcement variation.

at the effect of tariffs on evasion of customs duties and find evidence that this elasticity is affected by characteristics correlated with enforcement possibilities. Fack & Landais (2016) show that both the elasticity of reported income with respect to the tax rate and the elasticity of reported charitable contributions with respect to price are sensitive to the level of enforcement.<sup>5</sup> Much in line with Fack & Landais (2016) we show suggestive evidence that changes in the strictness of the enforcement environment affect estimated elasticities. However, we show such evidence in a situation where only minor changes are made to the reporting environment, where Fack & Landais (2016) use a large reform severely changing reporting standards. These results hence suggest that elasticities might be very sensitive to even small changes in enforcement. Further we include a new dimension of tax system design in this literature, by showing that the existence of reporting thresholds in the tax system directly affects elasticities. This tax design feature is not uncommon - both in the case of charitable contributions but also in other dimensions of tax systems. In the United Kingdom such a threshold exists for on-the-job mileage expenses, Germany has a reporting threshold on work-related expenses and the US has a threshold for charitable contributions.

Lastly, we contribute to a literature estimating the elasticity of charitable contributions with respect to the price of giving. This literature has been very focused on the US case and recently also a few other developed countries (Almunia et al. 2017; Duquette 2016; Fack & Landais 2010).<sup>6</sup> Given the importance of tax system design for elasticities, these reported elasticities might not be very informative for less developed economies or economies generally characterised by more informality or less enforcement. We contribute to this literature by providing estimates for a semi-formal economy with tax rules concerning charitable contributions which are representative of other semi-formal economies.

The rest of the paper is structured as follows: Section 3.2 describes the institutional environment and the data we use in the empirical analysis. Section 3.3 presents our results from the regression discontinuity analysis investigating behavioural responses to enforce-

---

<sup>5</sup>A related paper in this context is Doerrenberg et al. (2017) which documents responsiveness of total deductions to tax changes in Germany. Further they show a significant difference in the elasticity of taxable income and the elasticity of gross income.

<sup>6</sup>This is a non-exhaustive list of contributions to this extensive literature. Other contributions for instance include Adena (2014), Auten et al. (2002), Bakija & Heim (2011) and Karlan & List (2007).

ment. In Section 3.4 we use bunching techniques to separate real and reporting responses, Section 3.5 provides estimates of the tax price elasticity of charitable contributions, and Section 3.6 concludes.

## 3.2 Institutional Context and Data

In this section, we describe the institutional context of charitable giving in Cyprus, the associated reforms we exploit, and the administrative dataset we use to implement our identification strategies.

### 3.2.1 Institutional Details

As in most countries, charitable giving in Cyprus is subsidised through tax incentives. Specifically, the amount donated is deducted from taxable income, reducing the effective price of giving to  $(1 - \tau)$ , where  $\tau$  is the marginal tax rate. Due to feasibility constraints, there is no automatic third-party reporting by charities to the tax authority. Instead, tax filers are required to provide receipts of donations if claimed deductions exceed some specified level. To reduce both hassle and administrative costs, a threshold has been set, up to which no receipts are necessary. For any amounts claimed beyond this threshold, receipts must be provided.

We exploit several sources of exogenous variation in filing thresholds and marginal income tax rates, which allow us to examine how contributions respond to both the filing environment and to the price of giving. We first explain the reforms associated with the filing environment. The filing threshold schedule, determining at which donation level receipts are necessary, features several discontinuities and reforms between 1999-2010. This is illustrated in Figure 3.1.

Prior to 2003 the maximum amount one could declare without providing receipts was a function of salary income. For salary earnings above CYP 10,000, this threshold was CYP 150, for earnings between CYP 7,500-10,000 it was CYP 100, etc. These salary-dependent cutoffs, introduced in 1989, were abolished by the Regulatory Administrative Act No. 823 of 2003. No new law or regulatory act set any new threshold; rather, the tax authority

created a de-facto threshold at CYP 150 for everyone by clearly stating the following on the 2003 tax form: “*For donations above £150 please attach receipts*” (shown in Appendix 3B Figure 3B.1b). This is the first tax form that denotes a specific threshold; up to 2002 the tax form simply stated: “*attach relevant receipts*” (Figure 3B.1a). The 2003 wording was kept the same up to 2007. This threshold was changed again when Cyprus switched currency and adopted the Euro. The Euro was phased in during 2008, and the tax return for the 2008 fiscal year (which coincides with the calendar year) had to be filed in Euros. The tax return now stated “*attach receipts only for donations above €300*” (Figure 3B.1c). Given the locked exchange rate<sup>7</sup> of CYP 0.585274 = €1, this was equivalent to CYP 175. Tax returns are published after the end of the fiscal year and need to be submitted by the end of April. Therefore, this new threshold was published *after* the end of the 2008 fiscal year, precluding any real responses in contributions during 2008.

Besides threshold discontinuities and reforms, we are also able to exploit exogenous variation in the tax price of giving generated by marginal tax rate reforms. Figure 3.2 shows the income tax schedule in Cyprus between 1999-2010, where marginal tax rates were changed six times in total, affecting all parts of the income distribution.

Our empirical strategy draws on three sources of variation generated by the institutional setting. We start by focusing on the pre-2003 salary-based discontinuities in the filing threshold to establish that donations respond strongly to the filing environment. We then exploit the unique timing of the 2008 reform of this threshold to set bounds on the real and pure reporting components of the response. Lastly we draw on the variation in marginal tax rates to estimate the tax price elasticity of donations and examine whether this is sensitive to the filing environment.

### 3.2.2 Data

The data come from first-time access to the administrative records of the Tax Department of the Republic of Cyprus. It covers the universe of tax filers between 1999-2010, and includes information from the main fields of the I.R. 1A tax return, as well as basic demographic and firm-related characteristics. All employees are required to file taxes, unless

---

<sup>7</sup>This became legally binding by the Regulatory Administrative Act No. 311/2007.

they earn a gross amount below some tax free level. The self-employed are required to file regardless of the amount earned. Besides individuals with earnings from labour, the dataset also includes pensioners and individuals out of the labour force who may be filing because it is a requirement for accessing government welfare programmes.

To create our working dataset, we impose the following restrictions. First, we consider only individuals with a single employer, who report at least some positive salary income and are aged between 25-54.<sup>8</sup> Second, we drop individuals in the top 0.1% of donations. Our working dataset contains about 1.5 million observations and 225,000 unique individuals. Table 3.1 shows summary statistics for our sample.

It is important to note that due to the way the tax administration provides the tax data, our variable measuring donations also includes trade union subscription fees, which is another tax deduction that appears on the same section of the tax return. Prior to 2003, it also includes a so-called “professional” tax. This is a lump-sum tax that was a step-function of earnings (Figure 3B.2 in Appendix 3B shows the exact schedule). We can deal with the professional tax by simply removing the amount from our variable, as we know the exact amount individuals had to pay based on their salaries.

The only remaining issue is that of trade union membership fees. This does not affect our first two empirical strategies, because we use variation where adding noise to the measure should not affect any results. For instance, in the case of the regression discontinuity design this should only shift the level on both sides of the cut-off, but not affect the size of the discontinuity.<sup>9</sup> Similarly, for our bunching estimates, fees should affect both the level of bunching and the counterfactual, thereby leaving the bunching estimate unchanged. More importantly, we are not interested in the level of bunching per se, but in the changes in bunching across years. However, for the last part of the analysis where we consider elasticities, these union fees play a role and we need to correct for them. Since we don’t directly observe trade union membership in our data, we tackle this issue by using detailed sectoral information available in our dataset, which we combine with information on union

---

<sup>8</sup>As is explained below, we need to know workers’ sector and salary to determine their potential union membership fees. In our data, we can observe individual salaries, but not salaries per employer. We therefore drop the 3% of our sample having more than one job, to ensure we can do this accurately.

<sup>9</sup>Importantly, union fee rates are fixed across all salary ranges and did not change following any tax or threshold reforms.

fee rates we have collected directly from trade unions and from the Ministry of Labour and Social Insurance. Union fees are a fixed proportion of salaries, deducted every month from employers through the PAYE system. We use the information on union rates, salaries and sectors to residualise our donation measure from union fees in highly (or fully) unionised sectors where we can be sure we are correctly accounting for them.

To our advantage, union membership in Cyprus, which has about 50% coverage in our study period (Ioannou and Sonan 2014) is highly concentrated in just a few sectors: commercial banking, the public sector, hotel services and construction. Due to industry-wide agreements and automatic enrolment upon employment, both the public sector and commercial banking have nearly 100% coverage (Ioannou and Sonan 2014). They are also relatively large sectors, and together make up about 43% of our estimation sample (7% banks, 36% public sector). In comparison, hotel and construction have unionisation rates of about 75%, and make up 3% and 7% of our sample respectively. Besides these, union membership is very low for the remaining sectors. Our raw donation measure will therefore only be significantly affected for workers in the highly unionised sectors, for which we can accurately correct.<sup>10</sup> For each of our empirical strategies, we run a battery of robustness checks to show that our results are not affected by the way we deal with union fees.

### **3.3 Behavioural Responses to the Enforcement Environment**

As explained above, reforms changing the level of the reporting thresholds were implemented in 2003 and 2008. In both years thresholds were exclusively increased. In Figure 3.3 we clearly see the effect of these reforms in the evolution of average donations over time. The figure shows average donations among donors across the sample period. Aside from a clear increasing time trend we see jumps exactly in the two reform years. This initial time-profile of donations suggests that these reforms, which uniformly relaxed the

---

<sup>10</sup>In the few cases where we cannot, we simply drop them from our analysis. For instance, we exclude the public sector whenever we run specifications using our adjusted measure because our data does not distinguish between different types of workers in the public sector which can be subject to different union fee rates within the sector.

enforcement environment, caused a substantial increase in reported donations in Cyprus.

In this section, we look further into behavioural responses to enforcement by exploiting the salary-based discontinuities in the amount of donations tax filers can report without receipts.

### 3.3.1 Regression Discontinuity Design Estimates

As shown in Figure 3.1, before 2003 the reporting threshold was a function of gross salaries. This setup lends itself to a regression discontinuity design. For our main RD estimates, we use years 1999-2001 and restrict our sample to those with only salary income. We exclude 2002 because a reform in that year shifted the first income tax threshold to CYP 9,000, meaning that individuals with salaries below this level cease to be a reliable sample as they had no obligation to file a tax return and no tax incentive to claim deductions. We also exclude individuals with non-salary income because the threshold we want to exploit is a function of salary income, and we want to preserve the income trend in donations.<sup>11</sup> We focus on two discontinuities: the jump from CYP 100 to CYP 150 at the CYP 10,000 salary threshold, and the jump from CYP 60 to CYP 100 at the CYP 7,500 salary threshold. We do not consider lower thresholds because they are located at income levels where individuals have no tax filing obligation.

To motivate our approach, Figure 3.4 plots the average donation by salary bins of width 50 between 1999-2001. As is clearly seen, average donations jump at exactly the income thresholds associated with different reporting standards, but otherwise evolve smoothly. Note that our measure here also includes professional taxes and union fees. We do not remove these, since neither involve any discontinuities at our thresholds of interest. This is seen in Figure 3B.2 of Appendix 3B, which shows that the professional tax indeed evolves smoothly across these thresholds. Likewise, union fees are always set at fixed percentage of salary, and hence do not jump at different income levels.

Our aim is to estimate the jumps in reported donations using an RDD, treating individual salaries  $s_i$  as our assignment variable. Before doing so, we check that our identification

---

<sup>11</sup>For robustness, we also run our main specification including 2002, and including individuals with non-salary income, and find very similar results.

strategy is valid. The identifying assumption is that there is no precise manipulation of the assignment variable, i.e. workers cannot precisely choose their salaries in order to manipulate the different thresholds. For instance, if workers just to the right of a threshold strategically placed themselves there in order to be able to report more, then workers with salaries just below the threshold would not provide a good counterfactual. The possibility that workers specifically search for wage-hours packages in order to respond to the thresholds associated with charitable giving is however very unlikely. More importantly, even if this scenario were true, it is highly unlikely that they would be successful in doing so *precisely*. There are significant labour market frictions associated with searching for wage-hours packages. Indeed, a public finance literature on taxable income bunching (Chetty et al. 2011; Kleven and Waseem 2013; Gelber et al. 2017) and work hours constraints (Dickens and Lundberg 1993; Blundell et al. 2008) shows that there are significant frictions associated with precisely choosing earnings.

We check for evidence of manipulation by examining the density of salaries, as shown in Figure 3B.3 of Appendix 3B. A sign of manipulation would be significant amounts of bunching or sorting around the donation-related salary thresholds, but not elsewhere. In particular, we would expect individuals to sort just to the right of these cutoffs, in order to take advantage of the higher reporting thresholds. As we can see, this is not the case. Figure 3B.3a shows some bunching at our thresholds, but there is also (much larger) bunching at many other levels. Specifically, this comes from round-number bunching at *all* multiples of CYP 500, which is characteristic of the fact that salaries have a high propensity to be set at round numbers. This is also confirmed by Figure 3B.3b, which shows the density of salaries when we drop individuals with a salary that is an *exact* multiple of CYP 500, i.e.  $s_i \bmod 500 = 0$ . In this case, the bunching at our thresholds disappears, as it does at all other round numbers. We also use a McCrary test to formally test for the existence of any significant discontinuities in the density around each cutoff, the results of which are also reported in the figure. In line with visual evidence, the null of no discontinuity cannot be rejected, supporting our identifying assumption of no precise manipulation of  $s_i$ .

We proceed by dropping the rounders from our estimation sample, but show as a robustness check that their inclusion makes no difference. To estimate the jumps, we treat



individual salary  $s_i$  as our assignment variable, and run regressions of the form:

$$Y_i = \alpha_0 + \alpha_1 Treated_i + f(s_i, \beta) + Treated_i \times f(s_i, \gamma) + X_i' \delta + \epsilon_i \quad (3.1)$$

where we define, for each threshold  $T \in \{7500, 10000\}$  separately,  $Treated_i = \mathbb{1}\{s_i \geq T\}$ .  $f(s_i, \cdot)$  is a polynomial function with parameter vector  $\beta$  that controls for the salary trend and  $\gamma$  that controls for the interaction between the salary trend and treatment status.  $\alpha_1$  measures the jump in average donations due to the change in the reporting threshold. Some specifications also include a vector of controls  $X$  (sex, year and sector fixed effects). Lastly, we only consider bandwidths of up to 2000 to ensure that no estimation sample includes more than one threshold.

Table 3.2 shows our results split into two panels: A for the CYP 10,000 threshold and B for the CYP 7,500 threshold. To check the robustness of our estimates, we present results from eight different specifications of (3.1): with a first and second order polynomial of the assignment variable,<sup>12</sup> with and without controls, and with different bandwidths. Each panel reports two estimates: (a) the size of the discontinuity, and (b) the implied takeup, which scales our estimate by the size of the notch in the donation schedule.

Starting with the 10,000 cutoff, column (1) of Panel A shows that the effect of increasing the threshold from CYP 100 to CYP 150 leads to a CYP 36.64 increase in donations, with this effect estimated with very high precision. This estimate implies a takeup of 73%, i.e. that workers increase their donations by 0.73 for every unit increase in the amount of donations that can be reported without providing receipts. These results are highly robust to the choice of polynomial order, inclusion of controls, and bandwidth. As columns (2) - (8) show, the estimated effect is on average CYP 34.5 and the implied takeup is hence about 70%. Very similar results are found when we consider the second cutoff at CYP 7,500 (panel B). The increase in donations is estimated at about CYP 30. This is lower than the effect estimated at the higher cutoff, which is expected given that the discontinuity in reporting threshold is also smaller in magnitude. When we scale the effect by the size of the notch, we find a very similar implied takeup, estimated at about 74%, which suggests that

---

<sup>12</sup>As is clear from Figure 3.4 a linear specification allowing for different slopes on each side of each threshold should be sufficient. For robustness we also report results using second-order polynomials. We have tried running specifications with higher order polynomials, and this does not change any results.

the behavioural responses are highly comparable across the two cutoffs. Again, estimates are extremely robust to the alternative specifications across (1)-(8).

We next present robustness tests of our selected estimation sample. First, we include individuals with round-number salaries. Second, we include the year 2002. Third, we include individuals who also have non-salary income (conditional on having some salary income). Fourth, we use our variable cleaned from professional taxes and union fees.<sup>13</sup> Lastly, we run separate regressions for individuals working in highly and not highly unionised sectors. The results are shown in Appendix 3A Tables 3A.1-3A.7. Our estimates are extremely robust to every variation we consider.

As a last robustness test we also check for any discontinuities in covariates around our thresholds. Non-smoothness in covariates could suggest non-smoothness in the distribution of unobserved heterogeneity, thereby casting doubt on the validity of our method. We consider four covariates: age, the probability of being female, the probability of working in a highly unionised sector, and the level of other deductions. We also check the first and last price of giving, where these are defined as the prices before and after donations respectively.<sup>14</sup> Appendix 3B Figure 3B.4 shows plots of each of these cases as a function of our assignment variable. All plots confirm smoothness around our two cutoffs.

Finally we investigate heterogeneity in these responses by sex and age. Appendix 3A Table 3A.8 shows results for males, and Table 3A.9 for females. We do not find any notable heterogeneity around the 10,000 threshold, but do find a somewhat larger response around the 7,500 threshold among females. The implied takeup for females is around 0.82, compared to 0.66 among males.<sup>15</sup> Table 3A.10 shows results for taxpayers aged 25-39 and Table 3A.11 those aged 40-54. Younger taxpayers seem to respond slightly more to the lower threshold, but slightly less to the higher threshold, compared to older taxpayers. The implied takeup rate among the younger group is about 0.62 at the 10,000 and 0.77 at the 7,500 threshold. Among the older group, it is 0.75 and 0.68 respectively. Overall, these results do not reveal substantial heterogeneity, and importantly all groups seem to exhibit large responses independent of specification or threshold choice. Even the lowest implied

---

<sup>13</sup>In this case, we drop the public sector from Table 3A.4 because we cannot correct the trade union payments with high precision.

<sup>14</sup>These definitions are explained in detail in Section 3.5.

<sup>15</sup>These are the averages of the implied takeups across specifications.

takeups are quite substantial, confirming the large impact these reporting thresholds have on taxpayer behaviour.

### 3.4 Separating Real and Reporting Responses

Having established that reported donations respond strongly to the filing environment, we now focus on separating real and reporting responses by exploiting the timing of the 2008 reform. As explained, the threshold up to which no receipts were necessary was moved from CYP 150 to CYP 175 (€300) in 2008, but this change was only announced after the end of the 2008 fiscal year. Therefore, any response to the new threshold for 2008 can only be a change in reporting behaviour.

Our strategy exploits changes in bunching patterns, generated by the discontinuities associated with the filing threshold, around the 2008 reform. To do this, we implement standard bunching techniques (Saez 2010; Chetty et al. 2011; Kleven 2016) and estimate the excess mass of individuals located at each threshold between 2003-2010. We group donations in bins of width 5 and fit an 11th order polynomial to estimate the counterfactual mass of filers in the absence of these thresholds. The difference between the actual and counterfactual count is therefore the excess mass ascribed to the discontinuities in filing requirements. To assure comparability of bunching across different thresholds, we estimate the normalised excess mass  $b$  by scaling the excess mass by the height of the counterfactual. In our estimation, we also control for round number bunching in multiples of 50 and 100 (thereby flexibly allowing for different levels of roundedness for each).<sup>16</sup>

Our bunching results then allow us to indirectly estimate what proportion of observed responses to raising such thresholds are pure reporting effects by comparing the bunching in 2008 at the CYP 175 new threshold, which can only be driven by a pure reporting response, to the total bunching in 2007 at the CYP 150 threshold:

$$L_R = \frac{b_{175}^{2008}}{b_{150}^{2007}} \quad (3.2)$$

---

<sup>16</sup>For years 2008-2010, we have converted the currency from Euros to CYP using the official exchange rate. In this case, we control for round numbers by using the CYP converted amounts of the round numbers in Euros, since that was the actual currency used to file the tax return.

$L_R$  provides a lower bound, since bunching and responses in subsequent years can include both a real and a reporting dimension. In words,  $L_R$  reports the fraction of the excess mass at the previous threshold that moves to the new threshold before real responses are feasible.

### 3.4.1 Bunching Results

We start by focusing on our main sample, defined in the same way as in the RD section. We do not remove union fees because we want to preserve the raw patterns in the data. This of course means that we are identifying our effects from the non-unionised sample. We show in the next section that our results are highly robust to accounting for union fees.

Figure 3.5 shows the empirical density of reported donations between 2003-2010 for our main sample. Each sub-figure plots the density in bins of CYP 5, and reports the normalised excess mass at each threshold,  $b_{150}$  and  $b_{175}$ , with bootstrapped standard errors shown in parentheses. To highlight how the bunching moves across the two thresholds, each sub-figure also demarcates the threshold in place in a given year by a solid vertical line, and the other threshold by a dashed vertical line. We do not include the estimated counterfactuals here to avoid cluttering, but show these separately for each threshold in Appendix 3B Figures 3B.5 and 3B.6.

We find that bunching at the CYP 150 threshold is very large in magnitude and sharp (i.e. there is no diffuse bunching around the threshold). The normalised excess mass steadily increases between 2003 and 2007, starting from a level of 9.3 and peaking at 27.2 in the last year this threshold is effective. By 2007 therefore, there are 27 times as many individuals at CYP 150 than what there would be absent the filing discontinuity. At the same time, there is no excess bunching at CYP 175 throughout this period (marked with a dashed vertical line). These patterns are followed by a dramatic change in 2008. The bunching at CYP 150 stops growing and instead exhibits a large drop to 7.5, and continues decreasing thereafter. Bunching at CYP 175 now appears, producing a normalised excess mass of 9.9 in 2008. In a symmetrically opposite way to the bunching at CYP 150, bunching at the new threshold exhibits further growth in years 2009-2010. What is striking is that while there is no bunching in any year before 2007 at CYP 175, a very large spike appears

in 2008, even though there was no knowledge of this new threshold, and hence no real response possible during the 2008 fiscal year.

To visualise these patterns Figure 3.6 plots our bunching estimates at each threshold across time, with the shaded areas demarcating our 95% confidence intervals. The bunching dynamics suggests learning, as it seems to take time for individuals to understand the incentives created by the thresholds and respond to them over the years.

Given our estimates, we find that the lower bound of the reporting response is  $L_R = 36\%$ . This implies that at least 36% of the response to an increase in the reporting threshold is due to changes in reporting, rather than real changes in contributions. To the extent that individuals take time to learn about, and understand, the changes in the reporting environment, the responses in 2009 and 2010 may also capture reporting responses, which would increase our estimate.

Next, we exploit the panel dimension of our data to check whether the patterns we observe is indeed driven by individuals moving from the old threshold to the new one. In Appendix 3B Figure 3B.7, we plot the empirical density of reported donations between 2003-2010, for the sample bunching at CYP 150 in 2007. We find that the overwhelming majority of the 2007 bunchers are “repeat” bunchers, locating at CYP150 for up to 4 years earlier. They also overwhelmingly shift to the new threshold in 2008, while some take a further one or two years to complete the shift. Appendix 3B Figure 3B.8 repeats the analysis for those bunching at the new threshold in 2008. Again, we find the same patterns; these are individuals who were previously bunching at the old threshold for up to 5 years before, with the shift being nearly complete by 2010.

Figure 3B.9 plots the fraction of individuals with reported donations at (1) the 150 threshold, (2) the 175 threshold, (3) above 150 and (4) above 175 throughout the entire sample period. The trends in the fraction of individuals filing 150 and 175 is in line with our previous results. The fraction of individuals at 150 is increasing from 2003 and peaks at 2007, before dropping sharply in 2008. Conversely, there are very few at 175 until 2008, when it increases sharply. What is more interesting is the trend in the proportion of people filing more than 150 and more than 175. While they exhibit parallel trends up to 2008, there is a sharp increase in the fraction of individuals filing above 150, but no change in

the fraction of those filing above 175. This confirms that the movement is purely between these two thresholds, and emphasises how important the filing environment is for taxpayer behaviour.

### 3.4.2 Robustness Analysis

We next discuss the robustness of these results and conduct a battery of checks on our main results. First, it is very important for the validity of the lower bound that the 2008 threshold change was not anticipated or somehow made public before the end of the fiscal year. Before the introduction of the Euro there was a large government campaign informing citizens that during the transition they should simply use the official locked exchange rate to convert prices, salaries etc. Following this, filers should have expected the threshold to remain unchanged at a converted value of €250, not €300. Further, tax returns are not published before the end of the fiscal year. Even if it was published early through unofficial channels, it is highly unlikely that filers would be so keen to obtain their tax return before the end of the fiscal year that they would search for it, especially since the fiscal year ends four months in advance of the tax return submission deadline. Moreover, if we examine the distribution of tax return submission dates in our data, shown in Appendix 3B Figure 3B.10, we find that the vast majority of tax filers procrastinate and submit their return just before the deadline of April 30th of the following year. Most filers therefore leave this to the last week of the deadline, which is clearly at odds with active tax return search and filing behaviour.

Second, we consider whether our bunching patterns could be affected by the existence of union fees. This would be implied by two, extremely unlikely, scenarios. The first is that we are picking up bunching at thresholds that is driven purely by fees which coincidentally place individuals at the thresholds. The second is that we are picking up the sum of union fees and donations, which again happen to consistently sum to these thresholds. Both are implausible. The first case would require that the salaries of such bunchers were at a level prior to 2008 that, when the fixed % of salary paid as union fees was applied, would place them at the filing threshold. At the same time, their salary in 2008 would also have to grow by a rate exactly large enough to move them to the new threshold (a 16.67% increase).

Thereafter, they would again have to revert back to a zero growth rate in order to stay at the new threshold. We can in fact check this. Figure 3B.11 of Appendix 3B plots the salary growth rate of the 2007 bunchers and shows that this is not the case;<sup>17</sup> salary growth rates are far from what this extreme scenario would imply. The second case would require a very high level of sophistication and meticulous tax planning. Specifically, individuals would have to predict their exact yearly salary and union fees, and dynamically adjust their donations to crowd-out union fees one-for-one, so that these would always sum to the exact threshold by the end of each fiscal year. Not only is this implausible, but importantly taxpayers have no financial incentive to engage in such a form of planning in the first place.

Nevertheless, we conduct robustness checks where we repeat all of our previous analysis, but restrict the sample to only workers not in highly unionised sectors. In this case, the contribution of any union fees will be minimal because the proportion of unionised members in this sample will also be limited. Our full set of results is presented in Appendix 3B Figures 3B.12 - 3B.15. The bunching patterns are nearly identical to our main findings, both in terms of magnitudes and dynamics. As shown in Figure 3B.15, bunching at the CYP 150 threshold grows steadily from 13.6 in 2003 to 31.2 by 2007, falls sharply in 2008 to 7.5 and keeps decreasing thereafter. In the exact same way as before, a normalised bunching mass of 12.1 first appears at the new threshold in 2008, and grows thereafter. The lower bound estimate of the reporting response is now 39%, which is extremely close to our previous estimate (of 36%). One difference worth mentioning is that the height of the counterfactual to the right of the thresholds now seems lower than before (shown in Appendix 3B Figures 3B.17 and 3B.18). This is exactly what we would expect when we remove the highly unionised, if they are also bunching at the donation threshold because their union fees would thereby place them above it.

Our main results are therefore highly robust to the presence of union fees. This of course means that we are identifying the effect of the 2008 reform from non-union members, since union members would be scattered around these thresholds with their donations, and as the previous analysis revealed, their union fees would place them to the right of them. This does not affect our analysis, since systematic union fees would only potentially affect

---

<sup>17</sup>There is a large drop from 2008, driven by the financial crisis.

the *level* of bunching. We are not however interested in this level per se, but rather in the *ratio* of bunching, between 2008 and 2007. While this is identified from the non-unionised, our result is generalisable to the extent that union members behave in a similar way. We check this by next repeating our main analysis after removing union fees among the highly unionised sectors.<sup>18</sup> The results are shown in Appendix 3B Figures 3B.16 - 3B.19. Again, we find the exact same patterns. Our lower bound estimate of the reporting response is 42%, which is very similar to our previous estimates.

As a final robustness test, we check whether our main results are sensitive to the choices we have made regarding the polynomial order used to estimate the counterfactual and the bandwidth used. Figure 3B.20 of Appendix 3B shows the bunching estimates for different specifications of these, and shows that our results are highly robust.

### 3.4.3 Heterogeneity Analysis

We have also tested for heterogeneity by running our main specification on different subsamples by sex and age. Appendix 3B Figure 3B.21 shows the bunching dynamics by sex and Figure 3B.22 by age groups. Overall, we do not find any evidence of significant heterogeneity. Results are extremely similar regardless of the sample split we consider.

## 3.5 The Elasticity of Giving w.r.t. Price and Tax Policy Design

In this section we turn to our last source of quasi-experimental variation: changes in the price of giving generated by tax rate reforms and our second question: the importance of tax policy design for elasticities. Most importantly we look at how the thresholds studied extensively above impact the elasticity of reported donations with respect to price. The tax rate reforms in our sample period are illustrated in Figure 3.2. We can exploit these reforms as a source of variation in the price of giving, since the tax rate essentially determines the size of the tax subsidy to charitable donations. The typical approach in the literature on the elasticity of the tax price of giving is to run log-specifications of the form:

---

<sup>18</sup>In this case, we exclude the public sector.



$$\ln(d_{it}) = \beta_1 \ln(1 - \tau_{it}) + \beta_2 \ln(y_{it}) + \beta_3' X_{it} + \Gamma_i + \Gamma_t + \varepsilon_{it} \quad (3.3)$$

where  $d_{it}$  is the donation amount and  $y_{it}$  is disposable income before donations, of individual  $i$  at time  $t$ .  $\tau_{it}$  is the marginal tax rate and hence  $1 - \tau_{it}$  is the price of giving.  $X_{it}$  is a vector of other controls. Specifications estimated using panel data can also include individual and time fixed effects,  $\Gamma_i$  and  $\Gamma_t$ . The price elasticity is then given by  $\beta_1$ .

Estimating this equation using standard OLS leads to an endogeneity problem. Charitable donations can affect the price of giving because these may shift taxpayers to lower tax brackets, thereby reducing the tax price and causing an upward bias in the estimated elasticities. This is a well-known endogeneity problem in the literature, and has been typically dealt with by instrumenting the “last-pound” (observed) price of giving with the “first-pound” price of giving ( $1 - \tau_{it|d_{it}=0} \equiv 1 - \tau_{it}^*$ ). In words, this is the price a taxpayer faces for the first pound of charitable contribution. This removes any price variation due to charitable giving, and results in a very strong first-stage because the first- and last-pound prices are mechanically very highly correlated.

For the exclusion restriction to hold, the relationship between the first-pound price and the level of donations must solely go through the last-pound-price of giving. As argued by Almunia et. al (2017), this exclusion restriction is violated when using price variation from tax reforms because such reforms create a second source of endogeneity. Specifically, changes in marginal tax rates can also affect other choices such as individual labour supply and earnings more generally. Tax reforms therefore affect taxable earnings choices and hence can determine the first-pound price of charitable giving because they determine which tax bracket a taxpayer is in. In other words, tax reforms can affect choices which both enter the donation decision and directly affect the first-pound price, violating the exclusion restriction.

Almunia et al. (2017) propose a solution that leverages the availability of panel data, based on the Gruber and Saez (2002) IV strategy that is widely used in the literature on the elasticity of taxable income (for a review, see Saez et al. 2012). Their instrument

uses lagged income values to predict the reform-induced change in the price of giving. Concretely they propose estimating the following differenced equation:

$$\Delta \ln(d_{it}) = \beta_1 \Delta \ln(1 - \tau_{it}^*) + \beta_2 \Delta \ln(y_{it}) + \beta_3' \Delta X_{it} + \Delta \varepsilon_{it} \quad (3.4)$$

where  $\Delta \ln(x_{it}) = \ln\left(\frac{x_{it}}{x_{i,t-k}}\right)$  for  $x_{it} = d_{it}, 1 - \tau_{it}^*, y_{it}$ , and the log change in the first-pound price is instrumented by:

$$\ln\left(\frac{1 - \tau_{i,t}^* (y_{i,t-k}^*)}{1 - \tau_{i,t-k}^* (y_{i,t-k}^*)}\right) \quad (3.5)$$

The variable  $k$  determines the time horizon of the difference. In words, the instrument is the change in the price of giving from time  $t-k$  to time  $t$  if taxable income at zero donations ( $y_{it}^*$ ) remained unchanged. This instrument solves the endogeneity problem because it uses past (pre-tax) income which should not be affected by future reform-related choices. The only threat to this exclusion restriction is potential anticipation responses to tax reforms. For this reason, we consider lags of both  $k = 1$  and  $k = 2$ , as the longer the lag, the less realistic such a violation becomes.<sup>19</sup>

In our empirical application, we mainly follow Almunia et al. (2017) and implement the differenced-IV specification (3.4). To compare our findings with existing practice, we also report results for both OLS and IV estimates of specification (3.3).

As already mentioned, our measure of donations also includes a union membership fee for some subset of workers. We cannot directly observe the size of this fee or who is a member of a union. However, using information on sector we can to some extent back it out. The banking sector in Cyprus is a large sector, and it is (almost) fully unionised (99%) due to automatic enrolment upon employment. Consequently, for the banking sector we can completely separate donations from union membership fees. Apart from the banking sector a small number of other sectors, such as the construction and hotel sectors, are highly but not fully unionised. Consequently for these sectors we can correct for fee payments

---

<sup>19</sup>We have also tried longer lags and find very similar results, confirming the lack of anticipatory effects.

with some error coming from the small fraction of workers who are not in a union.<sup>20</sup> Lastly, Cyprus has a number of sectors with very low unionisation, and so in this case, the donation measure will be somewhat noisy due to the small number of unionised workers. Following these considerations we use the sample of workers from the banking sector as our main sample in the following analysis, but for all results we show robustness to using either the full sample or the sample of highly unionised sectors instead.

### 3.5.1 Elasticity Results for the Whole Period 1999-2010

In this section we estimate the elasticity of donations with respect to price using the methods described above and the full sample period. Table 3.3 reports results for different specifications using the sample of salary earners in the banking sector. We only include individuals with a tax filing obligation in the sample period. Since we are dealing exclusively with the banking sector we can correct for union fee payments and hence obtain an accurate measure of charitable contributions for each individual.

In Table 3.3 we begin with the benchmark case where we estimate specification (3.3) by OLS. Column (1) shows results without controls, while column (2) also includes age squared and a dummy for whether a worker switched job as controls. All specifications include individual and year fixed effects. When including controls the OLS specification produces an estimate of -0.38 for  $\beta_1$ . We next turn to the IV version of specification (3.3) in columns (3) and (4), where we instrument the last-pound price using the first-pound price. As expected,  $\beta_1$  becomes much larger in absolute value (-0.56), which is in line with the upward bias of OLS. Finally, the last four columns turn to our preferred differenced-IV specification. Column (6) and (8) shows that the estimate of  $\beta_1$  is -0.51 when using controls and  $k = 1$ , and nearly unchanged when using  $k = 2$ . Consequently the estimate of  $\beta_1$  is very stable across our three different IV specifications at a value around -0.5 when including controls. Estimates are slightly higher, but comparable when excluding controls.

These results are robust to including individuals for whom we cannot be sure that we can correct precisely for union fee payments in the donation measure. In Table 3A.12 of

---

<sup>20</sup>In this sample we do not include the public sector even though this sector is highly unionised. This is due to the existence of different rates for the union membership fee in the public sector, which we cannot accurately correct for.

Appendix 3A we include other sectors which are highly unionised. As already mentioned, in this case we can correct the donation measure with some noise coming from the small fraction which is not in a union. Results from this table are very much in line with what we found above. In Appendix 3A Table 3A.13 we include the whole sample and correct the donation measure only for those sectors with a high degree of unionisation. In this case the donation measure will be somewhat more noisy. The results from this sample are nevertheless still very much in line with our main results above. Reassuringly, all findings are not sensitive to the potential noise created by the union fees in our measure. This is also not surprising, given that the level of union fees is very low.

In the results above we do not explicitly deal with censoring coming from people reporting zero donations. Given the logarithmic specification such observations are excluded from the estimation and this can potentially create bias. However, in our setting this is not a serious concern since 87% of tax filers report positive donations.<sup>21</sup> Following previous literature we nevertheless check for robustness of our results to this margin by estimating a Poisson model of the form:

$$d_{it} = \exp(\theta_1 \log(1 - \tau_{it}^*) + \theta_2 \log(y_{it}) + \theta_3 X_{it} + \alpha_i + \delta_t + u_{it}) \quad (3.6)$$

This model can be estimated using a pseudo-maximum likelihood method. The benefit of this approach is that it circumvents the selection issue by allowing the inclusion of both positive and zero values of the dependent variable.<sup>22</sup>

Table 3A.14 in Appendix 3A shows the results from this approach applied to our main sample in this section (the banking sector). Both column (1) and (2) use the first pound price of giving as the regressor, and these results are hence comparable to columns (3) and (4) of Table 3.3. Especially when including controls the elasticity resulting from this approach is very similar in magnitude to our main estimates from Table 3.3. Hence, our

---

<sup>21</sup>In our main estimation sample of bankers, this is even higher (91%). In contrast, this number is only 10% of the estimation sample in Almunia et al. (2017) and 15% in Fack & Landaïs (2016).

<sup>22</sup>It is also preferred to the Tobit model as it does not suffer from the incidental parameters problem, relies on weaker functional form assumptions, and allows for constant elasticities, which the Tobit specification does not. For more details, see Almunia et al. (2017).

findings confirm that our main results are not biased by potential censoring issues.

### 3.5.2 Elasticities and Reporting Thresholds

The first sections of this paper investigated in detail the behavioural responses to thresholds determining reporting standards. We now investigate whether the existence of such thresholds affects people's responsiveness to price changes and hence affects the elasticity estimated above. We have shown that people respond strongly to these thresholds by bunching at the exact threshold values. If such bunching behaviour is sticky when prices change, then this design feature could potentially have a strong effect on elasticities and consequently the effectiveness of providing subsidies. This is an important issue as this feature is not uncommon across the world, nor is it restricted to the treatment of charitable contributions. Moreover, tax authorities may set such thresholds for a variety of reasons including concerns about administrative costs.

To get a picture of the sensitivity of elasticities to this tax design feature, we split our sample into two types of workers. The first group consists of those workers who at some point bunch, meaning that at some point in the sample period we observe them exactly at a threshold value. The second group consists of those workers who never bunch meaning that we never observe them at a threshold value. We then estimate elasticities separately for these two groups of workers using the methods introduced above. We look only at the period 2003-2007 since this is the longest period in the data for which the reporting environment remains constant. In both year 2003 and year 2008 we observe changes in both marginal tax rates and hence prices, and also in the placement of the reporting thresholds. Since we are specifically interested in the stickiness of bunching behaviour around these thresholds, we want to clearly separate responses to prices from movements caused by changes in thresholds.<sup>23</sup>

In Table 3.4 panel A, we first estimate the elasticity for both groups together in the period 2003-2007. In column (1)-(2) we report estimates using the first-pound price IV strategy, and in columns (3)-(4) we report results using the differenced IV strategy with

---

<sup>23</sup>Note that we cannot look exclusively at the last period 2008-2010 since we have no variation in tax rates in this period. Focusing instead on the period before 2003 does not change the main conclusions of this section.

a time horizon of one year. Since results in the previous section were very insensitive to the time horizon of the difference, we focus only on the case  $k = 1$  moving forward. The results from panel A using both groups are unsurprisingly very similar to the results from the previous section. In panel B we look only at the group of workers who never bunch at a reporting threshold. Results from both estimation strategies show that this group displays much more sensitivity to price changes than the collective group of workers. Estimated elasticities double in magnitude compared to panel A. In panel C we now look exclusively at the group of workers who at some point in the period 2003-2007 bunch at a threshold value. This group is much less responsive to price changes and estimated elasticities are around half the size of the estimates for the entire group and about a fourth of the size of the estimates for the non-bunchers. If we instead use the sample of all highly unionised sectors or the sample of all sectors we see exactly the same patterns (Appendix 3A Tables 3A.15 and 3A.16). In these samples the estimated differences are somewhat smaller consistent with the fact that we cannot precisely separate bunchers from non-bunchers given the noise in the donation measure.

These results are in line with a situation where bunching behaviour around reporting thresholds is sticky and hence workers who bunch at the thresholds need strong incentives to change behaviour. If these differences are driven by stickiness around reporting thresholds then the differences above should be driven by bunchers being less likely to react to price changes and not from bunchers reacting less to a fixed price change compared to non-bunchers. If we look at the average fraction of workers who change the amount they report in donations from one year to the next, then this fraction is 38% for those who bunched at a threshold the previous year, while it is 79% for those who did not bunch at a threshold the previous year. If we only look at workers in years where they experience a change in the first-pound price then these fractions increase to 44% and 91% respectively. These numbers support the idea of stickiness in behaviour around thresholds since both unconditional and conditional on a price change bunchers are much less likely to change reporting behaviour from one year to the next.

Beyond these non-parametric measures, we can implement this analysis in a regression framework and include controls. Table 3.5 report results from a simple linear probabil-

ity model where the dependent variable is an indicator for changing reporting from one year to the next. On the right hand side we include an indicator for being located at a threshold value i.e. being a buncher, a constant and potential controls. Column (1) simply shows the unconditional average fractions reported above - the constant shows that 79% of non-bunchers change reporting from one year to the next, while the coefficient on the buncher-indicator shows that this percentage is  $78.5\% - 40.3\% = 38.2\%$  for bunchers. In column (2) we control for year fixed effects, age squared, salary and a dummy for changing employer and results remain unchanged. Column (3) reports the conditional average fractions discussed above, where we only look at individuals with a change in price. Also in this case adding controls makes little difference to the results. Tables 3A.17 and 3A.18 in Appendix 3A report results using the alternative samples of either all highly unionised sectors or all sectors. Both tables show results very much in line with the main results discussed above. Consequently we observe a large and statistically significant difference between bunchers and non-bunchers in the probability of changing reporting behaviour and reacting to price-changes. This lends support to the idea that the large effect of these reporting thresholds on elasticities, comes from sticky bunching behaviour around the threshold values. Such behaviour should hence be taken into account when trying to estimate elasticities in an environment characterised by discontinuous changes in reporting standards around threshold values.

### 3.5.3 Elasticities and Enforcement

Next we look at the importance of the overall enforcement strictness for estimated elasticities. In short we have three different sub-periods characterised by differences in the enforcement environment. Before 2003 reporting thresholds were salary dependent and relatively low for many income groups. In 2003 this threshold was increased for individuals with income below CYP 10,000, and the threshold was made independent of income. Hence between 2003 and 2007 the enforcement environment was laxer compared to the earlier period because a subgroup experienced a higher threshold. In 2008 the threshold was moved again - this time it was increased for everyone and hence again the enforcement environment became laxer. To investigate the importance of these enforcement changes for

elasticities we divide our sample period into sub-periods and perform separate estimations in each period. In the last period 2008-2010 we don't have any marginal tax rate reforms and hence we do not have variation in the price of giving. Therefore we only estimate on two different periods: 1999-2002 representing the stricter enforcement environment and 2003-2010 representing the laxer enforcement environment with higher thresholds.

Results for workers in the banking sector are given in Table 3.6. We report results both for the first-price IV strategy and for the differenced IV strategy, but estimated elasticities are very similar across these specifications and hence we focus on our preferred strategy using differences. In column (4) we find an estimated elasticity of  $-0.29$  for the stricter period when including controls. In column (8) we find an estimated elasticity of  $-0.51$  for the laxer period when including controls. Consequently we see a much larger elasticity in the laxer enforcement environment suggesting that enforcement matters for the size of elasticities. Importantly this large difference is present in all four different specifications of Table 3.6. Note further that standard errors are small in both periods and hence these elasticities are precisely estimated.

These results are again robust to including individuals for whom we cannot be sure that we can correct precisely for union fee payments in the donation measure. In Table 3A.19 of Appendix 3A we show analogous results for the larger sample of all highly unionised sectors corrected for union fees. Again we see large differences across periods, with elasticities being larger in the laxer enforcement period. In Table 3A.20 of Appendix 3A we show results for the whole sample where we correct the highly unionised sectors, but not the sectors with low unionisation. Again we see the same differences across periods.

## 3.6 Conclusion

This paper studies behavioural responses to enforcement and the importance of tax policy design for elasticities, using the context of charitable contributions in the Republic of Cyprus. We use multiple sources of quasi-experimental variation in reporting requirements and tax price subsidies to present several policy-relevant results.

First, we show evidence of substantial reactions to changes in reporting environment.



Exploiting salary-dependent thresholds for third party reporting of charitable contributions and a regression discontinuity approach, we estimate that reported donations increase by about 0.7 pounds when taxpayers can report 1 pound more in charitable contributions without having to provide receipts. Second, we separate real from reporting responses by utilising a reform that retroactively shifted the location of the threshold up to which receipts are not required. Our bunching analysis reveals that at least 36 percent of the responses to changes in thresholds are purely due to changes in reporting.

Finally, using quasi-experimental variation in tax prices generated by income tax reforms, we find that the tax price elasticity of reported charitable donations is about  $-0.5$ . Importantly, we show that this policy parameter is highly dependent on the reporting environment, partly because the existence of thresholds generates strong bunching responses which mitigates responses to tax price changes. Price elasticities double if we remove individuals who at some point in the sample period bunch at a threshold. Further, we estimate elasticities across different sub-periods with varying enforcement strictness, and show suggestive evidence that this parameter is sensitive to even minor changes in reporting environment.

Our findings have important policy implications. The very strong behavioural responses to reporting thresholds imply that this policy instrument severely affects taxpayers' behaviour and consequently government revenue. Further, it affects people's responsiveness to prices and therefore potentially the effectiveness of providing subsidies such as for instance the globally widespread subsidy to charitable contributions. To the extent that the fiscal authority wants to incentivise certain forms of behaviours that generate positive externalities, it is crucial to understand the conditions under which tax subsidies cannot achieve this goal. Given that such thresholds are common in tax systems of many countries, this aspect of tax design warrants a more direct incorporation in optimal tax theory.

# References

- [1] Adena, M. (2014) ‘Tax-Price Elasticity of Charitable Donations - Evidence From the German Taxpayer Panel’, WZB Working Paper, 2014-302.
- [2] Alm, J., Deskins, J. and McKee, M. (2009) ‘Do Individuals Comply on Income not Reported by Their Employer?’ *Public Finance Review*, vol. 37(2), pp. 120-140.
- [3] Almunia, M., Lockwood, B. and Scharf, K. (2017) ‘More Giving or More Givers? The Effects of Tax Incentives on Charitable Donations in the UK’, CESifo Working Paper 6591.
- [4] Auten, G. E., Sieg, H. and Clotfelter, C.T. (2002) ‘Charitable Giving, Income, and Taxes: An Analysis of Panel Data’, *American Economic Review*, vol. 92(1), pp. 371-382.
- [5] Bakija, J. and Heim, B.T. (2011) ‘How Does Charitable Giving Respond to Incentives and Income? New Estimates From Panel Data’, *National Tax Journal*, vol. 64(2), pp. 615-650.
- [6] Blundell, R., Brewer, M. and Francesconi, M. 2008. ‘Job Changes and Hours Changes: Understanding the Path of Labor Supply Adjustment’, *Journal of Labor Economics*, vol. 26(3), pp. 421-453.
- [7] Carrillo, P., Pomeranz, D. and Singhal, M. (2017) ‘Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement’, *American Economic Journal: Applied Economics*, vol. 9(2), pp. 144-164.

- [8] Central Bank of Cyprus, 2012. *The Cyprus Economy: Historical Review, Prospects, Challenges*. Nicosia: Central Bank of Cyprus.
- [9] Chetty, R., Friedman, J. N., Olsen, T. and Pistaferri, L. (2011) ‘Adjustment Costs, Firm Responses, and Micro Vs. Macro Labor Supply Elasticities: Evidence From Danish Tax Records’, *The Quarterly Journal of Economics*, vol. 126(2), pp. 749-804.
- [10] Dickens, W.T and Lundberg, S.J (1993) ‘Hours Restrictions and Labor Supply’, *International Economic Review*, vol. 34(1), pp. 169-192.
- [11] Doerrenberg, P., Peichl, A. and Siegloch, S. (2017) ‘The Elasticity of Taxable Income in the Presence of Deduction Possibilities’, *Journal of Public Economics*, vol. 151, pp. 41-55.
- [12] Duquette, N. J. (2016) ‘Do Tax Incentives Affect Charitable Contributions? Evidence From Public Charities’ Reported Revenues’, *Journal of Public Economics*, vol. 137, pp. 51-69.
- [13] Fack, G. and Landais, C. (2010) ‘Are Tax Incentives for Charitable Giving Efficient? Evidence From France’, *American Economic Journal: Economic Policy*, vol. 2, pp. 117-141.
- [14] Fack, G. and Landais, C. (2016) ‘The Effect of Tax Enforcement on Tax Elasticities: Evidence From Charitable Contributions in France’, *Journal of Public Economics*, vol. 133, pp. 23-40.
- [15] Gelber, A.M., Jones, D. and Sacks, D.W. (2017) ‘Estimating Earnings Adjustment Frictions: Method & Evidence From the Earnings Test’, Goldman School of Public Policy Working Paper.
- [16] Gruber, J. and Saez, E. (2002) ‘The Elasticity of Taxable Income: Evidence and Implications’, *Journal of Public Economics*, vol. 84, pp. 1-32.
- [17] Ioannou, G. and Sonan, S. (2014) *Trade Unions in Cyprus: History of Division, Common Challenges Ahead*. Friedrich-Ebert-Stiftung Report.

- [18] Karlan, D., and List, J.A. (2007) ‘Does Price Matter in Charitable Giving? Evidence From a Large-Scale Natural Field Experiment’, *American Economic Review*, vol. 97(5), pp. 1774-1793.
- [19] Kleven, H.J. (2016) ‘Bunching’, *Annual Review of Economics*, vol. 8, pp. 435-464.
- [20] Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S. and Saez, E. (2011) ‘Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark’, *Econometrica*, vol. 79(3), pp. 651-692.
- [21] Kleven, H. J. and Waseem, M. (2013) ‘Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan’, *The Quarterly Journal of Economics*, vol. 128(2), pp. 669-723.
- [22] Kopczuk, W. (2005) ‘Tax Bases, Tax Rates and the Elasticity of Reported Income’, *Journal of Public Economics*, vol. 89, pp. 2093-2119.
- [23] Kopczuk, W. and Slemrod, J. (2002) ‘The Optimal Elasticity of Taxable Income’, *Journal of Public Economics*, vol. 84, pp. 91-112.
- [24] Mishra, P., Subramanian, A. and Topalova, P. (2008) ‘Tariffs, enforcement, and Customs Evasion: Evidence From India’, *Journal of Public Economics*, vol. 92, pp. 1907-1925.
- [25] Phillips, M. D. (2014) ‘Individual Income Tax Compliance and Information Reporting: What do the U.S. Data Show?’ *National Tax Journal*, vol. 67(3), pp. 531-567.
- [26] Saez, E., Slemrod, J. and Giertz, S. H. (2012) ‘The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review’, *Journal of Economic Literature*, vol. 50(1), pp. 3-50.
- [27] Slemrod, J. (1994) ‘Fixing the Leak in Okun’s Bucket - Optimal Tax Progressivity When Avoidance Can be Controlled’, *Journal of Public Economics*, vol. 55, pp. 41-51.

# Main Tables

Table 3.1: Summary Statistics

	Mean	Std. Dev.
Salary Earnings Only	0.845	0.362
Ratio of Salary to Total Earnings	0.960	0.141
Taxable Income	12646.498	11139.126
Job Switches	0.028	0.164
Marginal Tax Rate	0.182	0.133
Positive Donations	0.873	0.333
Donations (cond. positive)	171.711	118.089
Positive Donations (Net of Union Fees)	129.87	103.436
Donations (Net of Union Fees, cond. positive)	0.788	0.409
Age	40.437	8.070
Female	0.385	0.487
Agriculture	0.005	0.069
Mining	0.003	0.051
Manufacturing	0.086	0.281
Construction	0.074	0.262
Utilities	0.021	0.142
Trade	0.120	0.325
Hotel Services	0.033	0.180
Other Services	0.186	0.389
Commercial Banking	0.066	0.248
Other Financial Services	0.033	0.179
Public Sector	0.360	0.480
Other Sector	0.011	0.102
<b>Observations</b>	1,462,409	

*Notes:* This table displays summary statistics for our sample. We distinguish between positive donations, and positive donations net of union fees, where we residualise our measure in the latter case from union fees. Both measures have professional taxes already removed.

Table 3.2: RD Estimates at Notches in Donation Schedule

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	36.64*** (0.95)	33.80*** (1.41)	34.40*** (0.97)	32.51*** (1.45)	35.08*** (1.33)	35.12*** (2.04)	33.79*** (1.38)	34.15*** (2.12)
Implied takeup	0.73	0.68	0.69	0.65	0.70	0.70	0.69	0.69
Observations	84 255	84 255	68 697	68 697	43 095	43 095	35 404	35 404
$R^2$	0.21	0.21	0.28	0.28	0.13	0.13	0.20	0.20
<b>Panel B:</b>								
Above 7.5k	29.64*** (0.74)	31.77*** (1.08)	30.48*** (0.79)	30.91*** (1.14)	30.78*** (1.01)	26.50*** (1.51)	30.21*** (1.06)	26.30*** (1.58)
Implied takeup	0.74	0.79	0.76	0.77	0.77	0.66	0.76	0.66
Observations	92 433	92 433	71 095	71 095	51 850	51 850	40 123	40 123
$R^2$	0.29	0.29	0.37	0.37	0.17	0.18	0.27	0.27
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows the results from estimating specification (3.1) on our main sample pooled over 1999-2001. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3.3: Elasticity of Donations w.r.t. Price (1999-2010)

Price variable	(1) OLS	(2) OLS	(3) IV	(4) IV	(5) $IV\Delta_{k=1}$	(6) $IV\Delta_{k=1}$	(7) $IV\Delta_{k=2}$	(8) $IV\Delta_{k=2}$
$\ln(1 - \tau)$	-0.556*** (0.051)	-0.375*** (0.056)						
$\ln(1 - \tau^*)$			-0.716*** (0.057)	-0.556*** (0.065)				
$\Delta \ln(1 - \tau^*)$					-0.715*** (0.066)	-0.506*** (0.062)	-0.747*** (0.074)	-0.475*** (0.072)
Individual FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	-	✓	-	✓	-	✓	-	✓
Observations	50 789	50 789	50 783	50 783	43 094	43 094	37 598	37 598
$R^2$	0.50	0.51	0.50	0.51	0.13	0.14	0.18	0.19

*Notes:* This table shows estimates of the elasticity of donations w.r.t. the price of giving using different specifications. The sample is restricted to those in the banking sector. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 3.4: Elasticity of Donations w.r.t. Price - Bunchers vs. Non-bunchers (2003-2007)

Price variable	(1) IV	(2) IV	(3) IV $\Delta_{k=1}$	(4) IV $\Delta_{k=1}$
<b>Panel A: All workers</b>				
$\ln(1 - \tau^*)$	-0.642*** (0.169)	-0.658*** (0.180)		
$\Delta \ln(1 - \tau^*)$			-0.419*** (0.160)	-0.375*** (0.154)
Observations	17 955	17 955	15 957	15 957
$R^2$	0.62	0.62	0.20	0.21
<b>Panel B: Excluding bunchers</b>				
$\ln(1 - \tau^*)$	-1.462*** (0.335)	-1.517*** (0.372)		
$\Delta \ln(1 - \tau^*)$			-0.818*** (0.261)	-0.786*** (0.247)
Observations	5 336	5 336	4 586	4 586
$R^2$	0.72	0.72	0.27	0.28
<b>Panel C: Only bunchers</b>				
$\ln(1 - \tau^*)$	-0.294 (0.188)	-0.306 (0.191)		
$\Delta \ln(1 - \tau^*)$			-0.253 (0.203)	-0.207 (0.198)
Observations	12 619	12 619	11 371	11 371
$R^2$	0.52	0.52	0.17	0.18
Individual FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Controls	-	✓	-	✓

*Notes:* This table shows how our estimates of the elasticity of donations w.r.t. the price of giving depends on the inclusion of bunchers, using our main sample (the banking sector) between 2003-2007. Bunchers are defined as those who were located at the donation filing threshold in any year during 2003-2007. Panel A includes all workers, B excludes bunchers and C restricts to only the bunchers. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.



Table 3.5: Buncher Stickiness - Linear Probability Model

	(1) All	(2) All	(3) $\Delta(1 - \tau^*) \neq 0$	(4) $\Delta(1 - \tau^*) \neq 0$
$\mathbb{1}_{(\text{buncher})}$	-0.403*** (0.007)	-0.391*** (0.007)	-0.474*** (0.013)	-0.458*** (0.013)
Constant	0.785*** (0.004)	0.813*** (0.011)	0.913*** (0.004)	0.895*** (0.015)
Observations	20 621	20 621	6 567	6 567
$R^2$	0.16	0.20	0.27	0.28
Year FE	-	✓	-	✓
Controls	-	✓	-	✓

*Notes:* This table shows estimates of models where we regress an indicator for having adjusted donations in a given year on a dummy for being a buncher in the previous year, using our main sample (the banking sector) between 2003-2007. The first two columns do not impose any further sample restrictions, while the last two restrict the sample to those experiencing a change in the (first) tax price of giving. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.


Table 3.6: Elasticity of Donations w.r.t. Price by Period

	<b>1999-2002:</b>				<b>2003-2010:</b>			
Price variable	(1) IV	(2) IV	(3) IV $\Delta_{k=1}$	(4) IV $\Delta_{k=1}$	(5) IV	(6) IV	(7) IV $\Delta_{k=1}$	(8) IV $\Delta_{k=1}$
$\ln(1 - \tau^*)$	-0.235*** (0.088)	-0.229** (0.098)			-0.580*** (0.125)	-0.554*** (0.141)		
$\Delta \ln(1 - \tau^*)$			-0.288** (0.131)	-0.285** (0.130)			-0.613*** (0.127)	-0.513*** (0.124)
Individual FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	-	✓	-	✓	-	✓	-	✓
Observations	12 110	12 110	8 198	8 198	33 855	33 855	30 260	30 260
$R^2$	0.71	0.72	0.34	0.34	0.55	0.55	0.12	0.12

*Notes:* This table shows how our estimates of the elasticity of donations w.r.t. the price varies by different sub-periods, using our main sample (the banking sector). All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

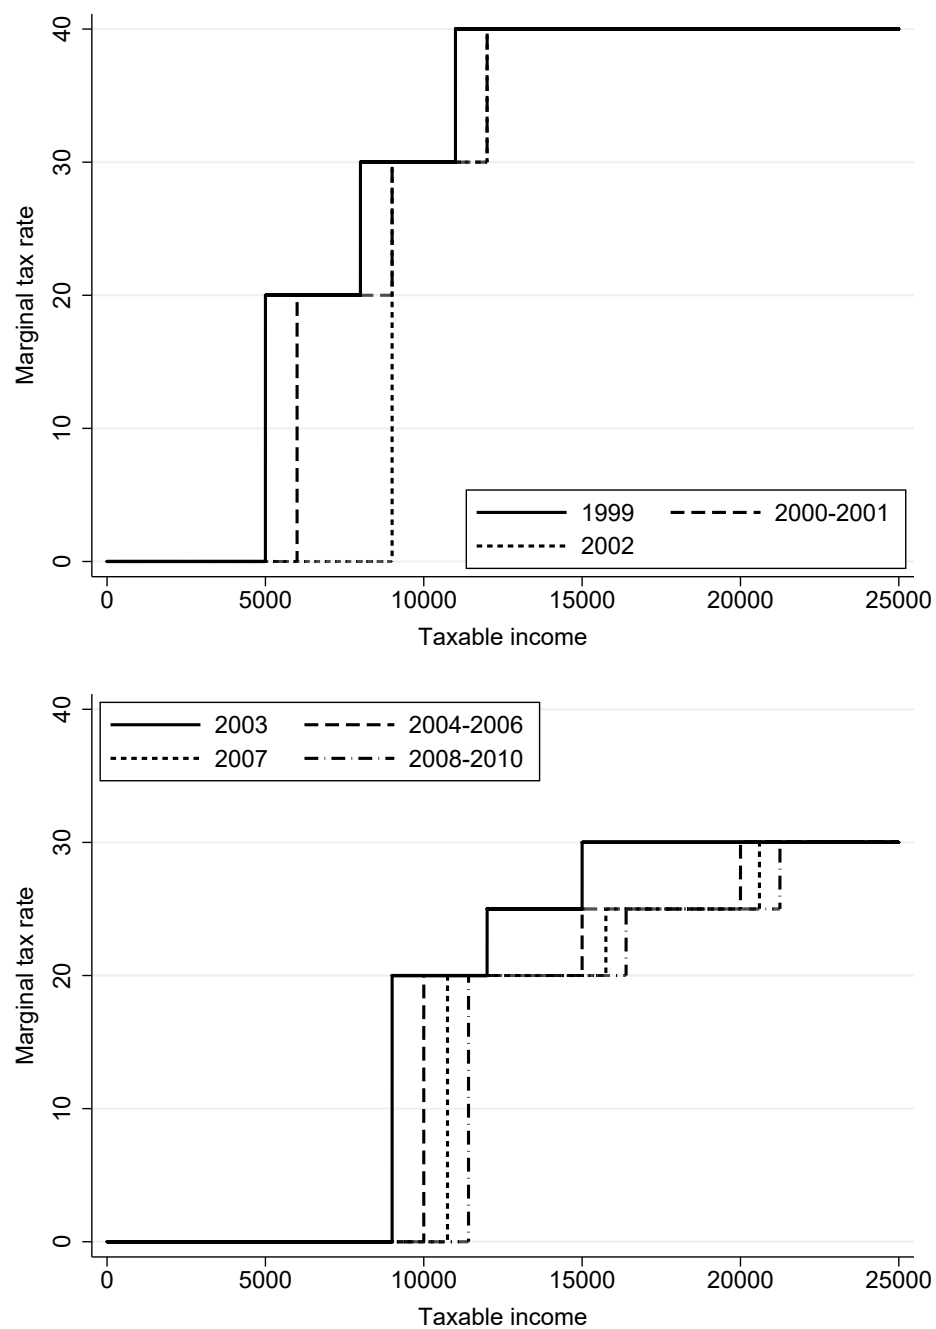
## Main Figures

Figure 3.1: Reporting Thresholds Over the Sample Period




*Notes:* This figure illustrates the thresholds up to which people could claim deductions for charitable contributions without providing receipts. In the years 1999-2002 this threshold was dependent on salary income with 5 different notches in the schedule. From 2003 the threshold became independent of income and was set at CYP 150. In 2008 this threshold was changed again to CYP 175.

Figure 3.2: Schedule of Marginal Tax Rates (Years 1999-2010)




*Notes:* This figure shows the schedule of marginal tax rates in place in the Republic Of Cyprus in the years 1999-2010.

Figure 3.3: Yearly Average Donations Among Donors



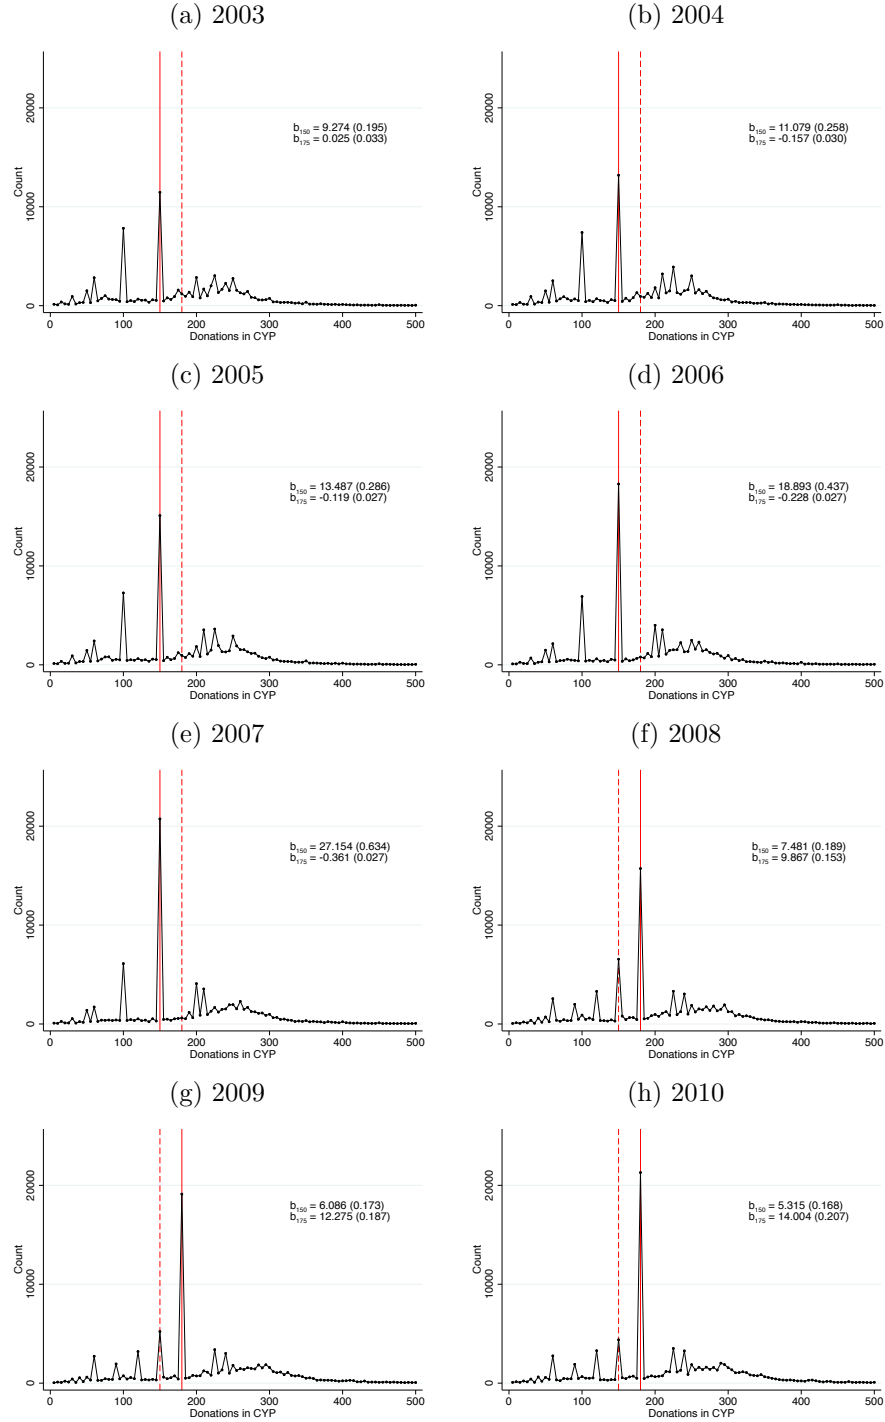
*Notes:* This figure shows the yearly average of positive donations. We remove the top 0.01% of donations within each year. The sample includes all tax filers in the age range 25-54, with some positive salary income and only one job within a given year.

Figure 3.4: Average Donations by Income 1999-2001




*Notes:* This figure shows average donations by income for our main sample, using income bins of width CYP 50 and pooled over the years 1999-2001.

Figure 3.5: Bunching Around Reporting Thresholds



*Notes:* This figure shows the bunching dynamics of donations among salary earners between 2003-2010, by plotting the yearly empirical distributions in bins of width CYP 5. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 150 and CYP 175 thresholds, with bootstrapped standard errors in parentheses. Vertical solid lines mark the relevant threshold that is in place in a given year (CYP 150 during 2003-2007 and CYP 175 during 2008-2010), while dashed lines mark the other threshold that has either been eliminated, or has not been yet introduced.

Figure 3.6: Bunching Estimates over Time



*Notes:* This figure shows for our main sample of salary earners, the estimates of the normalised excess mass around both the CYP 150 and 175 thresholds, between 2008-2010. The shaded areas demarcate 95% confidence intervals.



## Appendix 3A: Extra Tables

Table 3A.1: RD Estimates - Including Rounders

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	38.43*** (0.93)	36.35*** (1.39)	35.81*** (0.95)	34.10*** (1.42)	37.06*** (1.28)	37.40*** (2.00)	34.85*** (1.33)	35.29*** (2.09)
Implied takeup	0.77	0.73	0.72	0.68	0.74	0.75	0.70	0.71
Observations	91 768	91 768	74 064	74 064	46 676	46 676	38 012	38 012
$R^2$	0.26	0.26	0.32	0.32	0.20	0.20	0.26	0.26
<b>Panel B:</b>								
Above 7.5k	29.50*** (0.70)	31.79*** (1.02)	30.10*** (0.75)	30.82*** (1.09)	31.72*** (0.95)	27.66*** (1.42)	30.98*** (1.01)	26.56*** (1.50)
Implied takeup	0.74	0.79	0.75	0.77	0.79	0.69	0.77	0.66
Observations	107 684	107 684	81 197	81 197	60 945	60 945	46 272	46 272
$R^2$	0.33	0.33	0.39	0.39	0.22	0.22	0.30	0.30
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows robustness checks from estimating specification (3.1) when we include rounders to our main sample. In this case, the specification also includes round number fixed effects. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.2: RD Estimates - Years 1999-2002

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above10k	38.18*** (0.83)	34.43*** (1.25)	36.28*** (0.84)	33.41*** (1.27)	35.89*** (1.17)	34.91*** (1.80)	34.74*** (1.19)	35.09*** (1.84)
Implied takeup	0.76	0.69	0.73	0.69	0.72	0.70	0.69	0.70
Observations	114 588	114 588	93 830	93 830	59 023	59 023	48 459	48 459
$R^2$	0.20	0.20	0.29	0.29	0.13	0.13	0.22	0.22
<b>Panel B:</b>								
Above 7.5k	28.75*** (0.69)	30.70*** (1.00)	29.21*** (0.71)	29.50*** (1.04)	29.76*** (0.94)	26.76*** (1.41)	28.78*** (0.96)	26.37*** (1.44)
Implied takeup	0.72	0.77	0.73	0.74	0.74	0.67	0.72	0.66
Observations	116 837	116 837	90 641	90 641	64 941	64 941	50 592	50 592
$R^2$	0.26	0.26	0.36	0.36	0.15	0.15	0.27	0.27
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows robustness checks from estimating specification (3.1) when we also include year 2002 to our main sample. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.3: RD Estimates - Including Individuals with Some Non-salary Income

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	35.07*** (0.89)	31.84*** (1.35)	33.11*** (0.91)	31.07*** (1.39)	33.17*** (1.26)	32.36*** (1.95)	32.34*** (1.30)	31.78*** (2.03)
Implied takeup	0.70	0.64	0.66	0.62	0.66	0.65	0.65	0.64
Observations	96 826	96 826	79 271	79 271	49 475	49 475	40 784	40 784
$R^2$	0.20	0.20	0.27	0.27	0.12	0.12	0.20	0.20
<b>Panel B:</b>								
Above 7.5k	27.63*** (0.72)	29.46*** (1.05)	28.66*** (0.76)	28.80*** (1.11)	28.65*** (0.98)	24.49*** (1.48)	28.30*** (1.03)	24.79*** (1.54)
Implied takeup	0.69	0.74	0.72	0.72	0.72	0.61	0.71	0.62
Observations	105 388	105 388	81 401	81 401	59 031	59 031	45 849	45 849
$R^2$	0.27	0.27	0.35	0.35	0.16	0.16	0.25	0.25
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows robustness checks from estimating specification (3.1) when we include individuals with income from both salary and non-salary earnings to our main sample. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.4: RD Estimates - Professional Taxes and Union Fees Residualised

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	36.58*** (1.17)	35.19*** (1.72)	35.64*** (1.22)	33.45*** (1.79)	35.06*** (1.62)	35.51*** (2.48)	33.55*** (1.70)	32.96*** (2.58)
Implied takeup	0.73	0.70	0.71	0.67	0.70	0.71	0.67	0.66
Observations	62 960	62 960	47 402	47 402	32 269	32 269	24 578	24 578
$R^2$	0.11	0.11	0.24	0.24	0.08	0.08	0.22	0.22
<b>Panel B:</b>								
Above 7.5k	28.28*** (0.86)	27.25*** (1.25)	27.98*** (0.92)	25.68*** (1.34)	26.25*** (1.18)	23.67*** (1.78)	25.11*** (1.25)	22.60*** (1.88)
Implied takeup	0.71	0.68	0.70	0.64	0.66	0.59	0.63	0.57
Observations	75 557	75 557	54 219	54 219	41 914	41 914	30 187	30 187
$R^2$	0.14	0.14	0.26	0.26	0.10	0.10	0.23	0.23
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows robustness checks from estimating specification (3.1) when we net out professional taxes and union payments from our outcome variable. In this case, we exclude individuals working in the public sector. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.5: RD Estimates - Only Highly Unionised Sectors Using Residualised Outcome Variable

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	38.77*** (2.01)	33.67*** (2.81)	38.74*** (1.97)	33.88*** (2.75)	35.09*** (2.67)	34.97*** (3.94)	34.75*** (2.63)	34.43*** (3.87)
Implied takeup	0.78	0.67	0.77	0.68	0.70	0.70	0.70	0.69
Observations	13 805	13 805	13 805	13 805	7 496	7 496	7 496	7 496
$R^2$	0.13	0.13	0.16	0.16	0.11	0.11	0.14	0.14
<b>Panel B:</b>								
Above 7.5k	25.80*** (1.76)	25.70*** (2.67)	27.26*** (1.67)	25.25*** (2.54)	25.63*** (2.45)	20.40*** (3.88)	25.38*** (2.34)	21.82*** (3.70)
Implied takeup	0.65	0.64	0.68	0.63	0.64	0.51	0.63	0.55
Observations	13 531	13 531	13 531	13 531	7 638	7 638	7 638	7 638
$R^2$	0.11	0.11	0.19	0.19	0.09	0.09	0.17	0.17
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows robustness checks from estimating specification (3.1) when we net out professional taxes and union payments from our outcome variable, and restrict our sample to workers in highly unionised sectors (excluding the public sector). The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.6: RD Estimates - Only Highly Unionised Sectors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	34.81*** (1.29)	31.21*** (1.92)	35.52*** (1.27)	31.38*** (1.89)	32.72*** (1.82)	34.62*** (2.81)	33.38*** (1.80)	33.92*** (2.75)
Implied takeup	0.70	0.62	0.71	0.63	0.65	0.69	0.67	0.68
Observations	35 100	35 100	35 100	35 100	18 322	18 322	18 322	18 322
$R^2$	0.22	0.22	0.25	0.25	0.13	0.13	0.16	0.16
<b>Panel B:</b>								
Above 7.5k	31.52*** (1.17)	34.75*** (1.70)	32.44*** (1.15)	35.10*** (1.67)	33.07*** (1.57)	26.87*** (2.36)	33.19*** (1.53)	27.23*** (2.31)
Implied takeup	0.79	0.87	0.81	0.88	0.83	0.67	0.83	0.68
Observations	30 407	30 407	30 407	30 407	17 574	17 574	17 574	17 574
$R^2$	0.29	0.29	0.33	0.33	0.19	0.19	0.23	0.24
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows robustness checks from estimating specification (3.1) when we restrict our sample to those working in highly unionised sectors (including the public sector). The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.7: RD Estimates - Only Sectors with Low Unionisation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	33.08*** (1.54)	33.04*** (2.29)	33.55*** (1.49)	33.60*** (2.22)	33.88*** (2.17)	32.84*** (3.34)	34.57*** (2.11)	33.81*** (3.23)
Implied takeup	0.66	0.66	0.67	0.67	0.68	0.66	0.69	0.68
Observations	33 700	33 700	33 597	33 597	17 134	17 134	17 082	17 082
$R^2$	0.19	0.19	0.25	0.25	0.11	0.11	0.18	0.18
<b>Panel B:</b>								
Above 7.5k	28.99*** (1.12)	27.89*** (1.61)	28.76*** (1.08)	27.60*** (1.55)	27.79*** (1.51)	24.08*** (2.25)	27.51*** (1.45)	25.19*** (2.15)
Implied takeup	0.72	0.70	0.72	0.69	0.69	0.60	0.69	0.63
Observations	40 831	40 831	40 688	40 688	22 635	22 635	22 549	22 549
$R^2$	0.29	0.29	0.35	0.35	0.18	0.18	0.24	0.24
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows robustness checks from estimating specification (3.1) when we restrict our sample to those not working in highly unionised sectors. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.8: RD Estimates - Only Males

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	34.76*** (1.23)	33.04*** (1.80)	34.15*** (1.17)	32.74*** (1.71)	34.16*** (1.71)	32.76*** (2.56)	33.67*** (1.63)	32.78*** (2.44)
Implied takeup	0.70	0.66	0.68	0.65	0.68	0.66	0.67	0.66
Observations	47 705	47 705	47 632	47 632	25 220	25 220	25 184	25 184
$R^2$	0.21	0.21	0.28	0.28	0.13	0.13	0.21	0.21
<b>Panel B:</b>								
Above 7.5k	27.93*** (1.13)	28.64*** (1.64)	27.18*** (1.07)	27.76*** (1.54)	27.96*** (1.53)	22.78*** (2.29)	27.01*** (1.44)	22.91*** (2.13)
Implied takeup	0.70	0.72	0.68	0.69	0.70	0.57	0.68	0.57
Observations	43 398	43 398	43 316	43 316	24 150	24 150	24 097	24 097
$R^2$	0.30	0.30	0.37	0.37	0.18	0.18	0.27	0.27
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows the results from estimating specification (3.1) on our main sample restricting only to males. We pool years 1999-2001. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .



Table 3A.9: RD Estimates - Only Females

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	34.56*** (1.78)	30.41*** (2.81)	35.06*** (1.72)	32.43*** (2.69)	32.45*** (2.63)	33.08*** (4.39)	34.64*** (2.53)	37.30*** (4.21)
Implied takeup	0.69	0.61	0.70	0.65	0.65	0.66	0.69	0.75
Observations	21 095	21 095	21 065	21 065	10 236	10 236	10 220	10 220
$R^2$	0.21	0.21	0.29	0.29	0.12	0.12	0.21	0.21
<b>Panel B:</b>								
Above 7.5k	32.97*** (1.21)	34.20*** (1.73)	34.41*** (1.15)	34.14*** (1.66)	33.71*** (1.59)	28.53*** (2.39)	33.56*** (1.52)	29.40*** (2.30)
Implied takeup	0.82	0.86	0.86	0.85	0.84	0.71	0.84	0.74
Observations	27 840	27 840	27 779	27 779	16 059	16 059	16 026	16 026
$R^2$	0.32	0.32	0.40	0.40	0.21	0.21	0.30	0.30
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows the results from estimating specification (3.1) on our main sample restricting only to females. We pool years 1999-2001. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.10: RD Estimates - Ages 25-39

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	31.50*** (1.29)	29.63*** (1.97)	31.43*** (1.24)	29.97*** (1.90)	31.37*** (1.88)	30.97*** (2.90)	31.90*** (1.82)	32.17*** (2.77)
Implied takeup	0.63	0.59	0.63	0.60	0.63	0.62	0.64	0.64
Observations	42 068	42 068	42 014	42 014	21 096	21 096	21 069	21 069
$R^2$	0.19	0.19	0.26	0.26	0.11	0.11	0.18	0.18
<b>Panel B:</b>								
Above 7.5k	29.92*** (0.99)	33.50*** (1.43)	30.95*** (0.94)	33.01*** (1.37)	32.37*** (1.33)	26.22*** (1.99)	32.05*** (1.26)	26.93*** (1.90)
Implied takeup	0.75	0.84	0.77	0.83	0.81	0.66	0.86	0.67
Observations	45 534	45 534	45 446	45 446	26 164	26 164	26 106	26 106
$R^2$	0.29	0.29	0.36	0.36	0.18	0.19	0.27	0.27
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows the results from estimating specification (3.1) on our main sample restricting only to ages 25 to 39. We pool years 1999-2001. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.11: RD Estimates - Ages 40-54

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b>Panel A:</b>								
Above 10k	39.58*** (1.63)	36.19*** (2.36)	38.80*** (1.52)	37.24*** (2.23)	37.37*** (2.20)	35.88*** (3.40)	37.73*** (2.08)	36.95*** (3.25)
Implied takeup	0.79	0.72	0.78	0.74	0.74	0.72	0.75	0.74
Observations	26 732	26 732	26 683	26 683	14 360	14 360	14 335	14 335
$R^2$	0.24	0.24	0.32	0.32	0.16	0.16	0.27	0.27
<b>Panel B:</b>								
Above 7.5k	29.98*** (1.52)	27.48*** (2.16)	30.12*** (1.41)	26.80*** (2.01)	27.50*** (2.02)	23.57*** (3.01)	26.28*** (1.88)	24.38*** (2.77)
Implied takeup	0.75	0.69	0.75	0.67	0.69	0.59	0.66	0.61
Observations	25 704	25 704	25 649	25 649	14 045	14 045	14 017	14 017
$R^2$	0.31	0.31	0.42	0.42	0.18	0.18	0.30	0.30
Polynomial	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$	$p_1$	$p_2$
Controls	-	-	✓	✓	-	-	✓	✓
Bandwidth	2 000	2 000	2 000	2 000	1 000	1 000	1 000	1 000

*Notes:* This table shows the results from estimating specification (3.1) on our main sample restricting only to ages 40 to 54. We pool years 1999-2001. The *Implied takeup* is calculated as the parameter estimate divided by the notch size in the donation schedule (i.e. in panel A the parameter estimate is divided by 50 while in panel B the parameter estimate is divided by 40).  $p_s$  indicates that we fit a polynomial of order  $s$  on each side of the notch, while controls include sex, year and sector fixed effects. Robust standard errors clustered at the individual level in parentheses, \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.12: Elasticity of Donations w.r.t. Price (Sample: All Highly Unionised Sectors, 1999-2010)

Price variable	(1) OLS	(2) OLS	(3) IV	(4) IV	(5) IV $\Delta_{k=1}$	(6) IV $\Delta_{k=1}$	(7) IV $\Delta_{k=2}$	(8) IV $\Delta_{k=2}$
$\ln(1 - \tau)$	-0.402*** (0.040)	-0.279*** (0.041)						
$\ln(1 - \tau^*)$			-0.598*** (0.045)	-0.481*** (0.047)				
$\Delta \ln(1 - \tau^*)$					-0.554*** (0.050)	-0.499*** (0.050)	-0.556*** (0.059)	-0.459*** (0.059)
Individual FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	-	✓	-	✓	-	✓	-	✓
Observations	84 697	84 697	84 683	84 683	67 852	67 852	57 556	57 556
$R^2$	0.57	0.57	0.57	0.57	0.15	0.15	0.21	0.21

*Notes:* This table shows estimates of the elasticity of donations w.r.t. the price of giving using different specifications. The sample includes only workers in highly unionised sectors, between 1999-2010. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 3A.13: Elasticity of Donations w.r.t. Price (Sample: All, 1999-2010)

Price variable	(1) OLS	(2) OLS	(3) IV	(4) IV	(5) IV $\Delta_{k=1}$	(6) IV $\Delta_{k=1}$	(7) IV $\Delta_{k=2}$	(8) IV $\Delta_{k=2}$
$\ln(1 - \tau)$	-0.392*** (0.020)	-0.344*** (0.020)						
$\ln(1 - \tau^*)$			-0.606*** (0.024)	-0.557*** (0.024)				
$\Delta \ln(1 - \tau^*)$					-0.440*** (0.026)	-0.412*** (0.026)	-0.421*** (0.031)	-0.364*** (0.031)
Individual FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	-	✓	-	✓	-	✓	-	✓
Observations	219 732	219 732	219 706	219 706	168 999	168 999	139 480	139 480
$R^2$	0.65	0.65	0.65	0.65	0.16	0.16	0.23	0.23

*Notes:* This table shows estimates of the elasticity of donations w.r.t. the price of giving using different specifications. In this case, we do not impose any restrictions on our sample. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 3A.14: Poisson Regression (1999-2010)

	(1)	(2)
$\ln(1 - \tau^*)$	-0.557*** (0.052)	-0.470*** (0.057)
Observations	55 091	55 091
Individual FE	✓	✓
Year FE	✓	✓
Controls	-	✓

*Notes:* This table shows estimates from our Poisson specification, using our main sample (the banking sector) between 1999-2010. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 3A.15: Elasticity of Donations w.r.t. Price - Bunchers Vs. Non-bunchers  
(Sample: All Highly Unionised Sectors, 2003-2007)

Price variable	(1) IV	(2) IV	(3) IV $\Delta_{k=1}$	(4) IV $\Delta_{k=1}$
<b>Panel A: All workers</b>				
$\ln(1 - \tau^*)$	-0.428*** (0.131)	-0.426*** (0.148)		
$\Delta \ln(1 - \tau^*)$			-0.345*** (0.126)	-0.363*** (0.126)
Observations	25 822	25 822	22 644	22 644
$R^2$	0.69	0.69	0.26	0.27
<b>Panel B: Excluding bunchers</b>				
$\ln(1 - \tau^*)$	-0.725*** (0.208)	-0.774*** (0.213)		
$\Delta \ln(1 - \tau^*)$			-0.630*** (0.199)	-0.669*** (0.194)
Observations	9 110	9 110	7 708	7 708
$R^2$	0.75	0.75	0.32	0.32
<b>Panel C: Only bunchers</b>				
$\ln(1 - \tau^*)$	-0.361** (0.154)	-0.342** (0.160)		
$\Delta \ln(1 - \tau^*)$			-0.241 (0.162)	-0.232 (0.162)
Observations	16 712	16 712	14 936	14 936
$R^2$	0.64	0.64	0.23	0.24
Individual FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Controls	-	✓	-	✓

*Notes:* This table shows how our estimates of the elasticity of donations w.r.t. the price of giving depends on the inclusion of bunchers, using as our sample workers in all highly unionised sectors between 2003-2007. Bunchers are defined as those who were located at the donation filing threshold in any year during 2003-2007. Panel A includes all workers, B excludes bunchers and C restricts to only the bunchers. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 3A.16: Elasticity of Donations w.r.t. Price - Bunchers vs. Non-bunchers (Sample: All, 2003-2007)

Price variable	(1) IV	(2) IV	(3) IV $\Delta_{k=1}$	(4) IV $\Delta_{k=1}$
<b>Panel A: All workers</b>				
$\ln(1 - \tau^*)$	-0.432*** (0.064)	-0.409*** (0.065)		
$\Delta \ln(1 - \tau^*)$			-0.411*** (0.068)	-0.394*** (0.068)
Observations	55 416	55 416	47 574	47 574
$R^2$	0.74	0.74	0.27	0.27
<b>Panel B: Excluding bunchers</b>				
$\ln(1 - \tau^*)$	-0.555*** (0.105)	-0.523*** (0.107)		
$\Delta \ln(1 - \tau^*)$			-0.593*** (0.101)	-0.580*** (0.099)
Observations	23 604	23 604	20 729	20 729
$R^2$	0.81	0.81	0.30	0.30
<b>Panel C: Only bunchers</b>				
$\ln(1 - \tau^*)$	-0.360*** (0.080)	-0.344*** (0.081)		
$\Delta \ln(1 - \tau^*)$			-0.306*** (0.094)	-0.282*** (0.094)
Observations	31 812	31 812	26 845	26 845
$R^2$	0.64	0.64	0.25	0.25
Individual FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Controls	-	✓	-	✓

*Notes:* This table shows how our estimates of the elasticity of donations w.r.t. the price of giving depends on the inclusion of bunchers, using as our sample workers in both highly and non-highly unionised sectors between 2003-2007. Bunchers are defined as those who were located at the donation filing threshold in any year during 2003-2007. Panel A includes all workers, B excludes bunchers and C restricts to only the bunchers. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.



Table 3A.17: Buncher Stickiness - Linear Probability Model (Sample: All Highly Unionised Sectors)

	(1) All	(2) All	(3) $\Delta(1 - \tau^*) \neq 0$	(4) $\Delta(1 - \tau^*) \neq 0$
$\mathbb{1}_{(\text{buncher})}$	-0.381*** (0.006)	-0.370*** (0.006)	-0.428*** (0.011)	-0.421*** (0.011)
Constant	0.791*** (0.003)	0.816*** (0.009)	0.930*** (0.003)	0.885*** (0.011)
Observations	33 031	33 031	11 548	11 548
$R^2$	0.14	0.15	0.23	0.24
Year FE	-	✓	-	✓
Controls	-	✓	-	✓

*Notes:* This table shows estimates of models where we regress an indicator for having adjusted donations in a given year on a dummy for being a buncher in the previous year, restricting the sample to workers in highly unionised sectors between 2003-2007. The first two columns do not impose any further sample restrictions, while the last two also restrict the sample to those experiencing a change in the (first) tax price of giving. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 3A.18: Buncher Stickiness - Linear Probability Model (Sample: All)

	(1) All	(2) All	(3) $\Delta(1 - \tau^*) \neq 0$	(4) $\Delta(1 - \tau^*) \neq 0$
$\mathbb{1}_{(\text{buncher})}$	-0.425*** (0.004)	-0.423*** (0.004)	-0.473*** (0.007)	-0.468*** (0.007)
Constant	0.746*** (0.002)	0.777*** (0.006)	0.850*** (0.003)	0.863*** (0.009)
Observations	77 269	77 269	27 872	27 872
$R^2$	0.15	0.15	0.19	0.20
Year FE	-	✓	-	✓
Controls	-	✓	-	✓

*Notes:* This table shows estimates of models where we regress an indicator for having adjusted donations in a given year on a dummy for being a buncher in the previous year. The sample includes both highly and non-highly unionised sector workers between 2003-2007. The first two columns do not impose any further sample restrictions, while the last two also restrict the sample to those experiencing a change in the (first) tax price of giving. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

Table 3A.19: Elasticity of Donations w.r.t. Price by Period (Sample: All Highly Unionised Sectors)

	<b>1999-2002:</b>				<b>2003-2010:</b>			
Price variable	(1) IV	(2) IV	(3) IV $\Delta_{k=1}$	(4) IV $\Delta_{k=1}$	(5) IV	(6) IV	(7) IV $\Delta_{k=1}$	(8) IV $\Delta_{k=1}$
$\ln(1 - \tau^*)$	-0.118* (0.069)	-0.067 (0.072)			-0.343*** (0.093)	-0.342*** (0.106)		
$\Delta \ln(1 - \tau^*)$			-0.168* (0.097)	-0.168* (0.096)			-0.434*** (0.096)	-0.394*** (0.097)
Individual FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	-	✓	-	✓	-	✓	-	✓
Observations	25 298	25 298	15 732	15 732	50 461	50 461	44 091	44 091
$R^2$	0.70	0.70	0.27	0.27	0.63	0.63	0.18	0.18

*Notes:* This table shows how our estimates of the elasticity of donations w.r.t. the price varies by different sub-periods, using as our sample workers in all highly unionised sectors. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

Table 3A.20: Elasticity of Donations w.r.t. Price by Period (Sample: All)

	<b>1999-2002:</b>				<b>2003-2010:</b>			
Price variable	(1) IV	(2) IV	(3) IV $\Delta_{k=1}$	(4) IV $\Delta_{k=1}$	(5) IV	(6) IV	(7) IV $\Delta_{k=1}$	(8) IV $\Delta_{k=1}$
$\ln(1 - \tau^*)$	-0.131*** (0.037)	-0.119*** (0.037)			-0.421*** (0.041)	-0.403*** (0.042)		
$\Delta \ln(1 - \tau^*)$			-0.241*** (0.051)	-0.244*** (0.051)			-0.413*** (0.042)	-0.384*** (0.041)
Individual FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Controls	-	✓	-	✓	-	✓	-	✓
Observations	72 314	72 314	44 058	44 058	121 218	121 218	101 912	101 912
$R^2$	0.76	0.76	0.27	0.27	0.69	0.69	0.18	0.18

*Notes:* This table shows how our estimates of the elasticity of donations w.r.t. the price varies by different sub-periods, using as our sample workers in both highly and non-highly unionised sectors. All specifications control for log disposable income. Whenever indicated, specifications also include age squared and a dummy for changing employer as further controls. Robust standard errors clustered at the individual level are reported in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

## Appendix 3B: Extra Figures

Figure 3B.1: Information on Tax Returns Regarding Thresholds

(a) 2002

C MISCELLANEOUS DEDUCTIONS (attach necessary certificates )					
	1	2	1	2	
	DESCRIPTION	AMOUNT £	DESCRIPTION	AMOUNT £	
1	Professional licence / Tax		4	Donations to approved Charities	
2	Contributions to trade unions		5	Deposits under the specific savings scheme of the Housing Finance Corporation	
3	Subscriptions		6	Any other deduction	
TOTAL					

(b) 2003


B MISCELLANEOUS DEDUCTIONS (For donations over £150 please attach certificates / receipts. For donations of a lesser amount you should keep the certificates / receipts to submit when requested).		
	1	2
	DESCRIPTION	AMOUNT £
1	TRADE UNION CONTRIBUTIONS	
2	PROFESSIONAL SUBSCRIPTIONS	
3	DONATIONS TO APPROVED CHARITABLE ORGANISATIONS	
4	ANY OTHER DEDUCTION	
5	TOTAL	

(c) 2008

PART 5 – DEDUCTIONS / ALLOWANCES		
A MISCELLANEOUS DEDUCTIONS (Attach certificates / receipts <b>only</b> for donations over €300. For donations of a lesser amount you should keep the certificates / receipts to submit when requested).		
1	DESCRIPTION	2
		AMOUNT
1	TRADE UNION CONTRIBUTIONS	
2	PROFESSIONAL SUBSCRIPTIONS	
3	DONATIONS TO APPROVED CHARITABLE ORGANISATIONS	
4	ANY OTHER DEDUCTION	
5	TOTAL	


*Notes:* This set of figures shows the information provided on the tax return regarding filing thresholds for different years.

Figure 3B.2: Schedule of Professional Taxes (Before 2003)



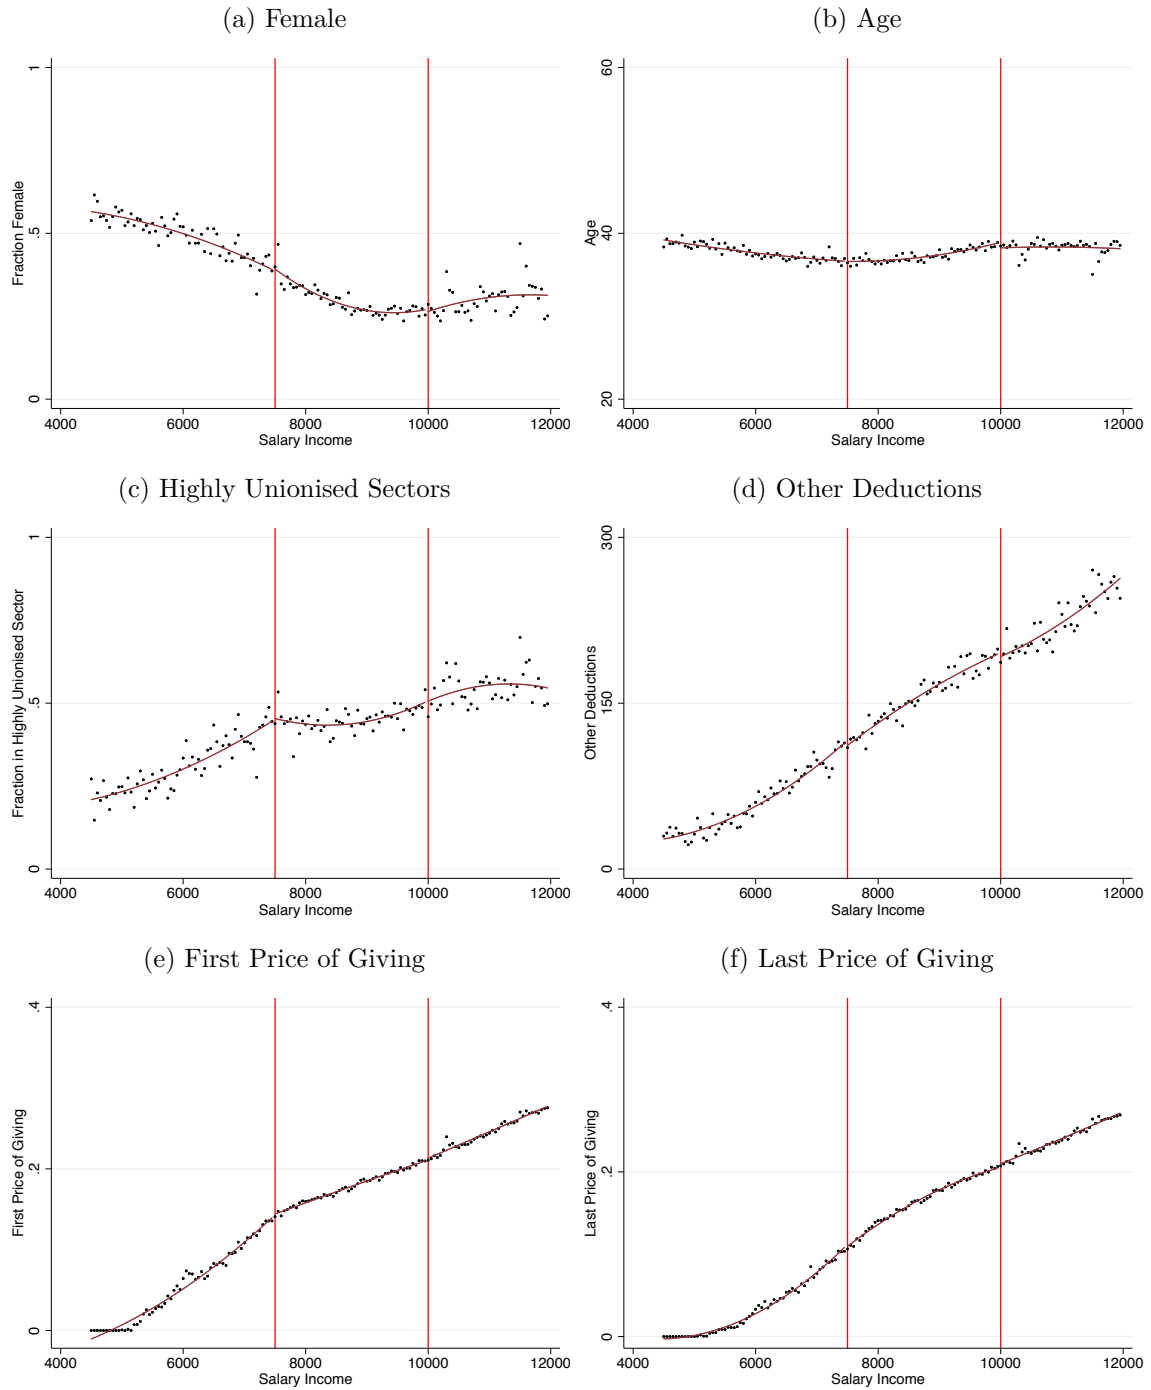
*Notes:* This figure shows the schedule of professional taxes in place in the years prior to 2003. The tax was dependent on salary income, with incremental steps until 30001. Vertical lines indicate the notches in the schedule for deductible donations without the provision of receipts.

Figure 3B.3: Density of Salary Income Between 1999-2001



*Notes:* This figure shows the density of salaries pooled over the years 1999-2001, for two samples: (a) all salary earners and (b) all salary earners excluding those earning at exact multiples of CYP500. In each case, the two earnings thresholds that we focus on (7,500 and 10,000) are marked with vertical lines. Panel (b) also reports the results from a McCrary test for discontinuities in the density of the assignment variable (the estimated log difference in height). The null of no discontinuity cannot be rejected at any of the two earnings cutoffs, in support of the assumptions of our RD design.

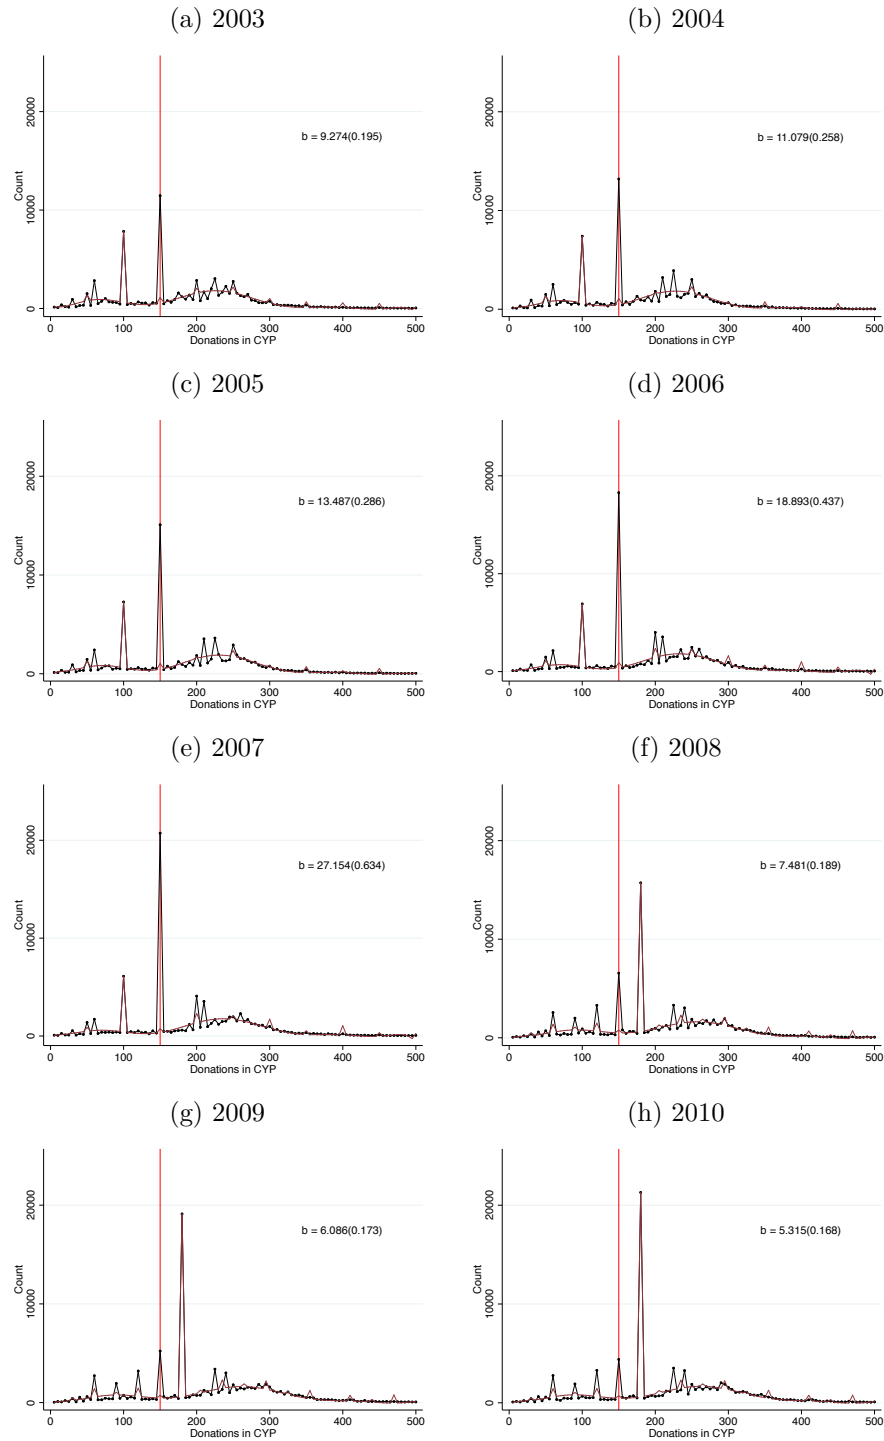
Figure 3B.4: Robustness Check: Smoothness of Covariates



*Notes:* This figure shows evidence in support of the RD identifying assumption. Each sub-figure shows the mean value of the given covariate in bins of width 50 of the assignment variable around each salary threshold. The sample is the same as in our main specification (pooled over 1999-2001).

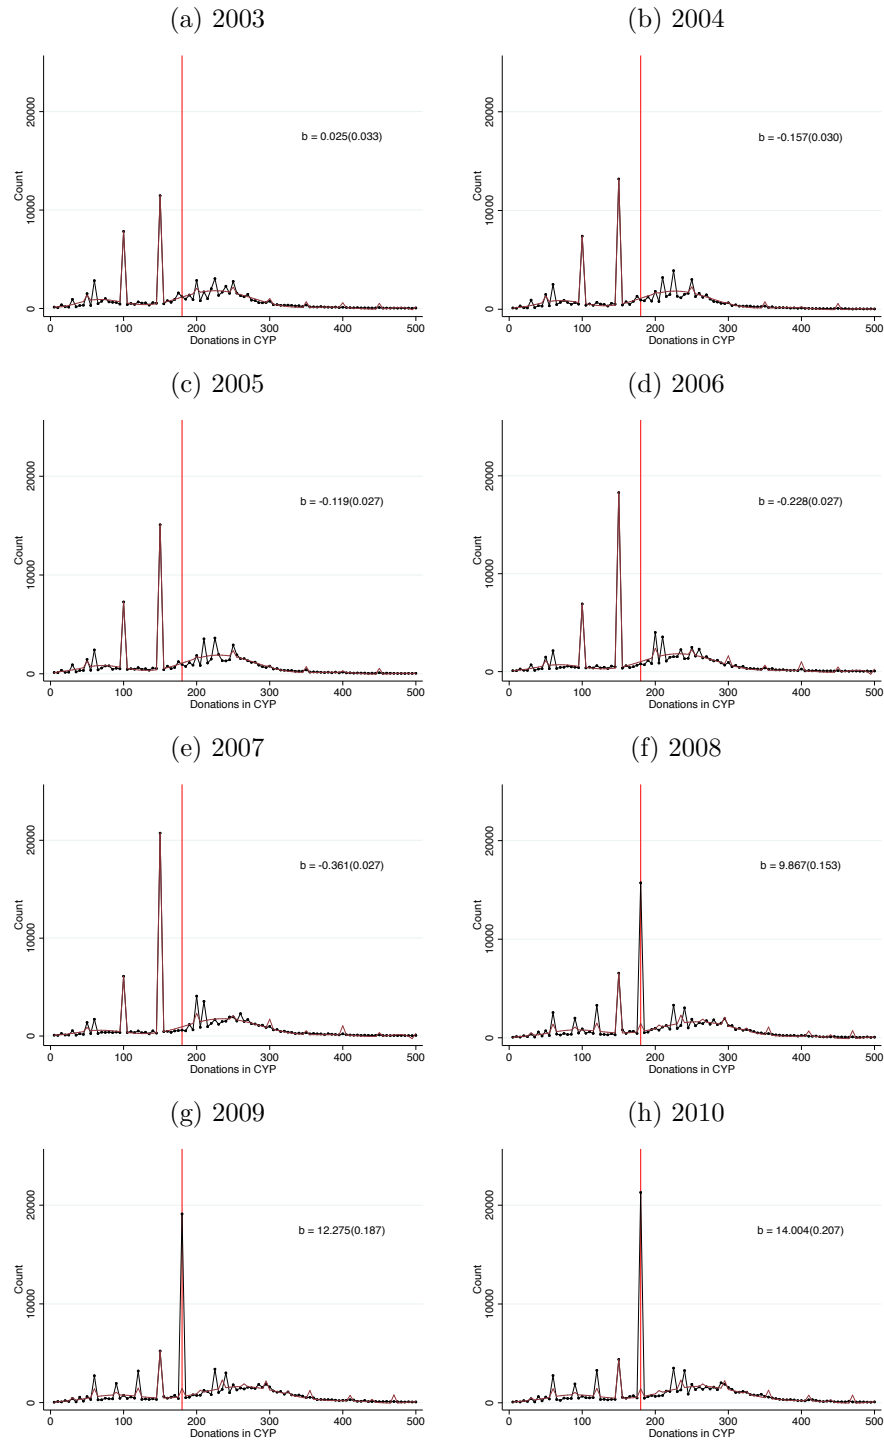


Figure 3B.5: Bunching at CYP 150 with Estimated Counterfactual, Main Sample




*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010. It plots the yearly empirical distribution in bins of width CYP 5, together with the estimated counterfactual. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 150 threshold. Bootstrapped standard errors are in parentheses.

Figure 3B.6: Bunching at CYP 175 with Estimated Counterfactual, Main Sample



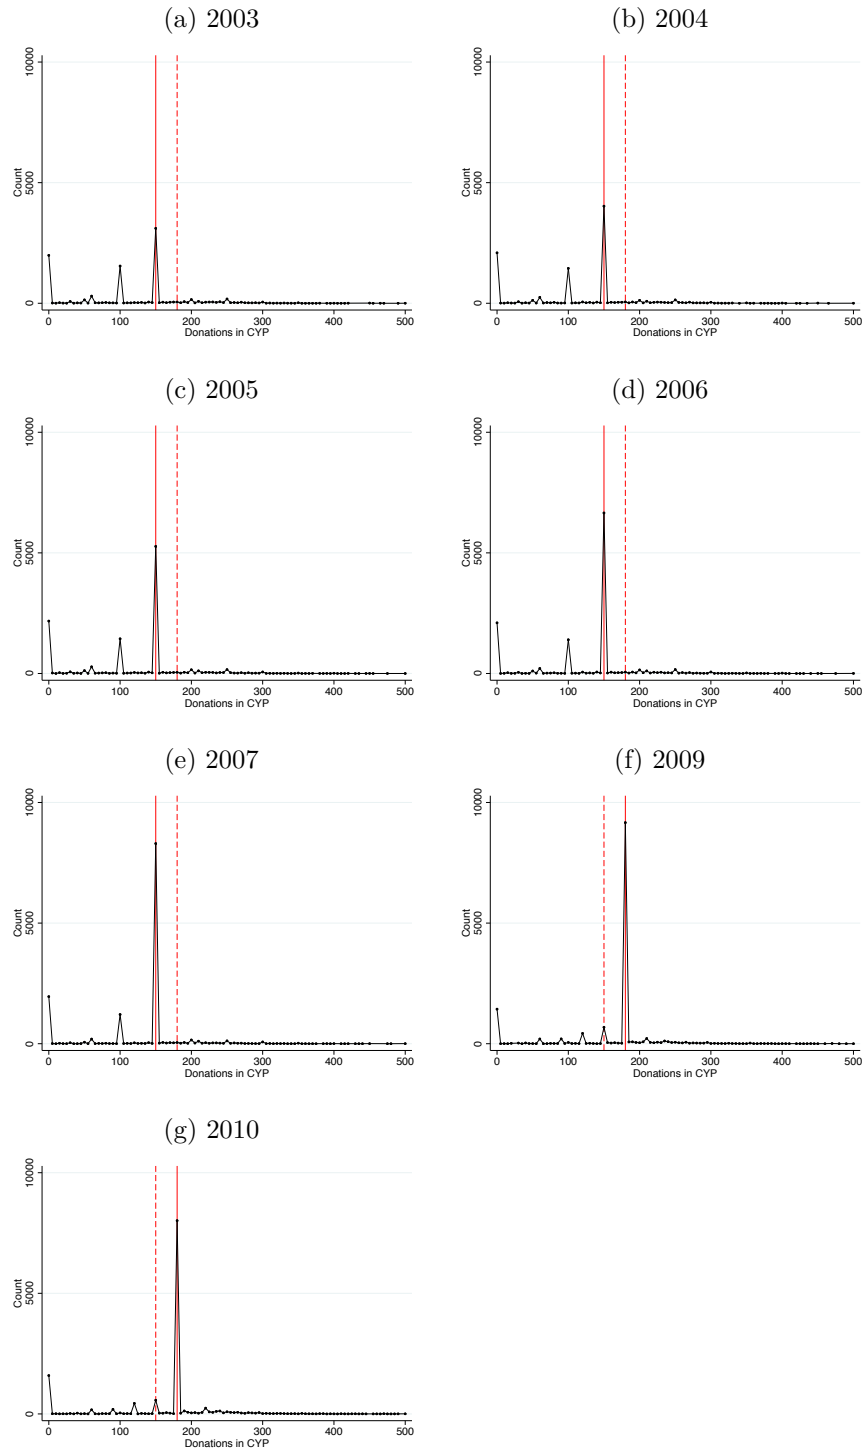
*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010. It plots the yearly empirical distribution in bins of width CYP 5, together with the estimated counterfactual. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 175 threshold. Bootstrapped standard errors are in parentheses.

Figure 3B.7: Donations Between 2003-2010 of Those Bunching at CYP 150 in 2007




*Notes:* This figure shows the empirical distribution of donations before and after 2007, for the sample of salary earners who bunched at CYP 150 in 2007. Vertical solid lines mark the relevant threshold that is in place in a given year (CYP 150 during 2003-2006 and CYP 175 during 2008-2010), while dashed lines mark the other threshold that has either been eliminated, or has not been yet introduced.

Figure 3B.8: Donations Between 2003-2010 of Those Bunching at CYP 175 in 2008




*Notes:* This figure shows the empirical distribution of donations before and after 2008, for the sample of salary earners who bunched at CYP 175 in 2008. Vertical solid lines mark the relevant threshold that is in place in a given year (CYP 150 during 2003-2007 and CYP 175 during 2009-2010), while dashed lines mark the other threshold that has either been eliminated, or has not been yet introduced.

Figure 3B.9: Fraction Reporting Specific Amounts of Donations Over Time



*Notes:* This figure shows the fraction of individuals reporting each of the following amounts of donations over time: over 175, over 150, 150 and 175.

Figure 3B.10: When Did People File Their Taxes for 2008?



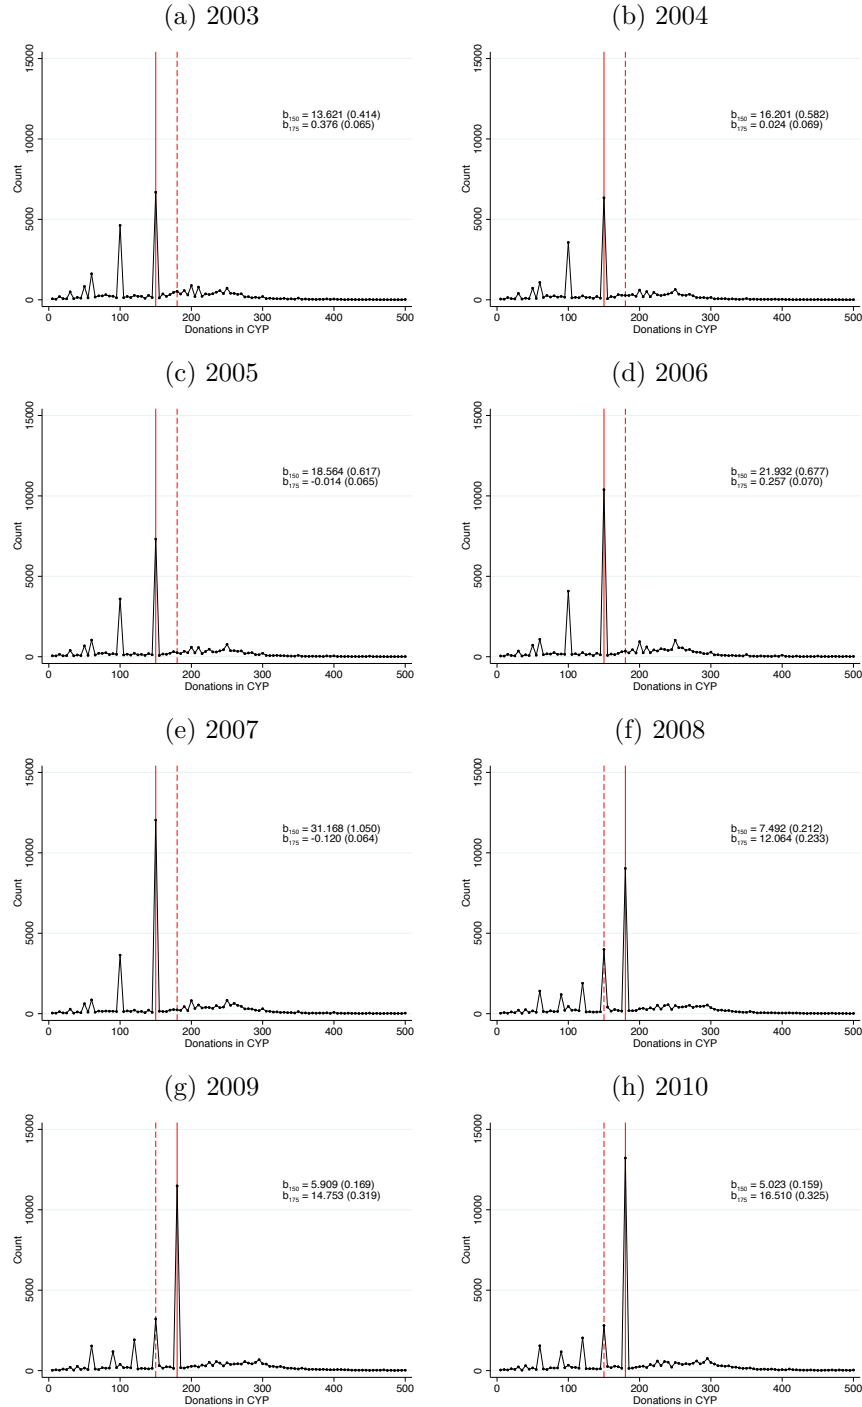
*Notes:* This figure shows, in weekly bins, the fraction of people filing their taxes for the fiscal year 2008. Vertical dashed and solid lines mark the end of the the fiscal year (31 December 2008) and filing deadline (30 April 2009) respectively.

Figure 3B.11: Salary Growth Rates of 2007 Bunchers



*Notes:* This figure shows the yearly salary growth rate between 2003-2010 of salary earners bunching at the CYP 150 threshold in 2007.

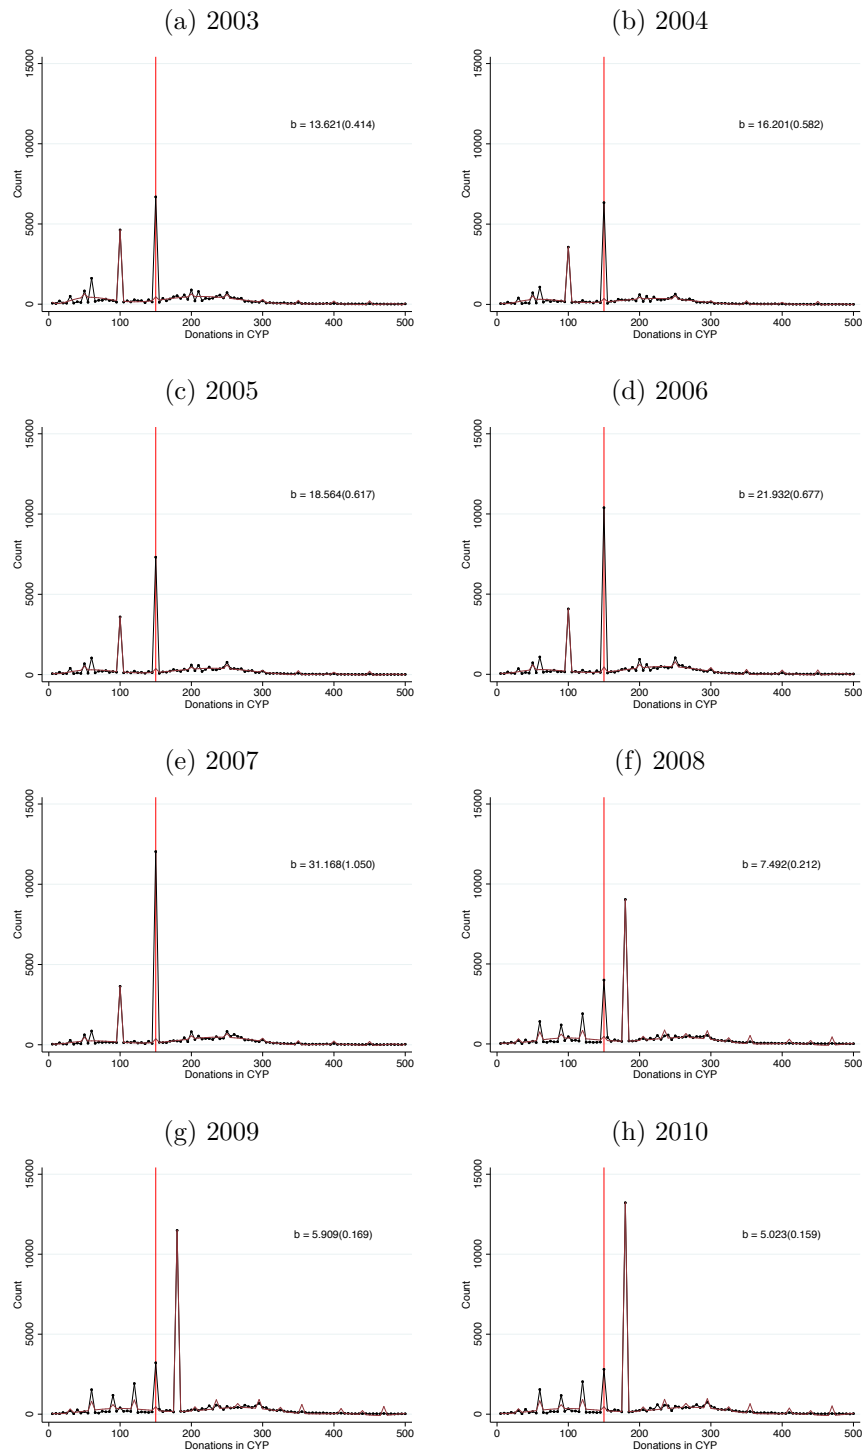
Figure 3B.12: Bunching Around Reporting Thresholds, Excluding Highly Unionised Sectors



*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010 by plotting the yearly empirical distributions in bins of width CYP 5. The sample is restricted to those not in highly unionised sectors and drops those whose sector is not observed. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 150 and CYP 175 thresholds, with bootstrapped standard errors in parentheses. Vertical solid lines mark the threshold that is in place in a given year, while dashed lines mark the other threshold that has either been eliminated, or has not been yet introduced.

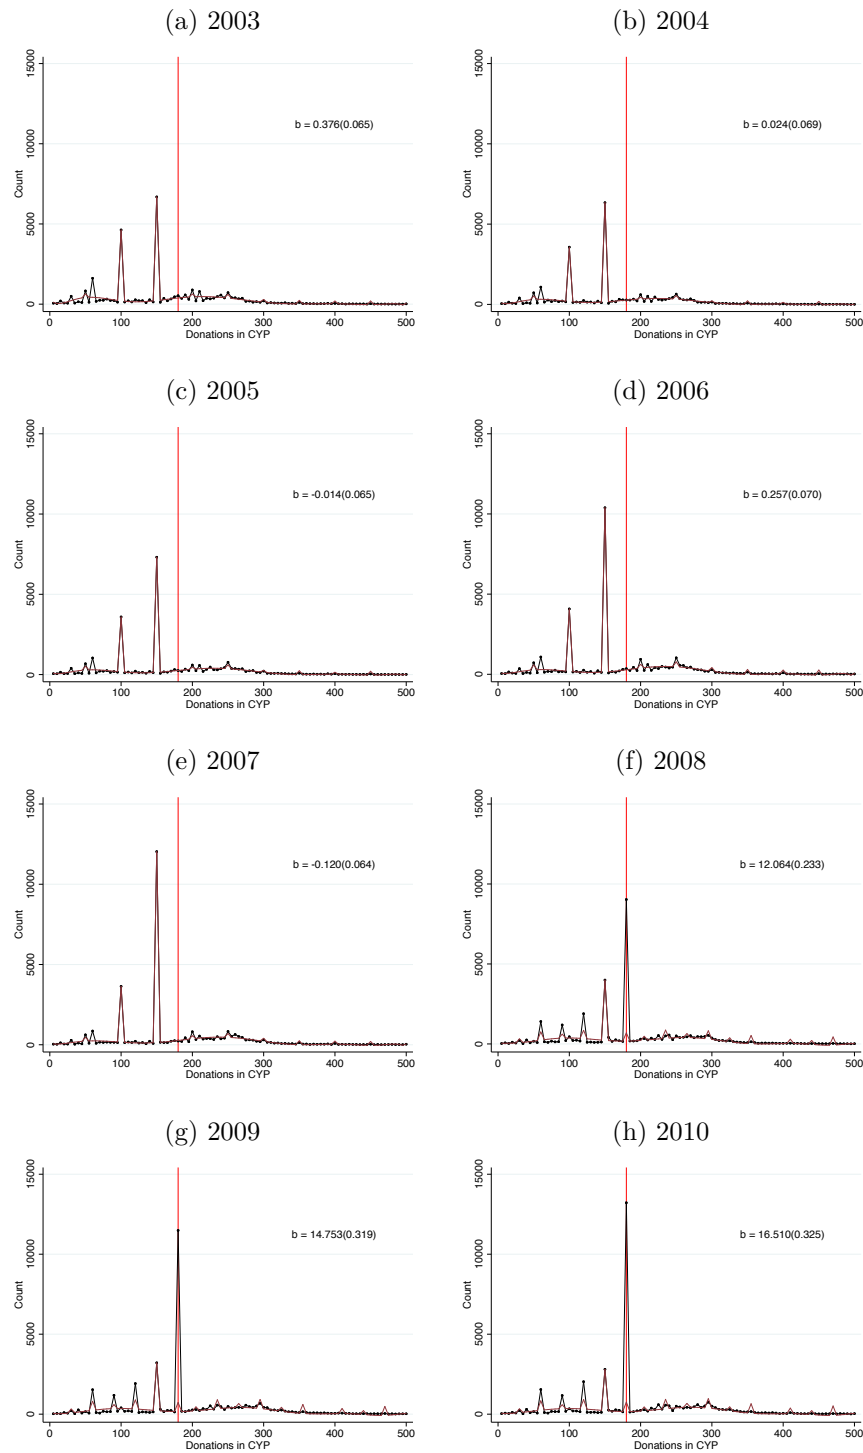


Figure 3B.13: Bunching at CYP 150 with Counterfactual, Excluding Highly Unionised Sectors




*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010, restricting the sample to those not in highly unionised sectors and dropping those whose sector is not observed. It plots the yearly empirical distribution in bins of width CYP 5, together with the estimated counterfactual. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 150 threshold. Bootstrapped standard errors are in parentheses.

Figure 3B.14: Bunching at CYP 175 with Counterfactual, Excluding Highly Unionised Sectors



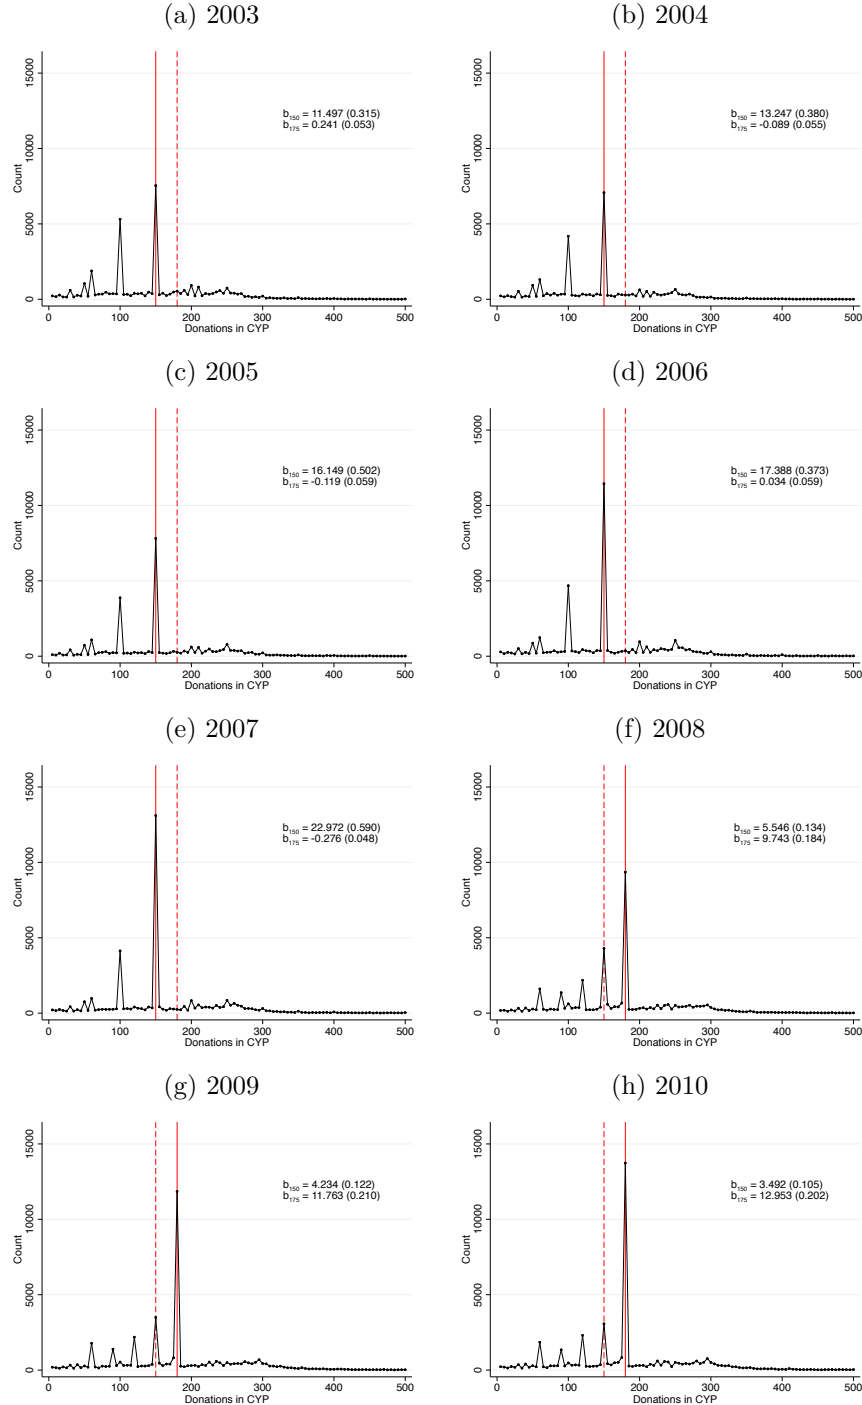
*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010, restricting the sample to those not in highly unionised sectors and dropping those whose sector is not observed. It plots the yearly empirical distribution in bins of width CYP 5, together with the estimated counterfactual. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 175 threshold. Bootstrapped standard errors are in parentheses.

Figure 3B.15: Bunching Estimates over Time, Excluding Highly Unionised Sectors



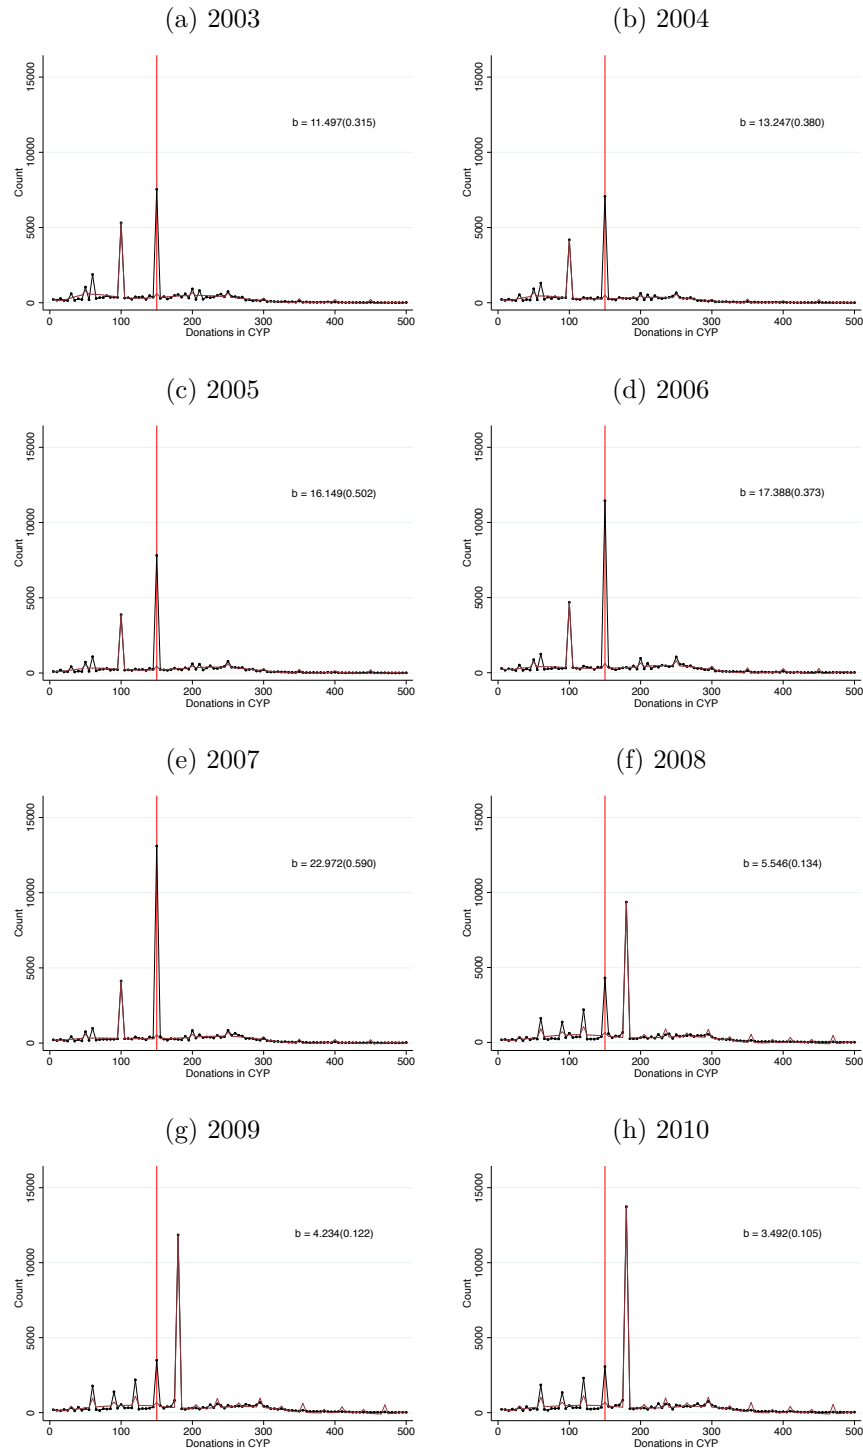
*Notes:* This figure shows the estimates of the normalised excess mass around both the CYP 150 and 175 thresholds between 2008-2010, restricting the sample to those not in highly unionised sectors. The shaded areas demarcate 95% confidence intervals.

Figure 3B.16: Bunching of Donations, Removing Union Fees



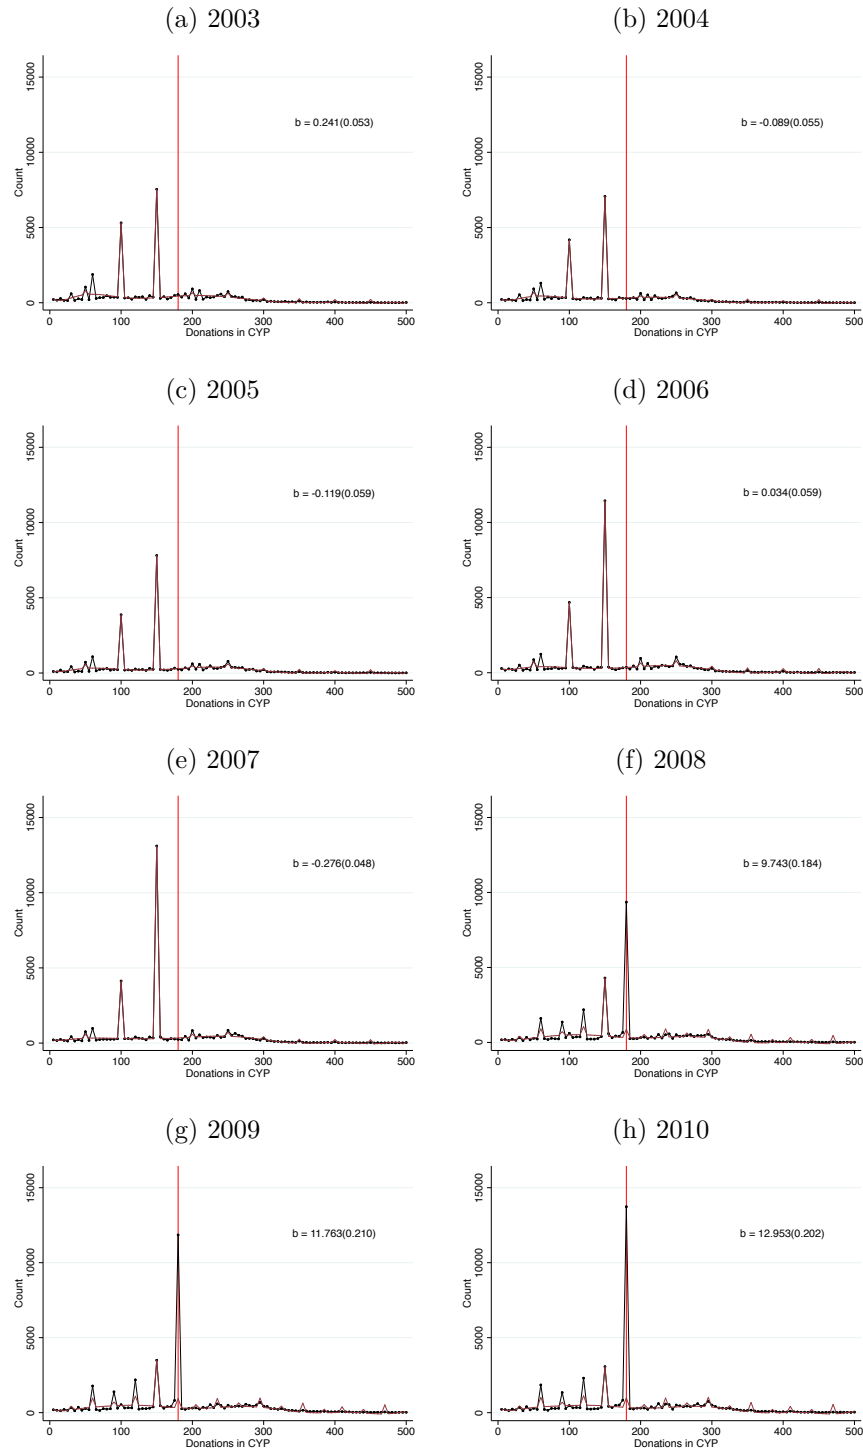
*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010, by plotting the yearly empirical distributions in bins of width CYP 5. The sample is restricted to only those whose sector can be observed. For those in highly unionised sectors, the union fees have been removed. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 150 and CYP 175 thresholds, with bootstrapped standard errors in parentheses. Vertical solid lines mark the threshold that is in place in a given year, while dashed lines mark the other threshold that has either been eliminated, or has not been yet introduced.

Figure 3B.17: Bunching at CYP 150 with Counterfactual, Removing Union Fees




*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010, for the main sample but restricted to only those whose sector can be observed. For those in highly unionised sectors, the union fees have been removed. It plots the yearly empirical distribution in bins of width CYP 5, together with the estimated counterfactual. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 150 threshold. Bootstrapped standard errors are in parentheses.

Figure 3B.18: Bunching at CYP 175 with Counterfactual, Removing Union Fees



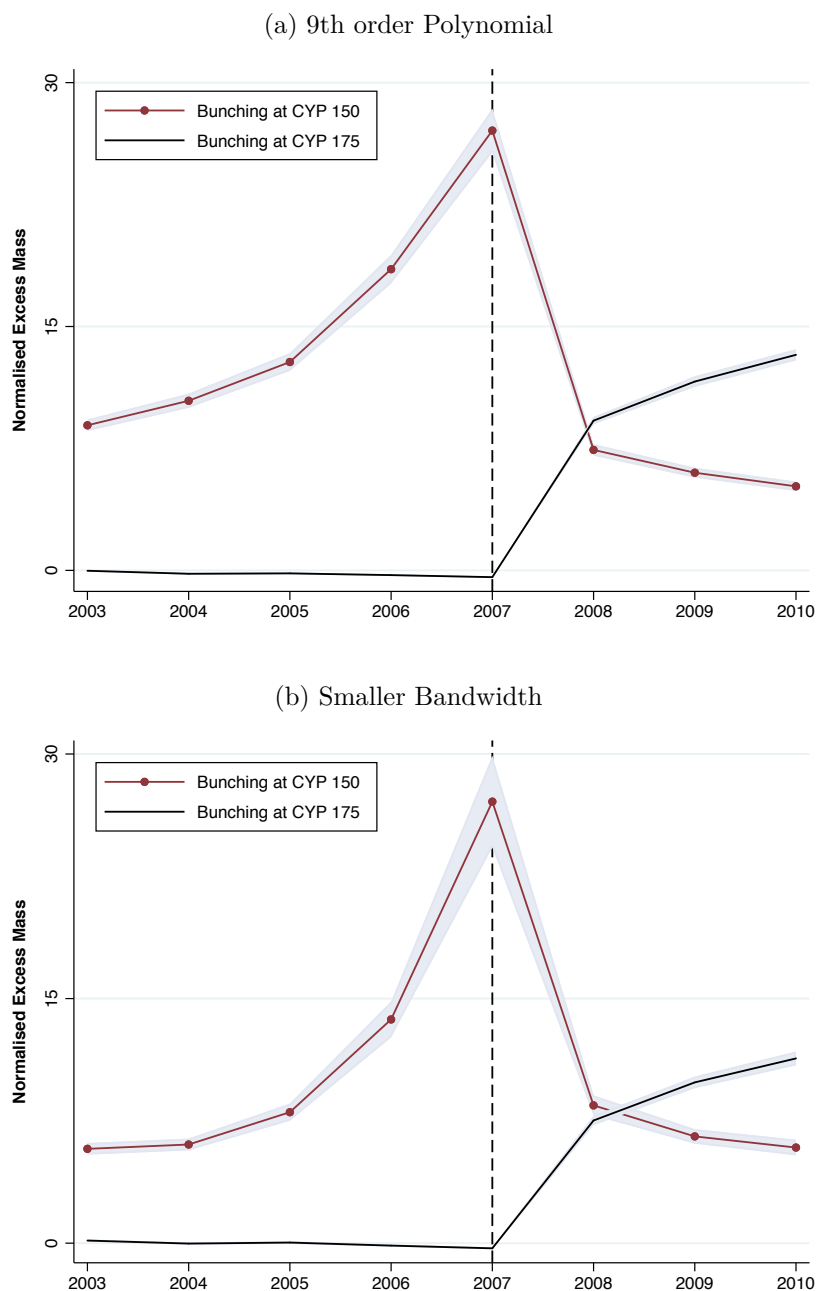
*Notes:* This figure shows the bunching dynamics of positive donations among salary earners between 2003-2010, for the main sample but restricted to only those whose sector can be observed. For those in highly unionised sectors, the union fees have been removed. It plots the yearly empirical distribution in bins of width CYP 5, together with the estimated counterfactual. Each sub-figure reports the normalised excess bunching mass  $b$  around the CYP 175 threshold. Bootstrapped standard errors are in parentheses.

Figure 3B.19: Bunching Estimates over Time, Removing Union Fees



*Notes:* This figure shows the estimates of the normalised excess mass around both the CYP 150 and 175 thresholds between 2008-2010. The shaded areas demarcate 95% confidence intervals. The sample is restricted to salary earners whose sector is observed (but excludes the public sector), and the outcome variable has been adjusted for union fees among those in highly unionised sectors.

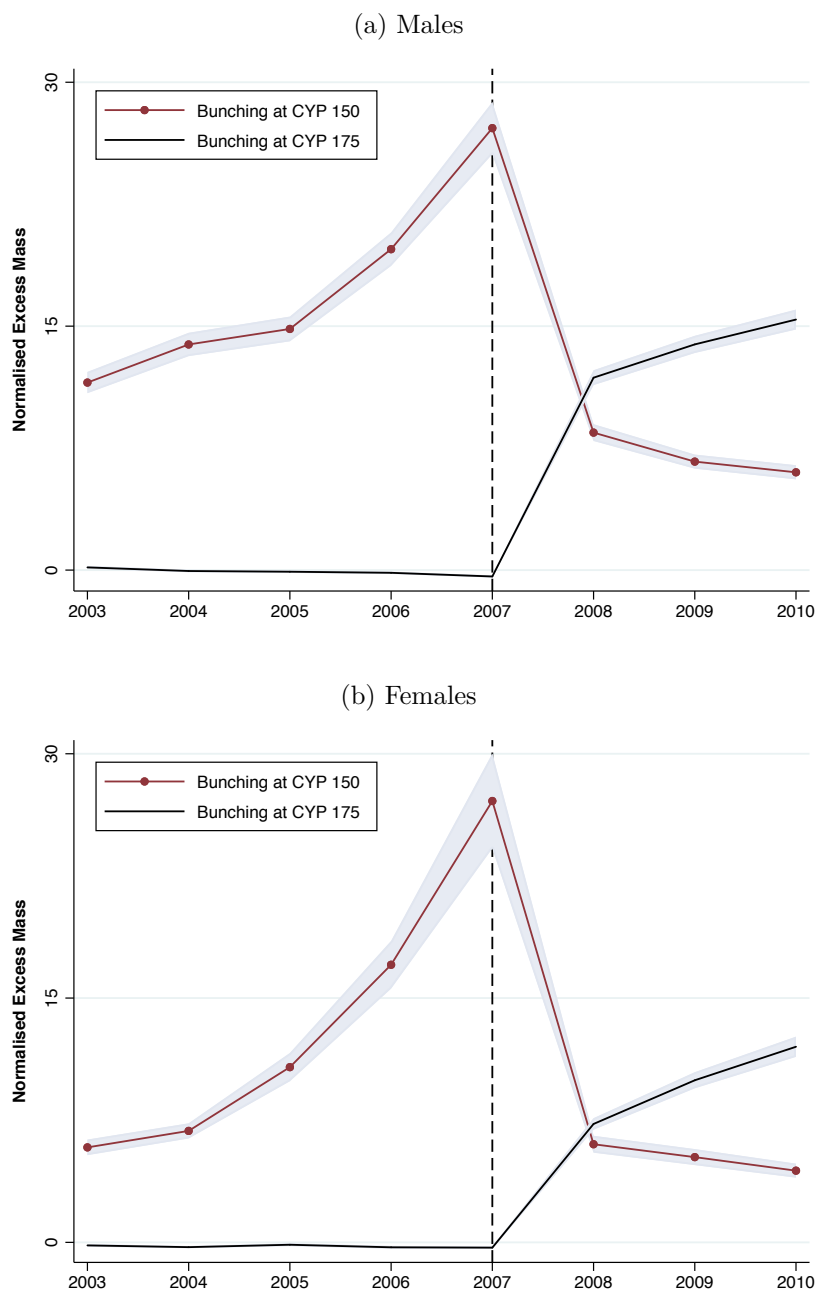
Figure 3B.20: Bunching Estimates over Time, Main Sample, Robustness Checks



*Notes:* This figure shows robustness checks for the estimates of the normalised excess mass around both the CYP 150 and 175 thresholds, between 2008-2010. The shaded areas demarcate 95% confidence intervals. Sub-figure (a) shows results using a 9th order polynomial and (b) restricting the bandwidth to between CYP 50-350.

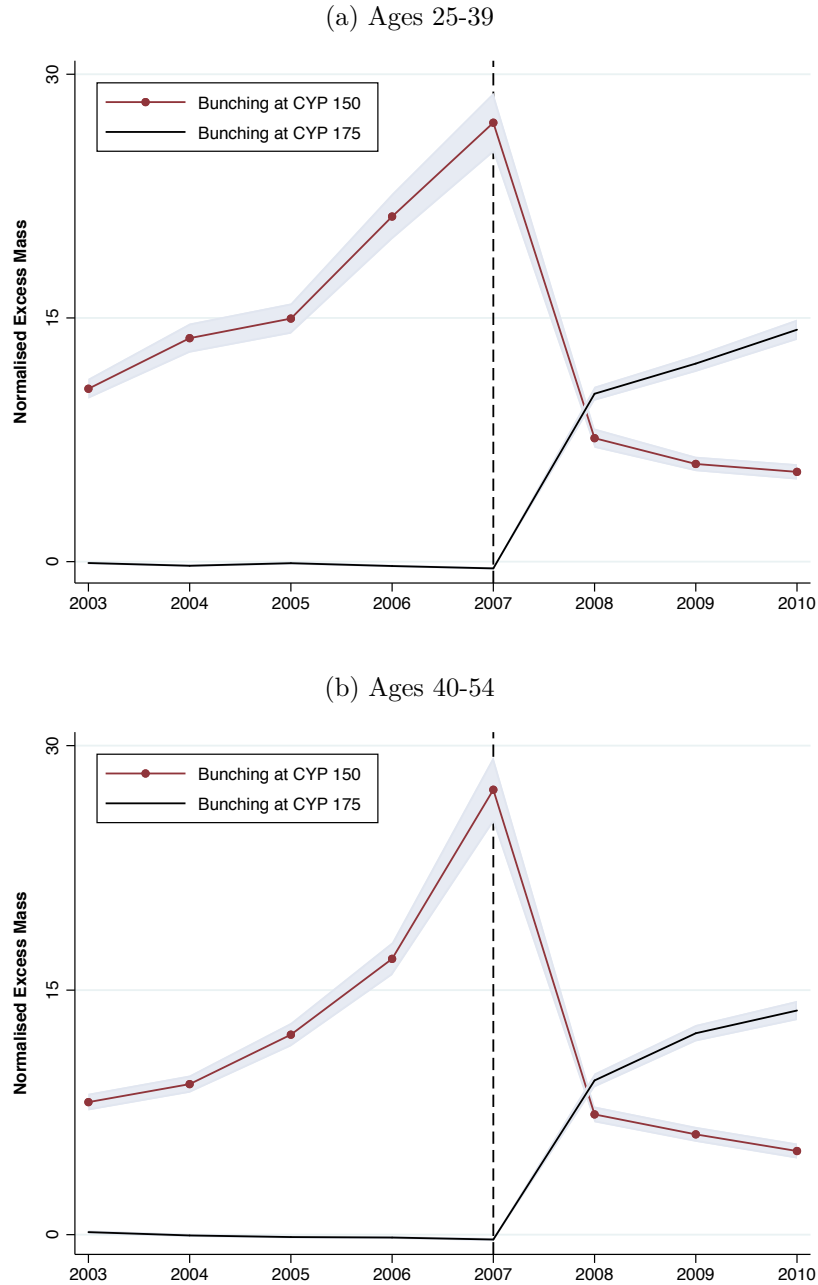


Figure 3B.21: Bunching Estimates over Time, Heterogeneity Analysis by Sex



*Notes:* This figure shows the estimates of the normalised excess mass around both the CYP 150 and 175 thresholds, between 2008-2010. The shaded areas demarcate 95% confidence intervals. Sub-figure (a) restricts our main sample of salary earners to males and (b) to females.

Figure 3B.22: Bunching Estimates over Time, Heterogeneity Analysis by Age



*Notes:* This figure shows the estimates of the normalised excess mass around both the CYP 150 and 175 thresholds, between 2008-2010. The shaded areas demarcate 95% confidence intervals. Sub-figure (a) restricts our main sample of salary earners aged 25-39 and (b) to those aged 40-54.