

DISCUSSION PAPER SERIES

IZA DP No. 12286

**Tax Bunching at the Kink in the Presence of
Low Capacity of Enforcement:
Evidence from Uruguay**

Marcelo Bergolo
Gabriel Burdin
Mauricio De Rosa
Matias Giacobasso
Martin Leites

APRIL 2019

DISCUSSION PAPER SERIES

IZA DP No. 12286

Tax Bunching at the Kink in the Presence of Low Capacity of Enforcement: Evidence from Uruguay

Marcelo Bergolo

Universidad de La Republica and IZA

Gabriel Burdin

The University of Leeds and IZA

Mauricio De Rosa

*Universidad de La Republica and Paris
School of Economics*

Matias Giacobasso

UCLA and Universidad de La Republica

Martin Leites

Universidad de La Republica

APRIL 2019

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Tax Bunching at the Kink in the Presence of Low Capacity of Enforcement: Evidence from Uruguay*

By using a bunching design on rich administrative tax records from Uruguay's tax agency we explore how individual taxpayers respond to personal income taxation in a context with high sheltering opportunities. We estimate a moderated elasticity of taxable income in the first kink point (0.16) driven by a combination of gross labor income and deductions responses. Taxpayers use personal deductions more intensively close to the kink point and underreport income unilaterally or through employer-employee collusion. Our results suggest that policy efforts should be directed at broadening the tax base and improving the enforcement capacities rather than eroding tax progressivity.

JEL Classification: H21, H24, H30, J22

Keywords: personal income taxation, tax bunching, elasticity of labor income, deductions behavior, misreporting, developing economies

Corresponding author:

Marcelo Bérgholo
Instituto de Economía (IECON)
Department of Economics
Universidad de La República
1375 Joaquín Requena
Montevideo 11200
Uruguay
E-mail: mbergholo@iecon.ccee.edu.uy

* We would like to thank Steven Hamilton, Mazhar Waseem, Carina Woodage and Eleanor Wilking, for comments on earlier drafts and presentations, as well as seminar participants at IECON seminar (December 2016), SEU Annual Conference (December 2016), ZEW Public Economics Conference (May 2017), UM seminar (April 2018), Zurich Development Economics and Policy Conference (June 2018), IIPF Annual Congress (August 2018), Phd seminars - Anderson School of Management, UCLA (February 2019). We also wish to thank the Dirección General Impositiva (DGI) by providing us with the administrative tax records, and specially to Gustavo Gonzalez and his staff for their comments to this work. This research was funded by the Agencia Nacional de Investigación e Innovación (ANII) - grant FCE_1_2014_104284.

1 Introduction

The state's capacity of collecting taxes is a critical determinant of long-run economic growth and one of the channels that could be used by democracies to enhance it (Dincecco and Katz, 2016; Acemoglu et al., 2018). Taxation plays a major role in boosting economic development and reducing income inequality in developing countries (Besley and Persson, 2013). Despite in the last years these economies have gradually increased their total tax collection measured as a percentage of the Gross Domestic Product (GDP), this share remains at relatively low levels - around 20%, compared to 35% on average for OECD countries (Besley and Persson, 2013; Gordon and Li, 2009). This has raised the question among policy and academic circles of what should developing countries do in order to increase their tax revenues.

Developing countries already rely heavily on indirect taxation and there is limited room to increase the tax burden associated with excise taxes like Value Added Tax (VAT) or Sales Tax.¹ On the other hand, direct taxes - in particular personal income taxation - play a minor role in the overall tax structure.² This has shifted the center of debate to the role of personal income taxation. Not only as a tool to increase tax revenues (ECLAC, 2013; Turnovsky and Basher, 2009) but also because it could be helpful to reduce the prevailing levels of wealth and income concentration (Duncan and Sabirianova Peter, 2016).

A major concern that should be addressed in this discussion are the efficiency costs of a strengthened personal income tax schedule. The standard economic argument indicates that a narrow income tax base - due to the presence of large informal sectors - and a limited capacity of enforcement should result in large behavioral responses (Bird and Zolt, 2004; Gordon and Li, 2009). Indeed, the natural prediction is that the elasticity of taxable income should be larger in developing countries where opportunities for tax evasion and avoidance are frequent. However, empirical studies that combine both credible identification strategies and high-quality data to gauge behavioral responses to income taxation in developing contexts is still scarce.³ This paper aims to fill this gap.

Using a bunching approach (Saez, 2010; Chetty et al., 2011), we exploit the discrete changes in economic incentives created by discontinuities in the marginal tax rates of a recently implemented progressive personal income tax in a middle income country:

¹OECD estimates that VAT and consumption taxes represent between 49% and 53% of total tax collection in Latin America and Africa, which is about 10% of their GDP.

²For instance, in 2016, the tax collection associated with personal income taxes represented less than 3% of the GDP in regions like Africa and Latin America. This is relatively low compared to OECD countries where on average it represented almost 9%.

³The growing literature that provides reliable empirical evidence about behavioral responses to personal income taxation has been carried out mostly for developed countries (Saez et al., 2012; Kleven et al., 2016).

Uruguay. Specifically, we study individuals' behavioral responses to the first kink point of the labor income tax schedule - i.e. the exemption threshold -. We focus the analysis on the first kink to maximize the credibility of our estimates (Chetty, 2012): it represents the point where the change in the net-of-tax rate is larger (10.5 percentage points) and more salient. Moreover, the first kink point is the kink point with the largest fraction of Uruguayan taxpayers. The identification assumption for this research design is that the labor income distribution in the absence of a discontinuity in the marginal tax rates (counterfactual distribution) is smooth around the kink point.

We use administrative records from the Uruguayan tax agency for the period 2010-2014 to construct a unique dataset that combines self-reported tax records from taxpayers, third-party reported earnings and deductions by the employers (Employer Statements) and firm-level records at the individual level. The result is a comprehensive dataset of almost 6 million observations in a period of five years. The large size and high-quality of the dataset (more than 1.2 million records per year) let us overcome the problems of attrition, measurement error and precision that typically arise when using survey data to identify causal effects (Card et al., 2010). The rich dataset we have available also allows us to explore in detail the mechanisms the taxpayers might use to adjust their reported income to tax incentives.

We divide the empirical analysis into two parts. In the first part, we analyze how individuals respond to labor income taxation and the magnitude of such response. We estimate the excess of mass in the observed distribution of labor income compared to the estimated counterfactual distribution at the first kink point and we calculate the implied elasticity of taxable income (ETI) with respect to the net-of-marginal tax change. We find a positive excess of mass of taxpayers around the first kink point. However, although notorious and statistically significant, the associated ETI is modest and approximately 0.16. Moreover, the observed excess of mass is hump-shaped and somewhat diffuse. Second, we divide the sample according to the different sources of income. Surprisingly, we find a similar pattern and magnitude of bunching for pure wage earners and self-employed individuals. We also analyze if taxpayers' responses follow a pattern consistent with a learning process. We find that the growth rate in the number of bunchers - i.e. people located at the bunching zone - grew by 6.1%, which is consistent with the learning hypothesis. Furthermore, we find that individuals located close to the kink point change their taxable income, but to a lesser extent than those located far from the kink. We also look for heterogeneous responses. Female workers and workers with only one job exhibit the largest responses. Individuals working in smaller firms seem to be more responsive compared to those employed in larger firms. Finally, we show that our results are remarkably robust to a number of different specification checks and definitions of the

sample. Hence, they are not driven by necessary - but partly arbitrary - methodological choices adopted as part of the analysis.

In the second part of our analysis, we investigate the mechanisms underlying bunching responses. Understanding the different margins of response matters because adjustments in real earnings (e.g., hours worked), evasion (e.g., misreporting of income) and avoidance (e.g., deduction behavior) have different consequences in terms of efficiency and policy implications. A limitation of our empirical setting is the lack of exogenous variation to identify the channels of response. In turn, we exploit some specific features of the Uruguayan tax code and different and independent sources of data (third party reported and self-reported).

We analyze separately the two components of taxable income: gross labor income and deductions. This analysis is restricted to the group of workers that only receives income from wages, limiting the number of potential simultaneous responses and simplifying the interpretation of the results. The fact that this group represents more than 95% of the total number of taxpayers reduces the concerns about a loss external validity as a result of restricting the sample.

We start the analysis of mechanisms by analyzing the use of deductions. First, we estimate behavioral responses using the taxable income distribution and the gross labor income distribution. We find a larger response in the former. This suggests that the observed bunching cannot be explained exclusively by the gross labor income. Second, we focus on one specific component of deductions: *itemized* deductions. These are the only ones that need to be explicitly claimed by the taxpayers and therefore they constitute the only margin of response within the deduction channel. We find a stronger response for taxpayers that used *itemized* deductions and we also report evidence of a disproportionate use of *itemized* deductions around the kink point.

Since claiming *itemized* deductions is strongly associated with filing a tax return, we explore whether stronger responses in itemizing taxpayers are associated to the use of deductions *per-se* or to a filing-a-tax-return channel. In order to test this, and because *itemized* deductions can also be claimed through the firm, we compare self reported data (tax returns) and third-party reported data (Employer Statements). We find a larger bunching response when using the self-reported information which is consistent with the filing-a-tax-return channel. However, we find that the additional behavioral response observed in the tax return is entirely driven by a sharp increase in the amount of deductions claimed. This suggests that, although filing a tax return is highly correlated with being a buncher, the main mechanism of response is associated with *itemized* deduction reporting.

Finally, we study behavioral responses in the gross labor income component. These

could be explained by income misreporting - i.e., underreporting the true gross labor income - and/or real earnings responses -e.g., reducing the hours worked. Misreporting can be done in three different ways: 1) unilaterally by the employee, 2) unilaterally by the employer, 3) collusion between both parties. A limitation of our empirical setting, is that it does not allow us to precisely identify each channel separately. In particular, it is especially difficult to disentangle misreporting responses due to collusion between employee and employer from real earnings responses.

We conduct a number of tests to explore if taxpayers bunch at the kink using the gross labor income and if so, how. In this part, we restrict the analysis exclusively on taxpayers that did not claimed *itemized* deductions to avoid conflating deduction and gross labor income responses. First, we explore differences in the size and nature of bunching between tax returns and third party reports for (non-itemizer) taxpayers that filed a tax return. A stronger response estimated with tax return data rather than with firm statements would suggest unilateral misreporting. Bunching at the kink using firm statements would suggest collusion/real response channel. We find evidence of both types of behavior. Then, we compare differences in the gross labor income reported in the tax returns and the gross labor income reported by firms. The difference gauges the extent of misreporting of gross labor income when taxpayers file a tax return. Consistent with the hypothesis of unilateral misreporting by the employee we find a sharp jump in the amount of gross income misreported in the tax return just at the kink point. Finally, we turn to explore the collusion/real response channels by analyzing heterogeneous responses by the number of jobs that a taxpayer has. Intuitively, collusion by employer-employee would be easier for those workers with only one source of income. In line with this argument, we report stronger bunching responses for taxpayers that have less coordination costs (i.e. workers with only one job). This result is consistent with the evidence of a larger bunch of taxpayers working in small-sized firms.

This paper contributes to four strands of literature. First, it adds to the growing empirical literature that analyzes local behavioral responses by exploiting discontinuities in the income tax schedule using the bunching approach (Kleven, 2016). Evidence on bunching responses to personal income taxation have been documented mainly for developed economies: e.g., Austria (Paetzold, 2018), Australia (Hamilton, 2018), Denmark (Chetty et al., 2011; le Maire and Schjerning, 2013), Germany (Schachtele, 2016), Ireland (Hargarden, 2015), Sweden (Bastani and Selin, 2014), and United States (Mortenson, Jacob and Whitten, Andrew, 2018). Most of these papers analyze behavioral responses on gross labor income or in deductions, but not in both components of the taxable income simultaneously. An advantage of our paper is that the empirical setting allows to consider both components of the taxable income at the same time. Indeed, we show that

behavioral responses to the first kink of the Uruguay’s labor income tax schedule can be explained by changes in both, gross labor income and deductions.

Our second contribution relates to the few recent quasi-experimental studies on behavioral responses to income taxes and tax design in developing countries using administrative data. Kleven and Waseem (2013), Best et al. (2015) and Best (2015) analyze responses to the personal income tax in Pakistan, Bohne and Nimczik (2017) analyzes bunching behavior in Ecuador and Bachas and Soto (2018) and Carrillo et al. (2017) do it for corporate taxation in Costa Rica and Ecuador, respectively. The closest paper to ours is Bohne and Nimczik (2017). The authors show how the dynamic learning process of the tax code yields to relevant adjustments in reported taxable labor income in the context of the Ecuador’s personal income tax. The main channel of adjustment in this paper is filling deductions behavior and they identify worker mobility across firms as the most important driver of information transmission. Our paper provides evidence that is consistent with their results but, in addition, we show that evasion behavior and the collusion/real response channels may be also relevant. In contrast with previous literature, we show that wage-earners are at least as responsive as self-employed individuals, even when they have less opportunities to adjust their reported income. However, in line with previous evidence on wage earners (e.g. Paetzold 2018; Chetty et al. 2011), we find that the observed local taxable income response for this population translates into modest elasticities. Indeed, on average the ETIs we estimate fall within the range of elasticities found for individuals in the bottom-middle part of the income distribution in advanced economies.

We also speak to the literature that analyzes the anatomy of taxable income elasticity. Distinguishing between “real” earnings responses (e.g., due to changes in hours worked) and sheltering responses, such as tax avoidance or evasion, has important consequences for the evaluation of income tax systems (Slemrod, 1995, 1998, 2019). As noted by Chetty (2009), the costs of sheltering income should account for transfers between different agents in the same economy. Additionally, how accurate is the elasticity of taxable income as a measure of the efficiency cost depends of the tax setting (e.g., Slemrod and Kopczuk 2002; Kopczuk 2005; Doerrenberg et al. 2017). Beyond efficiency considerations, determining whether behavioral responses are driven by real or “non-real” mechanisms may also affect tax policy prescriptions (Saez et al., 2012; Kopczuk, 2005). The empirical bunching literature, in turn, finds that sheltering responses have played a key role in different tax systems. le Maire and Schjerning (2013) show that intertemporal income-shifting accounts for a significant share of the observed short-run ETI among self-employed workers in Denmark. More recently Bohne and Nimczik (2017), Doerrenberg et al. (2017), Schachtele (2016) and Paetzold (2018) have found that wage earners also

react by changing their deduction behavior. Our paper provides evidence that supports two channels of response to labor income taxation: deductions and income misreporting behavior. The order of magnitude of each of these channels depends on the opportunities that different groups of taxpayers face.

Finally, our findings also relate to a small but growing literature exploring tax evasion behavior of taxpayers using discrepancies between different reports on the same tax base, either, in workers (e.g., Kleven et al. 2011; Kumler et al. 2013; Bergolo and Cruces 2014; Best 2015) or firms (e.g., Fisman and Wei 2004; Carrillo et al. 2017). We present direct evidence of gross income unilateral misreporting by wage earners, but most importantly we document how tax evasion behavior is affected by local tax incentives by showing a sharp increase in underreporting just at the kink point. In line with some previous studies (Bergolo and Cruces 2014; Best 2015), we also report evidence consistent with the hypothesis of collusion between employee and employer in a context of low coordination costs (Kleven et al. 2016). These results indicate the limits of third-party based enforcement systems in developing countries, and highlight the importance of improving the enforcement capacity of tax authorities.

The rest of this paper is organized as follows. Section 2 describes the institutional background details that are essential to understand our research design and empirical strategy. Section 3 describes the empirical strategy as well as some particular features that need to be addressed in the context of Uruguayan IRPF. Section 4 describes the data used as well as our sample restrictions. Section 5 presents the baseline results while Section 6 presents the analysis of the channels of response. Finally, Section 7 concludes and discusses some implications of our findings for tax design in developing countries.

2 Personal Labor Income Taxation in Uruguay (IRPF)

The progressive income tax schedule that we analyze in this paper was first introduced in 2007 in a context of a more general tax reform that aimed to increase the relative weight of direct taxation (Online appendix A.1 contains a more thorough description of the Uruguayan context and tax structure). The personal income tax is referred in Uruguay as *Impuesto a la Renta de las Personas Físicas* (IRPF in Spanish). While the labor income tax base is known as IRPF II, the capital tax base is referred as IRPF I. For the sake of simplicity, and since our study focuses only on the labor income tax base, henceforth we refer to the labor income taxation as IRPF.

Despite the changes introduced by the tax reform, the tax burden is still heavily

based on indirect taxes. According to the Uruguayan Internal Revenue Services (henceforth IRS), VAT constitutes the largest source of tax revenue representing about 50 % of the total tax collection. IRPF represented 12.52% of the whole tax revenue, positioning it as the second component of the tax structure in terms of total tax collection. Even though capital income is part of what is usually categorized as personal income, for simplicity and since we are only focusing on IRPF our definition of personal income will only include labor income.⁴ The IRPF's tax code contains both an individual and a household schedule with separate tax rates and labor income brackets. During the period covered by our analysis only 0.5% of the taxpayers decided to pay taxes as a household unit. For simplicity, these households were excluded from the analysis.

Figure 1 shows the gross labor income distribution in 2013. The vertical dashed lines depict the seven income tax brackets defined by the tax code. IRPF is levied on mid-high levels of the gross labor income distribution: individuals with annual labor income below USD 7,525 - the tax exemption threshold - were exempted of paying IRPF. As a reference, in the same period the yearly minimum wage was USD 3,316 while the median labor income was USD 6,038. According to IRS reports, only about 30% of registered employees paid labor income tax during the period 2010-2014. The total number of workers paying IRPF has been increasing systematically in the last years. Figure A.1.3. Panel a. in the online appendix A.1 shows the evolution of the total number of registered employees between 2010 and 2014. In this period, the total number of individuals registered in the social security - which are *potential* IRPF taxpayers according to their earnings - increased by 9.2%.⁵ Panel b. reports the number of tax returns presented to the IRS (left axis) between 2010 and 2014 and the percentage of taxpayers that made positive payments on behalf of IRPF in the same period (right axis). The total number of tax returns filed increased almost 70% during the period considered: from 102,448 to more than 171,232. Furthermore, the number of workers making positive payments of IRPF also increased, in this case from 27.5% to 37.2% in the five years considered. It is worth mentioning that during 2010-2014 - our period of analysis - the IRPF only suffered minor changes that do not affect our results.⁶

⁴In other words, our measure of taxable personal income do not consider capital sources of income that are taxed with proportional tax rates. Since we estimate the ETI using tax variation in the first kink of the IRPF tax schedule, only about 8% of taxpayers report positive capital income in our analysis sample.

⁵The IRS define potential taxpayers as the workers that could eventually pay IRPF regardless they were above or below the non taxable income threshold. In other words potential IRPF taxpayers are those with positive income from a source that is taxed by the IRPF.

⁶Since its creation, the IRPF schedule has remained virtually unchanged. The bottom brackets did not change at all in the period of analysis while the top bracket were slightly modified: originally, the IRPF schedule consisted only of six brackets with marginal rates ranging from 0% to 25%. In 2012, the last two brackets were rearranged such that a new top marginal rate was introduced: 30%. However, because of how brackets are located in the distribution of the taxable income, these changes only affected

Tax Schedule and Labor Income Components

IRPF relies on two components: a *tax part* and a *deduction part*. As most personal income tax structures around the world, the *tax part* is comprised by the total amount of gross labor income that is subject to taxation and a progressive rate structure. Left panel in Table 1, - Sources of Income - shows a detailed list of labor income concepts that are taxed by the IRPF. Broadly, all items related to wages and self employed income that a taxpayer receives during a fiscal year.

The *deduction part* includes all deductions allowed by the tax code and are also subject to a progressive deduction rate schedule (deduction rates). This is unusual compared to other modern tax system. The set of deductions allowed by the tax code is small and rigid and can be divided into two groups. The first group includes all social security contributions that are computed as a proportion of the gross labor income. This includes payroll taxes and health insurance mandatory contributions. In further sections, we will refer to them as “*non-itemized deductions*” (or “*mechanical deductions*”) since they are proportional to the amount of labor income received by each taxpayer and are automatically considered when calculating the final tax liability. The second group of deductions is formed by all deductions that are not proportional to the gross labor income. We will refer them as “*itemized deductions*” (or “*non-mechanical deductions*”) because they need to be explicitly and voluntarily claimed in order to be considered. This group of deductions can also be divided in two: personal deductions and housing expenses (Right panel - Deductions - in Table 1). Personal deductions include child care deductions and other non-proportional social security contributions (e.g. payroll taxes made by self-employed professionals to their own pension schemes). Housing expenses include both mortgage and rent expenses.

The final tax liability is the result of a three-step process. First, during the fiscal year, all income components detailed in Table 1 are added up to determine the annual gross labor income of each taxpayer.⁷ This amount is passed through the progressive tax rates and a *provisional tax liability* is computed. Panel a. in Table 2 shows the tax schedule for year 2013 with values expressed in United States Dollar (USD). IRPF is comprised by seven different brackets with marginal tax rates that range between 10% and 30%. The first bracket is taxed at a marginal rate of 0% and corresponds to the exemption threshold. Column (3) of Panel a. shows the change in the net-of-tax rate implied when moving from the n_{th} bracket to the $n + 1_{th}$. Figure 2, Panel a. depicts

a few number of taxpayers. These changes have no implications for our results since we are focusing only in the bottom bracket.

⁷In the case of self employed activity, a 30% deduction is automatically applied to the total amount from self employment on behalf of production costs and VAT

the marginal tax rates as a function of the gross labor income and the log-change in the marginal net-of-tax rates.

The second step is analogous. Deductions are added up and passed through the schedule reported in Table 2, Panel b. To the amount resulting of this step we will call it *tax deduction*. Panel b. in Table 2 and Figure 2 depict the marginal deduction rates. As marginal tax rates, marginal deduction rates are increasing with the total amount of deductions. Indeed, the size of the n_{th} deduction bracket is exactly the same as the size of $(n + 1)_{th}$ tax bracket and deductions rates also range from 10% to 30%.

The final step consists of subtracting the *tax deduction* computed in the second step from the *provisional tax liability* computed in the first step. These three steps yield the final tax liability corresponding to the fiscal year considered. Table A.2.1 in Appendix A.2 shows an example of how the final tax liability is computed for an hypothetical taxpayer.

Filing Rules and Withheld Tax

As is common across the world (Slemrod, 2008), the Uruguayan IRS collects part of the IRPF revenues by using monthly income tax withholding. For this reason, many employees are exempted of filling a tax return. Every month employers are required to remit to the tax authority part of their employees wages as in advance payments on behalf of IRPF. The *provisional tax liability* computed by each firm is calculated only with the information that firm's accountants have, i.e. income and deductions related to the employee's activity in the firm. Without further information, accountants can only compute the *deduction part* using proportional deductions because these are the only ones known by them. However, this is only a lower bound for the *deduction part*. Taxpayers can have a more accurate monthly tax withholding by providing further information to the firm's accountant through a 3100 form where they can report additional *itemized* deductions that otherwise are unknown for the firm (i.e. other non-proportional social security contribution or child care expenses). Nevertheless, if a taxpayer wants to claim housing expenses (i.e. rent or mortgage expenses) she necessarily needs to file an annual tax return.

Wage earners receiving labor income from only one source are exempted of filing a tax return but anyone can file a tax return if she wishes to. Most commonly, taxpayers voluntarily file a tax return if they are interested in claiming for *itemized* deductions that were not reported in the 3100 form or if they want to claim for housing expenses deductions. However, there are cases in which taxpayers file a tax return even when they are not required to do so and do not want to claim for additional deductions. Since IRPF is an annual base tax, this could be the case of taxpayers that worked only a fraction of the year but were withheld as if they worked the whole year.

There are two cases in which the tax agency requires taxpayers to file a tax return. First, when they receive labor income from multiple sources and the sum of these items is above a yearly defined threshold.⁸ Second, to all workers that receive at least one dollar from self-employment activities. In general, most workers do not file a tax return; e.g. in 2014 only 14% of all registered labor income earners did it. This is because most workers receive wages from only one source. Hence, unless they want to claim for *itemized* deductions, the amount of tax remitted by the firm to the IRS *via* tax withholdings is equal to the taxpayer's final tax liability. Therefore, filing a tax return does not provide any new or valuable information neither to the IRS or the taxpayer.

All tax returns must be submitted to the tax agency through an electronic form (1102 form). In this form, each taxpayer files both wage and self-employed annual earnings and also *itemized* and *non-itemized* deductions. With the inputs provided by the taxpayer, the 1102 form computes automatically the final tax liability and subtracting the third party withholdings, the remaining tax debit or credit.

In order to give a more precise idea to the taxpayer of how their IRPF tax liability is computed, the IRS developed a tax calculator that can be accessed through their website without any restrictions. (see section A.2.2 in online appendix for a detailed description of the tax calculator and a screenshot of the the spreadsheet provided by the IRS). Although this tax calculator makes IRPF kink points more salient, accessing to this information is costly. Taxpayers must be aware of the calculator and also need to invest some time in downloading this spreadsheet and learning how to use it. Therefore, it is reasonable to think that most taxpayers do not have an accurate idea of where is the exact kink point at which their marginal tax rate changes. Furthermore, since *non-itemized* deductions are tied to the gross labor income but *itemized* deductions are subject to a progressive schedule of marginal rates, calculating the exact point where the kink is located requires some degree of sophistication.

3 Empirical Strategy

In this study, we use a bunching approach to study the behavioral responses to labor income taxation in Uruguay. This methodological approach relies on the idea that kink (and notch) points can provide a credible source of identification to estimate behavioral parameters (Kleven, 2016). Next, we present the bunching formula and the key methodological decisions for estimation.

⁸In 2013, the Tax Agency defined that threshold in approximately 14,000 USD.

3.1 Bunching Formula and Estimation

Briefly, Saez (2010) proves that the elasticity of taxable income can be identified from the excess of mass of taxpayers located at the kink point.

Let $h_0(z)$ be a smooth density function that characterizes the taxable income distribution in an economy with a linear tax schedule and a marginal tax rate of τ_1 . Suppose that a new marginal tax rate ($\tau_2 > \tau_1$) is introduced creating a progressive tax schedule. The tax reform sets a threshold k such that the marginal tax rate applied to the part of the taxable income that exceeds k is τ_2 . Taxpayers whose taxable income is below k are not affected by the reform. Taxpayers whose taxable income is above k are negatively impacted by the new marginal rate since their after tax income is reduced as a consequence of the reform. This results in a new taxable income distribution $h(z)$ that satisfies $h_0(z) = h(z)$ for all taxpayers with $z < k$.

Motivated by the change in the incentives, a number of taxpayers will move to the left in the taxable income distribution to locate just at k , i.e. where the marginal tax rate is still τ_1 . Such mass of taxpayers can be written as:

$$B = \int_k^{k+\Delta z} h_0(z) dz = \Delta z h_0(\xi) \quad (1)$$

for some $\xi \in [k, k + \Delta z]$ where $[k, k + \Delta z]$ represents the interval in which taxpayers are better moving to k .

Feldstein (1999) defines the elasticity of taxable income at z ($\epsilon(z)$) as:

$$\epsilon(z) = \frac{1 - \tau}{z} \frac{\partial z}{\partial(1 - \tau)} \quad (2)$$

Therefore, at the kink k , $\epsilon(k)$ can be written as

$$\tilde{\epsilon}(k) = \frac{\Delta z}{\Delta(1 - \tau)} \frac{1 - \tau_1}{k} \quad (3)$$

Solving for Δz in equation (1) and substituting in equation (3), the elasticity of taxable income can be written as:

$$\tilde{\epsilon}(k) = \frac{B(\Delta z)}{k h_0(\xi) \times \frac{\Delta(1 - \tau)}{(1 - \tau_1)}} \quad (4)$$

where for small changes:

$$\lim_{\Delta\tau, \Delta z \rightarrow 0} \tilde{\epsilon}(k) = \epsilon(k) = \frac{dz}{d(1 - \tau)} \frac{(1 - \tau)}{z} = \frac{B(dz)}{k \times h_0(k) \times \log\left(\frac{1 - \tau_1}{1 - \tau_2}\right)} = \frac{b}{k \times \log\left(\frac{1 - \tau_1}{1 - \tau_2}\right)} \quad (5)$$

where ϵ is the structural compensated elasticity of taxable labor income, k is the kink point introduced by the tax reform, $B(dz)$ is the number of taxpayers that move towards k as a consequence of the reform, $h_0(k)$ is the value of the density function at k before the reform and $\frac{1-\tau_1}{1-\tau_2}$ is the percentage change in the net-of-tax marginal rate. Following Chetty et al. (2011) we define the excess of mass of taxpayers at the kink point relative to the counterfactual distribution as $b = \frac{B(dz)}{h_0(k)}$

All parameters in equation (5) are known or can be estimated from the taxable income distribution data. The kink point introduced by the tax reform (k) as well as the percentage change in the net-of-tax rate ($\ln(\frac{1-\tau_1}{1-\tau_2})$) are directly observable from the legislation. The excess of mass at k ($B(dz)$) can be defined as the difference between the observed number of taxpayers located at k after the reform and the number of taxpayers that would have been at the exact same point absent the reform. Therefore, what we need to estimate is this counterfactual distribution.

We follow the strategy proposed by Chetty et al. (2011). Intuitively, the counterfactual distribution is computed as a fitted polynomial of the actual distribution excluding a portion of the distribution that is nearly enough to the kink. The underlying identifying assumption is that the counterfactual distribution is a smooth function around the kink (Kleven, 2016). In order to estimate the counterfactual polynomial, we first group individuals in income bins and then we estimate:

$$C_j = \sum_{i=0}^q \beta_i^0 (Z_j)^i + \sum_{i=k^-}^{k^+} \gamma_i^0 \cdot 1\{Z_j = i\} + \varepsilon_j^0 \quad (6)$$

where C_j is the number of observations in the income bin j , q is the order of the polynomial and ε_j^0 is the error term. The function $1\{Z_j = i\}$ signals whether the income bin j lies on the bunching zone and controls for being located near the kink. Hence, the excess of mass at the kink point ($B(dz)$) is given by the difference between the observed and the estimated counterfactual distributions:

$$B = \sum_{j=k^-}^{k^+} (C_j - \hat{C}_j) \quad (7)$$

where \hat{C}_j is the counterfactual mass in bin j predicted with equation (6). Then, the estimated excess of mass at k between k^- and k^+ normalized by the counterfactual distribution can be expressed as:

$$\hat{b} = \frac{\sum_{j=k^-}^{k^+} (C_j - \hat{C}_j)}{\left[\frac{k^+ - k^-}{w}\right]^{-1} \sum_{j=k}^{k^+} \hat{C}_j} \quad (8)$$

where w is the bandwidth chosen to group the data and the denominator is the

average density of the counterfactual taxable income distribution between k^- and k^+ . Our analysis focuses only in the first kink of the Uruguayan IRPF schedule. There are a number of reasons that motivate this decision. First, it is here where we found the largest change in economic incentives in terms of marginal tax rate variation. Second, and arguably, the first kink - where an individual starts paying income taxes - is the most salient for the taxpayer. Third, because the bulk of IRPF falls on the lower tail of the gross labor income distribution, the first kink point is the one with the largest share of taxpayers.

In order to make the interpretation of results easier, we normalize the taxable income distribution by the kink point. Hence the kink point will be the $0th$ bin and all distances will be expressed as percentage of k . We group taxpayers such that the distance (d_j) between the j_{th} bin and the kink point is measured in steps of 2%. This should be interpreted as follows: being in the 4_{th} bin, either to the left or to the right of 0, is being 8% away of the kink point. The information used to estimate the counterfactual polynomial is restricted to observations within an 80% of distance relative to the kink point. Following Chetty et al. (2011), all fitted polynomials are degree 7 and standard errors are computed using bootstrap.⁹ The bunching zone is determined through visual inspection for each estimate. In sum, the two key equations can be written as:

$$e(k) = \frac{dz}{d(1-t)} \frac{(1-t)}{z} = \frac{B(dz)}{50 \times h_0 \times \ln(\frac{1}{0.9})} \approx \frac{B(dz)}{h_0} \frac{1}{5.27} \quad (9)$$

$$\frac{B(dz)}{h_0(k)} = \hat{b} = \frac{\sum_{j=k^-}^{k^+} (C_j - \hat{C}_j)}{\left[\frac{k^+ - k^-}{2}\right]^{-1} \sum_{j=k}^{k^+} \hat{C}_j} \quad (10)$$

Since the kink point k must be expressed in units of bandwidth (i.e. how many bandwidths does the kink measure), Equation (9) comes from substituting $k = 50$ in Equation (5). We also substituted $(\frac{1-\tau_1}{1-\tau_2})$ by $\ln(\frac{1}{0.9})$ which is the approximate percentage change in the net-of-tax rate associated with the kink point where the marginal tax rate changes from 0% to 10%. Equation (10) comes from substituting $w = 2$ in equation (8).

3.2 Distance to the Kink Point

As mentioned above, the effective location of the kink points is different for each individual because it depends not only of the gross labor income, but also of the total amount to deductions. Given the progressive nature of the deduction rates, one of the main challenges in our empirical setting is how to precisely calculate these individual kink

⁹In Section 5.4 we present robustness checks to show that our results are not driven by neither of these choices.

points. Next, we summarize the main steps used to calculate the percentage distance to the kink points. The process is explained thoroughly in the Appendix A.3.

The first step consists of combining both *tax and deductions parts* into a new mathematically equivalent structure with only one part. The new income brackets will be called *grossed-up* brackets. The distance to a kink point can be thought as how much additional income a taxpayer can earn (or resign) without changing her marginal tax rate. However, because some deductions are proportional to the gross labor income, receiving (resigning) an additional dollar does not only increase (decrease) the taxable income but also the amount of deductions that can be used. Therefore, measuring distances as the difference between the gross labor income and the *grossed-up* brackets would underestimate the actual margin that a taxpayer has to change their income without changing his marginal rate.

Thus, the second step corrects the distance between the *grossed-up* brackets and the gross labor income by a factor $\frac{1}{1-\tau}$, where τ is the proportional deduction rate associated with the *mechanical deductions*. This rate includes different components such as social security contributions, health insurance contributions and other deductible items associated with the gross labor income. In our data, we can only observe the actual τ , i.e. the current ratio of proportional deductions over taxable income. Therefore, assuming that this ratio does not change with an additional unit of income, we can use the observed $\hat{\tau}$ to correct the distance. Finally, after *grossing up* the brackets and correcting the distances by the $\hat{\tau}$ factor we can compute the actual distance to the kink point as the difference between the gross labor income and the corrected *grossed-up* brackets.

4 Data and Sample Restrictions

Our study relies on a large and high-quality administrative dataset that includes the universe of regular wage earners and self-employed workers for the period 2010-2014. This information was provided by the Uruguayan IRS and comes from three different administrative sources that can be linked at an individual level: 1) personal income tax returns (1102 form), 2) third-party reports of labor income and deductions from employer statements (equivalent of the W-2 form in the United States), and 3) firm level data. The personal income tax returns are only available for those individuals who actually filed a tax return, while third-party reports contains information of all workers registered with the Social Security Administration. The firm level data contains information of firms where registered workers are employed.

We merge the income tax returns with third-party reports and we obtain a dataset that includes gross labor income, tax withholdings as well as *itemized* and *non-itemized*

deductions for the universe of registered workers in Uruguay. For taxpayers who filed a tax return, we have both self and third party reported information on gross labor income and deductions. For taxpayers who did not file a tax return, this information comes from third-party reported records. This dataset also contains a limited number of socio-demographic variables including age and gender of the taxpayer. Finally, we merge the personal income tax data with the firm-level data that contains the total amount of sales and the industry activity code for 2010-2014 period at an individual level.

All the information provided by the tax agency is organized in a final dataset structured at taxpayer-year level. For each observation we have annual reports of labor income, deductions and personal and firm characteristics. Since we focus on responses to labor income taxation, we exclude some taxpayers for whom wage earnings are not their primary source of income. Hence, our sample is restricted to taxpayers who had between 21 and 60 years old at some point of the period. We also exclude taxpayers whose retirement income represented more than 40% of their total income and also those taxpayers who paid at least \$1 of retirement income tax. The total number of taxpayer-year observations with positive labor income during the 2010-2014 period is 6,098,483. After restricting the sample, we have 5,587,711 corresponding to 1,526,301 different taxpayers. This implies that a single taxpayer appears on average 3.7 times in our datasets of potential 5 appearances.

Table 3 provides summary statistics for the final sample of taxpayers pooling all the years included in the sample. The average annual income for the sample of workers considered is about USD 10,440. This can be decomposed in 10,120 USD (96.9%) coming from labor income, 140 USD (1.4%) from capital income and 180 USD (1.7%) from retirement income. At the same time, out of the 10,120 USD average labor income, 9,530 USD (about 94%) come from wages and only 590 USD come from self employed activity. This is a very illustrative picture of how the labor structure in Uruguay relies basically on salaried work. In general, the small share of income from self-employment is explained by the few individuals involved in self-employment. Note that only 4.2% of the taxpayers receive some self-employment income.

In terms of the personal income tax burden, the average tax liability among the whole sample is 580 USD which represents 6% out of the total average labor income. This low average burden is basically driven by the number of taxpayers that actually make positive payments: only 34.2% of the total number of observations in the sample include positive IRPF payments. The average IRPF payment among those who pay is about 1,710 USD. The deductions allowed by the tax code sample represent on average about 20% of the gross labor income for the whole sample. As expected, *Non-Itemized* deductions - i.e. deductions that are proportional to the wage earned - represent a larger

share of total deductions compared to the *itemized* deductions (60% vs 40%). The reason is simply mechanical: all workers that receive positive wages have this type of deductions. Last row in Panel b. shows that a very small fraction of the registered workers file a tax return - e.g. on average only about 11% of all formal labor income earners filed a tax return. In terms of the personal characteristics, 46% of the observations in our sample are women and most of the individuals are aged above 40 years old, roughly 50% are engaged in the services sector, and those who are employees work mostly in firms sized above the median of total sales.

Since the universe of workers is basically comprised by pure wage earners, the characteristics of the whole sample are basically driven by pure wage earners. Some differences arise when we compare pure wage earners (column (3)) with pure self-employed workers (column (5)). Although the differences in the total amount of deductions do not seem to be as large as in total income, because of the latter, the share of deductions relative to the total gross labor income is higher for the sample of pure wage earners (20.5% vs 12.2%). As a result, the average tax liability for pure self-employed is about 43% larger than the one for pure wage earners. In addition, as it was mentioned in Section 2, all workers with self employed income are required to file a tax return and therefore the share of self-employed workers who file a tax return is almost ten times larger than for pure wage earners. It is also important to note that more than 25% of pure self-employed workers are not fulfilling the tax agency requirement of filing a tax return since they did not submit a tax return even when it was required. Differences in the tax return filing rate may be the cause of the differences in the use of *itemized* deductions between pure wage earners and self-employed workers. If taxpayers are required to file a tax return they may be more likely to use *itemized* deductions because the marginal cost of reporting additional items should be low compared to workers that are not required to file one (the cost of filing a tax return is already sunk for workers that are required to file a tax return). In terms of personal characteristics, pure wage earners are less likely to be females, are younger, and are less likely to work in the service sector, relative to their counterparts working as pure self-employed.

5 Main Results

The model presented in Section 3 predicts that some taxpayers that in the absence of the kink would be located within $[k, k + \Delta z]$, will locate just at k as a response to the change in marginal tax rate. This section presents the main estimates for the excess of mass and the implied ETI pooling data from years 2010 to 2014. In all figures, the vertical red line at 0 represents the normalized individual effective kink point i.e. the exact value at

which each individual starts paying IRPF (calculated as explained in Section 3.2). Each blue dot represents the number of taxpayers in the bin. The position of a bin must be interpreted as follows: the first bin (value of 1 in the x-axis) includes taxpayers whose taxable income is between 0 and 2% larger than the effective kink, the second bin (value of 2 in the x-axis) includes taxpayers whose taxable income is between 2 and 4% larger than the effective kink, and so on. The red curve represents the counterfactual estimation of the distribution of the taxable income measured as a normalized percentage distance to the effective kink point. In other words, how the distribution of the taxable income would look like in the absence of the kink. The boundaries of the bunching zone are depicted with gray dashed lines and are defined by visual inspection. All bunching figures include a note with the estimated excess of mass relative to the counterfactual distribution at the kink (b), its standard error and the corresponding ETI.

5.1 Bunching Response Around the (Bottom) Kink

Panel a. in Figure 3 shows the observed and counterfactual taxable labor income distribution around the first IRPF kink point - i.e the exemption threshold, for all the population in our analysis sample. Compared to the predicted mass of taxpayers around the first kink, the excess of mass in the actual distribution is 84%, which is highly statistically significant. When considering the change in the net-of-tax rate, this excess of mass implies an ETI of 0.16, which is relatively modest. While the theoretical prediction is that taxpayers should sharply accumulate at the kink, the excess of mass observed in Panel a. is hump-shaped and somewhat diffuse.

There are many reasons that could explain this non-sharp bunching. For instance, information frictions can make difficult for taxpayers (and in particular wage earners) to have an accurate perception of the tax income schedule (Best, 2015; Chetty et al., 2011; Kleven and Waseem, 2013). Our discussion in Section 3, where we showed that it is not straightforward how to calculate the exact location of the kink point in the IRPF schedule, supports this idea. In addition, even after knowing the exact location of their individual kink point, there could be additional barriers such as rigid hours constraints (Rosen, 1976; Altonji and Paxson, 1988) or searching costs (Rogerson et al., 2005) that prevent them to adjust their income as they want. Even in a context with relative flexibility due to a sizable informal sector (Zenou, 2008), these costs may play a major role in preventing workers from manipulating their income in a precise manner.

In order to provide further empirical evidence for this discussion, Panels b., c. and d. of Figure 3 replicate the analysis presented in Panel a. for three different groups of taxpayers. We divide the sample in three sub-samples according to how they earn

their income: 1) pure wage earners, 2) pure self-employed and 3) both wage earners and self-employed. Panel b. in Figure 3 displays the taxable income distribution and the corresponding estimated counterfactual for pure wage earners. We define pure wage earners as those taxpayers who did not receive any income from self-employment activity within a given year. Since almost 95% of Uruguayan taxpayers fit this definition, Panel b. in Figure 3 confirms that the bunching observed in Panel a. is driven mainly by this type of workers. The graphs depicted in Panels a. and b. are almost identical both in shape and magnitude: the shape of the bunching is diffuse and hump-shaped, the excess of mass w.r.t. the counterfactual density is about 85%, and the implied ETI is 0.16. Although both figures show that these responses are statistically different from zero, an average implied elasticity of 0.16 is in the range of moderate bunching responses.

Adjustment frictions may be less relevant in self-employed workers since they have more scope for tax planning or tax avoidance and hence they may show a sharper bunching pattern (le Maire and Schjerning, 2013; Bastani and Selin, 2014). Since Panels a. and b. of Figure 3 are mostly the result of wage earners behavior, Panel c. displays the same estimate but in this case considering only the group of pure self-employed workers. It is worth noting, that we define pure self-employed workers as those workers whose only source of income comes from independent activity. The excess of mass around the first kink for this particular group of workers point is about 65% and statistically significant. After considering the change in the net-of-tax rate this excess of mass implies an ETI of 0.12. The size and shape of the bunching for pure self-employed workers observed for the Uruguayan IRPF contrast with the sharp responses of similar workers in other countries documented in previous literature. However, the magnitude of the estimated ETI is far from the lowest bound (see e.g. Bastani and Selin, 2014).

Finally, Panel d. in Figure 3 shows the observed and counterfactual distributions of taxable labor income for taxpayers that earn income both in salaried and self-employment activity. This group of taxpayers seems to be less sensitive to the change in the marginal tax rate at the first kink of the IRPF. In this case, the excess of mass is about 51% in the interval around the kink and the implied ETI is 0.10. Compared to the others, this is the group of workers that shows the smallest response.

Since as reported in Table 3, pure wage earners are more than 95% of the total number of workers, restricting our analysis to this group does not largely affects how generalizable our conclusions are. In addition, compared to the group of self-employed workers, restricting the analysis to pure wage earners makes easier to distinguish between the different adjustment channels that could explain the observed responses to taxation.¹⁰

¹⁰Self-employed workers have more access to extensive tax planning compared to wage earners. This could be done for instance via retained earnings, income shifting or deduction expenses ((le Maire and Schjerning, 2013; LaLumia et al., 2015; Harju and Matikka, 2016))

Since the real and avoidance adjustment channels are likely to be related, in the case of self-employed it is more intricate to “isolate” the role of the different mechanisms behind the observed responses of self-employed taxpayers. Namely, earnings supply responses, deduction behavior, and misreporting. Therefore, in what follows we will focus exclusively on the sub-sample of pure wage earners taxpayers.

5.2 Bunching Patterns Over Time: Learning and Persistence

Although the results presented in Section 5.1 show that there is a statistically significant response to the first kink point in the personal income tax schedule, its magnitude is relatively small and the shape of the bunching is somewhat diffuse. As already mentioned, one possible reason is the existence of optimization frictions that may prevent taxpayers to manipulate their income precisely (Chetty et al., 2011). In particular for the IRPF, a relevant kind of friction could be associated to the insufficient knowledge that taxpayers may have about the tax code that defines it.¹¹ Indeed, the IRPF is a relatively novel tax. More precisely, it started to be collected in January of 2008 and in 2014 - the last year of our data- the progressive income taxation in Uruguay was only six years old.¹² Because of its novelty, the period analyzed can be thought as a period of learning in which taxpayers were getting used to the new tax structure. In this context, bunchers may have first examined thoroughly the tax system and, only once they learn how to do it, they started to locate themselves close to the kink point. To explore for this possibility, we present two pieces of evidence.

First, for each year in our data we estimate the total number of taxpayers of taxpayers located at the bunch zone and the share of total labor force that they represent. If workers need time to learn how to locate just before the kink, the number of taxpayers located at the bunch zone should increase with years as more taxpayers become aware of how and where to bunch. It is important to note that even though this test is very simple, it is difficult to argue that there are other reasons that may explain the increase of the number of taxpayers located just at the kink, other than the learning channel. Figure 4, Panel a. depicts the evolution of the total number of *bunchers* (red line) for each year of available data (r.h.s axis).¹³ In addition, the blue line in this figure plots the ratio of *bunchers* over the total number of taxpayers (l.h.s axis). We do this to avoid a mechanical increase in the number of *bunchers* as a consequence of an increase in the overall

¹¹See for instance, Chetty and Saez (2013) and Chetty et al. (2013), for a discussion and evidence about the role of insufficient knowledge about the tax system on taxpayers behaviors.

¹²Before IRPF, the income taxation consisted of a flat tax rate on the personal labor income of workers.

¹³A taxpayer is defined as a “*buncher*” according to the bunching zone used in Panel b. of Figure 3, i.e. those taxpayers who are located in the bunching window (-20% to 6%).

number of registered workers. The total number of bunchers between years 2010 to 2014 increased by 34.7%. This implied an annual growth rate of 6.1% with a starting point of 107,471 *bunchers*. On the other hand, the share of *bunchers* relative to the total number of taxpayers also grew over the period considered from 10.9% to 13.4% although at a smaller rate (4.1%). Part of the increase in the number of *bunchers* could be explained by a mechanical effect resulting from an increase in the total number of taxpayers. However, the difference between the two growth rates can be mostly attributed to a learning process by which taxpayers improve their decision making process year to year.

Second, we measure differences in the rate of income persistence along the taxable income distribution. Specifically, we estimate the share of taxpayers that remain within a 4 bin window compared to the last year distance to the kink. If there is a learning process, and since there were no major tax changes in the analysis period, a taxpayer that accurately bunched in some period is likely to have learned the correct strategy, and it may be easier for him to repeat it in subsequent periods. If this holds, the rate of persistence around the kink should be substantially higher than the rate of persistence for taxpayers located at larger distances of the kink. Panel b. of Figure 4 reports the bunching persistence. The *x-axis* is the exact same used in Figure 3 though in this case in the *y-axis* we plot the proportion of taxpayers located on a window of 8% (or four bins) in consecutive years. This figure shows that the rate of persistence is increasing at the left of the kink point and peaks at the -3 bin (just before the kink) where the rate of persistence is 36.9%. This peak is followed by a sharp decline in the rate of persistence that averages 32.9% to the right of the kink. In other words, the persistence rate is disproportionately high when taxpayers locate just before the kink point compared to taxpayers that are relatively far away of k .

The evidence presented in this section suggests a dynamic component of the bunching behavior which is consistent with the presence of a learning process to respond the local incentives of the IRPF schedule.

5.3 Heterogeneous Responses

In this section, we explore differences in bunching responses for subgroups of individuals according to their own socio-demographic characteristics and the characteristics of the firms where they are employed. Specifically, we perform the same bunching analysis applied in Section 5.1 splitting the main sample into groups defined by gender (male/female), age (aged below/above 40 years old), number of jobs (one/multiple jobs), size of the firm where the taxpayer is employed (below/above median sales) and industry sector (services/goods). For the sake of conciseness we report the estimates of the heterogeneity

analysis (excess of mass and implied elasticities) summarized in Table A.4.1 in online appendix A.4 . The graphical visualization, analogous to our baseline graphics in Figure 3, is reported in Figures A.4.1 and A.4.2.

Left panel in Table A.4.1 reports behavioral responses according to personal characteristics of the taxpayers. Breaking down the analysis by gender, the estimated excess of mass at the first kink is more than three times larger in female workers when compared to male workers (about 150% in comparison to 50%). After accounting for the net-of-tax rate change, the implied elasticities are 0.297 and 0.090 for women and men, respectively. This result is consistent with previous evidence that suggest that womens' earnings are more responsive to changes in the tax rates compared to men (Best, 2015; Chetty et al., 2011; Paetzold, 2018).¹⁴ In terms of their age, we might expect larger responses to the first kink for younger workers compared to older workers if the former are more worried about career concerns, less attached to labor market, or change jobs more frequently. Contrary our expectation, the excess of mass for workers with more than 40 years is larger relative to those younger than 40 years old (roughly 100% vs 67%). A possible explanation for this result is that elderly people may have more experience in the labor market, and thus have a more comprehensive understanding of how tax system affect their earnings. However, these results are inconsistent with pre-existing evidence. However, due to the lack of additional socio-demographic variables, a more detailed study of this hypothesis cannot be performed.

As mentioned in Section 5.1 one of the key determinants of taxpayers' behavioral responses are the adjustment costs. In particular, these responses may be different depending on the type of labor market insertion that individuals have. For instance, the coordination costs of choosing the taxable income such to locate precisely at the kink may be different for workers with only one job compared to workers with more than one job. This includes the cost of managing different sources of income or colluding with employers to (mis)report wage earnings as shown by (Best, 2015). To test this hypothesis, we analyze the bunching behavior for wage earners with only one job and wage earners with multiple jobs separately. We find the excess of mass at the kink to be about 2.5 times higher for the group of wage earners with only one job (100% vs 35%). Furthermore, the observed bunching among the group of wage earners with only one job implies an ETI of 0.184, which is of an order of magnitude very close to elasticity estimated for the whole sample of pure wage earners. We will come back to discuss this point later in Section 6.2.

An additional source of heterogeneity could be related to the evasion opportunities that taxpayers experience in their work environment. The limited amount of variables

¹⁴Specifically, given the literature of labor supply, we might expect that women with children are more responsive to the tax rate changes than either, men or women without children. However, data restrictions do not allow us to test for this specific hypothesis.

included in our data prevent us to observe evasion opportunities directly. For this reason, and based on correlations found in previous literature, we use some firm-level characteristics as a proxy to evasion opportunities. In particular, we use industry-sector (e.g., Best, 2015) and size of the firm (Kleven et al., 2011). It is worth noting that a potential drawback for this analysis is that firm-level characteristics are missing for a sizable number of individuals. Thereby, the results that we obtain using this sample should be taken cautiously.

We first show the results for individuals who work in firms of the service sector and firms of the goods sector, respectively. The leading hypothesis is that firms in the goods sector are more enforced by the IRS compared to firms in the services sector. The reason is that the former is involved in transactions that usually include some type of third party report. Therefore, employers and employees have less margin of response because firm's activity is more enforced and there is less margin for evasion/avoidance. Our results suggest that wage earners in the services sector are more responsive around the kink than their counterparts working in firms in the goods sector. However, the observed differences between the two groups both in terms of the excess of mass and implied elasticity appear to be negligible. Second, we split the sample according to the firm's size in which workers are employed. We use the volume of sales in the fiscal year and divide the sample in two groups: above and below the median of sales. Here, we expect that smaller firms are less controlled by the tax agency and therefore workers on those firms may have more opportunities of evasion. In line with this hypothesis, the observed bunching response is significantly stronger at smaller firms relative to larger firms - an excess of mass of 2.164 *vs* 0.431 which translates into an implied elasticity of 0.412 *vs* 0.082 - . This evidence is consistent with previous literature that showed that earnings responses - mostly driven by evasion behavior - are negatively correlated with firm size (Best, 2015; Kleven et al., 2016).

5.4 Robustness Checks

In this section we report different robustness checks in order to show that our results are not driven by arbitrary decisions neither in the selection of the polynomial degree or the bunching zone or the selection and pooling of the data.

First, we replicate the results reported in Section 5.1 for the four groups reported in Figure 3: 1) all workers, 2) pure wage earners, 3) pure self-employed and 4) both self-employed and wage earners. In particular we replicate the estimates for two additional polynomial specifications: 6th and 8th degree and two additional bunching windows that increase (decrease) the bunching zone of our preferred specification in 1 bin for both

lower and upper bounds. In order to report the results concisely, rather than reporting the visual representation of each estimate, we summarize the results in Tables A.5.1 through A.5.4 in section A.5 in the online appendix. Each table reports the excess of mass, implied elasticity of taxable income and number of observations included in the estimation for each of the groups considered.

Overall, there are no significant differences in the estimates depending on the specific combinations of polynomial degree and bunching window. For all workers, the elasticities vary between 0.093 in the 8th degree polynomial in a bunching window of (-18%, 4%) and 0.178 in an 7th degree polynomial and between (-22%, 8%). The range of these estimates is about 0.080 and the average ETI of the nine estimates is 0.124. The results essentially hold when we look into pure wage earners since they represent more than 95% of the total population. When we analyze the results for the group of pure self employed workers the range of the estimates is narrower (0.041) and the magnitude of the average ETI is slightly smaller (about 0.109). In particular, the maximum ETI for pure self-employed workers is reached in the 7th degree polynomial with a bunching zone of (-12%, 6%) and the lowest ETI is found in an 8th degree polynomial where the bunching area is (-14%, 8%). As in the case of all workers and pure wage earners, our preferred specification lies in the upper bound of the range of estimates. Finally, when analyzing the set of estimates for workers that earn income from wages and self-employed activity, the conclusions are similar. The range of the ETI is about 0.02 being the minimum ETI 0.070 with an 8th degree polynomial estimated in the (-10%, 10%) bunching zone and a maximum of 0.097 that is achieved with a 7th degree polynomial in the window of (-8%, 8%). In general, from the analysis of this set of estimates it follows that our preferred estimates lie in the upper bound of the range of estimates for the ETI for each of the four groups considered. However, it is also worth considering that this is the best fit according to visual inspection.

Another potential concern about the validity of our results is that they could be driven by changes in the composition of the taxpayers included in the sample, i.e. if entries and exits of the sample are biased towards specific characteristics of the workers that are also related with their bunching behavior. We test this by replicating the analysis of our main results in a balanced panel, i.e. including taxpayers that were in the sample during the five years considered (Figure A.5.1 and Table A.5.5 in online appendix). There are no significant differences in our estimates when comparing the results using the balanced and unbalanced panel. The excess of mass for all workers in the balanced sample is 0.965 which implies an ETI of 0.184 while the excess of mass for the all workers in the unbalanced panel was 0.844 and the implied ETI 0.161. The same slightly larger response is observed for pure wage earners, pure self-employed and both wage earners and self-

employed. However, these differences are not economically relevant and do not affect the overall interpretation and conclusions of our main results.

6 Explaining Bunching Responses. Deductions and Gross Labor Income

So far we have documented a non-negligible bunching around the first kink point in the Uruguayan IRPF. We also presented some evidence that supports the hypothesis of a learning process and a preliminary analysis of heterogeneous responses which was very restricted by the quality of our socio-demographic and firms data. However, the mechanisms used by the taxpayers to locate close to the kink remain unclear. In this section we go one step further and we focus on the potential mechanisms behind the behavioral responses documented above.

Taxable labor income can be defined as a combination of gross labor income and deductions. Hence, the observed responses in taxable labor income to personal income taxation should be also a mix of gross labor income and deduction responses (Hamilton, 2018). We first explore the role of deductions in accounting for the observed behavioral responses to the IRPF. Next, we analyze the importance of the other component of the taxable income: gross labor income. In particular, we focus on the role of *real earnings* and evasion (*misreporting*) mechanisms. We do this by studying the behavioral responses of pure wage earners. As argued in Section 5.1, restricting the sample to this particular group of workers limits the alternative channels of response that may be operating simultaneously and makes the interpretation of results more straightforward.

It is worth noting that the evidence presented henceforth is not intended to be interpreted as causal. The reason is that we do not observe any source of exogenous variation in the tax schedule for this group of workers since the IRPF has remained essentially the same since it was first implemented. In spite of that, the institutional context - different ways of claiming deductions and different rules for filing a tax returns- and our empirical setting, allow us to provide valuable insights about the anatomy of these behavioral responses.

6.1 The importance of Deductions

Traditionally, empirical work has focused on the elasticity of taxable income by analyzing the income side of the response. However, in the last years a new strand of literature has claimed that when analyzing the mechanisms behind the taxable income responses, deduction responses are at least as important as gross labor income responses. In partic-

ular, recent evidence has shown that deduction may in some circumstances be the major channel of response. Based on this literature, the leading hypothesis for this section is that taxpayers indeed make use of deductions as one of the channels of response to locate just at the kink. Because of the lack of clear exogenous variation to identify directly the role of deductions, we provide evidence based on a number of alternative tests that leverage different features of the institutional background and test for the use of deductions as a driver of the behavioral response using different avenues.

Gross Income versus Taxable Income

Our first test consists of replicating the bunching analysis reported in Figure 3 using two different concept of income: gross labor income - i.e. labor income before considering deductions - and taxable labor income - the combination of income and deductions that represents the tax base-. This is similar to the tests performed by Chetty et al. (2011), Bastani and Selin (2014) and Schachtele (2016). By comparing these two distributions we can have preliminary evidence of the components that are involved the bunching response observed in Figure 3.

Figure 5, Panel a., depicts the gross labor income distribution. Since this definition of income does not include deductions, the x-axis represents the distance to the statutory kink point defined by the tax schedule (see Table 2). Unlike the distance to the first kink in the taxable income distribution measure, the distance to the first kink of the gross labor income is the same for all the workers considered in the analysis (7,525 USD in 2013). In this case, the value of the bin in which a taxpayer is located should be interpreted as follows: e.g. being in the second bin to the right of the kink point in the gross labor income distribution is having a gross labor income that is between 2 and 4% larger than 7,525 USD. The remaining features of the figure should be interpreted as usual. In order to make the comparison of gross labor income and taxable labor income distributions easier, Panel b. in Figure 5 depicts the same results that Panel b. in Figure 3.

On the one hand, Panel a. of Figure 5 shows that there is no sign of bunching around the first kink of the gross labor income distribution. In other words, before considering deductions, there is no evidence of a behavioral response to persona income taxation. On the other hand, as already mentioned in Section 5.1, after considering deductions (i.e. using the taxable income distribution instead of the gross labor income distribution), there is a clear bunching pattern around the first kink point. Since taxable income is defined as a combination of gross labor income and deductions, this piece of evidence does not allow us to claim that bunching is driven exclusively by deduction behavior because both components must be considered simultaneously. However, it allows us to claim

the deductions are at least as important as gross labor income to explain the bunching behavior since without them there is no sign of an excess of mass around the statutory kink. This first piece of evidence is consistent with our hypothesis that deductions play a major role in explaining the bunching behavior.

Itemizers versus Non-itemizers

Our second test compares the bunching patterns for taxpayers according to their use of *itemized* deductions. As explained in Section ??, the IRPF allows for two types of deductions. *Non-itemized* deductions are automatically included in the final tax liability and they do not need to be explicitly claimed by the taxpayer. Thereby, these type of deductions do not represent a margin of behavioral response for the taxpayers. On the contrary, *itemized* deductions are the result of a voluntary decision of taxpayers. Individual taxpayers could use them to adjust their taxable labor income in order to reduce their tax liability by locating just before the first kink point. Because claiming for *itemized* deductions gives the taxpayers an additional margin of response, our hypothesis for this section is that the excess of mass around the kink point should be larger for *itemizing* taxpayers. Although we still cannot interpret these results as causal, differences between these two groups would be suggestive of an intensive use of *itemized* deductions as a mechanism to locate just at the kink, supporting our leading hypothesis about the importance of the deduction channel to explain the bunching responses.

Figure A.6.1, Panel a., in Appendix A.6, shows the evolution the average amount of deductions considered as a whole and also its composition. It is worth noting that despite the sustained increase in the average amount of deductions, the share of *itemized* and *non-itemized* deductions remained constant (ca. 79% vs ca. 21%). Figure A.6.1, Panel b., shows the composition of *itemized* deductions. As mentioned in Section 2 and reported in Table 1, *itemized* deductions are comprised by personal and housing deductions. Overall, personal deductions are the most important component of personal deductions representing on average about 88% of total personal deductions for the five-year period considered. Despite housing deductions increased their participation as a share of the total amount of personal deductions between 2010 and 2014, they always represented less than 20% of total *itemized* deductions. We restrict the analysis to taxpayers that did not claim housing deductions in order to have a more direct interpretation of our results. Because they represent a relatively small fraction of the taxpayers, this restriction does not affect substantially the external validity of our results. Hence, when we refer to the group of *itemizing* taxpayers we mean individuals that only claimed personal deductions - i.e. non-proportional social security contributions and/or child care deductions-.

Figure 6 shows the taxable labor income distribution with its counterfactual distri-

bution breaking down the sample by *non-itemizers* (Panel a.) and *itemizers* taxpayers (Panel b.). In both cases the estimated excess of mass is statistically significant. However, there are clear differences in the magnitude and shape of the excess of mass. For *itemizing* taxpayers (Panel b.), although still not sharp, there is a clear hump around the kink point. This contrasts with the shape of the excess of mass for *non-itemizing* taxpayers which is relatively more diffuse. In terms of the magnitude, the excess of mass for *itemizing* taxpayer is almost 2.5 times larger than for the group of *non-itemizing* taxpayers (excess mass of 1.31 *vs* 0.55). This translates into an implied ETI of 0.25 and 0.11, respectively. These results support our hypothesis of a larger response for taxpayers that claim *itemized* deductions. This group of taxpayers seems to be much more elastic to the observed change in the marginal tax rate at the kink. However, the whole response cannot be attributed to *itemized* deductions since there is still a slightly significant response in *non-itemizers*. We will come back to this results later when we analyze additional channels of response (Section 6.2).

Figure 7 reports two additional tests to support the hypothesis of an intensive use of *itemized* deductions to target the first kink point in the IRPF schedule. Panel a. in Figure 7 depicts the intensity in the use of *itemized* deductions along the taxable labor income distribution measured as a proportion of the total amount of deductions. We expect that if bunching is explained by the use of *itemized* deductions, taxpayers should use them more intensively around the kink. Consistent with this idea, the ratio of *itemized* deductions has an inverted v-shape over the taxable income distribution: it increases until reaching k and then starts decreasing. Furthermore, the highest average intensity is reached just before the kink and it is about 11%.

In addition, we are also able to analyze the importance of *itemized* deductions by testing its correlation with being a *buncher*. If *bunchers* make an intensive use of this type of deductions, the share of taxpayers located at the bunching zone should be increasing with the intensity of *itemized* deductions. Panel b. in Figure 7 shows the probability of being a *buncher* by decile of intensity in the use of *itemized* deductions by fitting a simple linear regression. The results are consistent with our hypothesis: the share of *bunchers* among users of *itemized deductions* is increasing with the intensity of their use. Moving one decile further in the intensity-of-*itemized* deductions distribution increases the proportion of *bunchers* by 1.5 p.p.

In sum, even with the lack of an exogenous change that allow us to test directly the role of deductions, all these tests considered together constitute a second piece of evidence that strongly supports the hypothesis of the use of deductions as one of the main channels to explain bunching responses.

Tax Deduction or Filing Behavior?

As mentioned in Section 2 taxpayers can claim *itemized* deductions by filing a tax return. In this process, taxpayers do not only report *itemized* deductions but they also report all the other components of their taxable income for the period considered. Since filing a tax return gives the taxpayer the opportunity to self-report both gross labor income and *itemized* deductions, it is not clear whether the response observed in previous sections for itemizing taxpayers can be attributed exclusively to *itemized* deductions or could be also the result of changes in reported gross labor income. This latter mechanism is associated with misreporting and tax evasion behavior.

The correlation between claiming an *itemized* deduction and filing a tax return is strong but not one to one in the Uruguayan IRPF. Recall from Section 2 that instead of filing a tax return, employees can also claim *itemized* deductions *via* their employers. We exploit this specific feature of the IRPF administrative regulation to analyze the bunching pattern for taxpayers that claimed *itemized* deductions but did not file a tax return. This allow us to rule out the *filing-a-tax-return* channel and lead us to more direct evidence of the use of deductions as a mechanism to bunch. If non-fillers are using *itemized* deductions intensively to locate close to the kink the excess of mass around k cannot be attributed to filing behavior.

Figure 8 shows the bunching estimates for non-filers *itemizers*. This group of workers represents 15.7% of the total number of pure wage earners and about of the 83.36% of pure wage earners itemizers. Since most taxpayers claimed *itemized* deductions through the firm instead of filing a tax return, Figure 8 clearly resembles Panel b. of Figure 6 where we reported the bunching pattern for all *itemizers* taxpayers regardless how they claimed the *itemized* deductions. This suggests that the strong response observed for all *itemizers* is mostly driven by taxpayers that did not file a tax return and therefore this response is at least partially explained by an intensive use of deductions rather than by the *filing-a-tax-return* channel. The excess of mass in this case is of 120% and with a implied elasticity of about 0.23.

We can also test the deduction channel vs. the *filing-a-tax-return* channel for the group of workers that claimed *itemized* deductions and filed a tax return. To do this, we exploit the fact that for this group of taxpayers we count with two independent sources of information: third party reports from the Employers Statement and self-reported information from tax returns. This information includes not only the gross labor income paid by the employer but also the amount of *itemized* and *non-itemized* deductions claimed at the firm. This group represents 3.1% of the total number of pure wage earners and about 16.6% of taxpayers that claimed for *itemized* deductions.¹⁵

¹⁵It is worth mentioning that because of errors in administrative records, for about 9% of the pure

First, we analyze the timing of the response, i.e. we explore if taxpayers that claimed *itemized* deductions and filed a tax return started to bunch after filing the tax return or they already did at the firm. This allow us to have a first approximation to the use of the tax return as a mechanism for locating close to the kink point. Observing an excess of mass in the tax return but not in the firm would be suggestive of a bunching response that it is driven by the *filing-a-tax-return* channel. On the other hand, if the bunching is already observed at the firm we could rule out the *filing-a-tax-return* mechanism as a potential driver of the behavioral response.

Panel a. of Figure 9 shows the usual bunching analysis for taxpayers that claimed *itemized* deductions and filed a tax return using the self-reported information contained in the tax return. Panel b. of Figure 9 reports the same analysis but instead of using the tax return information, it is based on third party reports by the employer. The results are very suggestive. In the firm, (i.e. before the tax return - Panel b.) there is no sign of bunching around the kink and the excess of mass is negative and statistically insignificant. However, when we turn to the data reported in the tax return (Panel a.) there is a sharp excess of mass of about 360%, which implies an elasticity of 0.69. The immediate implication of this result is that *itemizing* taxpayers use the tax return as a mechanism to bunch.

At a first glance this seems to contradict the conclusions drawn from Figure 8 in which we attributed the behavioral response to the *itemized* deductions channel. However, for this specific group of taxpayers we still cannot rule out the *itemized* deduction channels since bunching in the tax return may be achieved both through adjustments in gross labor income reports or in the *itemized* deduction reports. Because our data allow us to test this directly, next we explore which are the items in the tax return that explain differences between the self and third party reported information. Hence, we look for discrepancies between gross labor income and amount of *itemized* deductions between tax returns and employer reports. Figure 10 depicts the average differences in gross labor income (Panel a.) and deductions (Panel .b) between the self and third-party reported information. Taxpayers are grouped and sorted using the tax return information and according to their distance to the taxable income kink point so that the x-axis remains the same that in previous figures. In both panels we excluded taxpayers that were already at the bunching zone in the firm (22.8%) because they do not have incentives to change their reports and they mechanically reduce the differences in the bunching zone. Panel a. shows that there is some evidence of income misreporting specially to the left of the kink point. However, this behavior is not associated with the bunching area and

wage earners who filed a tax return we do not have the third-party reported records in data. Henceforth, when we compare both sources of data in both cases we will include only those taxpayers for whom we observe both self and third-party reported data.

evolves relatively smoothly all over the taxable income distribution. Panel b. shows that taxpayers usually over-report deductions in the tax return compared to what they report to the firms. However, unlike income misreporting, these differences are strikingly large around the bunching zone. In particular, the differences between *itemized* deductions reported in the tax return compared to the ones reported in the firm jump from about 1,000 USD to 2,300 USD just before the bunching zone and then starts decreasing. This indicates that taxpayers that locate themselves in the bunching zone after they file a tax return make use of *itemized* deductions for this purpose. Overall, the bunching response observed in Figure 8 that could be first associated to a *filing-a-tax-return* behavior is indeed the result of filing a tax return but, specifically, through a more intensive use of deductions.

In a nutshell, this section provides multiple pieces of evidence to supports the major role of the deduction channel in order to explain the bunching behavior observed for Uruguayan taxpayers. As already mentioned, none of these results can be interpreted in a causal manner since they are all based on the analysis of different sub-samples and we do not count with a source of exogenous variation. Despite this caveat, the different tests reported robustly and consistently support our claim. To complement the analysis presented in this section we next focus on the second component of the taxable income: gross labor income.

6.2 So, No Responses in Gross Labor Income at All? Misreporting and Real Adjustments

The previous section analysis was motivated by the finding reported in Figure 6 in which we show that *itemizing* taxpayers have substantially larger responses than *non-itemizer* taxpayers. However, although smaller, we still observe a behavioral response for *non-itemizing* taxpayers. Because of the nature of this group, these responses can only be explained by the gross labor income component of the taxable income. This section aims to shed some light about channels of response related with gross labor income and therefore we will focus only in the sub-group of taxpayers that do not claim *itemized* deductions.

For *non-itemizer* taxpayers a gross labor income response can be explained by two different margins. The first one is associated with a traditional labor supply channel. Taxpayers may change their total amount of hours dedicated to labor market activities as a response to a change in the marginal tax rate. The second margin of response is through income misreporting, i.e. taxpayers may locate near the kink by reporting to the IRS only part of the gross labor income, leaving the remaining part outside the scope

of tax authorities. This last channel is particularly important in developing countries where capacity of enforcement is limited and taxpayers evade taxes relatively easily. Misreporting can be implemented in three different ways: 1) unilaterally by the employee, 2) unilaterally by the employer, 3) collusion by both parties. Our data does not allow us to precisely disentangle the two margins of response and the three alternatives for misreporting. However, we conduct some exploratory analysis to shed some light about which one is more likely to be operating.

As in the previous section, we first exploit the availability of two different and independent sources of income to analyze differences in the gross labor income reported by the firm and what the taxpayer self-reports in the tax return. There are two reasons to explain why a taxpayer that do not want to claim an *itemized* deduction may file a tax return. First, they may be required to do so by the IRS if they have multiple sources of income and the sum of those is above a yearly defined threshold (see Section 2). Second, they may voluntarily file a tax return even if they are not required to do it. This is perfectly reasonable for a taxpayer that has a refund to collect but it could also represent a signal of the intention to misreport/evade taxes. Our first tests explores differences in the size of bunching between tax returns and third party reports for *non-itemizer* taxpayers that filed a tax return. Panel a. in Figure 11 reports the usual bunching analysis for *non-itemizer* taxpayers using the self-reported information contained in the tax return. Panel b. replicates the analysis using the data reported by the employer. In both cases we observe evidence of bunching, although the size of the response is larger in the tax return data. This result has two implications regarding the potential channels of response. First, the difference in the magnitude of the bunching between the tax return and the Employer Statement can only be explained by unilateral employee misreporting. This can be tested directly by comparing differences in the gross labor income reported in the tax return and the gross labor income reported by the firms. Figure 12 shows that there is a sharp jump in the bunching zone in the differences between the gross labor income reported in the tax return and the gross labor income reported by the employer. This strongly supports our hypothesis that *non-itemizer* taxpayers are unilaterally misreporting to bunch around the kink point.

The second implication derived from Figure 11 (Panel b) is that the excess of mass observed at the firm could be the result both of real responses or misreporting collusion between employees and employers.¹⁶ This is confirmed by Figure 13 that reports the results for non-itemizer non-filer taxpayers. Although modest and diffuse (the implied ETI for this group is 0.119), there still evidence of bunching for this specific group of tax-

¹⁶It is worth noting that unilateral misreporting by the firm cannot be ruled out, but this does not seem plausible since there are no incentives for the firm to report their employees income just below the kink.

payers where the only alternatives are real responses or employer-employee collusion to misreport earnings. However, we still cannot say anything about which of the two mechanisms may be leading the behavioral response. An indirect way of testing the hypothesis of collusion to misreport consists of analyzing if responses are larger for taxpayers that have less coordination costs (see e.g., Kleven et al. 2016). We can compare the bunching pattern by breaking down the sample by the number of different jobs that a taxpayer has. The leading hypothesis is that if bunching in the firm is a result of employer-employee collusion, responses should be stronger for taxpayers with lower coordination costs. In terms of number of jobs, taxpayers receiving income from only one firm have less coordination costs than taxpayers receiving income from multiple sources since they only need to coordinate with one employer. Then, we expect to observe a larger bunching for single-firm employees. In order to focus on the collusion/real response channels in the following test we exclude all taxpayers that showed differences between self and third-party reported gross labor income.

Panels a. and b. in Figure 14 reports the bunching analysis performed with the data of the tax returns dividing the sample into workers that received income from only one firm (Panel a.) and workers that receive income from more than one firm (Panel b.). As expected, the bunching observed for *non-itemizer* wage earners is explained almost exclusively by taxpayers receiving income from only one source, i.e. taxpayers with less coordination costs. The implied ETI for single job workers in the tax return is 0.581 which is five times larger than the ETI for multiple job taxpayers. This is also true for non-itemizer non-filers. Figure 15 reports that the implied ETI for workers in a single firm is about 0.134 while for workers in multiple firms is substantially low (0.06). This evidence is consistent with previous studies that have shown the reported taxable income response to tax incentives is stronger in contexts with lower cost of coordination for employers and employees, like in smaller firms (Bergolo and Cruces 2014; Best 2015).

In sum, we find suggestive evidence of a gross labor income response through unilateral misreporting and also through the real response/collusion in the firms channels. Even though we cannot rule out that the second ones are associated with real response, the exploratory analysis and pre existing evidence about the lack of real responses in this type of setting may slightly lean the scale toward the *collusion-in-the-firm* channel.

7 Conclusion

This paper analyzes the behavioral responses of taxpayers to the first kink of the Uruguay personal income tax schedule. Our approach to identification rests on local variation in tax rates generated by the tax system. We present evidence of significant taxable labor

income responses around the kink where taxpayer start to pay taxes. We obtain two main set of findings. First, when the universe of taxpayers is considered, there is a statistically significant and robust excess of mass around the first kink point. This excess of mass translates into an ETI of about 0.16 which is diffuse and relatively modest. In contrast with previous studies, we show that wage-earners are at least as responsive as self-employed individuals. However, in line with previous evidence on wage earners (e.g. Paetzold 2018; Chetty et al. 2011), we find that the observed local taxable income response for this population translates into modest elasticities. This finding is suggestive in terms of the importance of frictions and optimization costs when analyzing behavioral responses to taxation (Chetty et al., 2011; Chetty, 2012). Our main hypothesis to explain the magnitude and shape of the response is that progressive taxation is relatively novel in Uruguay and taxpayers are still in a process of learning how the new tax schedule exactly works. This is supported by the dynamic evidence reported in which we show that the number of bunchers has steadily increased and at a larger rate than the number of registered taxpayers in general. Taken together, the responses we estimate on average fall within the range of elasticities found for individuals in the bottom-middle part of the income distribution in advanced economies.

When we deepen the analysis and explore the mechanisms that drive this response, we find that the two components of the taxable labor income are playing an important role. The order of magnitude of each of this channels depends on the opportunities that different groups of taxpayers face. In particular, a number of tests support the hypothesis that *itemized* deductions are a major channel of response. In addition, there is also evidence that supports the existence of gross labor income responses. In this regard, we document suggestive evidence of labor income misreporting through two mechanisms: unilateral underreporting by the employee and employer-employee collusion within the firm. A key limitation of our analysis of gross labor income responses is we cannot rule out the possibility that real earnings responses and misreporting may be operating simultaneously.

Our results contribute to the understanding of behavioral responses to personal income taxation in developing countries, for which existing evidence is comparatively scant, and inform the debate over tax policy in these countries. Calibrating the efficiency costs of taxation is crucial for many countries that need to raise tax revenue collection, and increase the provision and quality of certain public goods, without introducing major allocative distortions. First, we document non-negligible though moderate behavioral responses to personal income taxation. Second, our detailed investigation of the channels through which individuals respond to taxation is also an important input for tax design. Our results suggest that ETI may not be a sufficient statistics to account for

the efficiency costs of taxation. On the one hand, we find evidence of income misreporting. This suggests the existence of hiding transfers between agents in the same economy and, as pointed out by Chetty (2009), in the presence of fiscal externalities efficiency costs calculations based on the ETI may be misleading. Second, we document a more intensive use of deductions for individuals located just before the first kink, which may also affect the sufficiency of the ETI. As recently shown by Doerrenberg et al. (2017), the elasticity of taxable income is not a sufficient statistic for the welfare costs of taxation when deductions are responsive to tax-rate changes. This seems to be true even in the context of the Uruguayan personal income tax in which taxpayers enjoy limited deduction possibilities. In fact, the existence of few deductions may be an additional explanation for the relatively modest behavioral responses found in this paper. Our results show that policymakers should proceed with caution on reforms aimed to expand tax deduction opportunities and improve the quality of third-party reporting mechanisms in order to limit tax evasion. This is highly relevant in the context of developing countries characterized by low enforcement capacities. In the context of significant deduction or tax evasion responses to income taxation, broadening the tax base and improving the administrative capacities of tax authorities, rather than eroding the progressivity of the tax system, might be the appropriate policy options.

References

- Acemoglu, D., Naidu, S., Restrepo, P., and Robinson, J. A. (2018). Democracy Does Cause Growth. *Journal of Political Economy*, 127(1):47–100.
- Altonji, J. G. and Paxson, C. H. (1988). Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs. *Journal of Labor Economics*, 6(2):254–276.
- Amarante, Veronica and Gomez, Marcela (2016). El proceso de formalizacion en el mercado laboral uruguayo. Technical Report 20, Montevideo.
- Bachas, P. and Soto, M. (2018). Not(ch) Your Average Tax System: Corporate Taxation Under Weak Enforcement. Technical Report 8524, The World Bank.
- Bastani, S. and Selin, H. (2014). Bunching and non-bunching at kink points of the Swedish tax schedule. *Journal of Public Economics*, 109:36–49.
- Bergolo, M. and Cruces, G. (2014). Work and tax evasion incentive effects of social insurance programs: Evidence from an employment-based benefit extension. *Journal of Public Economics*, 117:211–228.
- Besley, T. and Persson, T. (2013). Chapter 2 - Taxation and Development. In Alan J. Auerbach, Martin Feldstein and Emmanuel Saez, R. C., editor, *Handbook of Public Economics*, volume 5 of *handbook of public economics*, vol. 5, pages 51–110. Elsevier.
- Best, M. C. (2015). The Role of Firms in Workers Earnings Responses to Taxes: Evidence from Pakistan. Technical report.
- Best, M. C., Brockmeyer, A., Kleven, H. J., Spinnewijn, J., and Waseem, M. (2015). Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan. *Journal of Political Economy*, 123(6):1311–1355.
- Bird, R. M. and Zolt, E. M. (2004). Redistribution via Taxation: The Limited Role of the Personal Income Tax in Developing Countries. *UCLA Law Review*, 52:1627.
- Bohne, A. and Nimczik, J. S. (2017). Learning Dynamics in Tax Bunching at the Kink - Evidence from Ecuador. Technical report.
- Card, D., Chetty, R., Feldstein, M., and Saez, E. (2010). *Expanding Access to Administrative Data for Research in the United States*. NSF White Paper 1.
- Carrillo, P., Pomeranz, D., and Singhal, M. (2017). Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement. *American Economic Journal: Applied Economics*, 9(2):144–164.
- Chetty, R. (2009). Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance. *American Economic Journal: Economic Policy*, 1(2):31–52.
- Chetty, R. (2012). Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica*, 80(3):969–1018.

- 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*, 126(2):749–804.
- Chetty, R., Friedman, J. N., and Saez, E. (2013). Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings. *American Economic Review*, 103(7):2683–2721.
- Chetty, R. and Saez, E. (2013). Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients. *American Economic Journal: Applied Economics*, 5(1):1–31.
- DGI (2013). Estimacion de la Evasion en el Impuesto al Valor Agregado Mediante el Metodo del Consumo: 2000-2012. Technical report, Direcci\textbackslash'on General Impositiva (DGI), Montevideo.
- Dincecco, M. and Katz, G. (2016). State Capacity and Long-run Economic Performance. *The Economic Journal*, 126(590):189–218.
- Doerrenberg, P., Peichl, A., and Siegloch, S. (2017). The elasticity of taxable income in the presence of deduction possibilities. *Journal of Public Economics*, 151:41–55.
- Duncan, D. and Sabirianova Peter, K. (2016). Unequal inequalities: Do progressive taxes reduce income inequality? *International Tax and Public Finance*, 23(4):762–783.
- ECLAC (2013). *Fiscal Panorama of Latin America and the Caribbean. Tax Reform and Renewal of the Fiscal Covenant*. Economic Commission for Latin America and the Caribbean (ECLAC), Santiago, Chile, united nations edition.
- Feldstein, M. (1999). Tax Avoidance and the Deadweight Loss of the Income Tax. *The Review of Economics and Statistics*, 81(4):674–680.
- Fisman, R. and Wei, S. (2004). Tax Rates and Tax Evasion: Evidence from Missing Imports in China. *Journal of Political Economy*, 112(2):471–496.
- Gordon, R. and Li, W. (2009). Tax structures in developing countries: Many puzzles and a possible explanation. *Journal of Public Economics*, 93(7-8):855–866.
- Hamilton, S. (2018). Optimal deductibility: Theory, and evidence from a bunching decomposition. Technical report.
- Hargarden, E. P. (2015). *Taxpayer Responses over the Cycle: Evidence from Irish Notches*.
- Harju, J. and Matikka, T. (2016). The Elasticity of Taxable Income and Income-Shifting: What is “Real” and What is Not? *International Tax and Public Finance*, 23(4):640–669.
- Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, 8(1):435–464.

- 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*, 79(3):651–692.
- Kleven, H. J., Kreiner, C. T., and Saez, E. (2016). Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries. *Economica*, 83(330):219–246.
- 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*, 128(2):669–723.
- Kopczuk, W. (2005). Tax bases, tax rates and the elasticity of reported income. *Journal of Public Economics*, 89(11):2093–2119.
- Kumler, T., Verhoogen, E., and FrÅas, J. A. (2013). Enlisting Employees in Improving Payroll-Tax Compliance: Evidence from Mexico. Working Paper 19385, National Bureau of Economic Research.
- LaLumia, S., Sallee, J. M., and Turner, N. (2015). New Evidence on Taxes and the Timing of Birth. *American Economic Journal: Economic Policy*, 7(2):258–293.
- le Maire, D. and Schjerning, B. (2013). Tax bunching, income shifting and self-employment. *Journal of Public Economics*, 107:1–18.
- Mortenson, Jacob and Whitten, Andrew (2018). Bunching to Maximize Tax Credits: Evidence from Kinks in the U.S. Tax Schedule. Technical report.
- OECD (2017). OECD.Stat [Data extracted on 25 Jul 2017 17:40 UTC (GMT) from OECD.Stat]. Technical report.
- Paetzold, J. (2018). How do taxpayers respond to a large kink? Evidence on earnings and deduction behavior from Austria. *International Tax and Public Finance*.
- Rogerson, R., Shimer, R., and Wright, R. (2005). Search-Theoretic Models of the Labor Market: A Survey. *Journal of Economic Literature*, 43(4):959–988.
- Rosen, H. S. (1976). Taxes in a Labor Supply Model with Joint Wage-Hours Determination. *Econometrica*, 44(3):485–507.
- Saez, E. (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3):180–212.
- 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*, 50(1):3–50.
- Schachtele, S. (2016). The Importance of Deductions in Response to the Personal Income Tax: Bunching Evidence from Germany. Technical report.

- Slemrod, J. (1995). Income Creation or Income Shifting? Behavioral Responses to the Tax Reform Act of 1986. *The American Economic Review*, 85(2):175–180.
- Slemrod, J. (1998). Methodological Issues in Measuring and Interpreting Taxable Income Elasticities. *National Tax Journal*, 51(4):773–788.
- Slemrod, J. (2008). Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance. *National Tax Journal*, 61(2):251–275.
- Slemrod, J. (2019). Tax Compliance and Enforcement. *Journal of Economic Literature*. *Forthcoming*.
- Slemrod, J. and Kopczuk, W. (2002). The optimal elasticity of taxable income. *Journal of Public Economics*, 84(1):91–112.
- Turnovsky, S. J. and Basher, M. A. (2009). Fiscal policy and the structure of production in a two-sector developing economy. *Journal of Development Economics*, 88(2):205–216.
- Zenou, Y. (2008). Job search and mobility in developing countries. Theory and policy implications. *Journal of Development Economics*, 86(2):336–355.

Tables

Table 1: Tax Schedule Components: Gross Labor Income and Deductions

Sources of Income	Deductions
Wages and Salaries	<i>Non - Itemized</i>
Self Employed Income (70%)	a. Payroll Taxes
Tips and Overtime	b. Health Insurance Contribution
Annual Complementary Salary and Holiday Bonus	<i>Itemized</i>
	<i>Personal Deductions</i>
	c. Other Social Security Contributions
	d. Child Care Deductions
	<i>Housing Expenses</i>
	e. Rent deductions
	f. Mortgage deductions

Notes: This table shows a detailed list of the components that are included in the definition of gross labor income and deductions used when computing IRPF. The left panel - Sources of Income - shows income components that are comprise the IRPF tax base while the right panel - Deductions - shows the type of deductions that are allowed by the tax code. Housing deductions are comprised by two types of deductions, rent deductions and mortgage deductions. Although both refer to housing expenses deductions, they are treated differently when computing the final tax liability. Rent deductions are treated as tax credits while mortgage deductions are included as any other in the deduction part of the schedule.

Table 2: Uruguayan Personal Income Tax (IRPF) Schedule

a. Tax rates			b. Deductions rates		
Bracket (2013 USD)	Mg. rate	$ln(\frac{1-\tau_i}{1-\tau_{i+1}})$	Bracket (2013 USD)	Mg. rate	$ln(\frac{1-\tau_i}{1-\tau_{i+1}})$
0-7,525	0		0 - 3,225	0.10	0.105
7,525 - 10,750	0.10	0.105	3,225 - 8,600	0.15	0.057
10,750 - 16,126	0.15	0.057	8,600 - 46,226	0.20	0.061
16,126 - 53,752	0.20	0.061	46,226 - 73,102	0.22	0.025
53,752 - 80,628	0.22	0.025	73,102 - 116,104	0.25	0.039
80,628 - 123,629	0.25	0.039	More than 116,104	0,30	0.069
More than 123,629	0,30	0.069			

Notes: This table shows the tax schedule for year 2013 with values measured in USD current dollars. Reported brackets are the current brackets in the last year of the period considered in this analysis (2010-2014). In Panel a. - Tax Rates - Column (1) depicts the annual income brackets, while column (2) shows the marginal tax rate applied to each of them. Column (3) shows the change in the net-of-tax marginal rate with respect to the previous bracket. This is computed as $ln(\frac{1-\tau_i}{1-\tau_{i+1}})$. For instance, all income between USD 10,750 and USD 16,126 is taxed at 15% rate. The marginal earned dollar above 16,126 is taxed at 20%, therefore, the change in the net-of-tax marginal rate from 3rd to 4th bracket is 6.1%. The result of passing the annual labor income through this *tax part* schedule is what we call “provisional tax liability”. In Panel b. - Deductions Rates - Column (1) depicts the annual deduction brackets, while column (2) shows the marginal deduction rate applied to each of them. Column (3) shows the change in the analogous change in deduction rates with respect to the previous bracket. This is computed as $ln(\frac{1-\tau_i}{1-\tau_{i+1}})$. For instance, deductions between USD 8,600 and USD 46,226 are deducted at a 20% rate. The marginal deduction rate above 16,126 is 25%, therefore, the change in the marginal deduction rate from 3rd to 4th bracket is 2.5%. The result of passing the annual total amount of deductions through this *deduction part* schedule is what we call “deduction amount”. The final tax liability consists of subtracting this deduction amount to the provisional tax liability that results from Panel a.

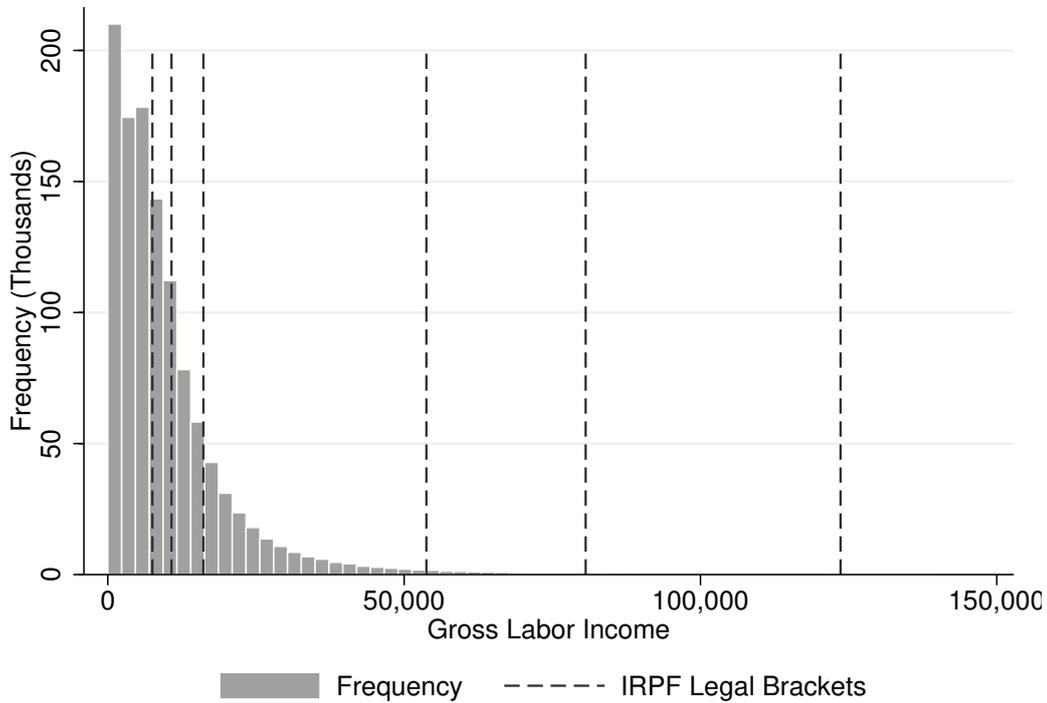
Table 3: Summary Statistics - Final Sample (2010-2014)

	All workers		Pure Wage Earners		Pure Self Employed (SE)		Both Wage Earners and SE	
	Mean (1)	SD (2)	Mean (3)	SD (4)	Mean (5)	SD (6)	Mean (7)	SD (8)
<i>Panel a. Income</i>								
Total income (USD in 1000s)	10.44	16.38	9.74	14.91	17.19	24.27	33.97	35.91
Total Labor Income (USD in 1000s)	10.12	14.07	9.45	12.60	16.08	20.47	32.98	32.82
Wage Income (USD in 1000s)	9.53	13.16	9.45	12.60	0.00	0.00	20.74	26.88
Self Employed income (USD in 1000s)	0.59	4.89	0.00	0.00	16.08	20.47	12.24	18.29
Retirement Income (USD in 1000s)	0.18	1.48	0.17	1.43	0.36	2.44	0.32	2.23
Capital Income (USD in 1000s)	0.14	7.57	0.12	7.29	0.75	12.13	0.67	12.35
<i>Panel b. Personal Income Tax (IRPF)</i>								
Tax liability (USD in 1000s) (labor income)	0.58	2.41	0.51	2.21	0.83	2.45	3.35	6.00
Tax liability > 0 (%)	34.21	47.44	33.04	47.04	39.58	48.90	78.15	41.32
Tax liability (among tax liability>0)	1.71	3.88	1.55	3.63	2.09	3.54	4.29	6.48
Total deductions (USD in 1000s)	2.10	2.76	2.00	2.49	2.10	5.59	6.12	5.48
Non-itemized (USD in 1000s)	1.65	1.96	1.63	1.89	0.53	0.75	3.60	3.57
Itemized (USD in 1000s)	0.44	1.49	0.37	1.15	1.57	5.44	2.53	3.41
Filing a tax return (%)	10.92	31.19	7.95	27.06	74.86	43.38	80.66	39.50
<i>Panel c. Socio-Demographics</i>								
Female (%)	46.12	49.85	45.81	49.82	50.38	50.00	55.40	49.71
Age	37.68	11.87	37.49	11.89	42.36	10.69	41.70	10.14
Under 40 years	41.83	49.33	41.22	49.22	57.07	49.50	54.36	49.81
Over 40 years	58.17	49.33	58.78	49.22	42.93	49.50	45.64	49.81
Source of income (%)								
One income source	.	.	75.61	42.94	100.00	0.00	.	.
Multiple income sources	.	.	24.39	42.94	0.00	0.00	.	.
Firm's activity (%)								
Services	50.48	50.00	49.74	50.00	67.93	46.68	82.55	37.96
Goods	49.52	50.00	50.26	50.00	32.07	46.68	17.45	37.96
Firm size (USD total sales) (%)								
Below median sales	19.41	39.55	19.52	39.63	.	.	13.87	34.56
Above median sales	80.59	39.55	80.48	39.63	.	.	86.13	34.56
Observations	5,587,711		5,351,080		107,594		129,037	
Individuals	1,526,301		1,471,560		29,383		25,358	

Notes: This table reports summary statistics of the final sample considered in the analysis for the period: 2010-2014. Panel a. reports statistics about income component, Panel b. about personal income tax items and Panel c. about socio-demographic characteristics as well as the place of work. All statistics are computed on a taxpayer-year level, i.e. the number of times that a taxpayer is included in the calculation of each variable mean is the number of years he had positive labor income. If a taxpayer do not have a positive labor income on a specific year, it is excluded from the sample in that particular year. Odd columns report the mean of the variable while even columns report the standard deviation. Columns (1) and (2) show summary statistics for the whole sample; columns (3) and (4) do the same for the pure wage earners group; columns (5) and (6) for the pure self-employed group and columns (7) and (8) for the group of workers receiving income both from wages and self-employment activity. Last two rows report the number of taxpayers and taxpayer-year included in the sample. All variables in monetary values are expressed in 2013 USD.

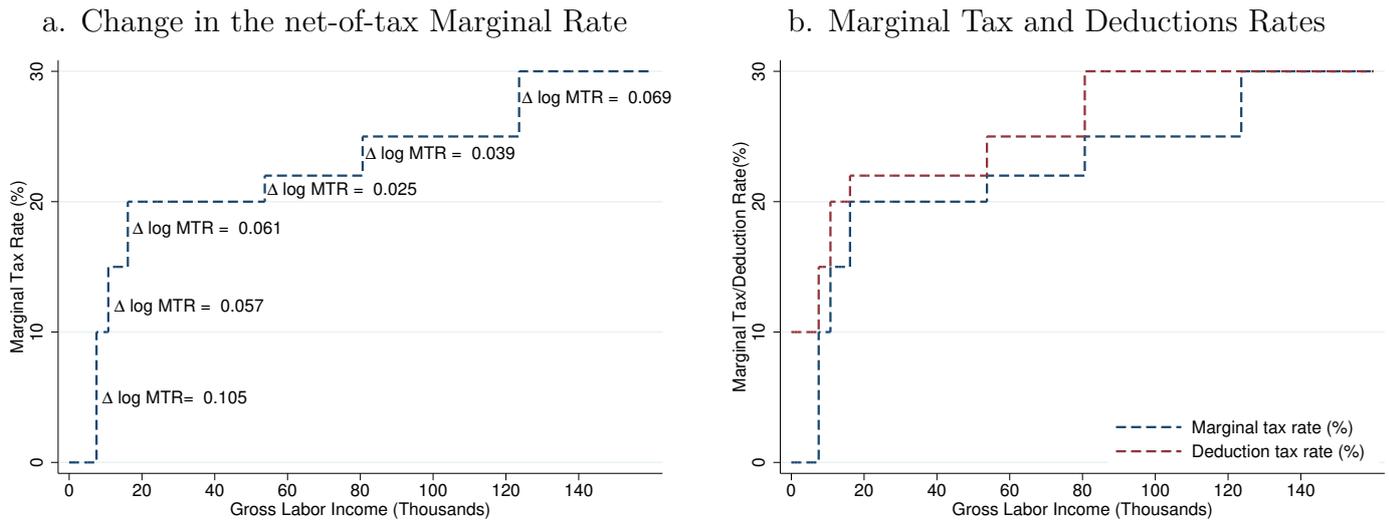
Figures

Figure 1: Tax Schedule and Gross Labor Income Distribution (2013)



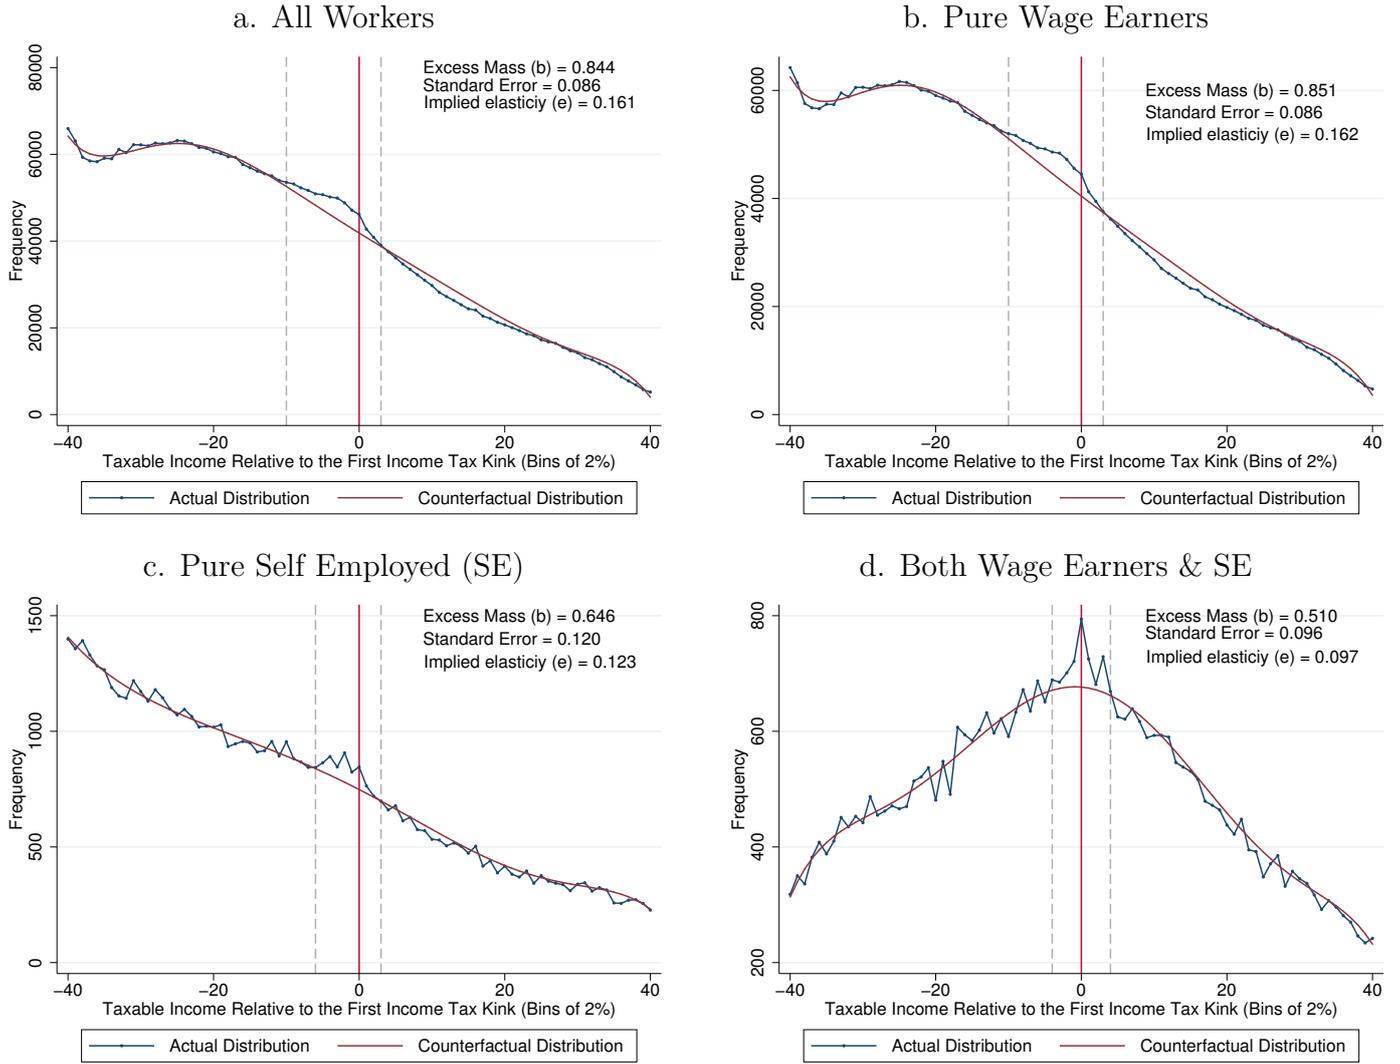
Notes: This figure shows gross labor income distribution (i.e., before deductions) and location of each income bracket in 2013 which is the valid tax schedule for the last year in our sample. Gross labor income is defined as all income earned by workers before considering any deduction. Data comes from the administrative records of 2013. Each bars represent the amount of people located at each a particular bin of the gross labor income measured in thousands (y-axis). Dashed lines show the tax income brackets at which the marginal tax rates changes.

Figure 2: Marginal Changes in Labor Income Tax and Deductions Schedule (2013)



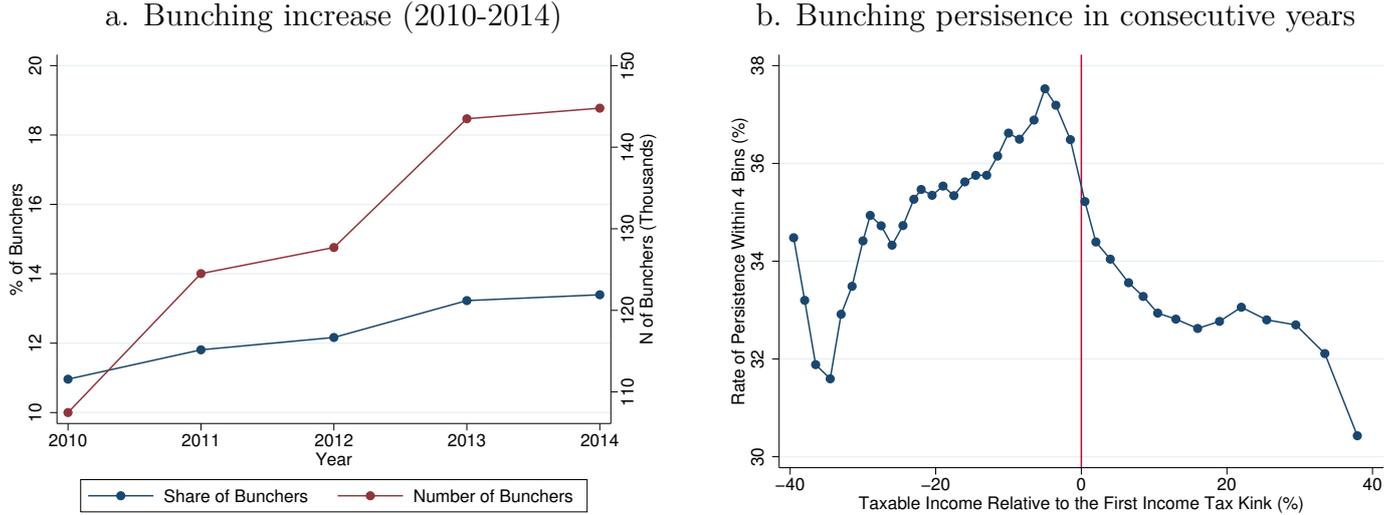
Notes: This figure depicts both tax and deductions parts of the IRPF schedule. The brackets are plotted for the 2013 USD income brackets which is the valid tax schedule in the last year of the period considered. Panel a. shows the *tax part* of the IRPF (Panel a. in Table 2). The blue dashed line represents the marginal tax rate (y-axis) as a function of the gross labor income (x-axis). We also include the change in the net-off tax rate of passing from one bracket to the next one as described in Table 2. Panel b. plots jointly the tax and deduction part. As it was explained in Section 2 the width of the n_{th} deduction bracket is exactly the same as the size of $n + 1_{th}$ tax bracket and this also happens with the marginal rate. As in Panel a. the blue dashed line shows the marginal tax rates for each income bracket while the red dashed line shows the marginal deduction rates also as a function of the gross labor income.

Figure 3: Taxable Labor Income Bunching by Source of Income, 2010-2014



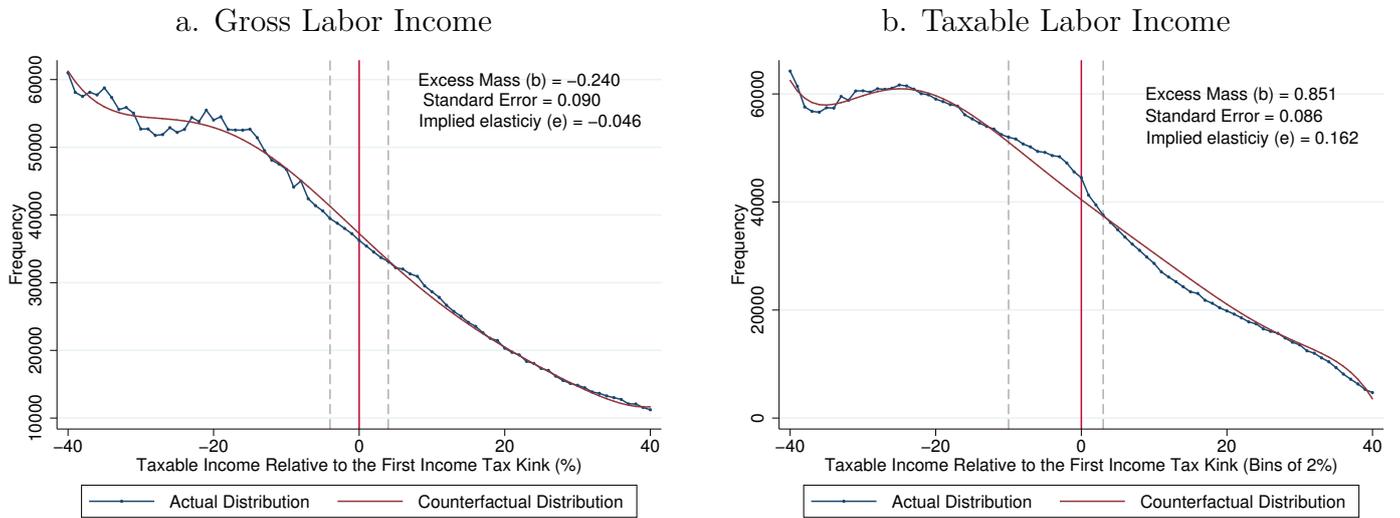
Notes: The figures in Panels a. through d. plot the empirical and counterfactual taxable income distributions normalized by the individual first kink point as explained in detail in Section 3.2 (See Appendix A.3 for a numerical example). All estimates are performed by using the pooled data for the 2010-2014 period. Panel a. shows the taxable income distribution for the main sample considering all taxpayers. Panel b. replicates the analysis for pure wage earners, – i.e. those workers who do not receive self-employed income –, Panel c. do the same with taxpayers that only earn income from self-employment work (pure self employed) while Panel d. reports the results for workers that receive both types of income. In all cases, the taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts paying IRPF, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. The counterfactual estimation is based on Chetty et al. (2011) and thoroughly detailed in Section ???. We use a 7 degree polynomial and bunching windows, identified with gray dashed lines, are defined in each particular case according to visual inspection. In each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 4: Evolution and Persistence of Bunching - Pure Wage Earners, 2010-2014



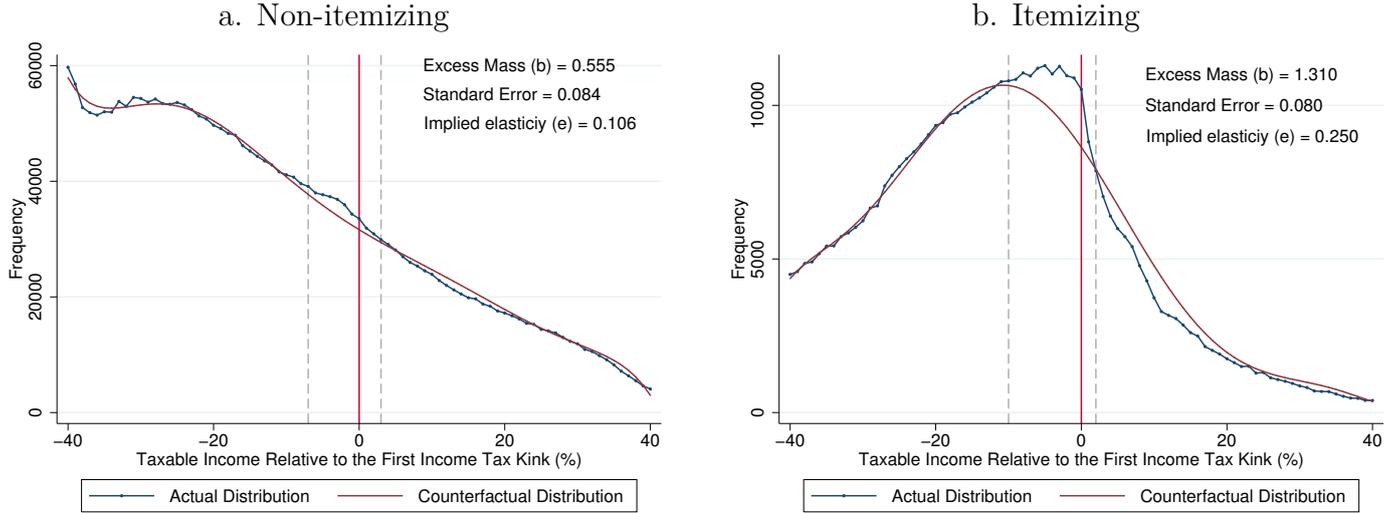
Notes: The analysis in this figure restricts the sample to the group of wage earners that do not receive self employed income, i.e., pure wage earners. Panel a. shows the evolution of bunchers between 2010 and 2014. Bunchers are defined as taxpayers located in the bunching zone depicted in Figure 3 and in particular, since we are showing the results for the pure wage earners sample, we use the zone defined in Panel b (-20%, 6%). The red line depicts the evolution of the absolute number of bunchers (right y-axis). Since the number of registered workers increased during the period, the blue line shows the evolution of the share of bunchers, i.e. the evolution of the number of bunchers as a percentage of the whole number of registered workers (left y-axis). Panel b. shows the ratio of persistence of a taxpayers as a function of the distance to the first kink point (The x-axis is exactly the same used in Figure 3). In this case, a worker is defined as persistent if its actual income remains within a window of 8% (or 4 bins) compared to last year's income.

Figure 5: Gross Labor Income versus Taxable Labor Income Distribution - Pure Wage Earners, 2010-2014



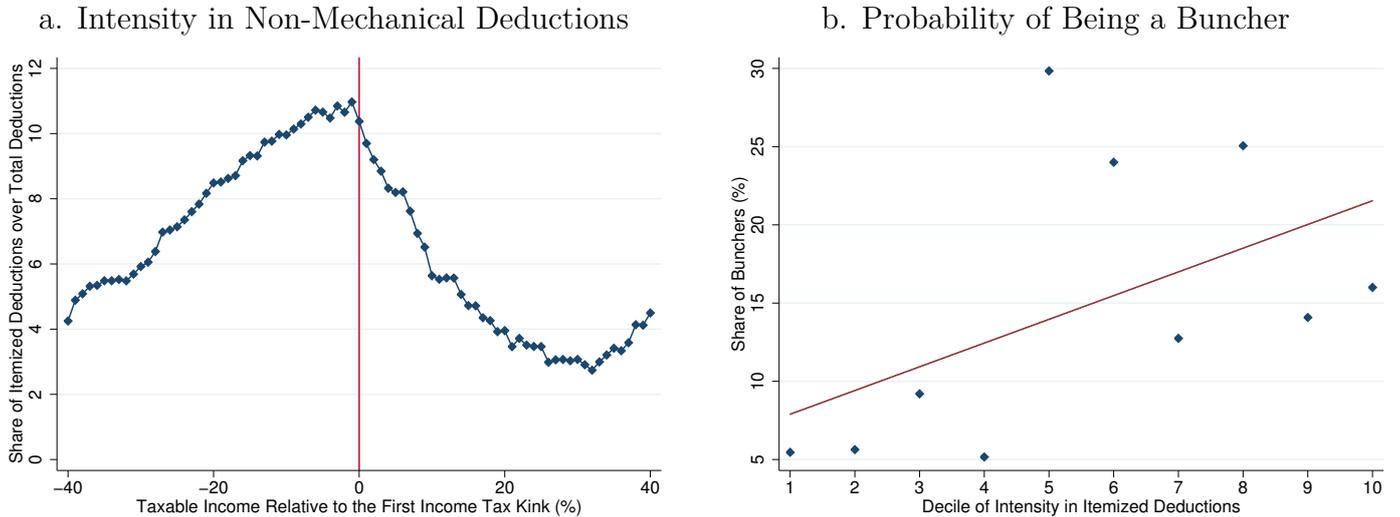
Notes: The figures in Panels a. and b. plot the empirical and counterfactual gross labor income and taxable labor income distributions respectively. The analysis restricts the sample to the group of wage earners that do not receive self employed income, i.e., pure wage earners. In Panel a. gross labor income distribution is normalized to the statutory income bracket (USD 7,525). In Panel b. taxable income distribution is normalized as usual (see Section 3.2 and Appendix A.3). All estimates are performed by using the pooled data for the 2010-2014 period. For Panel a. the vertical red line represents the statutory income bracket which is identical for each taxpayer. The values in the x-axis should be interpreted as the percentage distance to USD 7,525. For an easier comparison with the gross labor income distribution, Panel b. of this figure repeats the figure reported in Panel b. in Figure 3. In this figure the vertical red line represents the exact point at which each taxpayer starts paying IRPF, i.e. the first kink point. In both figures, Panel a. and b., the data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. The counterfactual estimation is based on Chetty et al. (2011) and thoroughly detailed in Section ???. We use a 7 degree polynomial and bunching windows, identified with gray dashed lines, are defined in each particular case according to visual inspection. In each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 6: Taxable Labor Income Bunching: Pure Wage Earners by Itemized Status



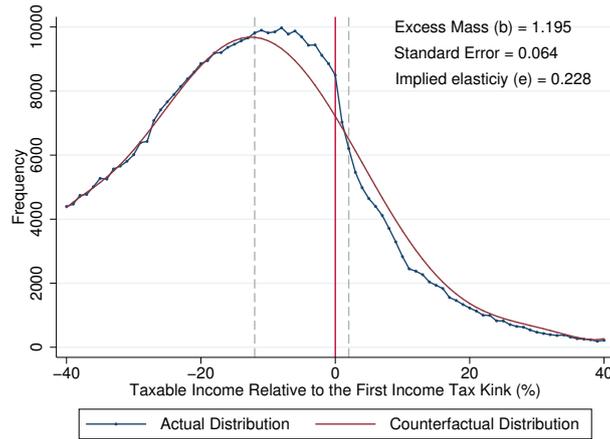
Notes: This figure plots the empirical taxable income distributions for two sub-groups of population for the 2010-2014 period: non-itemizing and itemizing taxpayers. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. Panel a. shows the taxable income distribution for non-itemizing taxpayers, i.e. taxpayers that did not claim for *itemized* deductions such as child care and housing deductions. Panel b. replicates the analysis for all taxpayers that make positive claims of *itemized* deductions. The taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts to pay IRPF II, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. This estimation was made following the approach in Chetty et al. (2011) with a 7 degree polynomial. Bunching windows are identified with gray dashed lines and are defined in each particular case according to visual inspection. In the upper right side of each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 7: Use of Non-Mechanical Deductions and Probability of Being a Buncher, 2010-2014



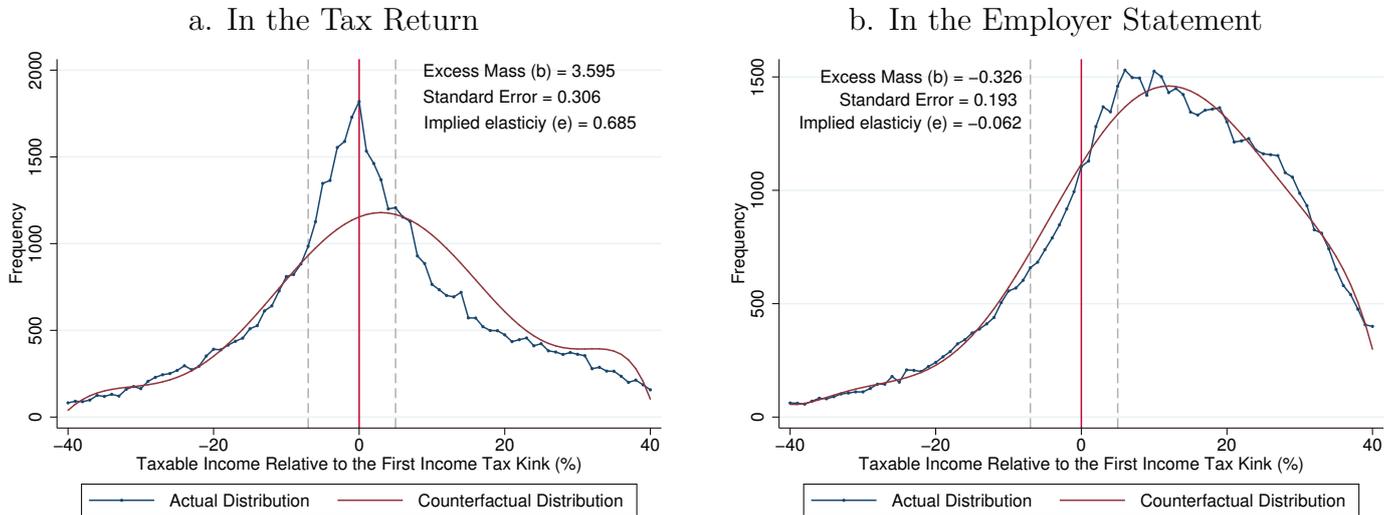
Notes: The analysis restricts the sample to the group of wage earners that do not receive self employed income, i.e., pure wage earners. Panel a. shows the intensity in the use of *itemized* deductions around the first kink point for taxpayers that do claim some of this type of deductions. The x-axis presents the taxable labor income distribution as the percentage distance to the first kink point. The data is grouped in bins of 2% wide. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots depicts the ratio of *itemized* deductions as a percentage of the total amount of deductions (*itemized* plus *non-itemized*). Panel b. depicts the share of bunchers according to the decile of the distribution of intensity in the use of *itemized* deductions. In order to do this, we construct the deciles of the intensity in the use of *itemized* deductions and we plot against the share of bunchers in each decile. The blue dots represent the correlation between these two variables, computed as the coefficient that results from regressing the share of bunchers in each decile on the decile of intensity in *itemized* deductions. Therefore, the β coefficient can be interpreted as the increase in the probability of being a buncher that results from moving from the n_{th} decile to the $n + 1_{th}$ decile.

Figure 8: Taxable Labor Income Bunching - Pure Wage Earners: Itemizers and Non-Fillers, 2010 -2014



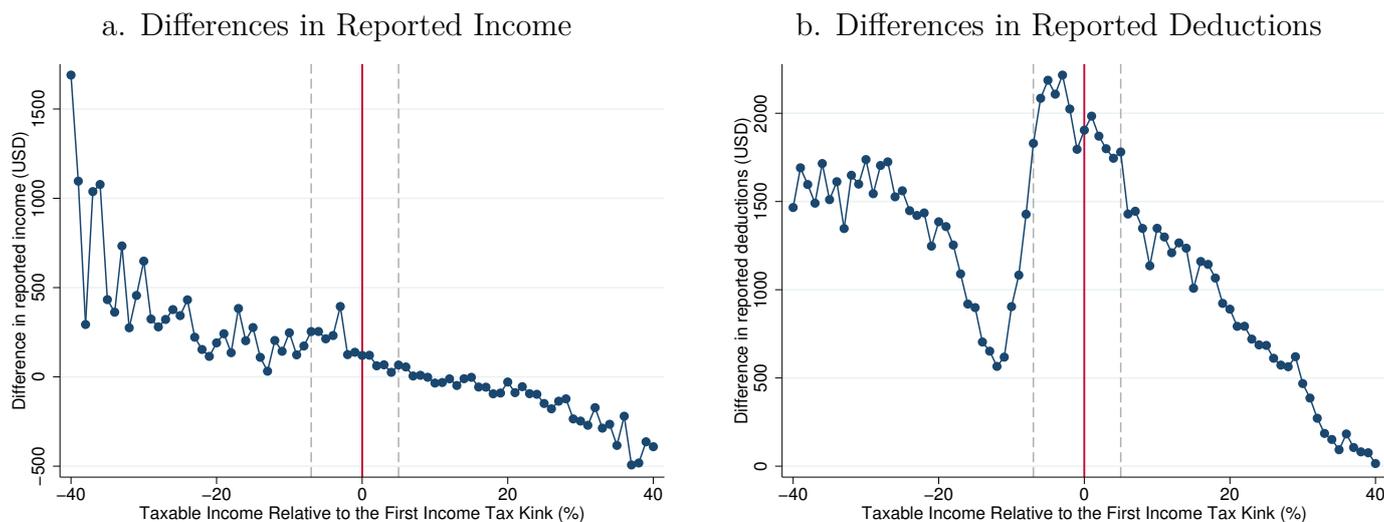
Notes: This figure plots the empirical taxable income distributions for the sub-group of itemizing taxpayers who do not file a tax return in the period 2010-2014. The analysis restricts the sample to the group of wage earners that do not receive self employed income, i.e., pure wage earners. The taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts to pay IRPF II, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. This estimation was made following the approach in Chetty et al. (2011) with a 7 degree polynomial. Bunching windows are identified with gray dashed lines and are defined in each particular case according to visual inspection. In the upper right side of each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 9: Taxable Labor Income Bunching: Pure Wage Earners Itemizers and Fillers - By Source of Information, 2010-2014



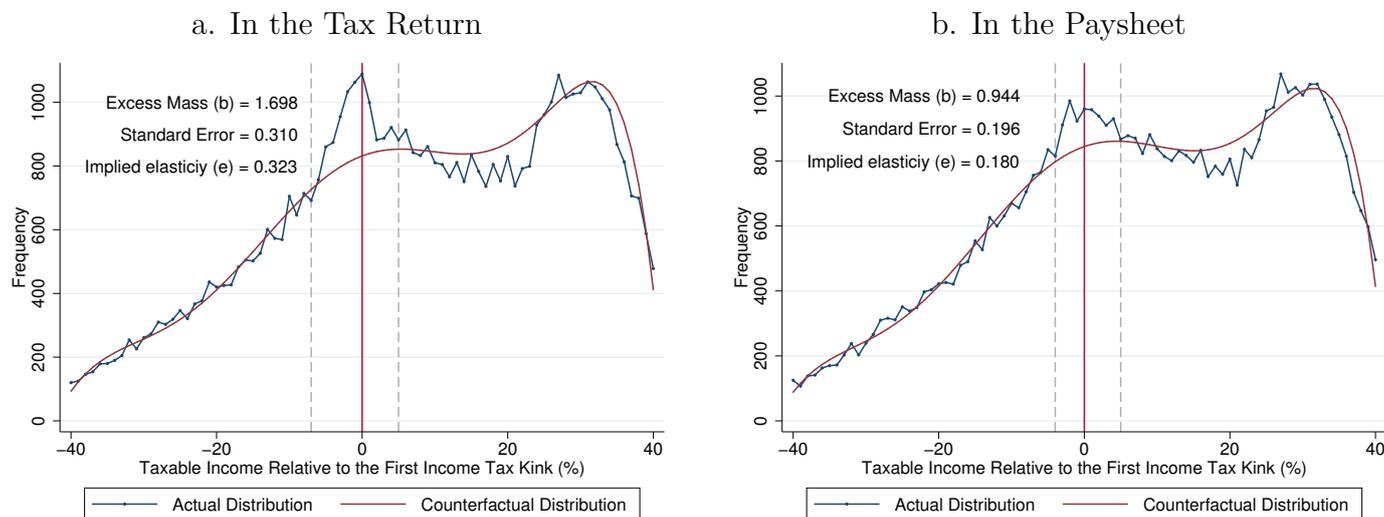
Notes: This figure plots the empirical taxable income distributions for two sub-groups of wage earners for the 2010-2014 period. Panel a. shows the taxable income distribution for itemizing taxpayers who file a tax return using the self-reported information contained in the tax return. Panel b. replicates the analysis but instead of using the tax return information, it is based on third party reports by the employer. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. The taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts to pay IRPF II, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. This estimation was made following the approach in Chetty et al. (2011) with a 7 degree polynomial. Bunching windows are identified with gray dashed lines and are defined in each particular case according to visual inspection. In the upper right side of each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 10: Differences in Reported Gross Labor Income and Deductions in Itemizers and Filers - Pure Wage Earners, 2010-2014.



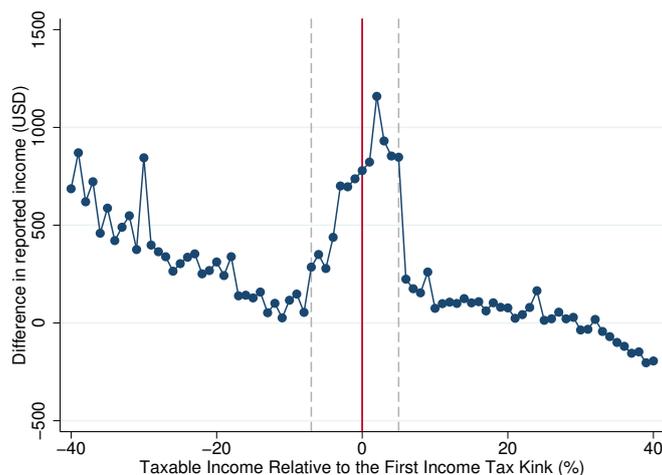
Notes: This figure plots the average differences in gross labor income (Panel a.) and deductions (Panel b.) between the self and third-party reported information for taxpayers who itemize deductions and file a tax return in the period 2010-2014. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. It also drop those taxpayers that were already at the bunching zone in the firm. Taxpayers are grouped and sorted using the tax return information and according to their distance to the taxable income kink point. The difference in reported income (amount of deductions) is grouped in bins of 2% wide who are represented by the blue dots. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window.

Figure 11: Taxable Labor Income Bunching: Wage Earners Non-itemizers and Filers - By Source of Information, 2010-2014



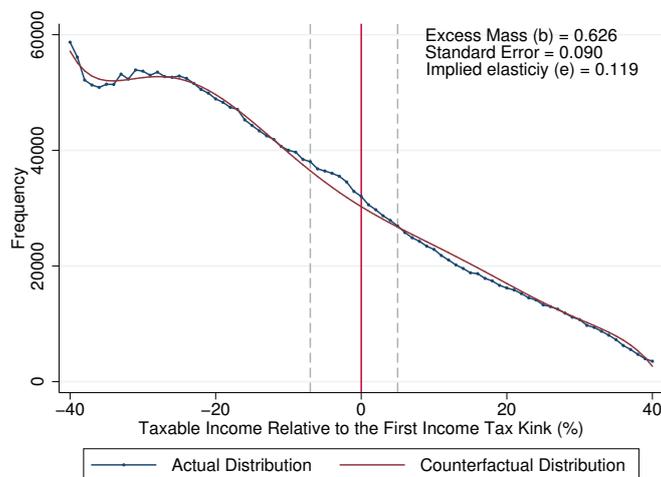
Notes: This figure plots the empirical taxable income distributions for two sub-groups of wage earners for the 2010-2014 period. Panel a. shows the taxable income distribution for non-itemizers taxpayers who file a tax return using the self-reported information contained in the tax return. Panel b. replicates the analysis but instead of using the tax return information, it is based on third party reports by the employer. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. The taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts to pay IRPF II, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. This estimation was made following the approach in Chetty et al. (2011) with a 7 degree polynomial. Bunching windows are identified with gray dashed lines and are defined in each particular case according to visual inspection. In the upper right side of each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 12: Differences in Reported Gross Labor Income for Non-Itemizers and Filers - Pure Wage Earners, 2010-2014



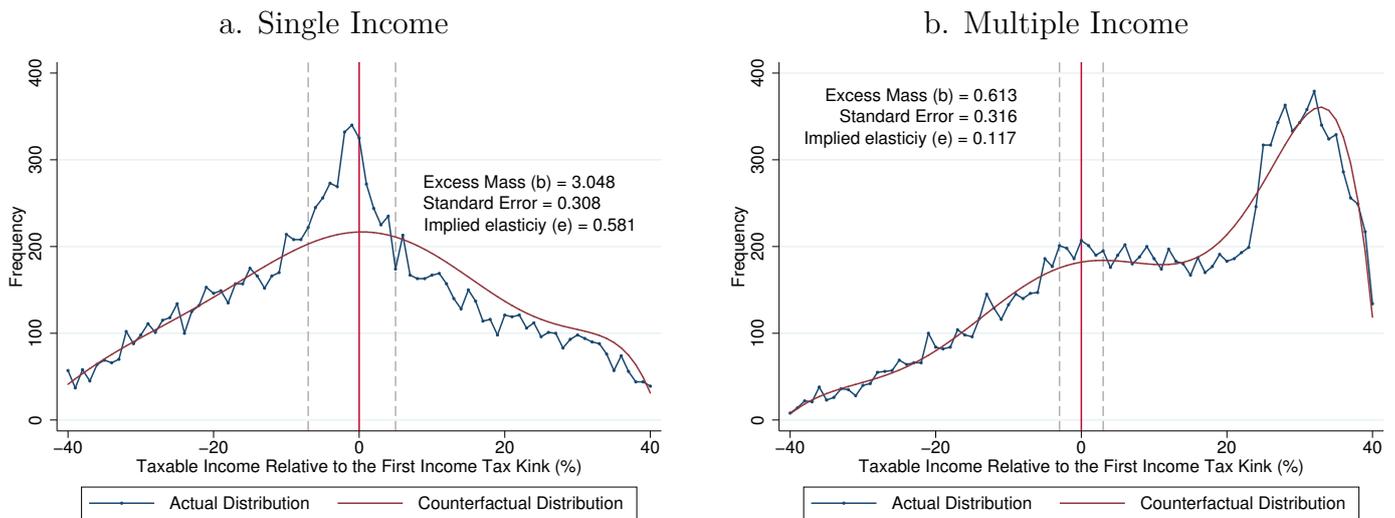
Notes: This figure plots the average differences in gross labor income between the self and third-party reported information for taxpayers who file a tax return but non-itemize deductions in the period 2010-2014. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. It also drop those taxpayers that were already at the bunching zone in the firm. Taxpayers are grouped and sorted using the tax return information and according to their distance to the taxable income kink point. The difference in reported income (amount of deductions) is grouped in bins of 2% wide who are represented by the blue dots. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window.

Figure 13: Taxable Labor Income Bunching: Pure Wage Earners Non-itemizers and Non-Filers, 2010-2014



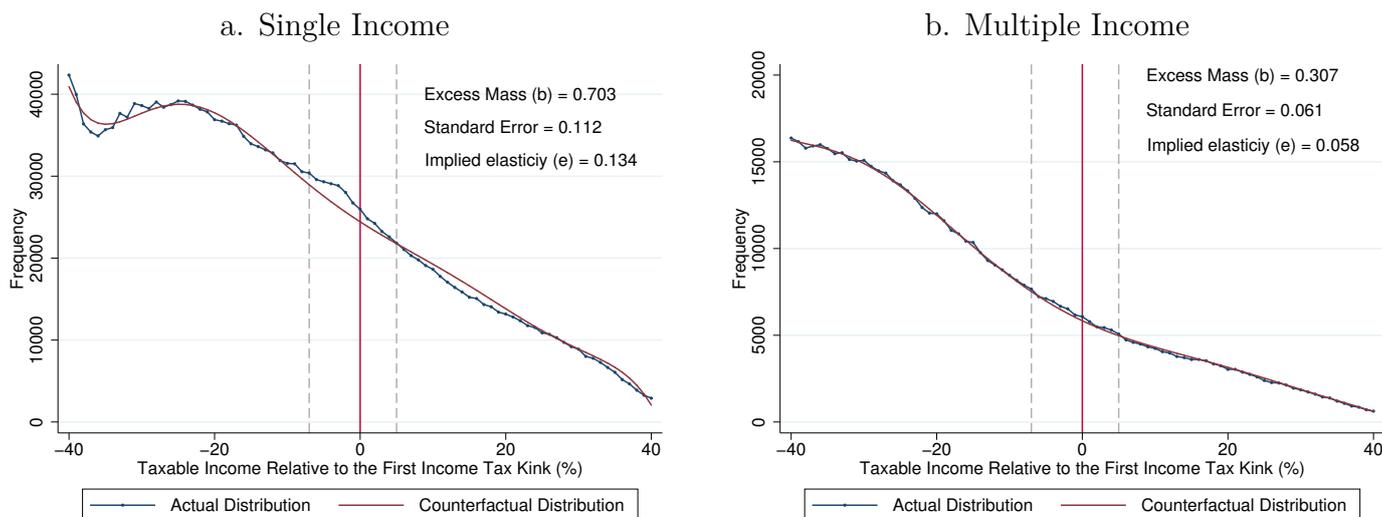
Notes: This figure plots the empirical taxable income distribution for non-itemizing and non-filers taxpayers. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. The taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts to pay IRPF II, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. This estimation was made following the approach in Chetty et al. (2011) with a 7 degree polynomial. Bunching windows are identified with gray dashed lines and are defined in each particular case according to visual inspection. In the upper right side of each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 14: Taxable Labor Income Bunching: Pure Wage Earners Non-itemizers with no differences between sources of data - By number of Income Sources



Notes: This figure plots the empirical taxable income distributions for two sub-groups of non-itemizer wage earners for the 2010-2014 period. Panel a. shows the taxable income distribution for workers that received income from only one firm. Panel b. replicates the analysis but using workers that receive income from more than one firm. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. It also excludes all taxpayers that showed differences between self and third-party reported gross labor income. The taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts to pay IRPF II, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. This estimation was made following the approach in Chetty et al. (2011) with a 7 degree polynomial. Bunching windows are identified with gray dashed lines and are defined in each particular case according to visual inspection. In the upper right side of each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Figure 15: Taxable Labor Income Bunching: Pure Wage Earners Non-itemizers and Non-Fileers with no differences between sources of data - By Income Sources



Notes: This figure plots the empirical taxable income distributions for two sub-groups of non-itemizers and non-filers wage earners for the 2010-2014 period. Panel a. shows the taxable income distribution for workers that received income from only one firm. Panel b. replicates the analysis but using workers that receive income from more than one firm. The analysis restricts the sample to the group of wage earners that did not receive self employed income, i.e., pure wage earners. The taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts to pay IRPF II, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. This estimation was made following the approach in Chetty et al. (2011) with a 7 degree polynomial. Bunching windows are identified with gray dashed lines and are defined in each particular case according to visual inspection. In the upper right side of each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Tax Bunching at the Kink in the Presence of Low Capacity of Enforcement: Evidence From Uruguay

Marcelo Bergolo, Gabriel Burdin, Mauricio De Rosa
Matias Giacobasso, and Martin Leites

April, 2019

A.1 Tax Structure in Uruguay

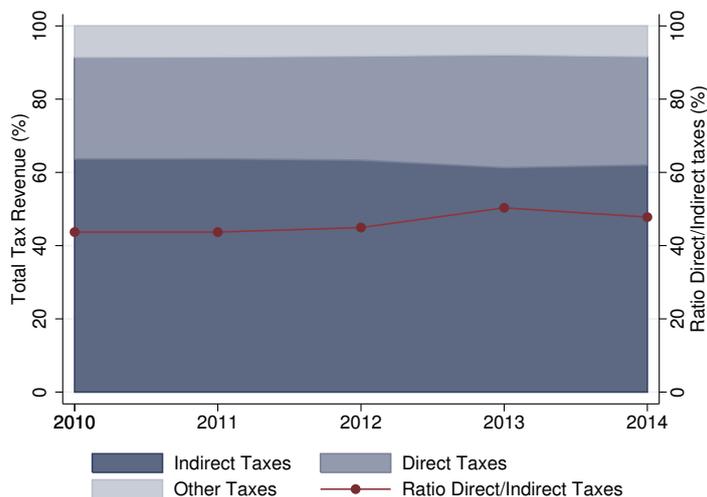
Uruguay is a South American country with an annual GDP per capita of USD 15,000 and a strong and solid institutional background. Like other South American countries, Uruguay performed extraordinary well (in terms of GDP growth) in recent years -a 4.3% average annual GDP growth for the period 2004-2017-. This was accompanied by an explicit strategy to enhance tax enforcement mechanisms which also led to a sharp increase in tax revenues. For instance, Value Added Tax (VAT) evasion contracted from 40% to 13.4% between 2002 and 2012 (DGI, 2013) and unregistered employment reduced from 35% to 25% between 2006 and 2014 (Amarante, Veronica and Gomez, Marcela, 2016). Several reforms were introduced during this period. One of the most important was the 2007 tax reform, which increased the relative weight of direct taxation. Indirect taxes like VAT and other specific taxes on consumption were cut and progressive forms of income taxation were introduced. In particular, a dual personal income tax, including a progressive tax on labor income and a proportional tax on capital income. The personal income tax is referred in Uruguay as *Impuesto a la Renta de las Personas Físicas* (IRPF in Spanish). While the labor income tax base is known as IRPF II, the capital tax base is referred as IRPF I. For the sake of simplicity, and since our study focuses only on the labor income tax base, henceforth we refer to the labor income taxation as IRPF.

Taxation in Uruguay can be broadly categorized into VAT, personal income tax, corporate tax, and some other small specific taxes to consumption and wealth. According to the Latin American module of OECDStat, overall tax burden in Uruguay was about 23.9% of the GDP in 2005 and reached 27% in 2015 (OECD, 2017). This ratio is lower than the OECD average (34%) and slightly larger than the Latin America average (23.1%). Despite the changes introduced by the tax reform, the tax burden is still heavily based on indirect taxes. According to the Uruguayan Internal Revenue Services (henceforth IRS), VAT constitutes the largest source of tax revenue representing about 50 % of the total tax collection, excluding social security contributions. Although significantly less important in terms of share of tax collection, direct taxation has increased its weight in the last years. Currently, it represents more than one third of total tax collection. Figure A.1.1 shows that between 2010 and 2014, the share of direct taxes over total tax collection increased from 28% to 30% of total tax revenues (l.h.s axis). In terms of indirect tax collection, direct taxes represented 43.7% in 2010 while at the end of 2014 raised to 47.8% (r.h.s axis).

Direct taxation in Uruguay can be divided into: 1) personal income tax (IRPF and capital income tax), 2) retirement income tax (IASS), 3) corporate income tax (IRAE), 4) and a non residents personal

income tax (IRNR). Figure A.1.2 describes the structure of direct tax collection. The two major components of direct taxation are IRPF and IRAE which combined explain more than 80% of total direct tax collection at the end of the period. In particular, IRPF represented 12.52% of the whole tax revenue, positioning it as the second component of the tax structure in terms of total tax collection. Even though capital income is part of what is usually categorized as personal income, for simplicity and since we are only focusing on IRPF our definition of personal income will only include labor income.¹⁷

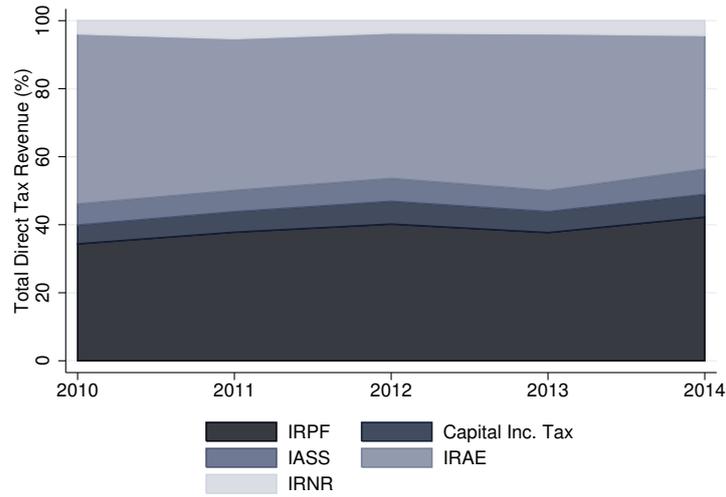
Figure A.1.1: Tax Structure in Uruguay (2010-2014)



Notes: This figure shows the composition of the Uruguayan tax structure between 2010 and 2014. This information comes from the annual series of tax revenues that are available in the IRS website. The total tax revenue is decomposed in direct and indirect taxes and a residual category that is formed by a group of other small items. Indirect taxes include the two major indirect taxes in Uruguay: VAT and IMESI: The VAT is the traditional Value Added Tax and it represents almost 80% of the total indirect tax revenue. IMESI is a tax that is collected from the first sale of a particular set of goods such as alcoholic beverages, tobacco and fuel, among others. Direct taxes include all personal income tax components (capital, labor and retirement) and the main corporate tax (see note in figure A.1.2 for a more detailed description of direct taxes). The left axis shows the percentage of total tax revenue. On the right axis we show the ratio of direct over indirect taxes. In the figure, this ratio is represented year by year by the black line.

¹⁷In other words, our measure of taxable personal income do not consider capital sources of income that are taxed with proportional tax rates. Since we estimate the ETI using tax variation in the first kink of the IRPF tax schedule, only about 8% of taxpayers report positive capital income in our analysis sample.

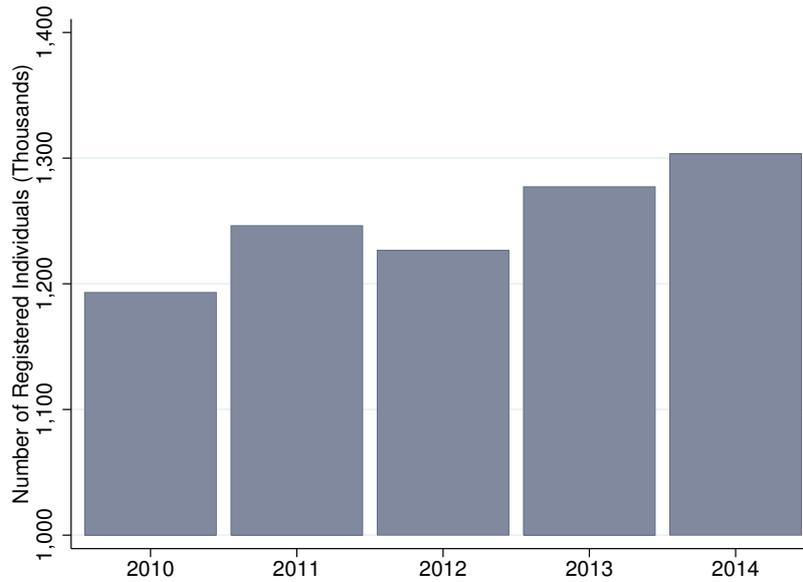
Figure A.1.2: Composition of Direct Taxes in Uruguay (2010-2014)



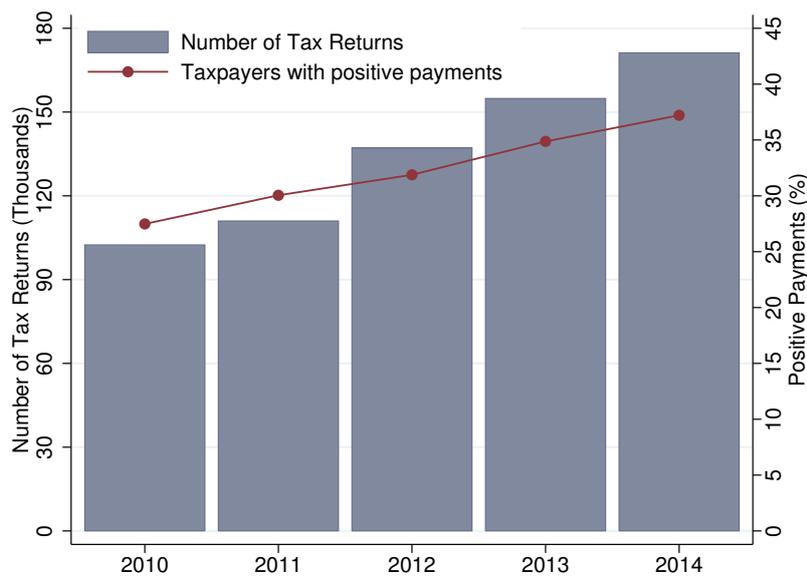
Notes: This figure shows the composition of the Uruguayan set of direct taxes between 2010 and 2014. This information comes from the annual series of tax revenues that are available in the IRS website. In this figure we depict all taxes included in the direct taxes category in figure A.1.1. Specifically, we include labor and capital components of personal income taxation (IRPF II and capital income tax respectively), retirement income tax (IASS), corporate tax (IRAE) and non residents personal income tax (IRNR). The left axis shows the percentage of total direct tax revenue.

Figure A.1.3: The Uruguayan Context

a. Evolution of Registered Employment (2010-2014)



b. Evolution of Tax Returns and Taxpayers with Positive Tax Payments (2010-2014)



Notes: Panel a. shows the evolution of the registered employment in the 2010-2014 period measured in thousands of individuals. In this category we include all workers that the IRS consider as potential IRPF taxpayers, i.e. all workers receiving income that is registered in the Social Security Agency. Information comes from the annual IRS tax return campaign reports that are available on the IRS website. In Panel b. we show the evolution of the number of tax returns submitted to the IRS and the evolution of people that paid at least 1 USD on behalf of IRPF in the corresponding year. Each bar represents the total number of tax returns filled every year from 2010 to 2014 measured in thousands (left axis). The solid line represents individuals that made positive IRPF payments measured as a percentage of the total registered employment. This percentage is depicted on the right y-axis.

A.2 Details of IRPF calculation

A.2.1 Example of how IRPF is computed

Table A.2.1 shows how the final tax liability would be computed for an hypothetical taxpayer whose annual earnings are UYU 900,000 (Equivalent to about USD 30,000). In this case, the UYU 900,000 are comprised by UYU 800,000 from wages and UYU 100,000 from self employed income. All self employed income is computed on a base of 70%, since 30% is discounted on behalf of VAT and other costs incurred in the independent activity. Thus, the final taxable income that is going to be passed through the tax schedule is UYU 870,000. Regarding deductions, this taxpayer has a total of \$223,899 where \$190,125 comes from the mechanical deductions associated to payroll taxes and other social security contributions and \$33,774 comes from the fixed amount that is defined for a single child in charge with no disabilities. From passing \$870,000 through the tax schedule and \$223,899 we have a *tax part* of \$113,207 and a *deduction part* of \$28,908. Subtracting the latter from the *tax part*, we have the *final tax liability*, which in this case is about \$84,299 (or about USD 2,800).

Table A.2.1: IRPF Calculation Example

Panel a. Summary of Income and Deduction Components

Annual Income	
Gross Wage Income	800,000
Gross Self - Employed Income	100,000
Mechanical Self Employed Income Deduction	-30,000
Total	870,000
Annual Deductions	
Social Security Contributions	135,000
Health Insurance Contribution	54,000
Other Social Sec. Contribution	1,125
Child Care Expenses	33,774
Total	223,899

Panel b. Tax Part Part Calculation

From(\$)	To(\$)	Part(\$)	Rate(%)	Tax(\$)
0	218,232	218,232	0	0
218,232	311,760	93,528	10	9,353
311,760	467,640	155,880	15	23,382
467,640	1,558,800	402,360	20	80,472
1,558,800	2,338,200	-	22	0
2,338,200	3,585,240	-	25	0
3,585,240	and more	-	30	0
Provisional Tax Liability		870,000		113,207

Panel c. Deduction Part Part Calculation

From	To	Part	Rate	Deduction
0	93,528	93,528	10	9,353
93,528	249,408	130,371	15	19,556
249,408	1,340,568	-	20	0
1,340,568	2,119,968	-	22	0
2,119,968	3,367,008	-	25	0
3,367,008	and more	-	30	0
Deductions		223,889		28,908
Final Tax Liability				84,299

A.2.2 Screenshot of the IRPF calculator

In order to give a more precise idea to the taxpayer of how their IRPF tax liability is computed, the IRS developed a tax calculator that can be accessed through their website without any restrictions. Taxpayers can compute their annual tax liability or their monthly tax withholding by introducing their annual labor income and additional information regarding *itemized deductions*. One feature of the tax calculator is that it provides a detailed description of all the steps used to compute the *provisional tax liability* and the *tax deduction amount*.

Figure A.2.1 shows a screenshot of the first page in the tax calculator exactly as it can be seen by any Uruguayan taxpayer. The shaded cells of the form represent the cells that should be filed by the taxpayer with its personal information. Rows (1) and (2) of the first section are the ones in which taxpayers should input their personal income. Row (1) refers to wage income while Row (2) refers to self-employed income. Row (3) computes the mechanic self-employed deduction which is also automatically considered. Section 2 contains all the fields associated with *itemized* and *non-itemized* deductions. For instance, taxpayers should enter the number of children in Row (1) of the deductions' panel and social security contributions in Row (8). In the third section, taxpayers are supposed to enter the amount of tax that was already withheld by employers and the last row of the page reports the final tax debt or credit according to the specific case.

Figure A.2.2 reports the second page of the tax calculator in which the taxpayer can access to the specific details of how the final tax liability is computed. Both panels represent each part of the tax schedule as reported in Table 1. Details in the calculation are analogous to the example reported in Table A.2.1. In this case, since the amount computed in the *deduction part* is larger than the amount calculated in the *tax part*, the final tax liability for this hypothetical taxpayer is 0.

Figure A.2.1: Tax Calculator - First Page: Input Taxpayer Data

 DIRECCIÓN GENERAL IMPOSITIVA		IMPUESTO A LA RENTA DE LAS PERSONAS FÍSICAS Categoría II: Rentas del Trabajo Estimación del IRPF individual - Período liquidación 2013	
Ingresos:			
Ingresos en relación de dependencia		\$	306,985
Ingresos fuera de la relación de dependencia (IVA excluido)		\$	0
Deducción legal 30%		\$	0
Importes facturados por los escribanos correspondientes a aportes a la Caja Notarial		\$	0
Deducción por incobrables (ingresar, si corresponde, valores positivos)		\$	0
Total Ingresos:		\$	306,985
Deducciones:			
Cantidad de personas a cargo:	Deducción de 100%		
	Sin Discapacidad (hijos menores)		1
	Con Discapacidad		0
	Importe a Deducir	\$	33,774
Si es profesional:	Fondo de solidaridad	0	\$ 0
	Adicional Fondo de solidaridad	No	\$ 0
	Aporte CJPPU o Caja Notarial		\$ 0
Aportes a la Seguridad Social del período:			
	Aportes Jubilatorios	\$	57,390
	Aportes FONASA, Seguros de salud	\$	
	Aporte FRL	\$	
Otras Deducciones (Ingresar importe del período)		\$	0
Importes reintegrados AFAP (ingresar si corresponde, valores positivos)		\$	0
Importes reintegrados FONASA (ingresar si corresponde, valores positivos)		\$	0
Cuotas préstamos hipotecarios pagadas (ingresar importe del período)		\$	0
Total Deducciones:		\$	91,164
Monto estimado IRPF PERIODO LIQUIDACIÓN 2013		\$	0
Total retenciones en relación de dependencia		\$	0
Total retenciones fuera de la relación de dependencia		\$	0
Total anticipos rentas de trabajo		\$	0
Crédito arrendamiento vivienda permanente			
Arrendamiento pagado \$ (importe del período)	0	Imputación 6% vivienda permanente	\$ 0
Saldo IRPF a pagar 2013		\$	0

Desarrollado por D.G.I.

Figure A.2.2: Tax Calculator - Second Page: Computing Details



IMPUESTO A LA RENTA DE LAS PERSONAS FÍSICAS

Categoría II: Rentas del Trabajo

Estimación del IRPF individual - Período liquidación 2013

Total Ingresos: **306,985**

Total Deducciones: **91,164**

Cálculo del impuesto/deducciones según escala de rentas

Rangos BPC	Desde	Hasta	Ingresos	Tasa	Impuesto
0 A 84 BPC	0	218,232	218,232	0.00%	0
84 A 120 BPC	218,232	311,760	88,753	10.00%	8,875
120 A 180 BPC	311,760	467,640	0	15.00%	0
180 A 600 BPC	467,640	1,558,800	0	20.00%	0
600 A 900 BPC	1,558,800	2,338,200	0	22.00%	0
900 A 1380 BPC	2,338,200	3,585,240	0	25.00%	0
+ DE 1380 BPC	3,585,240		0	30.00%	0
			306,985		8,875

Rangos BPC	Desde	Hasta	Deducciones	Tasa	A Deducir
0 A 36 BPC	0	93,528	91,164	10.00%	9,116
36 A 96 BPC	93,528	249,408	0	15.00%	0
96 A 516 BPC	249,408	1,340,568	0	20.00%	0
516 A 816 BPC	1,340,568	2,119,968	0	22.00%	0
816 A 1296 BPC	2,119,968	3,367,008	0	25.00%	0
+ DE 1296 BPC	3,367,008	-	0	30.00%	0
			91,164		9,116

Monto estimado de IRPF ANUAL **0**

Desarrollado por D.G.I.

A.3 Example of IRPF Grossed-Up Brackets

In this appendix we show a detailed example of the gross-up process to compute the distance of taxable income to the exact point in which marginal tax rate changes. This example is based on data from a real taxpayer of the dataset that earns annually \$306,895 and have child care deductions corresponding to one children in charge. As we explained in Section 2, the first step in computing IRPF was passing the labor income through the progressive tax rates schedule. In this case, the annual labor income (\$306,895) exceeds the first income bracket of the *tax part* (\$218,232) and therefore the *provisional tax liability* is positive and \$8,875 (last row of Panel b. in Table A.3.1). The second step consists of passing the total amount of deductions (\$91,164) through the deduction schedule. In this case, when we compute the *deduction part* from the mechanical and itemized deductions (i.e. child care), the amount of the deduction to be subtracted is \$9,116 (Panel c. of Table A.3.1) which is greater than the *provisional tax liability* (\$8,875) . Therefore, from the third step, it results that the annual labor income from this taxpayer (\$306,985) is above the non taxable income threshold (\$218,232) but because of his deductions (\$91,164) he will not pay IRPF.

To calculate the grossed-up brackets, what we do is to split the total amount of deductions according to the *deduction part* schedule (column 3 of Panel c. in Table A.3.1) and sum these “pieces” of deductions into the income brackets from the *tax part*. Each piece is added to the upper bound of the immediately previous bracket and therefore also to lower bound of the income bracket that has the same marginal rate. In other words, we stretch the bounds of the *tax part* brackets by including the amount of deductions that corresponds to each one of them.

In our example, we take the \$91,164 that are deducted at a rate of 10% and sum this amount to the upper bound of the immediately previous income bracket (0%) and to the lower bound of the corresponding income bracket (10%). Thus the new individual non taxable income threshold raises up to \$309.396 which is the result of adding \$218,232 and \$91,164 (See Table A.3.1 of Appendix A.3). In addition, the second bracket is also modified since its lower bound changed and the upper bound remain unchanged. This new narrower income bracket told us the income part that is going to be taxed at a 10% rate, and it is smaller than the one in the *tax part* because of the simple action of deductions.

Since we want to know how far a taxpayer is from a change in his marginal tax rate, we also need to take into account that *mechanical deductions* generate that changes in income do not only affect income but also deductions. Therefore, we need to apply a correction factor to the grossed-up bracket. As it was explained in Section 3.2, this factor is $\frac{1}{1-\tau}$ where τ is the quotient of *mechanical deductions* over taxable income. In our case, the mechanical deductions of this taxpayer are \$57,390 (which is \$91,164 minus the amount corresponding to the one child in charge \$33,774). This implies that $\tau = 18.69$ and we need to multiply the grossed up brackets by 1.23. Now, we can compute the distance as the difference between the taxable income and the grossed-up brackets corrected by the *mechanical deduction* factor.

Table A.3.1: Example of the Gross-up Process

Panel a. Summary of Income and Deduction Components

Annual Income	
Total	306,985
Annual Deductions	
Total	91,164

Panel b. Tax Part Part Calculation

From(\$)	To(\$)	Part(\$)	Rate(%)	Tax(\$)
0	218,232	218,232	0	0
218,232	311,760	88,753	10	8,875
311,760	467,640	-	15	0
467,640	1,558,800	-	20	0
1,558,800	2,338,200	-	22	0
2,338,200	3,585,240	-	25	0
3,585,240	and more	-	30	0
Provisional Tax Liability		306,985		8,875

Panel c. Deduction Part Part Calculation

From	To	Part	Rate	Deduction
0	93,528	91,164	10	9,116
93,528	249,408	-	15	0
249,408	1,340,568	-	20	0
1,340,568	2,119,968	-	22	0
2,119,968	3,367,008	-	25	0
3,367,008	and more	-	30	0
Deductions		91,164		9,116
Final Tax Liability				0

Panel d. Grossed-Up Brackets

From(\$)	To(\$)	Part(\$)	Rate(%)	Tax(\$)
0	309,396	306,985	0	0
309,396	311,760	-	10	0
311,760	467,640	-	15	0
467,640	1,558,800	-	20	0
1,558,800	2,338,200	-	22	0
2,338,200	3,585,240	-	25	0
3,585,240	and more	-	30	0
Final Tax Liability				0

A.4 Heterogeneous Responses

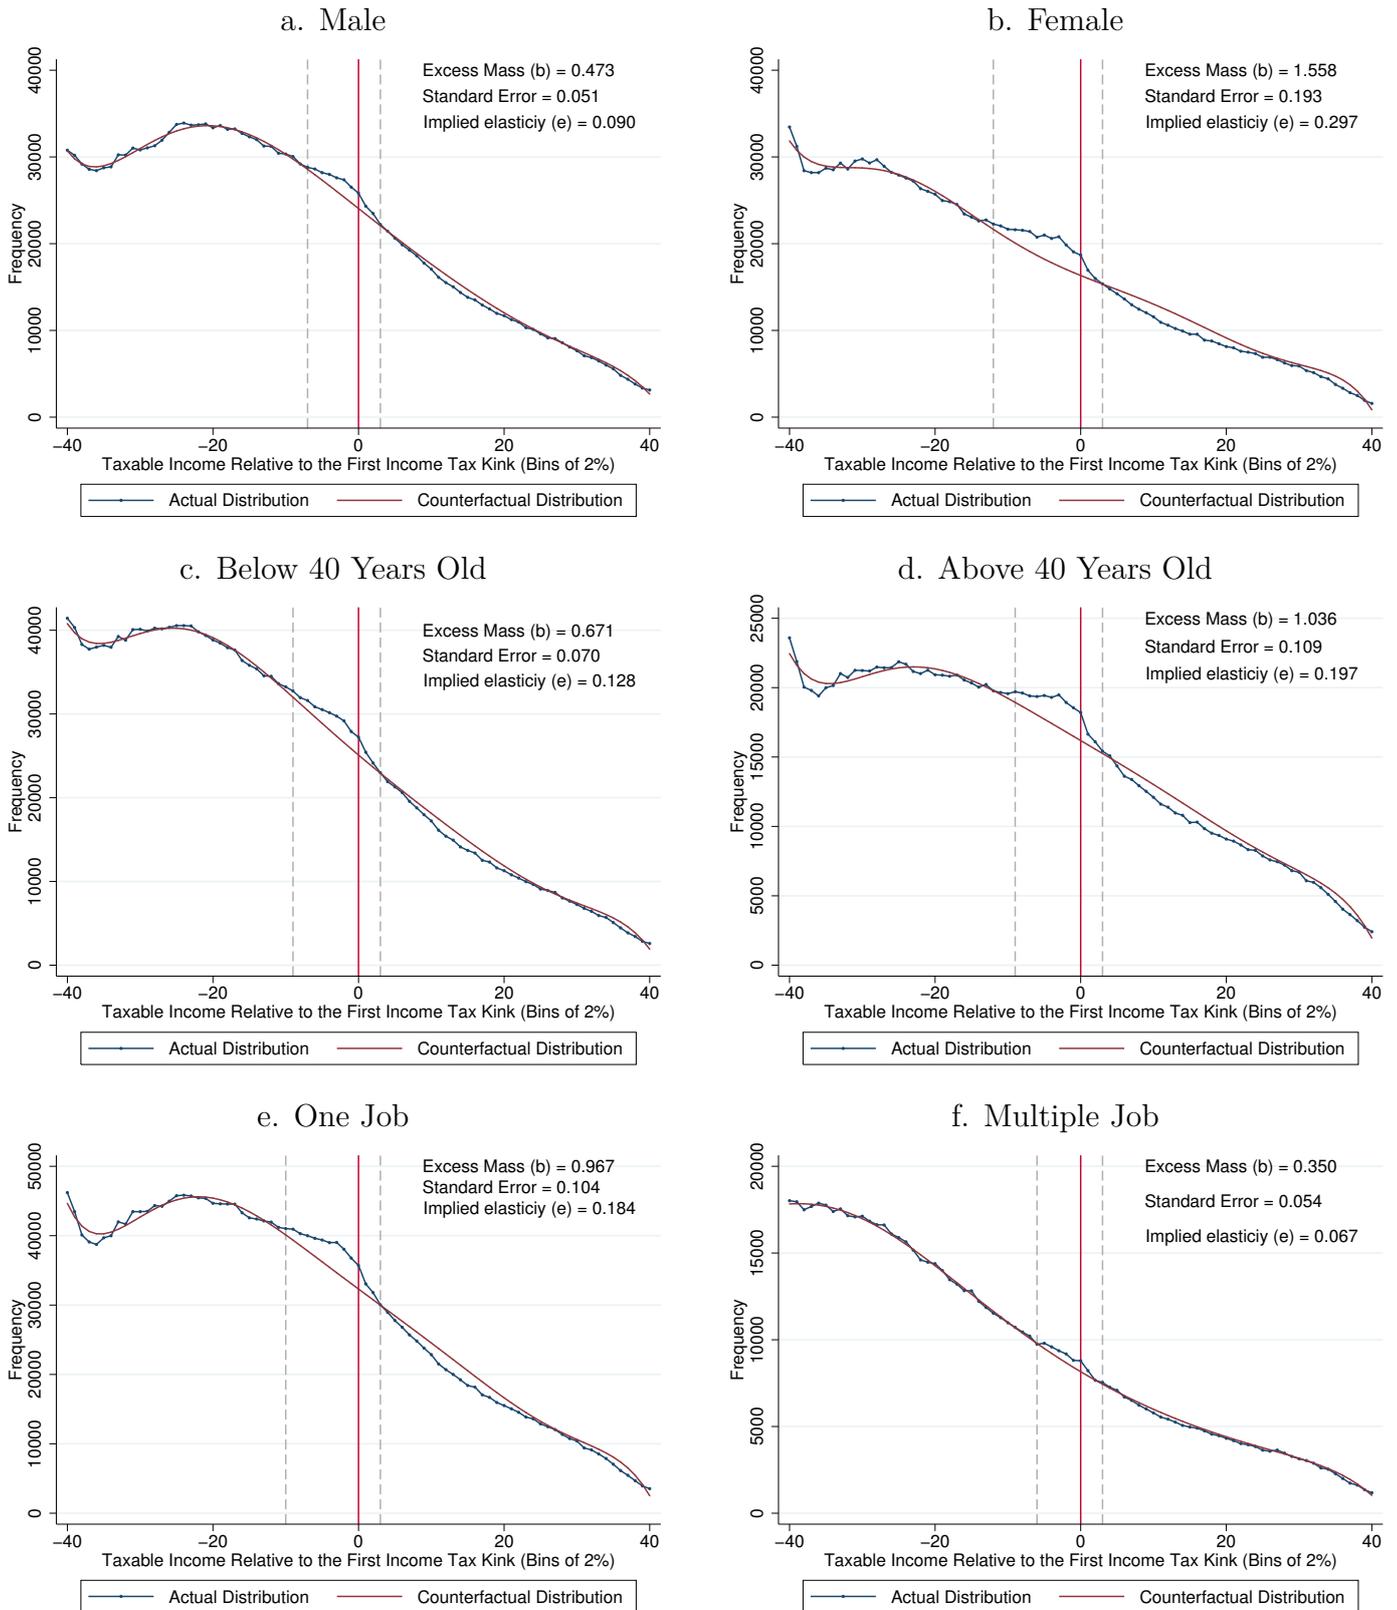
Table A.4.1: Taxable Labor Income Bunching - Heterogeneous Responses in Pure Wage Earners, 2010-2014

	Workers Characteristics			Firms Characteristics	
	Excess of mass (1)	Implied Elasticity (2)		Excess of mass (3)	Implied Elasticity (4)
Male	0.473 (0.051) [290,908]	0.090	Services	0.740 (0.061) [306,363]	0.141
Female	1.558 (0.193) [319,574]	0.297	Goods	0.681 (0.102) [215,178]	0.130
Below 40 years	0.674 (0.071) [363,606]	0.128	Below Median Sales	2.164 (0.264) [96,216]	0.412
Above 40 years	1.036 (0.109) [240,127]	0.197	Above Median Sales	0.431 (0.046) [369,714]	0.082
Multiple employments	0.350 (0.054) [88,704]	0.067			
Only one employment	0.967 (0.104) [524,682]	0.184			

Notes: This table summarizes the heterogeneity analysis reported in Figures A.4.1 and A.4.2 in the Appendix A.4 for the sample of pure wage earners. Estimates are computed for different groups according to personal and job characteristics. Columns (1) and (3) report the excess of mass estimates, (bootstrapped standard errors) and the [number of observations included in the estimation]. Columns (2) and (4) show the implied elasticity estimated with the equations (9) and (10). Left panel shows the estimates of breaking down the sample by personal variables (gender, age and number of jobs): Row (1) shows the taxable income distribution for males while row (2) replicates the analysis for females. Rows (3) and (4) break down the sample by age of the worker dividing the sample into above/below median groups. Rows (5) and (6). report the results for workers with only one job and workers with multiple jobs respectively. Right panel do the same with firms characteristics (size and activity sector): row (1) shows the taxable income distribution for individuals working in firms that perform their economic activity in the Services sector, row (2) is analogous but for individuals who work in the Goods sector. Rows (3) and (4) divide the sample in terms of the size of the firm in which workers are employed. Row (3) reports the results for individuals working in firms whose sales are below the median of sales for the total number of firms included in the sample. Row (4) replicates the analysis for workers that are employed in the largest 50% of the firms. In all cases, estimates are depicted in the exact same way that Figure 3 and details of visual representation can be consulted in the corresponding note.

By Taxpayer's Characteristics

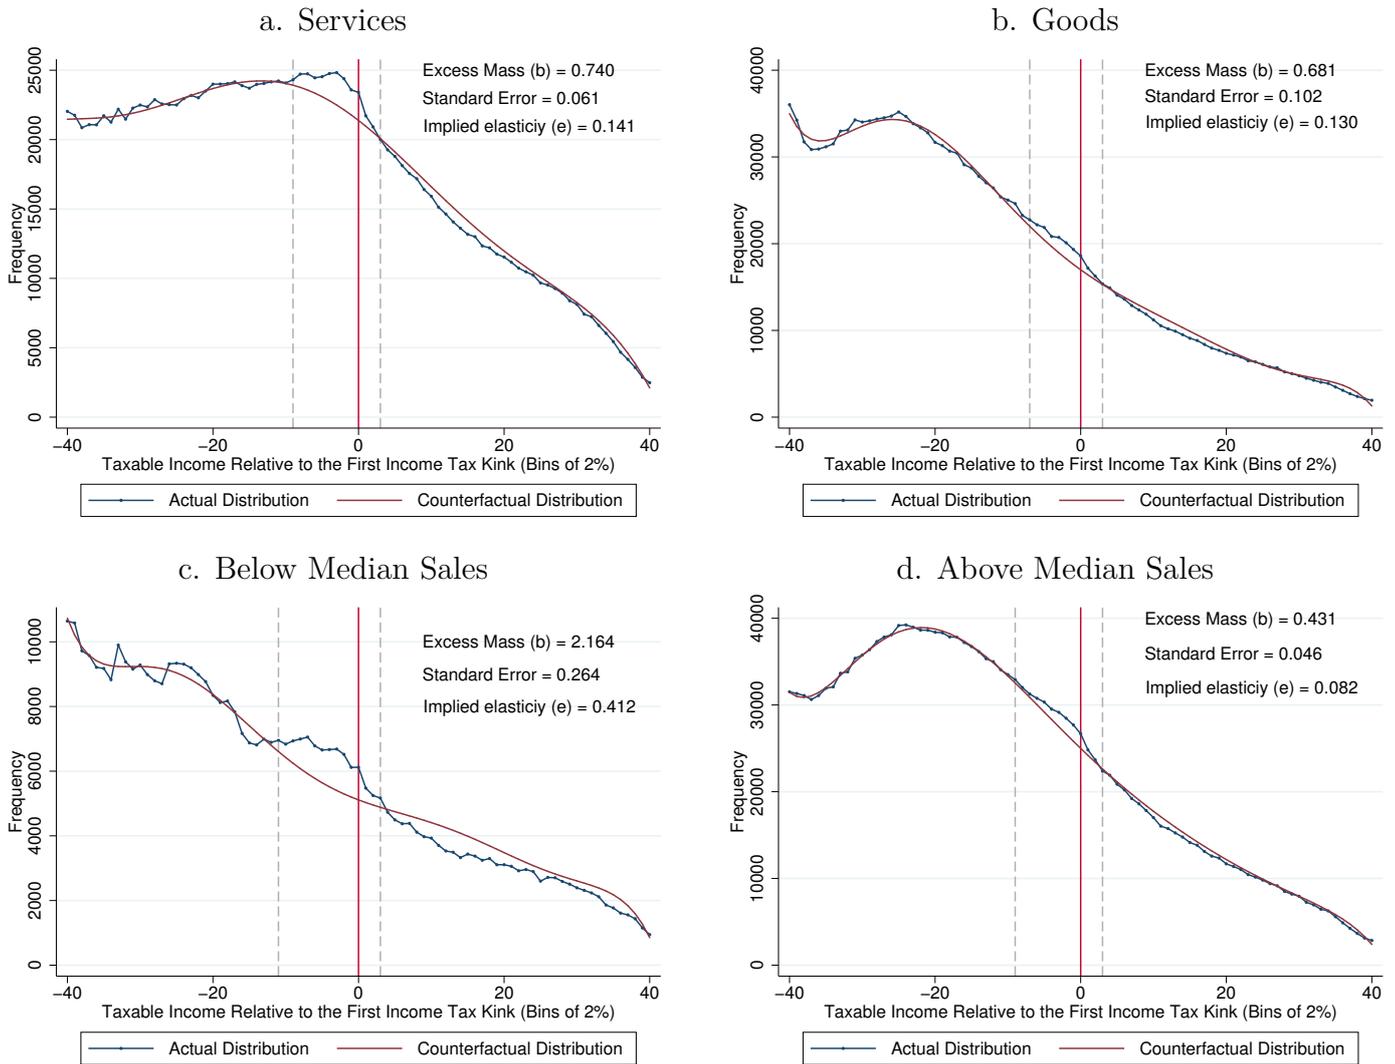
Figure A.4.1: Heterogeneous Responses by Personal Characteristics - Pure Wage Earners (2010-2014)



Notes: The figures in Panels a. through f. plot the empirical and counterfactual taxable income distributions normalized by the individual first kink point as explained in detail in Section 3.2 (See Appendix A.3 for a numerical example) for the sample of pure wage earners. All estimates are performed by using the pooled data for the 2010-2014 period. Panel a. shows the taxable income distribution for males while Panel b. replicates the analysis for females. Panels c. and d. break down the sample by age of the worker dividing the sample into above/below median groups. Panel e. and f. report the results for workers with only one job and workers with multiple jobs respectively. In all cases, estimates are depicted in the exact same way that Figure 3 and details of visual representation can be consulted in the corresponding note.

By Firm's Characteristics

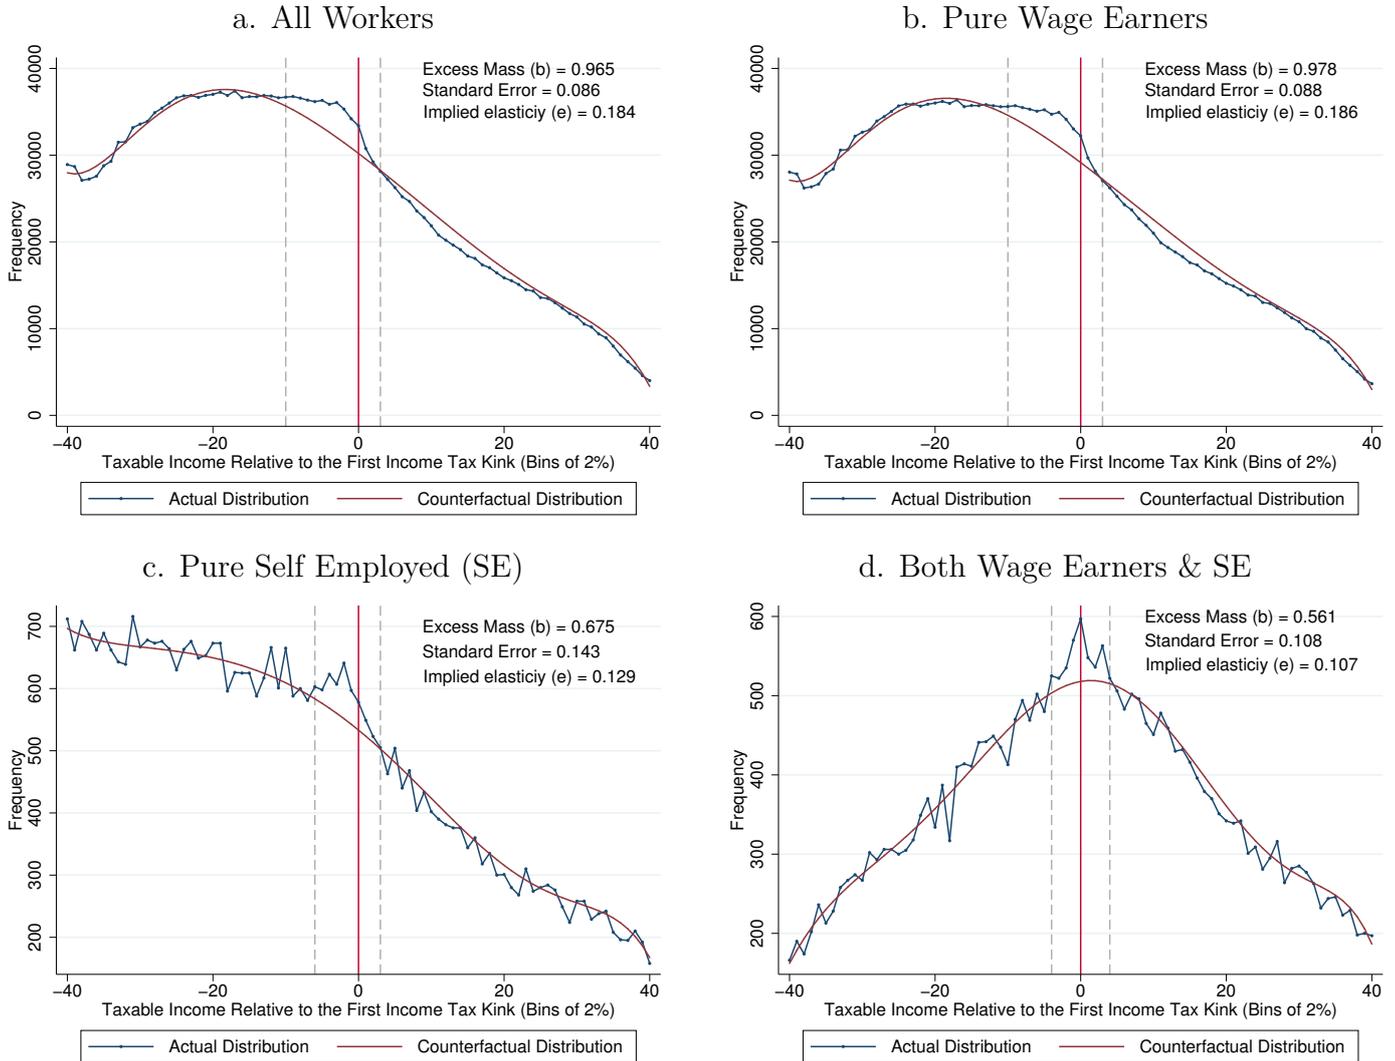
Figure A.4.2: Heterogeneous Responses by Firm's Characteristics - Pure Wage Earners, 2010-2014



Notes: The figures in Panels a. through d. plot the empirical and counterfactual taxable income distributions normalized by the individual first kink point as explained in detail in Section 3.2 (See Appendix A.3 for a numerical example) for the sample of pure wage earners. All estimates are performed by using the pooled data for the 2010-2014 period. Panel a. shows the taxable income distribution for individuals working in firms that perform their economic activity in the Services sector. Panel b. is analogous but for individuals who work in the Goods sector. Panels c. and d. divide the sample in terms of the size of the firm in which workers are employed. Panel c. reports the results for individuals working in firms whose sales are below the median of sales for the total number of firms included in the sample. Panel d. replicates the analysis for workers that are employed in the largest 50% of the firms. In all cases, estimates are depicted in the exact same way that Figure 3 and details of visual representation can be consulted in the corresponding note.

A.5 Robustness Checks

Figure A.5.1: Taxable Labor Income Bunching by Source of Income, 2010-2014: Balanced Sample of Taxpayers



Notes: The figures in Panels a. through d. plot the empirical and counterfactual taxable income distributions normalized by the individual first kink point as explained in detail in Section 3.2 (See Appendix A.3 for a numerical example). Unlike Figure 3 all estimates are performed by using the balanced pooled sample for the 2010-2014 period, i.e. including taxpayers that were in the sample for the five years. Panel a. shows the taxable income distribution for the main sample considering all taxpayers. Panel b. replicates the analysis for pure wage earners, – i.e. those workers who do not receive self-employed income –, Panel c. do the same with taxpayers that only earned income from self-employment work (pure self employed) while Panel d. reports the results for workers that receive both types of income. In all cases, the taxable income distribution is represented by the percentage distance to the first kink point. The vertical red line represents the exact point at which each taxpayer starts paying IRPF, i.e. the first kink point. The data is grouped in bins of 2% wide who are represented by the blue dots. This means that the first bin at the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis to a window of 40 bins at each side of the kink, which is an 80% window. The connected blue dots represent the observed taxable income distribution while the counterfactual estimation is depicted by a dark red line. The counterfactual distribution represents the estimated taxable income distribution in the absence of kinks in the tax schedule. The counterfactual estimation is based on Chetty et al. (2011) and thoroughly detailed in Section ???. We use a 7 degree polynomial and bunching windows, identified with gray dashed lines, are defined in each particular case according to visual inspection. In each figure we include a note with the excess of mass estimates relative to the counterfactual distribution at the kink (b), which is proportional to the magnitude of the labor income response to tax rates (Saez 2010), its standard error and the corresponding implied elasticity. Standard errors are estimated by bootstrapping using a N of 1000. The implied elasticity is computed using equations (9) and (10).

Table A.5.1: Robustness Checks - All Workers

Excess of mass	Implied elasticity	Polynomial Degree	Bunching Window (%)
(1)	(2)	(3)	(4)
0.634 (0.104) [768,619]	0.121	6	(-22, 8)
0.609 (0.092) [677,084]	0.116	6	(-20, 6)
0.568 (0.084) [584,524]	0.108	6	(-18, 4)
0.935 (0.095) [768,619]	0.178	7	(-22, 8)
0.844 (0.086) [677,084]	0.161	7	(-20, 6)
0.744 (0.080) [584,524]	0.142	7	(-18, 4)
0.547 (0.052) [768,619]	0.104	8	(-22, 8)
0.526 (0.044) [677,084]	0.100	8	(-20, 6)
0.489 (0.040) [584,524]	0.093	8	(-18, 4)

Notes: This table shows estimates of alternative specifications of the main results reported Panel a. in Figure 3, i.e. all workers. We vary the polynomial degree between 6 and 8 and the width of the bunching window by +/- 1 bin (or equivalently 2%). Column (1) contains information of the excess of mass estimates, (bootstrapped standard errors) and the [number of observations included in the estimation]. Column (2) of each panel shows the implied elasticity estimated with equations (9) and (10). Column (3) show the polynomial degree used for the estimation, while column (4) shows the bunching zone defined for the counterfactual estimation.

Table A.5.2: Robustness Checks - Pure Wage Earners

Excess of mass	Implied elasticity	Polynomial Degree	Bunching Window (%)
(1)	(2)	(3)	(4)
0.637 (0.103) [744,430]	0.121	6	(-22, 8)
0.611 (0.097) [655,739]	0.116	6	(-20, 6)
0.572 (0.088) [566,151]	0.109	6	(-18, 4)
0.945 (0.100) [744,430]	0.180	7	(-22, 8)
0.851 (0.088) [655,739]	0.162	7	(-20, 6)
0.753 (0.081) [566,151]	0.143	7	(-18, 4)
0.549 (0.052) [744,430]	0.105	8	(-22, 8)
0.527 (0.046) [655,739]	0.100	8	(-20, 6)
0.493 (0.039) [566,151]	0.094	8	(-18, 4)

Notes: This table shows estimates of alternative specifications of the main results reported Panel b. in Figure 3, i.e. pure wage earners. We vary the polynomial degree between 6 and 8 and the width of the bunching window by +/- 1 bin (or equivalently 2%). Column (1) contains information of the excess of mass estimates, (bootstrapped standard errors) and the [number of observations included in the estimation]. Column (2) of each panel shows the implied elasticity estimated with equations (9) and (10). Column (3) show the polynomial degree used for the estimation, while column (4) shows the bunching zone defined for the counterfactual estimation.

Table A.5.3: Robustness Checks - Pure Self Employed

Excess of mass	Implied elasticity	Polynomial Degree	Bunching Window (%)
(1)	(2)	(3)	(4)
0.576 (0.138) [9,708]	0.110	6	(-14, 8)
0.627 (0.124) [8,204]	0.119	6	(-12, 6)
0.605 (0.100) [6,662]	0.115	6	(-10, 4)
0.597 (0.147) [9,708]	0.114	7	(-14, 8)
0.646 (0.124) [8,204]	0.123	7	(-12, 6)
0.619 (0.099) [6,662]	0.118	7	(-10, 4)
0.432 (0.164) [9,708]	0.082	8	(-14, 8)
0.529 (0.137) [8,204]	0.101	8	(-12, 6)
0.529 (0.107) [6,662]	0.101	8	(-10, 4)

Notes: This table shows estimates of alternative specifications of the main results reported Panel c. in Figure 3, i.e. pure self-employed workers. We vary the polynomial degree between 6 and 8 and the width of the bunching window by +/- 1 bin (or equivalently 2%). Column (1) contains information of the excess of mass estimates, (bootstrapped standard errors) and the [number of observations included in the estimation]. Column (2) of each panel shows the implied elasticity estimated with equations (9) and (10). Column (3) show the polynomial degree used for the estimation, while column (4) shows the bunching zone defined for the counterfactual estimation.

Table A.5.4: Robustness Checks - Both Self Employed and Wage Earners

Excess of mass	Implied elasticity	Polynomial Degree	Bunching Window (%)
(1)	(2)	(3)	(4)
0.402 (0.110) [7,670]	0.077	6	(-10, 10)
0.509 (0.096) [6,394]	0.097	6	(-8, 8)
0.445 (0.079) [5,036]	0.085	6	(-6, 6)
0.403 (0.109) [7,670]	0.077	7	(-10, 10)
0.510 (0.093) [6,394]	0.097	7	(-8, 8)
0.446 (0.080) [5,036]	0.085	7	(-6, 6)
0.368 (0.129) [7,670]	0.070	8	(-10, 10)
0.505 (0.108) [6,394]	0.096	8	(-8, 8)
0.429 (0.088) [5,036]	0.082	8	(-6, 6)

Notes: This table shows estimates of alternative specifications of the main results reported Panel d. in Figure 3, i.e. workers earning both wage and self-employed income. We vary the polynomial degree between 6 and 8 and the width of the bunching window by +/- 1 bin (or equivalently 2%). Column (1) contains information of the excess of mass estimates, (bootstrapped standard errors) and the [number of observations included in the estimation]. Column (2) of each panel shows the implied elasticity estimated with equations (9) and (10). Column (3) show the polynomial degree used for the estimation, while column (4) shows the bunching zone defined for the counterfactual estimation.

Heterogeneous Responses Using a Balanced Sample

Table A.5.5: Taxable Labor Income Bunching - Heterogeneous Responses in Pure Wage Earners, 2010-2014. Balanced Sample

	Workers Characteristics			Firms Characteristics	
	Excess of mass (1)	Implied Elasticity (2)		Excess of mass (3)	Implied Elasticity (4)
Male	0.549 (0.056) [202,301]	0.105	Services	1.012 (0.070) [237,940]	0.193
Female	1.519 (0.162) [252,060]	0.289	Goods	0.659 (0.088) [160,488]	0.126
Below 40 years	0.736 (0.074) [258,196]	0.140	Below Median	2.491 (0.256) [81,372]	0.474
Above 40 years	1.140 (0.098) [186,902]	0.217	Above Median	0.458 (0.047) [269,942]	0.087
Multiple sources of income	0.491 (0.060) [66,352]	0.094			
One source of income	1.072 (0.095) [386,248]	0.204			

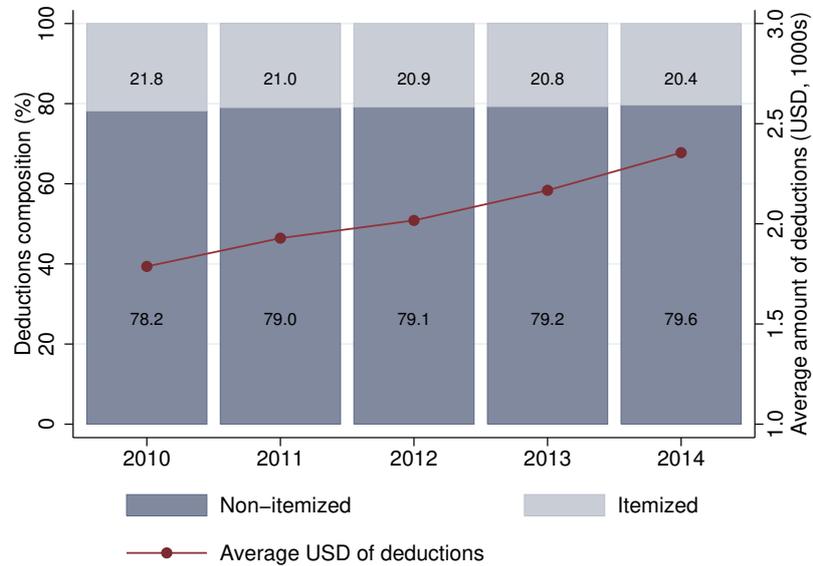
Notes: This table replicates Table A.4.1 in Section A.4 but for a balanced panel of pure wage earners. Estimates are computed for different groups according to personal and job characteristics. Columns (1) and (3) report the excess of mass estimates, (2) (bootstrapped standard errors) and (3) the [number of observations included in the estimation]. Columns (2) and (4) show the implied elasticity estimated with the equations (9) and (10). Left panel shows the estimates of breaking down the sample by personal variables (gender, age and number of jobs): Row (1) shows the taxable income distribution for males while row (2) replicates the analysis for females. Rows (3) and (4) break down the sample by age of the worker dividing the sample into above/below median groups. Rows (5) and (6). report the results for workers with only one job and workers with multiple jobs respectively. Right panel do the same with firms characteristics (size and activity sector): row (1) shows the taxable income distribution for individuals working in firms that perform their economic activity in the Services sector, row (2) is analogous but for individuals who work in the Goods sector. Rows (3) and (4) divide the sample in terms of the size of the firm in which workers are employed. Row (3) reports the results for individuals working in firms whose sales are below the median of sales for the total number of firms included in the sample. Row (4) replicates the analysis for workers that are employed in the largest 50% of the firms. In all cases, estimates are depicted in the exact same way that Figure 3 and details of visual representation can be consulted in the corresponding note.

A.6 Composition of Deductions

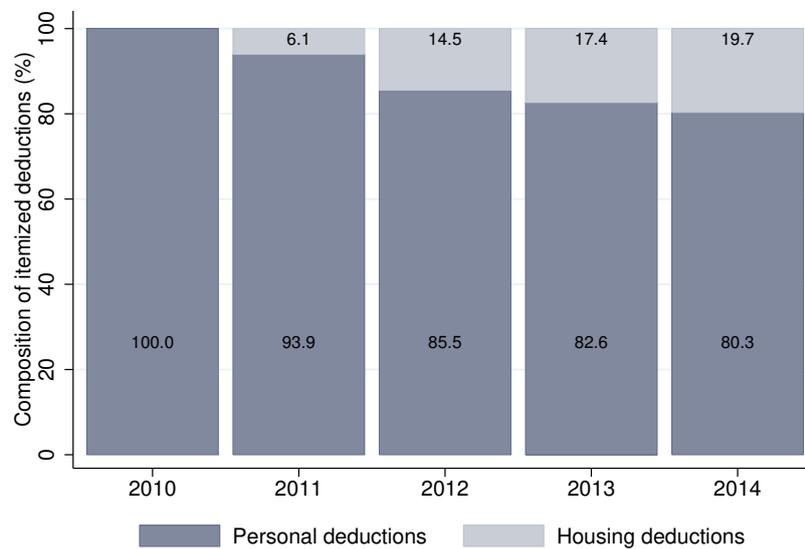
A.6.1 Deductions: Growth and Composition Over Time

Figure A.6.1: Evolution and Composition of IRPF's Deductions, 2010-2014

a. Growth of Total Deductions and Composition Over Time



b. Composition of Itemized Deduction Over Time



Notes: Panel a. shows the composition of total IRPF's deductions between 2010 and 2014. This information comes from the administrative data. The total deductions are decomposed in non-itemized and itemized components. Non-itemized deductions include payroll taxes and health insurance contributions. Itemized deductions include personal deductions and housing expenses. Panel b. shows the composition of itemized deductions between 2010 and 2014. The itemized deductions are decomposed in personal deductions and housing deductions. Personal deductions include Other social security contributions and child care deductions. Housing deductions include rent deductions and mortgage deductions.