

## DIGGING INTO THE CHANNELS OF BUNCHING: EVIDENCE FROM THE URUGUAYAN INCOME TAX\*

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

Based on detailed administrative tax records, we implement a bunching design to explore how individual taxpayers respond to personal income taxation in Uruguay. We estimate a very modest elasticity of taxable income at the first kink point (0.06) driven by a combination of gross labour income and deduction responses. Taxpayers use personal deductions more intensively close to the kink point and underreport income to the tax authority. Our results suggest that the efficiency costs of taxation are not necessarily large in contexts characterised by limited deduction opportunities. Policy efforts should be directed at broadening the tax base and improving enforcement capacity.

Enhancing state capacity and the ability to collect taxes is a major concern for many governments, particularly in developing countries. Taxation can play an important role in boosting economic development in the long run and in reducing income inequality (Besley and Persson, 2013; Dincecco and Katz, 2016; Acemoglu *et al.*, 2018). While total tax collection in developing countries has increased in recent years, it remains relatively low—around 20% of GDP in 2018 (22.3% in Latin America and 17.2% in African countries), as compared to an average of 34.3% for OECD countries in the same year.<sup>1</sup> Developing countries rely heavily on indirect taxation and have almost no margin to increase the tax burden associated with excise taxes.<sup>2</sup> By contrast, direct taxes—in particular, personal income taxation—play a relatively minor role in the overall tax structure.<sup>3</sup> This affects the overall level of total tax revenues, but also the ability to curb prevailing levels of wealth and income concentration (Turnovsky and Basher, 2009; ECLAC, 2013; Duncan

\* Corresponding author: Gabriel Burdin, Leeds University Business School—The University of Leeds, Maurice Keyworth Building, Leeds, LS2 9JT, UK. Email: [G.Burdin@leeds.ac.uk](mailto:G.Burdin@leeds.ac.uk)

*This paper was received on 9 June 2020 and accepted on 4 January 2021. The Editor was Barbara Petrongolo.*

The authors were granted an exemption to publish their data because access to the data is restricted. However, the authors provided the Journal with temporary access to the data, which allowed the Journal to run their codes. The codes are available on the Journal website. The data and codes were checked for their ability to reproduce the results presented in the paper.

We thank the Dirección General Impositiva (DGI) for providing us with the administrative tax records, and especially Gustavo Gonzalez and his staff for their comments on this work. We also wish to thank the editor, Barbara Petrongolo, and two anonymous referees for their substantial comments. We are grateful to 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). An earlier version of this paper circulated under the title “Tax Bunching at the Kink in the Presence of Low Capacity of Enforcement: Evidence from Uruguay”. This research was funded by the Agencia Nacional de Investigación e Innovación (ANII) under grant number FCE.1.2014.104284.

<sup>1</sup> Data extracted on 28 April 2020 16:04 UTC (GMT) from OECD.Stat.

<sup>2</sup> OECD 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.

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

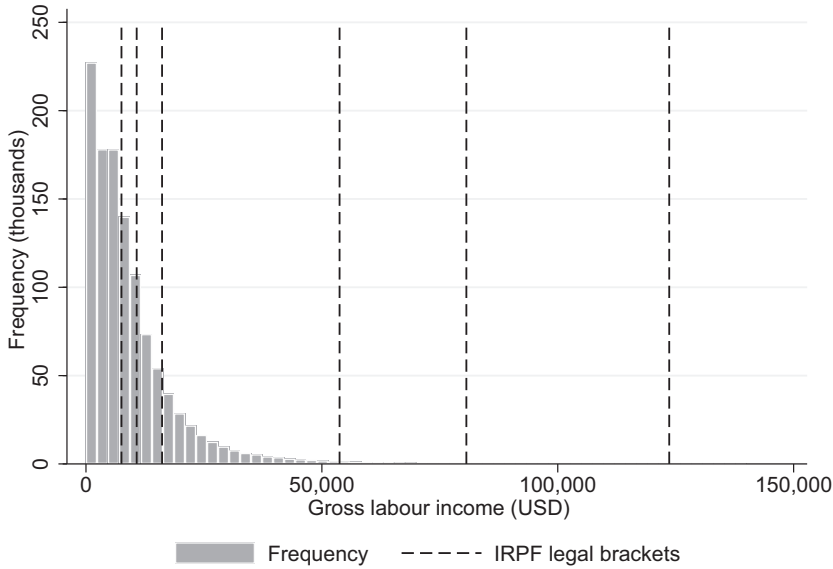


Fig. 1. *Tax Schedule and Gross Labour Income Distribution (2013).*

*Notes:* In this figure we show the gross labour income distribution (i.e., before deductions) and income brackets in 2013. Gross labour income is defined as all income earned by workers before considering any deductions. Each bar represents the number of people located in a particular gross labour income bin measured in thousands (y axis). Dashed lines show the income level at which the marginal tax rates change.

and Sabirianova Peter, 2016). In this context, one natural question is how to collect taxes more efficiently, without sacrificing the redistributive component of taxation.

One obvious concern when considering whether and how to increase tax capacity and the progressivity of personal income taxation is efficiency costs. Developing countries usually have large informal sectors. With narrow tax bases and limited enforcement capacity, the standard economic argument suggests that one should observe large behavioural responses to taxation in such countries (Bird and Zolt, 2004; Gordon and Li, 2009). However, empirical studies aimed at understanding how individuals respond to taxation in developing countries, and using both credible identification strategies and high-quality data, remain scarce.<sup>4</sup> In this paper we aim to fill this gap. We estimate the elasticity of taxable income using a bunching approach that exploits discrete changes in the marginal tax rates of a recently implemented progressive personal income tax in Uruguay (Saez, 2010; Chetty *et al.*, 2011). Specifically, we study individuals' behavioural responses to the first kink point of the tax schedule, where the marginal tax rate changes from 0% to 10%, i.e., the exemption threshold. We focus on this kink to obtain the most reliable estimates of bunching responses to taxation (Chetty *et al.*, 2009; Chetty, 2012): the largest change in economic incentives in terms of net-of-tax rate at the whole schedule occurs at the first kink (see Figure 2(a)); it is where an individual starts paying income taxes and thus is arguably the most salient kink point for taxpayers; and, finally, the first bracket is the one affecting the largest percentage of Uruguayan taxpayers (Figure 1).<sup>5</sup> Our underlying identification assumption is that,

<sup>4</sup> Recent empirical research regarding behavioural responses to personal income taxation has mostly analysed rich countries (Saez *et al.*, 2012; Kleven *et al.*, 2016).

in the absence of the progressive tax schedule, the labour income distribution would be smooth around the first kink point.

We use administrative records from the Uruguayan tax agency for the period 2010–2016 to construct a unique dataset. We combine tax returns containing self-reported information, employer-reported information about individuals' earnings and deductions, and firm-level records at the individual level. The result is a comprehensive dataset covering more than 1.6 million workers over a period of seven years. Both the large size and high quality of the information contained in the dataset allow us to overcome problems of attrition, measurement error and precision that typically arise in studies using survey data (Card *et al.*, 2010). The richness and granularity of the data allow us to examine in great detail the mechanisms used by taxpayers to adjust their reported income.

We divide the empirical analysis into two parts. In the first part, we examine whether individuals respond to labour income taxation and estimate the magnitude of such a response. In particular, we estimate the excess mass in the observed labour income distribution and we compare it to the estimated counterfactual distribution around the first kink point. To do this, we use a bunching design as developed by Saez (2010) and Chetty *et al.* (2011).<sup>6</sup> This estimate allows us to identify the elasticity of taxable income (ETI) with respect to the marginal net-of-tax rate. In the second part of the analysis, we investigate the underlying mechanisms behind the observed responses. For this purpose, we conduct the main analysis in different sub-samples of individuals and examine deduction behaviour by taking advantage of specific features of the Uruguayan tax code (e.g., the existence of itemised and non-itemised deductions). We also look at within-individual discrepancies between self- and third-party-reported gross labour income at the kink point to explore the extent to which income underreporting behaviour reacts to tax rate incentives.

Our analysis yields three key findings. First, there is a positive excess mass of taxpayers around the first kink point. Although statistically significant, the associated ETI is very modest, roughly 0.06. Interestingly, we document an increase in the number of bunchers—i.e., people located in the bunching zone—over time, a dynamic pattern consistent with a learning process. Second, the comparison of gross labour income (before deductions) and taxable labour income (combination of gross labour income and deductions) responses rules out that bunching behaviour is exclusively driven by the gross labour income component. The fact that there is no sign of bunching for the gross labour income distribution highlights the critical role played by deduction behaviour. In other words, before considering deductions, we find no evidence of a behavioural response to personal income taxation. Furthermore, our findings show that individuals who claim *itemised* deductions are more responsive to changes in the marginal net-of-tax rate than are individuals who do not claim this type of deduction. Distinguishing between different types of deductions is informative because, while *non-itemised deductions* are automatically considered for the purpose of calculating individuals' tax liability (taxpayers have no discretion in relation to this type of deduction), *itemised deductions* need to be explicitly and voluntarily claimed by the taxpayer. In addition, we find a relatively greater use of *itemised* deductions around the kink point and suggestive evidence that this response is related to a deduction behaviour *per se* rather than to the act of filing a tax return. This distinction is relevant because there is a

<sup>5</sup> Chetty (2012) showed that larger changes in the tax rates are better suited to estimate behavioural responses to taxation in environments where frictions can play a major role. Moreover, Chetty *et al.* (2009) showed that it is expected for individuals to underreact to taxation in contexts where taxes are less salient. Nonetheless, we also investigated the potential existence of behavioural adjustments at the other kinks of Uruguay's income tax schedule without finding any evidence of bunching responses (see the discussion in Subsection 4.3).

<sup>6</sup> Kleven (2016) provided a review of the bunching approach and its application.

strong, though not perfect, correlation between claiming an itemised deduction and filing a tax return. Hence, the two mechanisms can be easily confounded. We show, however, that part of the response observed in deduction behaviour is explained by taxpayers who did not file a tax return. Finally, consistent with individuals underreporting their income, we document a sharp jump in the observed discrepancies between self- and third-party-reported gross labour income for *non-itemising* taxpayers just at the kink point. There is no evidence of underreporting occurring via collusive employer-employee agreements at the firm level.

The ETI of 0.06 we determined for Uruguayan taxpayers falls within the range of elasticities estimated for individuals in the bottom-middle part of the income distribution in advanced economies. This result contrasts with the hypothesis that larger responses will occur in contexts with narrow tax bases and limited enforcement capacity. We argue that two factors help explain this finding. First, wage earners constitute more than 95% of the total number of workers in Uruguay. Since wage earners are usually subject to more frictions in the labour market, the opportunities for tax planning are fewer than for self-employed workers. This is consistent with previous evidence about wage earners' responses (e.g., Chetty *et al.*, 2011; Paetzold, 2019). Second, the tax schedule in Uruguay offers a very limited set of deductions, which, in turn, restricts the margins for tax avoidance. This is consistent with Slemrod and Kopczuk (2002) and Kopczuk (2005), who showed that the ETI is considerably larger when extensive deductions are allowed.

Our paper contributes to four strands of literature. As a first contribution, we add to the growing empirical literature that analyses local behavioural responses by exploiting discontinuities in the income tax schedule using the bunching approach (Kleven, 2016). Our contribution to this literature is two-fold. First, most prior research analyses behavioural responses either in the reporting of gross labour income or in deductions.<sup>7</sup> Our empirical setting allows us to consider both components of the taxable income simultaneously. Indeed, we show that the behavioural responses involve changes in both gross labour income *and* deductions. Failing to account for both channels of response might lead to an incomplete understanding of how taxpayers react to local incentives in tax schedules. In addition, we rely on a new method for implementing the bunching approach. Typically, researchers using bunching techniques use visual inspection to define both the bunching zone and the polynomial degree for constructing the predicted counterfactual labour income distribution. We depart from this visual approach and use instead a more systematic rule following Bosch *et al.* (2020). Using this data-driven approach to select key parameters of our estimating equation, we restrict our margin of discretion. To our knowledge, except for the original paper that proposed this approach, our paper is the first to implement this procedure. We believe that “tying our hands” when defining these two parameters enhances the validity and robustness of our results.

Our second contribution relates to the few recent quasi-experimental studies of behavioural responses to taxation in developing countries using administrative data and, in particular, to the very few that study the phenomenon using a bunching design. Bohne and Nimczik (2017) is the closest paper to ours.<sup>8</sup> The authors showed that Ecuadorian taxpayers increase their knowledge of the tax code over time, and this results in significant adjustments in reported taxable labour

<sup>7</sup> For instance, Australia (Hamilton, 2018), Austria (Paetzold, 2019), Denmark (Chetty *et al.*, 2011; le Maire and Schjerning, 2013), Germany (Schachtele, 2016), Ireland (Hargarden, 2020), Sweden (Bastani and Selin, 2014) and United States (Mortenson and Whitten, 2020).

<sup>8</sup> Other contributions include Kleven and Waseem (2013), Best *et al.* (2015), Best (2015), Carrillo *et al.* (2017) and Bachas and Soto (2018).

income. The primary margin of response is associated with the deductions component and driven, the authors claim, by information transmission between individuals moving across firms. Our paper provides evidence consistent with responses in the deduction margin being a key driver of the overall response, but we show, in addition, that income underreporting behaviour is also relevant.

We also speak to the literature on the anatomy of taxable income elasticity. It has been shown that, under some circumstances, distinguishing between “real” earnings responses (e.g., due to changes in hours worked) and sheltering responses (such as tax avoidance or evasion) could be very relevant to gauge the welfare costs of taxation. These circumstances usually include the existence of fiscal externalities or when sheltering/evasion behaviour is associated with transfer costs rather than resource costs (Slemrod, 1998; Chetty, 2009; Gillitzer and Slemrod, 2016; Doerrenberg *et al.*, 2017).<sup>9</sup> Beyond efficiency considerations, determining whether behavioural responses are driven by real or “non-real” mechanisms may also affect tax policy prescriptions (Kopczuk, 2005; Saez *et al.*, 2012). The empirical bunching literature, in turn, finds that sheltering responses have played a key role in different tax systems (le Maire and Schjerning, 2013; Schachtele, 2016; Bohne and Nimczik, 2017; Doerrenberg *et al.*, 2017; Paetzold, 2019). We contribute to this literature by disentangling two different sheltering responses to labour income taxation: deductions and income underreporting. The relative importance of these channels depends on the actual adjustment opportunities that different groups of taxpayers have at their disposal.

Finally, our findings also relate to the literature that explores tax evasion responses to changes in the marginal tax rate.<sup>10</sup> More specifically, we add to a small but growing literature that explores the extent of tax evasion using discrepancies between different reports of the same tax base, either among workers (e.g., Kumler *et al.*, 2020; Bergolo and Cruces, 2014; Best, 2015) or firms (e.g., Fisman and Wei, 2004; Carrillo *et al.*, 2017). By using quasi-experimental variation and comparing self- and third-party-reported information, we show that tax evasion can also be a relevant margin of response to tax rate changes for wage earners. Our evidence casts new doubts regarding the effectiveness of third-party reporting in contexts characterised by relatively weak enforcement.

The rest of the paper is organised as follows. In Section 1 we describe the institutional background. In Section 2 we present the empirical strategy as well as some very specific features of the Uruguayan setting that need to be addressed in the empirical strategy. In Section 3 we describe the data used for the analysis. In Section 4 we report the overall estimates for the ETI while in Section 5 we analyse the mechanisms that explain the estimated response. In Section 6 we conclude and discuss some implications of our findings for tax design in developing countries.

<sup>9</sup> Feldstein (1995; 1999) showed that, under some conditions, the ETI is a sufficient statistic for the efficiency costs of taxation. This approach assumes that the marginal social cost of reducing taxable income by \$1 is equal to the marginal tax rate that, at the optimum, is the same for every possible margin of response. In this framework, for instance, the marginal cost of sheltering \$1 equals the net marginal cost of reducing income by \$1 (see Saez *et al.*, 2012 for a discussion). However, as shown by Chetty (2009), in some cases the type of response matters a lot for efficiency cost calculations. For instance, in the presence of fiscal externalities—e.g., when the costs of sheltering are completely transferred to a different agent in the economy—individuals’ sheltering responses to taxation do not generate efficiency losses at all. Sheltering would be only a reallocation of resources across agents.

<sup>10</sup> Some early contributions include Clotfelter (1983) and Feinstein (1991). More recently, Kleven *et al.* (2011) exploited kinks in the Danish income tax schedule combined with audit data to show that bunching behaviour is partly explained by underreported self-employment income.

## 1. Personal Labour Income Taxation in Uruguay

In Uruguay, the progressive labour income tax (*Impuesto a la Renta de las Personas Físicas* or IRPF in Spanish) was first introduced in 2007 in the context of a more general tax reform that aimed at increasing the relative weight of direct taxation.<sup>11</sup> Despite the introduction of the IRPF, the tax burden still relies heavily on indirect taxation. According to the Uruguayan Tax Agency (*Dirección General Impositiva*; henceforth DGI), value-added tax (VAT) constitutes about 50% of total tax collections while IRPF, at 12.5%, is second in importance.

In Figure 1 we show the gross labour income distribution in 2013 and the seven IRPF income brackets. Bracket cutoff points are depicted by vertical dashed lines. The burden of the IRPF relies mainly on individuals in the middle and upper parts of the gross labour income distribution: individuals with annual labour income below 7,525 United States dollars (USD)—the tax exemption threshold—were exempted from paying IRPF. As a reference, the minimum annual wage for 2013 was USD 3,316, while the median labour income was USD 6,038. According to DGI reports, only about 30% of registered employees paid labour income tax during the 2010–2016 period (DGI, 2016). Nevertheless, the number of people affected by the tax has been increasing since it was first implemented. In panel a of Figure A.3 of Online Appendix A we show that the total number of registered employees—those who are *potential* IRPF taxpayers depending on their earnings—increased by 9.3% between 2010 and 2016. In panel b we show the number of tax returns filed (left axis) and the share of taxpayers making positive IRPF payments (right axis). The number of tax returns filed more than doubled in the period considered (from 102,448 to 209,353) while the number of workers making positive IRPF payments also increased, although by a smaller percentage (from 27.5% to 33.3%).

It is worth mentioning that over our period of analysis (2010–2016) the IRPF marginal rates only changed for the highest income brackets. These changes have no implications for our main results, which focus on the bottom bracket.

### 1.1. Tax Schedule and Labour Income Components

The IRPF relies on two components: a *tax part* and a *deduction part*. The difference between these components determines the IRPF's total tax liability ( $L$ ). A simplified version of  $L$  can be expressed as

$$L(y, d_i, d_n) = T(y) - D(d_i + d_n) \quad \text{with } d_n = c \times y, \quad (1)$$

where  $y$  is the gross labour income,  $d_i$  and  $d_n$  represent *itemised* and *non-itemised* deductions, respectively,  $T$  is a progressive set of tax rates and brackets that determines the *tax part* and  $D$  is also a progressive set of rates and brackets but in this case for deductions. It is worth noting that  $d_n = c \times y$ , where  $c$  is a proportional rate.

As with most personal income tax structures around the world, the *tax part* is determined by function  $T$  of the gross labour income  $y$ . The left panel in Table 1 shows a detailed list of labour income components that are taxed by the IRPF. In general, it includes all items related to wages and self-employment income received by a taxpayer during a fiscal year.

<sup>11</sup> Our definition of labour income does not include any type of capital income that is taxed in a separate schedule at proportional tax rates. The IRPF also contains a different schedule for married individuals filing jointly with their spouse. During the period covered by our analysis, only 0.5% of taxpayers chose to pay taxes as a household unit, i.e., jointly. For simplicity, these households were excluded from the analysis. We include a detailed description of the context and the tax structure in Online Appendix A.

Table 1. *Tax Schedule Components: Gross Labour Income and Deductions.*

Sources of income	Deductions
Wages and salaries	<i>Non-itemised</i>
Self employed income (70%)	(a) Payroll taxes
Tips and overtime	(b) Health insurance contribution
Annual complementary salary	<i>Itemised</i>
Holiday bonus	<i>Personal deductions</i>
	(c) Other social security contributions
	(d) Child care deductions
	<i>Housing expenses</i>
	(e) Rent deductions
	(f) Mortgage deductions

*Notes:* In this table we show a detailed list of the components included in the definition of gross labour income and deductions used when computing the IRPF. The left column, sources of income, shows income components that comprise the IRPF tax base while the right column, deductions, shows the type of deductions that are allowed by the tax code. Rent deductions and mortgage deductions are treated differently when computing the final tax liability. Rent deductions are treated as tax credits, while mortgage deductions are treated the same as any other item in the deduction part of the schedule.

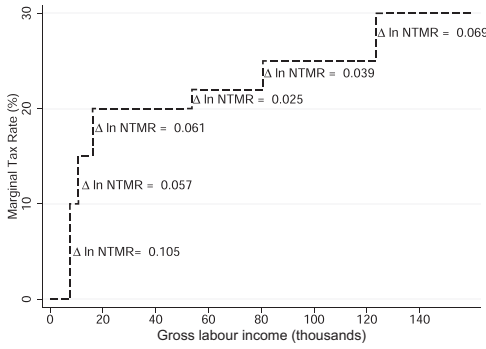
The *deduction part* includes all deductions allowed by the tax code, which are also subject to a progressive deduction rates schedule  $D$ . This feature of the Uruguayan tax system is unusual compared to other modern tax systems. The set of deductions allowed by the Uruguayan tax code is relatively limited and can be divided into two groups. The first group includes all social security contributions that are proportional (at a fixed rate  $c$ ) to the gross labour income (e.g., payroll taxes and mandatory health insurance contributions) and can be written as  $d_n = c \times y$ . These are automatically considered for the purpose of calculating the final tax liability. Throughout this paper, we refer to these as *non-itemised* deductions. The second group of deductions ( $d_i$ ) includes all deductions that are not proportional to gross labour income. We call them *itemised* deductions because they need to be explicitly and voluntarily claimed by the taxpayer. They can be divided into two types: personal deductions and housing expenses (see the right column of Table 1). Personal deductions include deductions for underage children and other non-proportional social security contributions, such as payroll taxes paid by self-employed professionals to their own pension plans. 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 each taxpayer's annual gross labour income  $y$ .<sup>12</sup> The progressive tax rates  $T$  are applied to this amount and a *provisional tax liability* is computed. In Table 2(a) we show the tax schedule  $T$  for year 2013 with values expressed in USD. The IRPF comprises seven different brackets with marginal tax rates that range from 10% to 30%. The first bracket, which includes income up to USD 7,525, is taxed at a marginal rate of 0% and corresponds to the amount of income exempt from taxation. The third column of Table 2(a) shows the change in the net-of-tax rate implied when moving from the  $n$ th bracket to the  $(n + 1)$ th. In Figure 2(a) we depict the marginal tax rates as a function of gross labour income and include notes with the log-change in the marginal net-of-tax rates.

The second step is analogous. Deductions  $d_n$  and  $d_i$  are added up and passed through the deduction schedule  $D$  shown in Table 2(b). The resulting amount is what we call *tax deduction*.

<sup>12</sup> In the case of self-employed activity, a 30% deduction is automatically applied to total earnings to account for production costs and VAT.

(a) Change in the net-of-tax marginal rate



(b) marginal tax and deductions rates

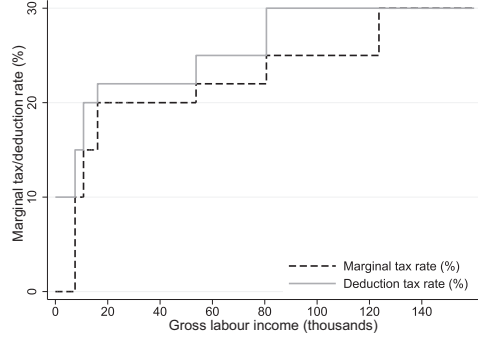


Fig. 2. Marginal Changes in Labour Income Tax and Deductions Schedule (2013).

Notes: In this figure we depict the tax and deduction components of the IRPF schedule. The brackets are plotted for the 2013 USD income brackets. In (a) we show the tax part of the IRPF (see Table 2(a)). The dashed line represents the marginal tax rate (y axis) as a function of gross labour income (x axis). We also include the log-change in the marginal net-of-tax rate, i.e.  $\Delta \ln \text{NTMR}$ , as described in Table 2. In (b) we jointly plot the tax and deduction components. The width of the  $n$ th deduction bracket is exactly the same as that of the  $(n + 1)$ th tax bracket; the same applies to the marginal rate. As in (a), the dashed line depicts the marginal tax rates as a function of the gross labour income while the solid line depicts the marginal deduction rates.

Table 2. Uruguayan Personal Income Tax (IRPF) Schedule (2013).

(a) Tax rates			(b) Deduction rates		
Bracket (2013 USD)	Marginal rate	$\ln\left(\frac{1-\tau_i}{1-\tau_{i+1}}\right)$	Bracket (2013 USD)	Marginal rate	$\ln\left(\frac{1-\tau_i}{1-\tau_{i+1}}\right)$
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: In this table we show the tax schedule for year 2013 with values measured in current USD. In (a), tax rates, the first column depicts the annual income brackets, the second column shows the marginal tax rate applied to each bracket and the third column shows the change in the net-of-tax marginal rate with respect to the previous bracket. This is computed as  $\ln((1 - \tau_i)/(1 - \tau_{i+1}))$ . In (b), deduction rates, the first column depicts the annual deduction brackets, the second column shows the marginal deduction rate applied to each bracket and the third column shows the change in deduction rates with respect to the previous bracket calculated in the same way as with for the tax rates.

In Table 2(b) and Figure 2(b) we show the marginal deduction rates. As with marginal tax rates, marginal deduction rates increase with the total amount of deductions. It is also worth noting that the size of the  $n$ th deduction bracket is exactly the same as the size of the  $(n + 1)$ th tax bracket and deduction rates also range from 10% to 30%.<sup>13</sup>

<sup>13</sup> It is worth noting that, while our research design exploits variation generated by the progressive tax structure  $T$ , the deduction rates schedule  $D$  adds an additional source of discontinuities that our approach does not take into account.



The final step consists of subtracting the tax deduction  $D(d_i + d_n)$  computed in the second step from the *provisional tax liability*  $T(y)$  computed in the first step. These three steps yield the final tax liability  $L$  given in (1).<sup>14</sup>

To illustrate more precisely how tax liabilities are computed, the DGI developed a tax calculator that is available on their website and that any taxpayer can freely use.<sup>15</sup> Although this tax calculator makes the IRPF kink points more salient, accessing this information is costly. Taxpayers must be aware of the calculator and also need to invest some time in downloading the spreadsheet and learning how to use it. Therefore, it is reasonable to suppose that most taxpayers do not actually use this calculator and therefore do not know the exact location of the kink point at which their marginal tax rate changes. Furthermore, since *non-itemised* deductions are tied to gross labour income but *itemised* deductions are subject to a progressive schedule of marginal rates, calculating the exact location of the kink requires some degree of sophistication.

### 1.2. Filing Rules and Withholdings

As is common in many countries, the DGI collects part of the IRPF liabilities using monthly withholdings for dependent workers, i.e., workers who are not self-employed (Slemrod, 2008). Withholdings are estimated for each employee in a firm based exclusively on the income and deductions related to her activity within the firm. In the simplest and typical case, i.e., a dependent worker with only one job who does not claim any deductions other than the automatic proportional ones, the sum of monthly withholdings will exactly match the annual tax liability. However, because of the progressivity of the tax and deduction rates, if individuals have multiple sources of income, or want to claim additional itemised deductions, the *tax* and *deduction parts* calculated within the firm are only a portion of the individual's total tax liability. For this reason, individuals with multiple sources of income that add up to more than the designated annual threshold are required to file a tax return.<sup>16</sup> Similarly, because self-employed workers are not subject to withholding, this group of workers is also required to file a tax return. To have one's monthly withholdings better align with one's expected tax liability, taxpayers can provide additional information to their employer using a 3100 form. However, this is restricted only to deductions for dependent children. If a taxpayer wants to claim housing expenses (i.e., rent or mortgage expenses), she must file an annual tax return.

While these rules dictate who *must* file a return, everyone has the option of filing a return if they wish to. Most commonly, taxpayers choose to file a return if they want to claim *itemised* deductions that were not reported in the 3100 form or if they want to claim deductions for housing expenses. However, there are cases in which taxpayers file a tax return even when they are not required to do so and do not claim additional deductions. Since the IRPF is an annual base tax, these filers could be individuals who worked only a fraction of the year but had taxes withheld as if they worked the whole year. In practice, most workers do not file a tax return—e.g., in 2016, only 16% of all registered labour income earners filed—because most workers receive wages from only one source.

<sup>14</sup> Table B.1 of Online Appendix B shows an example of how the final tax liability is computed for a hypothetical taxpayer.

<sup>15</sup> In Section B.2 of the Online Appendix, we provide a detailed description of the tax calculator and a screenshot of the spreadsheet provided by the DGI.

<sup>16</sup> In 2013, the tax agency defined the threshold as approximately USD 14,000.

Tax returns must be submitted to the tax agency via an electronic form (1102 form). In this form, each taxpayer provides information on her annual earnings from both wage labour and self-employment, as well as *itemised* and *non-itemised* deductions. With the inputs provided by the taxpayer, the 1102 form automatically computes the person’s final tax liability and, after subtracting the third-party withholdings, the tax owed or refund due.

## 2. Empirical Model

### 2.1. Bunching Formula and Estimation

Since Feldstein (1995; 1999), the modern empirical literature in public economics has focused on estimating the compensated ETI as a sufficient statistic for all types of behavioural response to income taxation. The ETI is defined as the percentage change in taxable income ( $z$ ) in response to a one percentage increase in the net-of-tax rate:  $\epsilon(z) = [(1 - \tau)/z]\partial z/[\partial(1 - \tau)]$ , where  $\tau$  is the marginal tax rate.

Consider a progressive tax schedule with two marginal tax rates  $\tau_2 > \tau_1$  and a kink point  $k$  that indicates the point at which the marginal tax rate changes. The theoretical prediction in a model with such a tax schedule, convex preferences and a smooth income distribution is that a number of taxpayers should bunch just before  $k$ . Saez (2010) showed that the compensated ETI can be identified from the excess mass of taxpayers located at the kink point, which is defined as the difference between the observed and a counterfactual income density function at  $k$ . In this setting, the compensated ETI is identified locally such that  $\epsilon(k) = b/[k \ln((1 - \tau_1)/(1 - \tau_2))]$ , where  $(1 - \tau_1)/(1 - \tau_2)$  is the percentage change in the marginal net-of-tax rate at  $k$ , and  $b$  represents the excess mass ( $B$ ) at  $k$  expressed as a share of the counterfactual density function ( $h_0$ ):  $b = B/[h_0(k)]$ .<sup>17</sup>

It is worth noting that all the parameters in this equation are known or can be estimated using the taxable income distribution data. Here  $k$  and  $(1 - \tau_1)/(1 - \tau_2)$  are directly observed from the tax code, while  $b$  is the only parameter we need to estimate. To do this, we follow the strategy proposed by Chetty *et al.* (2011), i.e.,

$$\hat{b} = \frac{\hat{B}}{[(k^+ - k^-)/w]^{-1} \sum_{j=k}^{k^+} \hat{C}_j} \quad \text{with } \hat{B} = \sum_{j=k^-}^{k^+} (C_j - \hat{C}_j),$$

where the estimated excess mass ( $\hat{B}$ ) is the sum of the differences between the observed ( $C_j$ ) and the predicted ( $\hat{C}_j$ ) numbers of taxpayers for each bin  $j \in [k^-, k^+]$  and  $w$  is the bin width chosen to group the data. We obtain  $\hat{C}_j$  by grouping individuals in income bins of size  $w$  and estimating:

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

Here  $q$  is the order of the polynomial and  $\varepsilon_j^0$  is the error term. The function  $1\{Z_j = i\}$  signals whether the income bin  $j$  lies in the bunching zone and controls for being located near the kink. Once we estimate  $\hat{C}_j$ , we compute  $\hat{B}$  and the relative excess mass  $\hat{b}$ . Substituting these values into the ETI expression derived in Saez (2010), we obtain the estimated ETI. Following Chetty *et al.*

<sup>17</sup> In Online Appendix C we present a detailed derivation of the bunching formula.

(2011), we use bootstrapping to calculate standard errors. We determine the optimal polynomial order using the Bayesian information criterion (BIC).

In order to facilitate the interpretation of results, we normalise the taxable income distribution by the kink point. Hence, the kink point is defined as the zeroth bin and all distances will be expressed as a 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%. The information used to estimate the counterfactual polynomial is restricted to observations within 60% of the distance to the kink point.

## 2.2. Distance to the Kink Point

Our empirical analysis will focus on the first kink, because it represents the largest and most salient change in the net-of-tax rate of the Uruguayan IRPF schedule. However, for completeness, in Online Appendix E we report the analysis at all other kink points of the tax schedule.

As mentioned above, the effective location of the kink points is different for each individual because it depends not only on one's gross labour income, but also on the total amount of 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 points. Next, we summarise the main steps used to calculate the percentage distance to the kink points. The process is explained in detail in Online Appendix D.

The first step consists of combining both the tax and deduction components into a single new mathematically equivalent structure. The new income brackets will be called *adjusted* brackets. The second step corrects the distance between the *adjusted* brackets and gross labour income by a factor  $1/(1-d)$ , where  $d$  is the proportional deduction rate associated with the *non-itemised deductions*. This is because the distance to a kink point can be thought as the additional income that a taxpayer can earn (or forfeit) without changing her marginal tax rate. Since some deductions are proportional to one's gross labour income, receiving (forfeiting) an additional dollar increases (decreases) both the taxable income and the amount of deductions. In this case, measuring distances to the kink point simply as the difference between gross labour income and the *adjusted* brackets would underestimate the actual margin by which a taxpayer can change his income without changing his marginal rate. The rate  $d$  includes different deduction components, such as social security contributions, health insurance contributions and other deductible items associated with one's gross labour income. In our data, we can only observe the actual  $d$ , i.e., the current ratio of proportional deductions to taxable income. Therefore, assuming that this ratio does not change with an additional unit of income, we can use the observed  $\hat{d}$  to correct the distance.

After adjusting the brackets and correcting the distances by the  $\hat{d}$  factor we can compute the actual distance to the kink point as the difference between gross labour income and the corrected *adjusted* brackets.<sup>18</sup>

<sup>18</sup> In some cases, when individuals are to the right of a kink but deductions are large enough to fully offset the income brackets to the left, we code the distance to the kink as missing. By doing this, we avoid the inclusion of such individuals in the analysis of kinks that are irrelevant in their decisions. These individuals are not completely excluded of the analysis, but are included in the analysis of the subsequent kinks.

### 2.3. Bunching Window Selection

Estimating  $b$  requires defining a bunching zone  $[k^-, k^+]$ . The usual practice in the empirical bunching literature is to select this zone by visual inspection (Kleven *et al.*, 2016). This introduces a large degree of discretion for the researcher that may bias the estimates.<sup>19</sup> To minimise arbitrariness and potential biases, we follow the data-driven approach proposed by Bosch *et al.* (2020). To our knowledge, this is the first paper, other than the paper that proposes this approach, to implement this method.

Figure C.1 of Online Appendix C, taken from Bosch *et al.* (2020), illustrates the basic intuition for this method. This figure plots a binned taxable income distribution for a given excluded region, in this example  $[-10, 10]$ . The solid line represents a fitted linear regression that was estimated using all the points outside the excluded zone. The bunching window is then defined as the zone in which contiguous bins fall outside the confidence interval for the predicted model, in this case  $[-5, 2]$ . To avoid arbitrariness in the selection of the excluded region, the method iterates through different limits for the excluded region and uses the mode of all estimates to define the “optimal” bunching zone.

The algorithm proposed by Bosch *et al.* (2020) can be summarised as follows.<sup>20</sup>

- (1) Start with the excluded region  $(x^-, x^+) = [0, 0]$ .
- (2) Run a local regression using all information outside  $(x^-, x^+)$  and obtain the fitted model.
- (3) Estimate a confidence interval around that prediction.
- (4) The bunching zone for this excluded area  $(x^-, x^+) = [0, 0]$  will be defined as the set of contiguous bins around the kink point that lie outside the confidence interval.
- (5) Repeat steps (2) to (4) for variations of the excluded regions  $(x^-, x^+)$  with  $x^- \in \{-X, (-X + 1), \dots, 0\}$  and  $x^+ \in \{0, 1, \dots, X\}$ , where  $X$  denotes an arbitrary but large enough limit for the excluded region. This generates a distribution of lower ( $k^-$ ) and upper ( $k^+$ ) bounds for the bunching window.
- (6) After obtaining a distribution of lower and upper bounds for the bunching zone from step (5), select the limits of the bunching zone from the modes of the distributions for  $k^-$  and  $k^+$ . This will be the bunching zone used for estimation.

Four points regarding the above procedure are worth mentioning. First, the iterative procedure to construct a distribution for upper and lower bounds of the excluded region aims to minimise concerns about possible arbitrariness in the selection of the excluded region. The number of iterations to obtain that distribution depends on the value  $X$ . In our case, we set  $X = 10$ , which implies that steps (2) to (4) are repeated 121 times for each estimate. Second, the confidence interval in step (2) is an additional parameter whose value is left to the researcher’s discretion. We use a typical 95% confidence interval level. Third, step (6) suggests using the modes of the distributions for  $(k^-, k^+)$ . While other statistics could be used, e.g., mean or median, Bosch *et al.* (2020) reported no major differences between these alternatives. Fourth, unlike the standard practice in the bunching literature, the method does not restrict the bunching window to be

<sup>19</sup> If we select a too small (high) bunching window, we would underestimate (overestimate) the excess mass and thereby underestimate (overestimate) the implied elasticity.

<sup>20</sup> See Appendix 1 of Bosch *et al.* (2020) for a formal derivation of the bunching window procedure.

symmetric around the kink. Indeed, most of the bunching windows we obtain by using this procedure are asymmetric.<sup>21</sup>

### 3. Data

#### 3.1. Administrative Data and Sample Restrictions

Our study relies on large, high-quality administrative records provided by the Uruguayan Tax Agency. These datasets contain information about all wage earners and self-employed workers in the 2010–2016 period. We are able to link each dataset using (masked) national identification numbers. The resulting dataset is a combination of three data sources linked at the individual level: (a) personal income tax returns (1102 form); (b) third-party reports of labour income and deductions from employer statements (equivalent to the W-2 form in the United States); and (c) firm-level data. While third-party reports and firm-level data are available for the whole universe of workers, we observe tax returns only for those individuals who actually filed one (e.g., 16% in 2016). If a taxpayer filed a tax return, we use the 1102 form as our primary data source; for those who did not, we use the third-party reports (W-2 equivalent). It is worth noting that, for taxpayers who filed a tax return, we have both self- and third-party-reported information. We exploit these alternative sources when we analyse in depth the mechanisms of response. The key variables included in the data are: gross labour income, *itemised* and *non-itemised* deductions and tax withholdings. We also have a limited number of socio-demographic variables, including age and gender of the taxpayer. Since we have the identification number of all firms in which a taxpayer is employed, we can link individual data with firm-level data containing sales during the fiscal year and the industry activity code.

In the original tax data, the total number of taxpayer-year observations with positive labour income during the 2010–2016 period is 8,714,279. However, to perform the analysis, we introduce some restrictions to the data. First, to focus on the active population in the labour market, we include only taxpayers aged between twenty-one and sixty years old. Second, because we focus on labour income responses, we exclude taxpayers for whom labour income is not their primary source of income. For instance, we exclude taxpayers with retirement income that is more than 40% of their total income and also those taxpayers who paid at least USD 1 of retirement income tax. We also exclude taxpayers who have capital income as their main source of income. All results are robust to the inclusion of these sub-groups. After restricting the sample, we end with 7,789,808 observations, corresponding to 1,623,468 different taxpayers. This implies that a single taxpayer appears in our datasets, on average, 4.8 times out of a potential maximum of seven appearances.

In addition to the administrative tax data, we use an auxiliary database that contains all minimum wages negotiated in the collective bargaining process in Uruguay for most of the analysis period (2010–2014) (Ceni *et al.*, 2020). This dataset consists of the minimum wages set

<sup>21</sup> If the optimal bunching zone estimated using this method results in a lower or upper bound larger than 25% of the kink (in absolute value), we set the corresponding bound for the bunching zone to 1 (in absolute value). Two reasons motivate this decision. First, even when considering that adjustment frictions may play an important role, it is hard to argue that someone more than 25% away from the kink should be considered as a *buncher*. Second, modes for  $k^-$  and  $k^+$  distributions that are very far away from  $k$  are most likely due to relatively uniform distributions, where the estimated mode has a very strong random component. In turn, this suggests that there is no clear concentration of bins outside the confidence interval for the fitted model in a neighbourhood relatively close to  $k$ . We interpret this to indicate a lack of bunching close to  $k$ . Empirically, this does not affect any of our results and nothing changes by defining alternative bunching zones for these very specific cases.

for ninety-two sector groups that are currently defined by the law.<sup>22</sup> Researchers at the Instituto de Economía (*Universidad de la República*, Uruguay) generated this database by scrapping the documents associated with the agreed union contracts from records by the Ministry of Labour, and made it available for use in our study. We employ this data to complement the analysis from the bunching estimates in Subsection 5.2.

### 3.2. Descriptive Statistics

In Table 3 we provide summary statistics for the final sample used in the analysis. The average annual income is USD 10,440, which can be split into USD 10,000 (95.8%) coming from labour income, USD 240 (2.3%) from capital income and USD 210 (2.01%) from retirement income. On average, 94.3% of labour income comes from wages, while only USD 570 comes from self-employed activity. This shows that the labour market structure in Uruguay is based predominantly on wage earners. In turn, only 3.4% of taxpayers receive earnings from self-employment and only 2% receive self-employment income exclusively. As to the IRPF burden, the average tax liability for the whole sample is USD 610, which represents 6.1% of labour income. This low burden is mostly attributable to the fact that only 33% of the total number of observations in the sample correspond to taxpayer years with positive IRPF payments. Restricting the sample to those who actually paid at least USD 1 of IRPF, the average payment is three times larger (USD 1,830). Deductions allowed by the tax code represent on average about 21.2% of gross labour income. As expected, *non-itemised* deductions—i.e., deductions that are proportional to the wage earned and do not need to be explicitly claimed—represent a larger share of total deductions than do to the *itemised* deductions (80% versus 20%). The last row in panel B shows that fewer than 12% of the observations in the sample correspond to workers who filed tax returns. In terms of demographic characteristics, 47% of the observations in our sample are women and most of the observations are individuals aged above forty years old. Roughly 53% are individuals engaged in the services sector and most of those who are employees work in firms with total sales above the median, i.e., relatively large firms.

The general picture depicted for the whole sample is almost exclusively explained by the group of pure wage earners (96.3% of the sample). We define a pure wage earner (purely self-employed) as a taxpayer who did not receive any income from self-employment (salary/wages) within a given year. There are some differences between pure wage earners (column (3)) and purely self-employed workers (column (5)). The income of purely self-employed workers doubles that of pure wage earners (USD 18,860 versus 9,340) and it is much more variable. Although the differences in the total amount of deductions do not seem to be as large as the differences in total income, the share of deductions relative to total gross labour income is higher for the sample of pure wage earners (21.9% versus 9.7%). As a result, the average tax liability for purely self-employed workers (USD 1,300) is more than twice that observed for pure wage earners (USD 530). The share of self-employed workers who file a tax return is more than seven times greater than the share of pure wage earners who do so (61.4% for purely self-employed workers and 78.6% for those who are both wage earners and self-employed). This is explained by the DGI requirement that any individual with positive self-employment income must file a tax return. The

<sup>22</sup> These groups do not include rural, public sector and domestic work, whose minimum wages are regulated by a different process.

Table 3. *Summary Statistics—Final Sample (2010–2016).*

	All workers		Pure wage earners		Pure 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>								
<b>Total income (USD in 1,000s)</b>	10.44	22.56	9.74	20.30	20.62	62.63	34.61	40.01
Total labour Income (USD in 1,000s)	10.00	17.23	9.34	14.48	18.86	58.54	33.25	35.26
Wage income (USD in 1,000s)	9.43	14.87	9.34	14.48	0.00	0.00	20.86	26.82
SE income (USD in 1,000s)	0.57	8.66	0.00	0.00	18.86	58.54	12.44	22.08
Retirement income (USD in 1,000s)	0.21	1.68	0.20	1.64	0.41	2.64	0.34	2.36
Capital income (USD in 1,000s)	0.24	13.75	0.20	13.58	1.34	20.06	1.02	15.51
<i>Panel B: personal income tax</i>								
<b>Tax liability (USD in 1,000s)</b>	0.61	3.56	0.53	3.11	1.30	11.45	3.55	7.02
<b>Tax liability &gt; 0 (%)</b>	33.19	47.09	32.02	46.65	43.43	49.57	79.13	40.64
<b>Tax liability (among tax liability &gt; 0)</b>	1.83	6.00	1.66	5.32	2.99	17.23	4.48	7.63
<b>Total deductions (USD in 1,000s)</b>	2.12	3.14	2.04	2.98	1.79	5.02	6.15	5.05
Non-itemised (USD in 1,000s)	1.71	2.29	1.67	2.22	0.80	1.41	4.25	3.94
Itemised (USD in 1,000s)	0.41	1.76	0.36	1.62	1.00	4.73	1.90	2.54
<b>Filling a tax return (%)</b>	11.17	31.50	8.84	28.39	61.36	48.69	78.57	41.03
<i>Panel C: socio-demographics</i>								
<b>Female (%)</b>	46.56	49.88	46.38	49.87	45.93	49.83	55.20	49.73
<b>Age</b>	39.69	62.22	39.06	58.52	74.14	182.08	41.66	27.05
Under forty years	41.64	49.30	41.12	49.20	57.94	49.36	52.93	49.91
Over forty years	58.36	49.30	58.88	49.20	42.06	49.36	47.07	49.91
<b>Source of income (%)</b>								
One income source	–	–	75.03	43.28	100.00	0.00	–	–
Multiple income sources	–	–	24.97	43.28	0.00	0.00	–	–
<b>Firm's activity (%)</b>								
Services	52.78	49.92	52.07	49.96	–	–	86.19	34.50
Goods	47.22	49.92	47.93	49.96	–	–	13.81	34.50
<b>Firm size (USD total sales) (%)</b>								
Below median sales	20.04	40.03	20.15	40.11	–	–	14.04	34.74
Above median sales	79.96	40.03	79.85	40.11	–	–	85.96	34.74
Observations	7,789,808		7,498,913		127,192		163,702	
Individuals	1,623,468		1,584,939		47,057		55,428	

*Notes:* In this table we report summary statistics of the final sample considered in the analysis for the period 2010–2016. In panel A we report statistics regarding the various income components, in panel B we report statistics regarding personal income tax items and in panel C we present statistics regarding taxpayers' socio-demographic characteristics as well as characteristics of the firm that represents the main income source. All statistics are computed based on taxpayer years, i.e., the number of times that a taxpayer is included in the calculation of each variable mean is the number of years he earned positive labour income. If a taxpayer does not have positive labour income in a given year, the individual is excluded from the sample in that particular year. Odd-numbered columns list the variable means while even-numbered columns list the SD. 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) show summary statistics for the purely self-employed (SE) group and columns (7) and (8) show them for the group of workers who receive income both from wages and self-employment activity. All monetary values are expressed in 2013 USD. The number of observations and unique individuals is based on the total labour income variable.

fact that about 35% of the observations of self-employed individuals with positive income are not associated with a tax return, even though filing a return is mandatory in such cases, indicates the limited enforcement capacity of the DGI. In terms of demographic characteristics, pure wage earners are less likely to be female, are younger, and are less likely to work in the service sector than are their purely self-employed counterparts.

## 4. Main Results

In this section we report the main estimates for the excess mass and the implied ETI pooling data from years 2010 to 2016. In all figures, the vertical solid line at 0 represents the normalised individual effective kink point. Each data point on the connected curve represents the number of taxpayers in a given bin. The data are grouped into 2%-wide bins: the first bin corresponds to taxpayers whose taxable income lies between 0% and 2% higher than the threshold, the second bin includes taxpayers whose taxable income lies between 2% and 4% higher than the threshold, and so on. The solid and smooth curve represents the counterfactual taxable income distribution estimated as explained in Section 2. The boundaries of the bunching zone, defined as explained in Subsection 2.3, are indicated by dashed vertical lines. The estimated excess mass ( $b$ ), bootstrapped standard error, estimated ETI and optimal polynomial degree according to the BIC criterion are indicated in each figure.

### 4.1. Bunching Response Around the (Bottom) Kink

In Figure 3(a) we show the observed and counterfactual taxable labour income distributions around the first IRPF kink point for the full sample. The estimated excess mass ( $\hat{b}$ ) is 29.1%, which is highly statistically significant. This  $\hat{b}$  implies a modest ETI of 0.055. While the theoretical prediction is that taxpayers should sharply bunch at the kink, the excess mass observed in panel (a) is hump shaped and somewhat diffuse. There are several possible explanations for this finding. For example, information frictions can make it difficult for taxpayers—and, in particular, wage earners—to accurately determine the tax income schedule (Chetty *et al.*, 2011; Kleven and Waseem, 2013; Best, 2015). This possibility is consistent with our discussion in Subsection 2.2, where we showed that calculating the exact location of the kink point is not straightforward. In addition, even if an individual was able to determine the exact location of her individual kink point, she may encounter additional barriers, such as rigid constraints on hours worked (Rosen, 1976; Altonji and Paxson, 1988) or searching costs (Rogerson *et al.*, 2005) that prevent individuals from adjusting their income as they wish. Even in contexts in which individuals have considerable flexibility due to the presence of a sizeable informal sector (Zenou, 2008), these costs may play a major role in preventing workers from manipulating their income in a precise manner. Finally, the lack of sharp bunching may be attributable, in part, to the fact that Uruguay's tax code only allows a reduced set of deduction possibilities (Kopczuk, 2005).

Panels (b), (c) and (d) of Figure 3 replicate the analysis reported in panel (a) for three different sub-samples of taxpayers: pure wage earners, purely self-employed workers and taxpayers who have both wage income and self-employment, respectively. In panel (b) we show that the behaviour of pure wage earners is almost identical to that in panel (a) for the whole sample: the excess mass with respect to the counterfactual density is about 28.7%, and the implied ETI is also 0.055, with a diffuse concentration of taxpayers around the kink point. This is explained by the heavy predominance of wage earners in Uruguay's labour market structure (96.2% of the total sample). Although both figures show that these bunching responses are statistically different from zero, an ETI of 0.055 is in the range of the moderate-to-low bunching responses found in previous studies. Adjustment frictions may be less relevant for self-employed workers. The latter have more scope for tax planning or tax avoidance and they are expected to show a sharper bunching pattern (le Maire and Schjerning, 2013; Bastani and Selin, 2014). In panel (c) we present the same estimates as panels (a) and (b) but for purely self-employed workers.



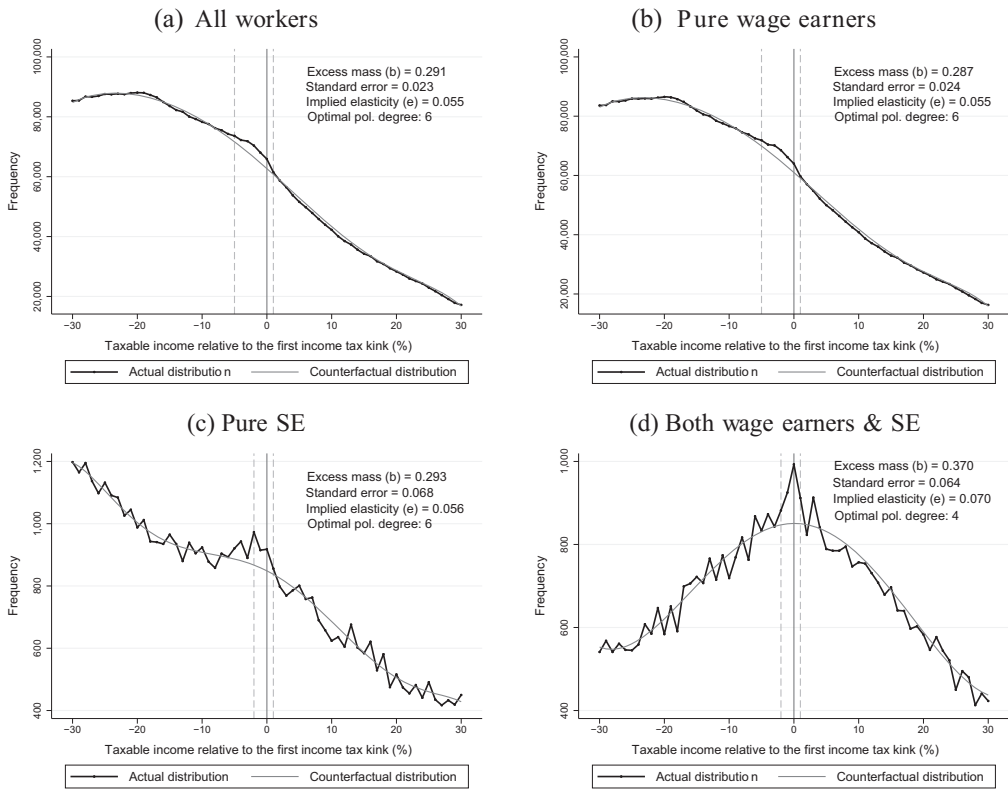


Fig. 3. Taxable Labour Income Bunching by Source of Income, 2010–2016.

*Notes:* In panels (a) through (d) we plot the empirical and counterfactual taxable income distributions normalised by the individual's first kink point, as explained in detail in Section 2 (see Online Appendix D for a numerical example). All estimates are performed by using the pooled data for the 2010–2016 period. In panel (a) we show the taxable income distribution for the main sample considering all taxpayers. For panel (b), we repeated the analysis for pure wage earners, i.e., those workers who do not receive self-employment income. In panel (c) we show the distribution for taxpayers who only earn income from self-employment activity (purely self-employed) while in panel (d) we present the results for workers who receive both types of income. In all cases, the taxable income distribution is expressed as the percentage distance to the first kink point. The vertical solid line indicates the exact point at which each taxpayer starts to pay income tax, i.e., the first kink point. The data are grouped into 2%-wide bins, which are represented by the marked data points (dots) in the connected line. This means that the first bin to the left of the kink point is 2% away from it, while the second is 4%, and so on. We restrict the analysis sample to individuals located up to thirty bins (60%) above or below the kink. The solid line connecting dots represents the observed taxable income distribution while the counterfactual estimation is depicted by the thinner and lighter 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 detailed in Subsection 2.1. We use the BIC criterion to determine the optimal degree of the polynomial, and the Bosch *et al.* (2020) procedure to determine the bunching windows (vertical dashed lines), as described in Subsection 2.3. In each plot, we provide the excess mass estimates relative to the counterfactual distribution at the kink ( $b$ ), its standard error, the corresponding implied elasticity, and the polynomial degree. Standard errors are estimated by bootstrapping using an  $N$  of 1,000.

Purely self-employed workers are defined as workers whose only source of income is independent activity. The excess mass around the first kink for this group is very similar to that observed for all workers and for pure wage earners (29.3% and statistically significant). This excess mass implies an ETI of 0.056. We further investigate differences between pure wage earners and self-employed workers in Subsection 4.4. Finally, in Figure 3(d) we show that, for taxpayers who earn both wage and self-employment income, the excess mass is 37% and the implied ETI is 0.070. Although the ETI for this group seems to be slightly larger, and the bunching pattern sharper, the magnitude of the response is still in the range of moderate to low.

Since pure wage earners represent more than 95% of the total number of workers, restricting our analysis to this group does not affect the generality of our conclusions. In addition, restricting the analysis to pure wage earners makes it easier to distinguish between the different adjustment channels that could explain the observed responses to taxation.<sup>23</sup> Therefore, in what follows, we focus exclusively on the sub-sample of taxpayers who are pure wage earners.

#### 4.2. *Bunching Patterns Over Time: Learning and Persistence*

The modest ETI reported in the previous section may reflect optimisation frictions, which may prevent taxpayers from manipulating their income precisely (Chetty *et al.*, 2011). In particular, taxpayers' lack of knowledge about the tax code could be important in the context of the IRPF.<sup>24</sup> It is worth recalling that the IRPF was first implemented in January of 2008 and, thus, in 2016—the last year of our data—progressive income taxation in Uruguay was only eight years old.<sup>25</sup> In this context, the period analysed can be considered a period of learning during which taxpayers familiarised themselves with the new tax structure. To explore this possibility, we present two pieces of evidence.

First, we estimate the total number of taxpayers located in the bunch zone and the share of the total labour force that they represent for each year in our data. If workers need time to learn how to locate the area just before the kink, the number of taxpayers located in the bunch zone should increase over time as more taxpayers become aware of how and where to bunch. While this is a very raw test of the learning hypothesis, it is hard to imagine alternative explanations for an increase in the number of taxpayers located just at the kink over time. In Figure 4(a) we depict the evolution of the total number of *bunchers* (solid connected line) for each year of available data (right-hand-side axis).<sup>26</sup> In addition, the dashed connected line in this figure plots the number of *bunchers* as a percentage of the total number of taxpayers (left-hand-side axis). We do this to control for an increase in the number of *bunchers* attributable simply to an increase in the overall number of registered workers. The total number of *bunchers* from 2010–2016 increased by 50%. This implies an annual growth rate of about 7%, with a starting point of roughly fifty thousand *bunchers*. On the other hand, the share of *bunchers* relative to the total number of taxpayers also

<sup>23</sup> Self-employed workers have more access to extensive tax planning than wage earners do. This could be done, for instance, via retained earnings, income shifting or by deducting 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, it is harder to isolate the role of the different mechanisms contributing to the observed responses in the case of self-employed taxpayers.

<sup>24</sup> See, for instance, Chetty and Saez (2013) and Chetty *et al.* (2013) for a discussion and evidence regarding the role of insufficient knowledge of the tax system on taxpayers' behaviour.

<sup>25</sup> Before the IRPF, income taxation in Uruguay consisted of a flat tax rate on the personal labour income of workers.

<sup>26</sup> A taxpayer is defined as a "*buncher*" according to the bunching zone used in Figure 3(b).

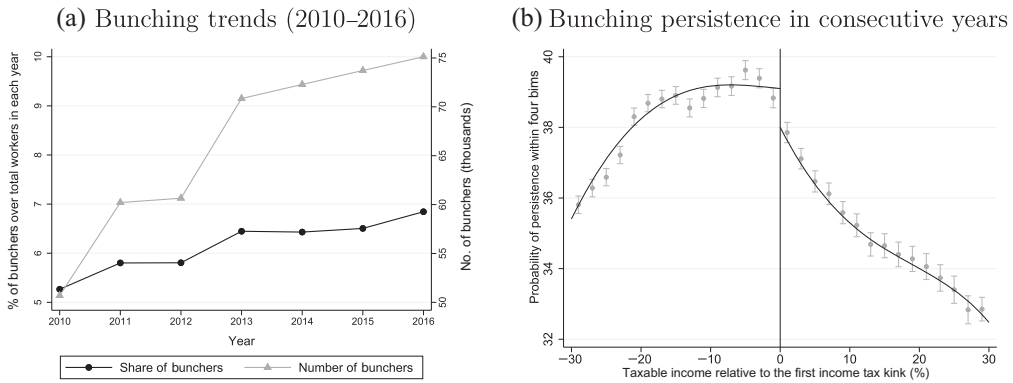


Fig. 4. *Evolution and Persistence of Bunching—Pure Wage Earners, 2010–2016.*

*Notes:* The sample used in this figure is restricted to pure wage earners. In panel (a) we show the evolution of the number of bunchers between 2010 and 2016 as well as the share of bunchers in the whole sample. Bunchers are defined as taxpayers located in the bunching zone depicted in Figure 3, i.e., individuals located between bins  $-5$  and  $1$ . The line connected by triangles depicts the evolution of the absolute number of bunchers (right-hand side  $y$  axis). Because the number of registered workers increased during the period, the line connected by circles shows the evolution of the share of bunchers, i.e., the evolution of the number of bunchers as a percentage of the total number of registered workers (left-hand side  $y$  axis). In panel (b) we show the ratio of taxpayer persistence as a function of the distance to the first kink point (the  $x$  axis is exactly the same as that used in Figure 3). In this case, a worker is defined as persistent if her actual income remains within 8% (or four bins) of her previous year's income.

grew over the period considered, although at a smaller rate, from 5.1% to 6.9%. An increase in the number of bunchers, both in absolute and relative terms, is consistent with the learning hypothesis and suggests that individuals incorporate information over time about their optimal income-reporting strategy.

As a second test, we measure differences in the rate of income persistence along the taxable income distribution. Specifically, we estimate the share of taxpayers that remained within a four-bin window as a function of their distance to the kink in the previous year. Since there were no major tax changes during the period analysed, if taxpayers learn from their own actions, someone who accurately bunched in a given period should be more likely to repeat her own strategy in subsequent periods. The rate of persistence just before the kink should be substantially larger than the rate for taxpayers who are still learning how to bunch. In Figure 4(b) we present the rate of persistence within the defined window. The  $x$  axis is the exact same as that used in Figure 3. On the  $y$  axis we plot the proportion of taxpayers located in an 8% window (or four bins) in consecutive years with corresponding 95% confidence intervals. We also add a polynomial fit estimated separately for observations on both sides of the kink. The rate of persistence is increasing as we move closer to the kink and peaks at the  $-4$  bin where it is 38.6%. Just at the kink we observe a discontinuous change in the rate of persistence, which decreases to 37.4% in the first bin to the right of the kink. This  $-1.2\%$  change is statistically significant at the 1% level. It is important to note that confidence intervals for consecutive kinks overlap for almost all points of the distribution, except from 0 to 1 and  $-20$  to  $-21$ . Consistent with the evidence reported for the evolution in the number of bunchers, a disproportionately high persistence rate just to the left of the kink is consistent with the learning hypothesis.

### 4.3. *Heterogeneous Responses*

In this section, we explore differences in bunching responses for sub-groups of individuals based on socio-demographic attributes and characteristics of the firms where they are employed. Specifically, we replicate the analysis in Subsection 4.1 for sub-samples defined by gender, age (below/above forty years old), number of jobs (one/multiple jobs), size of the firm where the taxpayer is employed (below/above median sales volume) and industry sector (services/goods). We also replicate our baseline model used to estimate the ETI at the first kink point at all other kinks in the tax schedule. All estimates related to the analysis by individual characteristics are reported in Table E.1 of Online Appendix E. Estimates for the additional kink points are reported in Figure E.3.

The graphical visualisation, analogous to our baseline graphics in Figure 3, is presented in Figures E.1 and E.2. The left-hand panel of Table E.1 presents behavioural responses according to individual characteristics of the taxpayers. Regarding gender differences, the estimated excess mass at the first kink is larger for female taxpayers (0.372) than for male taxpayers (0.191). The larger bunching response among women is consistent with previous evidence that showed that this group—mostly married women—are more responsive to tax rate changes than men (see, e.g., Chetty *et al.*, 2011; Paetzold, 2019). The reason could be related to women having more margin to adjust hours/income because they are more likely to be employed in part-time jobs or being the secondary earner in the household. Unfortunately, we cannot test these hypotheses directly because information such as marital status, hours worked and type of jobs (part versus full time) is not available in our dataset. Another possibility is that women are more likely to use deductions to bunch at the kink. In Subsection 5.1, we further explore differences between men and women by analysing deduction behaviour.

In terms of taxpayer age, we might expect smaller responses to the first kink for younger workers as compared to older workers. For instance, career concerns may induce younger workers to adjust hours worked to a lesser extent than their older peers. In addition, older individuals may have more experience in the labour market, and thus have a more comprehensive understanding of how the tax system affects their net-of-tax earnings. Contrary to these speculations, we do not observe any age-related differences in the ETI. As mentioned in Subsection 4.1, adjustment costs can strongly affect taxpayers' behavioural responses. In particular, these responses may be different depending on the type of labour market participation. For instance, the coordination costs of manipulating taxable income may be smaller for workers with only one job compared to workers holding multiple jobs because of the additional costs of managing multiple sources of income.<sup>27</sup> To test this hypothesis, we separately analyse the bunching behaviour for wage earners with only one job and wage earners with multiple jobs. We find the excess mass at the kink to be about 70% higher for the group of wage earners with only one job (30.7% versus 18.1%). Furthermore, the observed bunching among wage earners with only one job entails an ETI of 0.058, which is similar to the elasticity estimated for the whole sample of pure wage earners. Panels e and f of Figure E.1 also suggest differences in the bunching pattern of these two groups.

Evasion opportunities that taxpayers experience in their work environment might also affect the magnitude and patterns of response. The limited number of variables included in our data prevents us from observing evasion opportunities directly. For this reason, and guided by correlations found

<sup>27</sup> Multiple jobholders may find it difficult to engage in income planning in general and not due to coordination costs or ease of collusion. This, however, would be hard to reconcile with observed differences in behavioural responses between the two groups just at the first kink.

in previous studies, we use various firm-level characteristics as proxies for evasion opportunities, such as industry (e.g., Best, 2015) and firm size (Kleven *et al.*, 2011).

We first show the results for individuals who work in service sector firms and those who work in goods-producing firms. The leading hypothesis is that firms in the goods sector are more subject to enforcement by the DGI compared to firms in the services sector because the former's transactions are more likely to be subject to some sort of third-party reporting. Our results suggest that wage earners in the services sector are slightly more responsive near the kink than are their counterparts working in firms in the goods sector (see panels a and b of Figure E.2). However, the observed differences between the two groups, both in terms of the excess mass and implied elasticity, appear to be small (0.277 versus 0.209). Second, we split the sample according to the size of the firm in which workers are employed. We use the volume of sales in the fiscal year and divide the sample in two groups corresponding to those with sales volume above and below the median, respectively. Here, we expect smaller firms to experience less monitoring by the tax agency and therefore workers in those firms may have more opportunities for evasion. In line with this hypothesis, the observed bunching response seems to be stronger for employees of smaller firms relative to those at larger firms—an excess mass of 0.309 versus 0.192, which translates to an implied elasticity of 0.059 versus 0.037 (see panels c and d of Figure E.2). This evidence is consistent with previous research showing that earnings responses—driven mostly by evasion behaviour—are negatively correlated with firm size (Best, 2015; Kleven *et al.*, 2016). However, estimates for small firms are based on a much smaller sample and are less precise, so we cannot rule out that both groups have similar responses.

Finally, for completeness, we calculate our main estimates for the whole sample at all the other kinks in the personal income tax schedule. These results are reported in Figure E.3 of Online Appendix E. Overall, we do not observe any sign of behavioural response at any of these other kinks. This could be attributable to, first, the fact that the economic incentives are much smaller at upper kinks compared to the change in the marginal tax rate at the first kink point and, second, that these upper kinks are much less salient as compared to the first kink.

#### 4.4. Robustness Checks

In this section we report different specifications for the models estimated in the previous sections to analyse how sensitive our results are to the choice of the bunching zone and polynomial degree. In particular, departing from the parameters defined using the Bosch *et al.* (2020) approach, we calculate the estimates while introducing small variations of the bunching zone as well as with different polynomial degrees. Finally, we also report the main results using a balanced sample of taxpayers to show that our results are not driven by changes in the composition of the labour force during the years studied.

First, we calculate the results reported in Subsection 4.1 for the four groups reported in Figure 3: (1) all workers, (2) pure wage earners, (3) purely self-employed workers and (4) workers who are both self-employed and wage earners, using alternative definitions for the bunching zone and polynomial degrees. It is important to recall that, for the whole sample, our baseline results are based on a bunching zone defined by  $[-10\%, 2\%]$  and a sixth polynomial degree. Table F.1 of Online Appendix F shows nine different estimates using different combinations of bunching zones and polynomial degrees. Specifically, we replicate our results for all combinations of the bunching zones  $[-12\%, 4\%]$ ,  $[-10\%, 2\%]$  and  $[-8\%, 2\%]$  and fifth, sixth and seventh degree polynomials. The ETI estimated in our preferred specification (0.055) lies in the middle of the

distribution of the nine estimated ETIs. The average ETI for this set of specifications is 0.057, while the range is [0.047, 0.072]. Table F.2 focuses on pure wage earners and, as expected, yields essentially the same results. Table F.3 reports the robustness of our baseline estimates for the group of purely self-employed workers. Our preferred specification for this group uses a bunching zone of  $[-4\%, 2\%]$  and a sixth degree polynomial. We report results for different combinations of the bunching zones  $[-6\%, 4\%]$ ,  $[-4\%, 2\%]$  and  $[-2\%, 2\%]$  and fifth, sixth and seventh degree polynomials. Again, our preferred specification lies approximately in the middle of the distribution. The average ETI for the nine specifications is 0.044 and estimates range from 0.027 to 0.073. Finally, Table F.4 presents the results for both wage earners and self-employed workers. The baseline specification for this group of taxpayers consists of a bunching zone of  $[-4\%, 2\%]$  and a fourth degree polynomial. We estimate nine specifications iterating over the bunching zones  $[-6\%, 4\%]$ ,  $[-4\%, 2\%]$  and  $[-2\%, 2\%]$  and third, fourth and fifth degree polynomials. As with previous groups, our preferred estimate lies in the middle of the nine estimated ETIs. The average estimate is in this case 0.078 and the range is [0.062, 0.116]. Overall, these results suggest that the data-driven approach used to select both the bunching zone and the polynomial degree do not systematically over- or underestimate the excess mass, which alleviates the concerns about possible bias in our estimates.

Another potential concern regarding 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 taxpayers to and from the sample are biased towards specific worker characteristics that are also related to their bunching behaviour. We test this by repeating the analysis of our main results in a balanced panel, i.e., including taxpayers who were in the sample during the seven years considered (see Figure F.1 of Online Appendix F). Although the results remain qualitatively unchanged when using a balanced sample, there are some differences in the magnitude of the response when compared to the estimates using the unbalanced panel. In particular, the excess of mass and the implied estimated ETI seems to be larger when using workers who remained in the formal labour market during seven consecutive years. In particular, we observe that the ETI for the overall sample is now 0.065, compared to 0.055 in the unbalanced sample. Splitting the sample by the different sub-groups, the estimates are generally larger, being 0.064 (compared to 0.055) for pure wage earners, 0.124 (compared to 0.056) for purely self-employed workers and 0.062 (compared to 0.070) for those who are both wage earners and self-employed workers. These differences accord with our previous results. Moreover, they are consistent with the learning hypothesis presented in Subsection 4.2 because they suggest that taxpayers with a more stable connection to the labour market, who are exposed year after year to the tax schedule, are more responsive than are those who fluctuate between formal employment and informal employment or unemployment.

One interesting result is that, when using the balanced sample, the ETI for the purely self-employed workers almost doubles that observed for the pure wage earners, which is more in line with the existing literature. However, the response of self-employed individuals may still be considered surprisingly modest. One possibility is that our focus on individuals with relatively low earnings is masking heterogeneous responses amongst self-employed individuals. Self-employed individuals near the lowest kink may differ from the typical self-employed individual in the upper part of the income distribution. For example, employers may rely on false self-employment contracts to hire workers under more flexible conditions and without paying payroll taxes. Individuals in this situation usually have more precarious working conditions and are

probably less sophisticated compared to more affluent individuals who choose self-employment for convenience.

## 5. Explaining Bunching Responses

In this section we analyse the potential mechanisms behind the behavioural responses documented in previous sections. We do this by studying the behavioural responses of pure wage earners. As argued in Subsection 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. Taxable labour income is a combination of gross labour income and deductions and taxpayers might respond in either of these two components (Hamilton, 2018). We first explore the role of deductions in accounting for the observed behavioural responses to the IRPF. Next, we analyse changes in gross labour income and, in particular, we focus on the role of tax evasion (*underreporting*). The evidence presented in this section is not intended to be interpreted as causal as we do not manipulate any source of exogenous variation in the tax schedule. The IRPF has remained virtually unchanged since its implementation. Instead, we exploit the granularity of the data and some very specific features of the institutional setting to characterize taxpayers' behavioural responses.

### 5.1. Deduction Behaviour

Traditionally, empirical research regarding the elasticity of taxable income has focused on the income side of the response. However, in recent years, a new strand of literature has claimed that deduction responses are at least as important as gross labour income responses. In fact, recent evidence shows that, under some circumstances, changes in deduction behaviour is the major channel of response (see, e.g., Kopczuk, 2005; Doerrenberg *et al.*, 2017). Accordingly, in this section we analyse in detail whether taxpayers use deductions as one of the main mechanisms to locate themselves right at the kink.

#### 5.1.1. Gross income versus taxable income

Our first test follows the main strategy used in Figure 3 using two different income concepts: gross labour income—i.e., labour income before considering deductions—and taxable labour income, which is the combination of income and deductions that defines the actual tax base. This test is similar to those reported in Chetty *et al.* (2011), Bastani and Selin (2014) and Schachtele (2016). By comparing these two distributions, we can rule out whether the bunching observed in Figure 3 is caused exclusively by the gross labour income component or if, instead, deductions play a role.

In Figure 5(a) we depict the gross labour income distribution. Since gross labour income does not include deductions, the  $x$  axis depicts the distance to the statutory kink point defined by the tax schedule (7,525 USD in 2013; see Table 2 for details). In contrast to the distance of one's taxable income to the first kink point, the distance of one's gross labour income to the statutory kink point is the same for all taxpayers. In this case, being in the first bin to the right of the kink point means that your gross labour income is between 0% and 2% larger than USD 7,525, and so on. The remaining features of the figure should be interpreted as usual.<sup>28</sup>

<sup>28</sup> For comparison purposes, Figure 5(b) includes the figure reported in Figure 3(b).

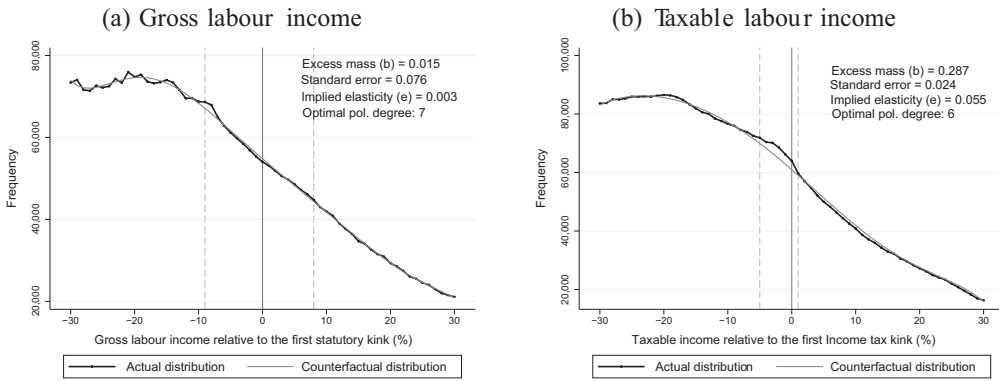


Fig. 5. *Gross Labour Income Versus Taxable Labour Income Distribution—Pure Wage Earners, 2010–2016.*

*Notes:* In panels (a) and (b) we plot the empirical and counterfactual gross labour income and taxable labour income distributions, respectively. The sample used in this figure is restricted to pure wage earners. In panel (a), the gross labour income distribution is normalised to the statutory income bracket (USD 7,525). In panel (b), the taxable income distribution is normalised as usual (see Subsection 2.2 and Online Appendix D). All estimates are performed by using the pooled data for the 2010–2016 period. In panel (a), the vertical solid line represents the statutory income bracket that is identical for each taxpayer. The values on the  $x$  axis should be interpreted as the percentage distance to USD 7,525. For an easier comparison with the gross labour income distribution, in panel (b) we reproduce the plot shown in panel (b) of Figure 3. In this plot, the vertical solid line represents the exact point at which each taxpayer starts to pay the IRPF, i.e., the first adjusted kink point. See Figure 3 for a more detailed explanation of each element in the figure.

Figure 5 reveals two things. On the one hand, panel (a) shows that there is no sign of bunching around the first kink for the gross labour income distribution. In other words, before considering deductions, there is no evidence of a behavioural response to personal income taxation. On the other hand, panel (b) reports our baseline result for pure wage earners and shows evidence of modest bunching around the first kink point when using the taxable income. Although the only difference between panels (a) and (b) are deductions, we cannot claim that this modest bunching is driven exclusively by deduction behaviour. What we can claim, however, is that bunching is only observed when both components are considered simultaneously. This finding is consistent with our hypothesis that deductions are a relevant piece of the puzzle for explaining the observed bunching behaviour.

### 5.1.2. *Itemiser versus non-itemiser taxpayers*

Our second test compares taxpayers' bunching patterns depending on their use of *itemised* deductions. As explained in Section 1, the IRPF allows for two types of deductions. *Non-itemised* deductions are automatically included in the final tax liability. Thereby, they do not represent a channel of response for the taxpayers. On the contrary, claiming *itemised* deductions is the result of a voluntary and explicit choice on the part of taxpayers. More specifically, a taxpayer might claim an *itemised* deduction in a way that allows her to fall just to the left of the kink point. Our hypothesis in this section is that the excess mass around the kink point should be larger for *itemising* taxpayers. Although we still cannot interpret these results as causal, differences between these two groups would suggest an intensive use of *itemised* deductions as a mechanism for bunching.



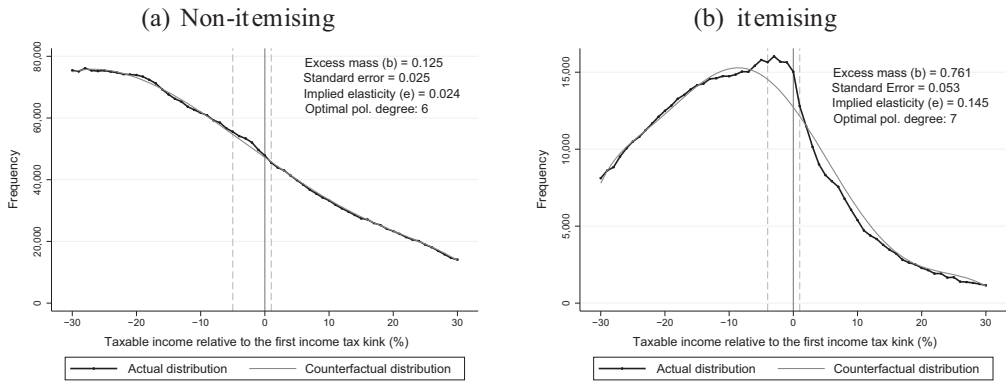


Fig. 6. Taxable Labour Income Bunching: Pure Wage Earners by Itemised Status, 2010–2016.

Notes: Plots of the empirical taxable income distributions for two sub-groups of the population for the 2010–2016 period: non-itemising and itemising taxpayers. The analysis restricts the sample to the group of wage earners who did not receive self-employment income, i.e., pure wage earners. In panel (a) we show the taxable income distribution for non-itemising taxpayers, i.e., taxpayers who did not claim *itemised* deductions such as child care and housing expenses. For panel (b), we repeated the analysis for all taxpayers who claimed *itemised* deductions. See Figure 3 for a more detailed explanation of each element in the figure.

For descriptive purposes, in Figure G.1, panel a, of Online Appendix G we show the evolution of the annual average deduction between 2010–2016 as well as its composition in terms of *itemised* and *non-itemised* deductions. Despite the fact that the average deduction increased continuously between 2010–2016 (from about USD 1,700 to 2,400), the share of *itemised* and *non-itemised* deductions has remained constant (ca. 80% versus ca. 20%). In turn, in Figure G.1, panel b, we show the composition of *itemised* deductions. Overall, personal deductions (see Table 1 for a detailed description) are the most important component of *itemised* deductions, representing on average about 84% of the total for the seven-year period considered. Although housing-related deductions increased between 2010 and 2016, they always represented less than 26%. In what follows, we restrict the analysis to taxpayers who did not claim housing deductions. We do this to achieve a more direct interpretation of our results. Because workers who claimed housing deductions represent a relatively small fraction (between 1.2% of all pure wage earners in 2010 and 6% in 2016), this restriction does not substantially affect the external validity of our results.

Figure 6 disaggregates the sample into *non-itemising* (panel (a)) and *itemising* taxpayers (panel (b)). In both cases, the estimated excess mass is statistically significant, although there are clear differences in the magnitude and shape of the excess mass. For *itemising* taxpayers (panel (b)), there is a clear hump around the kink point. The excess mass is 75%, implying an ETI of 0.145, which is about three times our baseline result for pure wage earners. This observation contrasts strongly with the evidence for *non-itemising* taxpayers. For the latter group, the excess mass is less than one sixth of the value for *itemising* taxpayers and the implied elasticity is very modest (0.024). These results support our hypothesis of a larger response for taxpayers that claim *itemised* deductions.

In Figure E.4 of Online Appendix E, we dissect taxable income responses of itemising taxpayers by gender, documenting a larger excess of mass in the case of female taxpayers. This suggests that heterogeneous responses by gender documented in Subsection 4.3 reflect a more intense

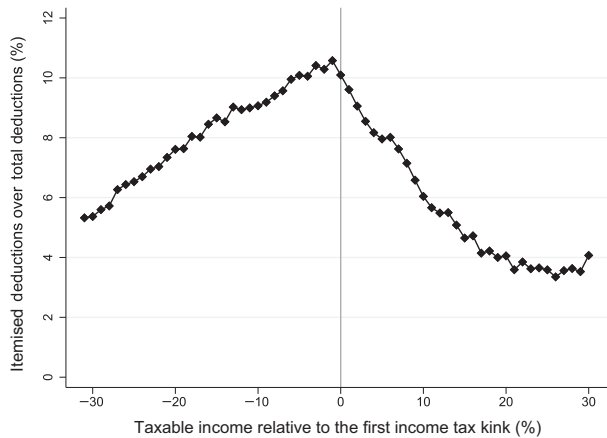


Fig. 7. *Intensity in the Use of Itemised Deductions, 2010–2016.*

*Notes:* The analysis restricts the sample to the group of wage earners who do not receive self-employment income, i.e., pure wage earners. The figure shows a relative increase in the use of *itemised* deductions around the first kink point for taxpayers who claim itemised deductions. Along the  $x$  axis we present the taxable labour income distribution as a percentage of income at the first kink point. The data are grouped into 2%-wide bins, represented by squares. The curve connecting marked data points indicates how large a percentage of a taxpayer's total deduction amount (i.e., *itemised* plus *non-itemised*) *itemised* deductions account for.

deduction behaviour of female taxpayers around the kink.<sup>29</sup> Differential real labour supply responses could also explain the observed differences between male and female taxpayers. However, this explanation seems to be unlikely in our context as we do not find evidence for gross labour income (see Figure E.5).

Our final test to analyse the extent of deduction behaviour is presented in Figure 7. In this figure, we show intensity in the use of *itemised* deductions as a function of distance to the first kink point. We define intensity as the ratio of *itemised* deductions to the total amount of deductions. If bunching behaviour is related to the use of *itemised* deductions, we would expect to see a more intensive use of these deductions close to the kink. Consistent with this hypothesis, the ratio of *itemised* deductions to the total amount of deductions has an inverted “V” shape that reaches its maximum (11%) just before  $k$  and then starts decreasing.

<sup>29</sup> To our knowledge, while there is some evidence on gender differences in evasion responses, there are no previous studies documenting such differences in deduction responses. For instance, Kleven *et al.* (2011), using randomised audits data, found that women are less likely than men to evade income tax. Given the lack of detailed information about each component comprising itemised deductions and also about household characteristics, we can only speculate about what causes the difference in deduction behaviour we find. It could be the case that in households with two wage earners women are more likely to claim child-care deductions than men. It is worth recalling that child-care deduction is one of the few deduction possibilities available in the Uruguayan system. However, as explained in Section 1, deductions are subject to a progressive deduction rates schedule. If households want to minimise the total tax burden, child-care deductions should be claimed by the main earner. As women are more likely to be secondary earners, this would be inconsistent with our finding. We hypothesize that other behavioural frictions may be at work. For instance, women may be more likely than men to recall the identity card number of their children, which is required at the moment of claiming the child-care deductions. This would be consistent with evidence that shows gender differences in the time spent in parenting and domestic responsibilities. Moreover, households may rely on a heuristic extrapolated from other social benefits (e.g., family allowances) in which women are the preferred beneficiary within the household.

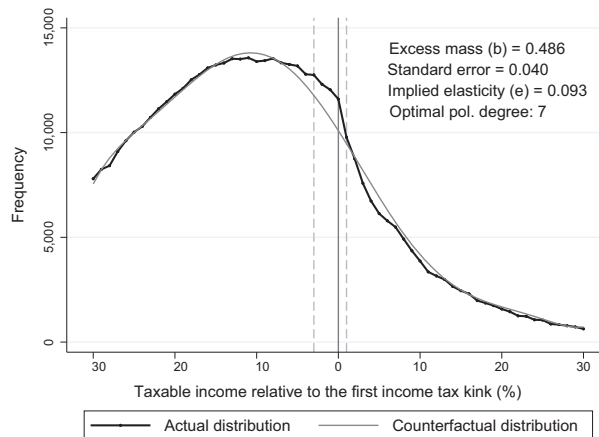


Fig. 8. *Taxable Labour Income Bunching—Pure Wage Earners: Itemisers and Non-Fillers, 2010–2016.* Notes: Plot of the empirical taxable income distributions for the sub-group of itemising taxpayers who did not file a tax return in the period 2010–2016. The analysis restricts the sample to the group of wage earners who do not receive self-employment income, i.e., pure wage earners. See Figure 3 for a more detailed explanation of each element in the figure.

Taken together, the various pieces of evidence reported in this section support the existence of a deduction behaviour channel. These tests reinforce the idea that deduction behaviour is partly channelling bunching responses.

### 5.1.3. Tax deduction or filing behaviour?

Taxpayers can claim *itemised* deductions by filing a tax return. In this process, taxpayers report *itemised* deductions and also all the other components of their taxable income. Since filing a tax return gives the taxpayer the opportunity to self-report both gross labour income and *itemised* deductions, it is not clear whether the response observed in the previous section for *itemising* taxpayers can be attributed exclusively to *itemised* deductions; it may also be associated with changes in reported gross labour income. However, even though the correlation between claiming an *itemised* deduction and filing a tax return is strong, it is not perfect. As explained in Section 1, employees can also claim *itemised* deductions via their employers (3100 form). We exploit the existence of alternative ways to claim deductions to analyse whether bunching for *itemising* taxpayers relates more to the deduction channel or is instead driven by filing a tax return. If taxpayers who do not file tax returns are using *itemised* deductions intensively—via their employers' tax reporting—to locate themselves close to the kink, the bunching cannot be attributed to filing behaviour.

In Figure 8 we present the analysis for non-filer *itemisers*. This group of workers constitutes 15% of the total number of pure wage earners and about 82% of the pure wage earners that claim *itemised* deductions. Because most taxpayers claimed *itemised* deductions via their employers instead of filing a tax return, Figure 8 resembles Figure 6(b) although with a less pronounced hump. The excess mass in this case is 49% with an implied elasticity of 0.093. This result suggests that part of the response observed for all *itemisers* is explained by taxpayers who did not file a tax return. Therefore, at least part of the bunching observed for *itemising* taxpayers must be driven by an intensive use of deductions. We can also test for the deduction channel

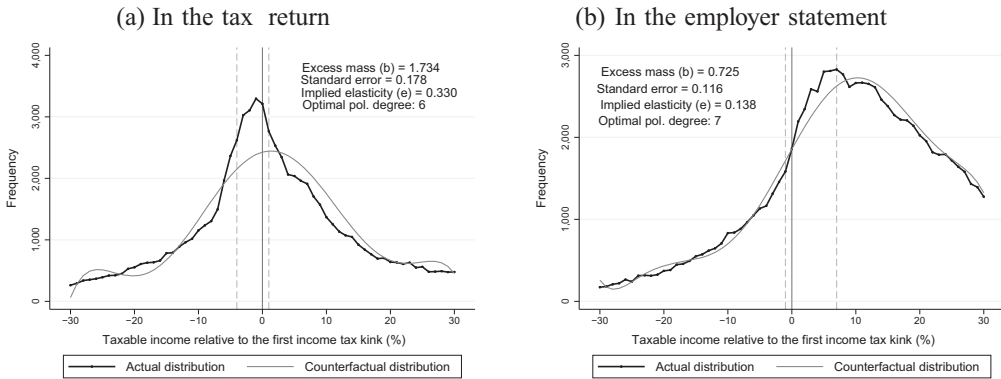


Fig. 9. Taxable Labour Income Bunching: Pure Wage Earner Itemisers and Filers—by Source of Information, 2010–2016.

Notes: Plots of the empirical taxable income distributions for two sub-groups of wage earners for the 2010–2016 period. In panel (a) we show the taxable income distribution for itemising taxpayers who file a tax return using the self-reported information contained in the tax return. For panel (b), we repeated the analysis but instead of using the tax return information, we used third-party reports submitted by employers. The analysis restricts the sample to the group of wage earners who did not receive self-employment income, i.e., pure wage earners. See Figure 3 for a more detailed explanation of each element in the figure.

versus the *filing-a-tax-return* channel in the group of workers that claimed *itemised* deductions and filed a tax return. To do this, we exploit the fact that, for this group of taxpayers, we have two independent sources of information: third-party reports from the employers' tax statement and self-reported information from tax returns. This information includes not only the gross labour income paid by the employer but also the amount of *itemised* and *non-itemised* deductions the worker claimed at the firm. This group represents 3.3% of the total number of pure wage earners and about 18.2% of taxpayers that claimed *itemised* deductions.<sup>30</sup>

First, we analyse the timing of the response. In particular, we explore whether taxpayers who claimed *itemised* deductions and filed a tax return started to bunch after filing the tax return or if that pattern was already evident in the employer's tax statement. Observing an excess of mass in the tax return but not in the information reported by the employer would suggest that the observed bunching is driven, at least partly, by the *filing-a-tax-return* channel. On the other hand, if bunching is already happening at the firm, we could rule out the possibility of the *filing-a-tax-return* mechanism being the only driver of the behavioural response.

In Figure 9(a) we show our baseline analysis for taxpayers who claimed *itemised* deductions and filed a tax return using the self-reported information contained in the tax return. For Figure 9(b), we repeated the analysis using the third-party reports from the firm. At the firm (i.e., before the tax return; panel (b)), there are some signs of bunching just after the kink point, which would be a “sub-optimal” response. However, when we turn to the self-reported information in the tax return (panel (a)), there is a pronounced excess mass of about 173%, which implies an elasticity of 0.330. This suggests that *itemising* taxpayers may be relatively limited in their ability to optimise their tax liability at the firm level and therefore use the tax return as the main channel to adjust it.

<sup>30</sup> Because of errors in administrative records, we do not have the third-party reported records in the database for about 9% of the pure wage earners who filed a tax return. Henceforth, our comparisons of the two sources of data will include only those taxpayers for whom we observe both self- and third-party-reported data.

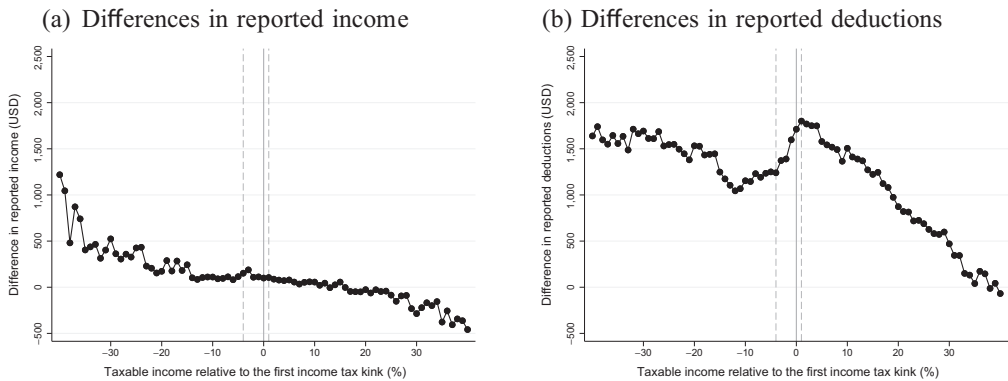


Fig. 10. *Differences in Reported Gross Labour Income and Deductions in Itemisers and Filers—Pure Wage Earners, 2010–2016.*

*Notes:* Plots of the average differences in gross labour income (panel (a)) and deductions (panel (b)) between the self- and third-party-reported information for taxpayers who itemised deductions and filed a tax return in the period 2010–2016. The analysis restricts the sample to the group of wage earners who did not receive self-employment income, i.e., pure wage earners. It also excludes those taxpayers who were already in the bunching zone based on information workers provided to the firm. Taxpayers are grouped and sorted using the tax return information and according to their distance to the taxable income kink point. The differences in reported income (amount of deductions) are grouped into 2%-wide bins, represented by the dots. We restrict the analysis sample to individuals located up to thirty bins (60%) above or below the kink.

At first glance, this seems to contradict the conclusions drawn from Figure 8 where we attributed the behavioural response to the *itemised* deductions channel. However, adjustments to tax liability performed in the tax return may be achieved both through adjustments in gross labour income and through *itemised* deductions. Hence, we still cannot rule out the deduction channel. Because our data allow us to test this directly, in Figure 10 we present discrepancies between individuals' tax returns and employer reports in gross labour income (panel (a)) and *itemised* deductions (panel (b)) by distance to the taxable income kink point. For deductions, positive differences mean that self-reported deductions are higher than third-party-reported deductions. For gross labour income, positive differences mean that labour income reported in the employer's tax statement is higher than the corresponding self-reported item in the tax return. In both panels we excluded observations that were already in the bunching zone in the firm (8.7%) because these individuals have no incentive to report different information on their tax return and including them dilutes potential differences between the two sources of information. Panel (a) shows some general evidence of income underreporting, specially to the left of the kink point. However, this behaviour is not related to a bunching pattern because the differences are smoothly distributed. Panel (b) shows that taxpayers usually report deductions in the tax return, i.e., they claim more deductions on their tax return than at the firm. However, unlike income underreporting, these differences are disproportionately large around the bunching zone. In particular, taxpayers located in the thirtieth bin overreport about USD 1,500 in *itemised* deductions. This amount decreases until the fifteenth bin, where it is around USD 1,000. Here, the differences start to grow again until they peak just around the kink point at USD 1,500 and then start decreasing again. This suggests that taxpayers who locate themselves in the bunching zone when they file a tax return use *itemised* deductions for this purpose. These figures suggest that the results

reported in Figure 9 that seemed to contradict the deduction channel hypothesis are, in fact, consistent with it because tax return filers get close to the kink point by using *itemised* deductions intensively.

Overall, in this section we have provided additional evidence that the deduction channel plays a major role in explaining bunching behaviour. To complement the analysis, we next focus on the second component of taxable income: gross labour income.

## 5.2. Gross Labour Income and Underreporting Behaviour

In this section, we investigate potential responses along the gross labour income margin. To that end, we focus on *non-itemising* taxpayers who, by design, are not able to respond using the deduction margin. Despite the fact that the ETI estimated for this group was very small (0.024)—though statistically significant—this average response may be masking potentially informative changes in gross labour income. For *non-itemising* taxpayers, gross labour income responses can be explained by responses in two margins. The first margin is associated with a traditional labour supply response. Taxpayers might change the total amount of hours they dedicate to labour market activities as a response to a change in the marginal tax rate. The second margin is through income underreporting, which can be implemented in two different ways: unilaterally by the employee and through collusive employer-employee agreements.<sup>31</sup>

### 5.2.1. Unilateral underreporting by employees

Here we focus on unilateral income underreporting. In this case, taxpayers may locate near the kink by reporting only part of their gross labour income to the DGI. This is particularly important in developing countries where enforcement capacities are limited and taxpayers can evade taxes more easily. Our data does not allow us to precisely distinguish these two margins but the results of an exploratory analysis are suggestive.

As in the previous section, we exploit the availability of self- and third-party reports of individuals' gross labour income.<sup>32</sup> In Figure 11(a) we present our baseline analysis for *non-itemising* taxpayers using the self-reported information contained in the tax return. For panel (b), we repeated the analysis using the information reported by the employer. In the tax return (panel (a)), we observe strong evidence of a pronounced bunch just at the kink point (with an excess mass of 62% and an implied ETI of 0.120).

Since we focus on *non-itemising* taxpayers, the differences in the magnitude of the bunching between the tax return and the Employer Statement can only be explained by income underreporting. This can be tested more directly. In Figure 12 we show the difference in the gross labour income reported by the firm and the gross labour income self-reported by the taxpayer on the tax return. This figure shows two interesting results. First, the positive differences observed along the whole distribution show that, on average, people report less income on their tax return than the amount the firms report paying them, which suggest taxpayers care for reducing their total tax liability. Second, and more directly related to our analysis of mechanisms, the difference in

<sup>31</sup> A third possibility would be unilateral misreporting by the firm. This does not seem to be incentive compatible in our context as there are no incentives for the firm to report their employees' income just below the kink.

<sup>32</sup> There are two reasons why a taxpayer that does not want to claim an *itemised* deduction may nonetheless file a tax return. First, they may be required to do so by the DGI if they have multiple sources of income and their total earnings are above a defined annual threshold (see Section 1). Second, they may voluntarily file a tax return even if they are not required to do so. This is perfectly reasonable for a taxpayer who has a refund to collect, but it could also signal an intention to misreport/evade taxes.

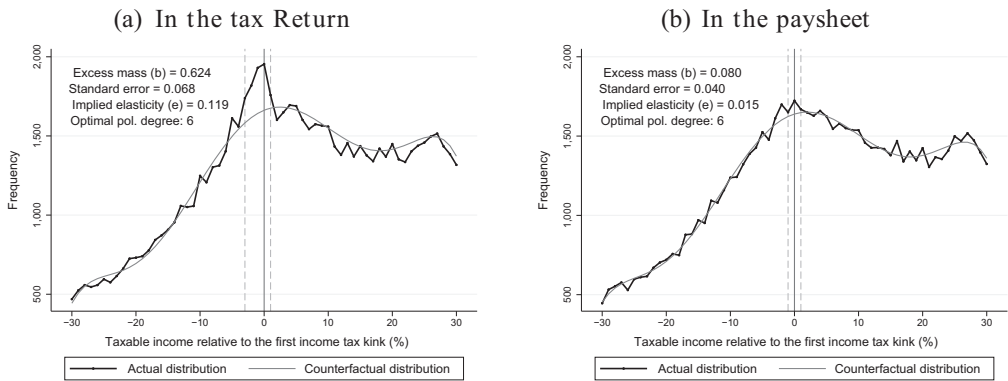


Fig. 11. *Taxable Labour Income Bunching: Wage Earner Non-itemisers and Filers—By Source of Information, 2010–2016.*

Notes: Plots of the empirical taxable income distributions for two sub-groups of wage earners for the 2010–2016 period. In panel (a) we show the taxable income distribution for non-itemising taxpayers who filed a tax return using the self-reported information contained in the tax return. For panel (b), we repeated the analysis but instead of using the tax return information, we used third-party reports submitted by employers. The analysis restricts the sample to the group of wage earners who did not receive self employment income, i.e., pure wage earners. See Figure 3 for a more detailed explanation of each element in the figure.

the self- and third-party-reported gross labour income increase sharply just at the bunching zone. This shows that taxpayers who do not claim *itemised* deductions use income underreporting as a mechanism to adjust their taxable income and locate themselves to the left of the kink point.

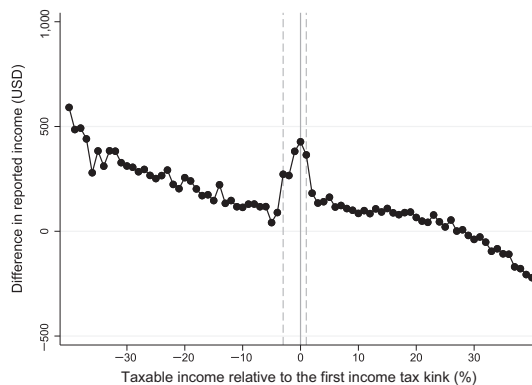


Fig. 12. *Differences in Reported Gross Labour Income for Non-Itemisers and Filers—Pure Wage Earners, 2010–2016.*

Notes: Plot of the average differences in gross labour income between the self-reported and third-party-reported information for taxpayers who filed a tax return but did not claim itemised deductions in the period 2010–2016. The analysis restricts the sample to the group of wage earners who did not receive self-employment income, i.e., pure wage earners. It also excludes those taxpayers who were already in the bunching zone based on the information the firm reported. Taxpayers are grouped and sorted using the tax return information and according to their distance to the taxable income kink point. The differences in reported income (amount of deductions) are grouped into 2%-wide bins, represented by the dots. We restrict the analysis sample to individuals located up to thirty bins (60%) above or below the kink.

Furthermore, this finding provides direct evidence that tax evasion behaviour—through unilateral underreporting—for wage earners is responsive to changes in the tax rate, and complements previous findings in the literature that found this behaviour for self-employed taxpayers (e.g., Kleven *et al.*, 2011).

### 5.2.2. Underreporting via collusive employer-employee agreements

Figure 11(b) shows no conclusive evidence of bunching at the firm. The implied ETI is very small (0.015). While gross labour income adjustments via collusive agreements between employers and individual employees are a theoretical possibility, the *collusion-in-the-firm* channel does not appear to be economically relevant in the Uruguayan context.

An alternative hypothesis could be that collective bargaining agreements between employers and trade unions take income tax brackets and associated tax rates into consideration (Chetty *et al.*, 2011).<sup>33</sup> This is theoretically plausible in the Uruguayan context as wages are negotiated at the industry level and collective bargaining coverage is very high, comparable with many European countries. If this is the case, personal coordination between individual taxpayers and employers would be unnecessary. To check whether these features of wage-setting institutions in Uruguay may be inducing particular behavioural responses, we link the administrative tax records with a complementary dataset that contains (almost) all minimum wages negotiated in the collective bargaining process. We link this data with the tax records using the activity sector code at a four-digit level. The results from this exercise, presented in Figures 13 and 14 (see also Figures H.1 and H.2 of Online Appendix H), do not provide support for the idea that the first tax bracket would be a focal point for wage bargaining by industry in Uruguay. Negotiated wage floors are well below the first kink of the personal income tax. Moreover, there is no relationship between the distance of negotiated minimum wages to the first kink and sector-specific implied elasticities of taxable income.

In sum, we document a gross labour income response through unilateral misreporting among *non-itemiser* pure wage earners. Moreover, we find little support for the existence of income underreporting occurring via collective bargaining agreements or employer-employee collusive behaviour at the firm level.

## 6. Conclusion

In this paper we analyse the behavioural responses of taxpayers to the first kink of Uruguay's personal income tax schedule. Our identification approach relies on local variation in tax rates generated by the progressivity of the tax system. We obtain two main sets of findings. First, when we consider the universe of taxpayers, there is a statistically significant and robust—although somewhat diffuse—excess mass around the first kink point. This translates to an ETI of about 0.06, which is relatively modest and falls within the lower tier of elasticities found for individuals in the bottom-middle part of the income distribution in advanced economies. In contrast to previous studies, we show that wage earners are at least as responsive as are self-employed individuals. However, in line with previous evidence on wage earners (e.g., Chetty *et al.*, 2011; Paetzold, 2018), we find that the observed local taxable income response for this group of workers translates to modest elasticities. The magnitude and shape of our estimated responses may reflect the fact that the environment in which taxpayers operate determines the strength of

<sup>33</sup> We thank an anonymous referee for drawing our attention to this point.



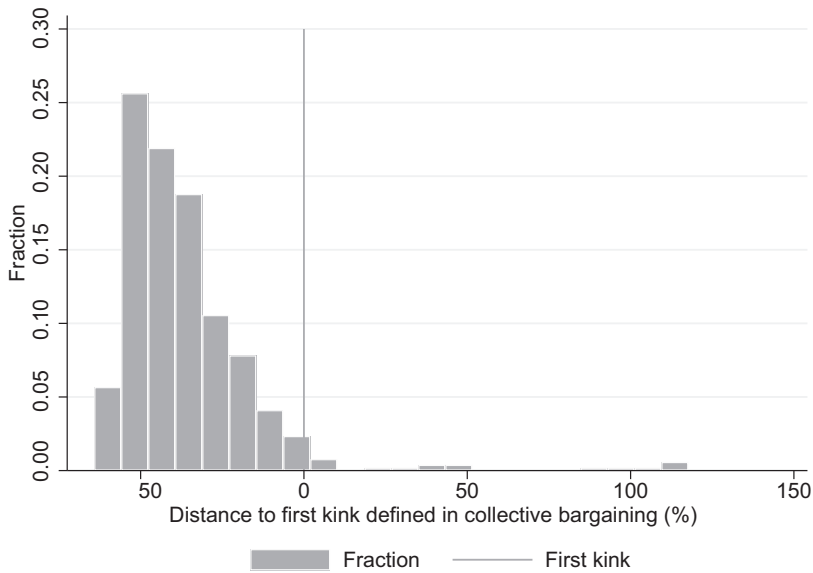


Fig. 13. *Distribution of Minimum Wage by Sub-Sector (2010–2016).*

*Notes:* In this figure we depict the distribution of the distance between the minimum wages negotiated in collective bargaining to the statutory first kink of the personal income tax schedule. Each data point represents one of the ninety-one sub-groups defined by the collective bargaining legislation in a given year for the period 2010–2015. Within a sub-group year, we considered the average minimum wage for different jobs. On the  $x$  axis we measure the distance to the exemption threshold in thousands of USD adjusted by the CPI. The  $y$  axis represents the share of minimum wages in a given interval of the distance. The vertical line represents the exemption threshold. Each observation included in the histogram represents a sub-group year.

their behavioural response (Slemrod and Kopczuk, 2002). Progressive taxation is relatively novel in Uruguay and taxpayers are still figuring out how the new tax schedule works. Our findings regarding learning dynamics are consistent with this hypothesis. Moreover, a modest elasticity is also consistent with the restricted set of deductions allowed by Uruguay's tax code, which limits the tools available to taxpayers to adjust (precisely) their taxable income (Kopczuk, 2005). In addition, modest responses may also be explained by frictions and optimisation costs that might constrain behavioural responses to taxation, particularly for wage earners (Chetty *et al.*, 2011; Chetty, 2012).

The second set of main findings concerns the mechanisms taxpayers use to adjust their behaviour. We find that the two components of the taxable labour income (gross labour income and deductions) both play an important role. Furthermore, their relative importance depends on the opportunities available to different groups of taxpayers. In particular, a number of tests support the hypothesis that *itemised* deductions are a major channel of response. However, we also report evidence showing that, when taxpayers do not make use of *itemised* deductions, they underreport their gross labour income in the tax return as a mechanism of adjustment.

Our results advance our understanding of behavioural responses to personal income taxation in developing countries and also inform the policy debate. Assessing the efficiency costs of taxation

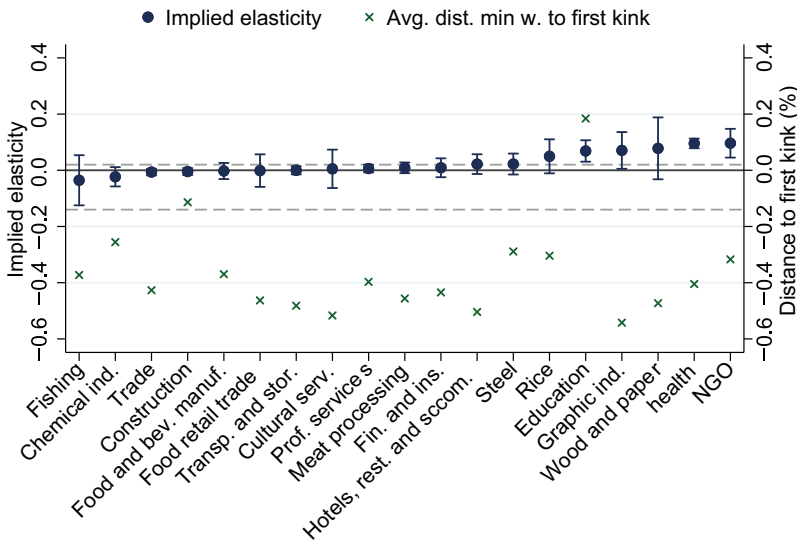


Fig. 14. *Bunching Elasticities and Minimum Wage Distance to the Kink by Sub-Sectors, 2010–2014.*

*Notes:* In this figure we report the estimated ETI for each group defined by the collective bargaining legislation. We do this for all non-itemising and non-filler pure wage earners. We merge the administrative records to the collective bargaining dataset using ISIC codes at a four-digit level. Each blue circle (measured on the left-hand side y axis) represents the estimated ETI for one of the groups, while crosses (measured on the right-hand side y axis) represent the average distance of the average minimum wage in the group to the statutory first kink point. Vertical dashed lines represent the bunching area defined as in Figure 3(b), which is the estimated optimal bunching area for the whole sample of pure wage earners. We exclude rural work and housing service activities since they are regulated by a different legislation.

is crucial for countries that need to raise tax revenue to increase the supply of public goods, without introducing major distortions. We document that taxpayers respond to personal income taxation to a very limited extent. This is a significant finding because it shows that, in this context, the efficiency costs of taxation are not necessarily large. Moreover, our detailed analysis of the channels of response suggests that using the estimated ETI as a “sufficient statistic” in welfare analysis may lead to an overestimation of the dead-weight loss of income taxation (Slemrod, 1998; Chetty, 2009; Doerrenberg *et al.*, 2017). As shown by Chetty (2009), this could be the case if some of the private costs of income underreporting are transfer costs rather than resource costs. While resource costs are costs incurred by agents that are not recovered by anyone in the society, transfer costs are associated to a transfer of resources between different agents. In our context, part of the increase in income underreporting could have been recovered by the tax agency through an increase in revenue from fines to tax evaders or simply represent hidden transfers across private agents. In addition, it is also worth noting that deduction behaviour is usually associated with externalities (Doerrenberg *et al.*, 2017). For instance, allowing for child-care deductions usually incentivise expenses that could be welfare enhancing. The type of deductions allowed by the Uruguayan system, even when very restrictive, seems to be consistent with this type of externality. Notwithstanding the above, our results show that policymakers should proceed cautiously with reforms aimed at expanding tax deduction opportunities and focus on improving the quality of third-party reporting mechanisms to limit tax evasion. Such reforms are especially

important in the context of developing countries that are usually characterised by low enforcement capacities. In a context of significant misuse of deductions and/or tax evasion, broadening the tax base and improving the administrative capacity to enforce tax collection seems an appropriate policy response (Kopczuk, 2005).

*Universidad de la República, Uruguay, IZA, Germany and GLO  
Leeds University Business School, University of Leeds, UK, Universidad de La República & IZA,  
Germany  
Universidad de la República and Paris School of Economics, France  
Anderson School of Management, UCLA, USA and Universidad de La República, Uruguay  
Universidad de La República, Uruguay*

Additional Supporting Information may be found in the online version of this article:

### Online Appendix Replication Package

### References

- Acemoglu, D., Naidu, S., Restrepo, P. and Robinson, J.A. (2018). 'Democracy does cause growth', *Journal of Political Economy*, vol. 127(1), pp. 47–100.
- Altonji, J.G. and Paxson, C.H. (1988). 'Labor supply preferences, hours constraints, and hours-wage trade-offs', *Journal of Labor Economics*, vol. 6(2), pp. 254–76.
- Bachas, P. and Soto, M. (2018). 'Not(ch) your average tax system: corporate taxation under weak enforcement', Policy Research Working Papers, 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*, vol. 109, pp. 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*, vol. 117, pp. 211–28.
- Besley, T. and Persson, T. (2013). 'Chapter 2—Taxation and development', in (A.J. Auerbach, R. Chetty, M. Feldstein and E. Saez, eds.), *Handbook of Public Economics*, pp. 51–110, North Holland: Elsevier.
- Best, M.C. (2015). 'The role of firms in workers earnings responses to taxes: evidence from Pakistan', Working Paper, London School of Economics.
- 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*, vol. 123(6), pp. 1311–55.
- 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*, vol. 52, p. 1627.
- Bohne, A. and Nimczik, J.S. (2017). 'Learning dynamics in tax bunching at the kink—evidence from Ecuador', VfS Annual Conference 2017 (Vienna): Alternative Structures for Money and Banking, Verein für Socialpolitik/German Economic Association.
- Bosch, N., Dekker, V. and Strohmaier, K. (2020). 'A data-driven procedure to determine the bunching window: an application to the Netherlands', *International Tax and Public Finance*, pp. 951–79.
- Card, D., Chetty, R., Feldstein, M. and Saez, E. (2010). 'Expanding access to administrative data for research in the United States', *Political Economy: Structure & Scope of Government eJournal*.
- Carrillo, P., Pomeranz, D. and Singhal, M. (2017). 'Dodging the taxman: firm misreporting and limits to tax enforcement', *American Economic Journal: Applied Economics*, vol. 9(2), pp. 144–64.
- Ceni, R., Carrasco, P., Parada, C. and Santín, S. (2020). 'Wage distribution effects and displacement on a centralized wage bargaining scheme', *Universidad de la República, Departamento de Economía, Uruguay*.
- Chetty, R. (2009). 'Is the taxable income elasticity sufficient to calculate dead-weight loss? The implications of evasion and avoidance', *American Economic Journal: Economic Policy*, vol. 1(2), pp. 31–52.
- Chetty, R. (2012). 'Bounds on elasticities with optimization frictions: a synthesis of micro and macro evidence on labor supply', *Econometrica*, vol. 80(3), pp. 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*, vol. 126(2), pp. 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*, vol. 103(7), pp. 2683–721.
- Chetty, R., Looney, A. and Kroft, K. (2009). 'Salience and taxation: theory and evidence', *American Economic Review*, vol. 99(4), pp. 1145–77.
- Chetty, R. and Saez, E. (2013). 'Teaching the tax code: earnings responses to an experiment with EITC recipients', *American Economic Journal: Applied Economics*, vol. 5(1), pp. 1–31.
- Clotfelter, C.T. (1983). 'Tax evasion and tax rates: an analysis of individual returns', *The Review of Economics and Statistics*, vol. 65(3), pp. 363–73.
- DGI. (2016). 'Informe de Resultados de IRPF IASS—Ejercicio 2016', Dirección General Impositiva (DGI), Montevideo.
- Dincecco, M. and Katz, G. (2016). 'State capacity and long-run economic performance', *Economic Journal*, vol. 126(590), pp. 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*, vol. 151, pp. 41–55.
- Duncan, D. and Sabirianova Peter, K. (2016). 'Unequal inequalities: do progressive taxes reduce income inequality?', *International Tax and Public Finance*, vol. 23(4), pp. 762–83.
- ECLAC. (2013). *Fiscal Panorama of Latin America and the Caribbean. Tax Reform and Renewal of the Fiscal Covenant*, Santiago, Chile: United Nations.
- Feinstein, J.S. (1991). 'An econometric analysis of income tax evasion and its detection', *The RAND Journal of Economics*, vol. 22(1), pp. 14–35.
- Feldstein, M. (1995). 'The effect of marginal tax rates on taxable income: a panel study of the 1986 tax reform act', *Journal of Political Economy*, vol. 103(3), pp. 551–72.
- Feldstein, M. (1999). 'Tax avoidance and the dead-weight loss of the income tax', *The Review of Economics and Statistics*, vol. 81(4), pp. 674–80.
- Fisman, R. and Wei, S. (2004). 'Tax rates and tax evasion: evidence from missing imports in China', *Journal of Political Economy*, vol. 112(2), pp. 471–96.
- Gillitzer, C. and Slemrod, J. (2016). 'Does evasion invalidate the welfare sufficiency of the ETI?', *The BE Journal of Economic Analysis & Policy*, vol. 16(4).
- Gordon, R. and Li, W. (2009). 'Tax structures in developing countries: many puzzles and a possible explanation', *Journal of Public Economics*, vol. 93(7–8), pp. 855–66.
- Hamilton, S. (2018). 'Optimal deductibility: theory, and evidence from a bunching decomposition'.
- Hargarden, E.P. (2020). 'Taxpayer responses in good times and bad', *Journal of Economic Behavior & Organization*, vol. 176, pp. 653–90.
- 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*, vol. 23(4), pp. 640–69.
- Kleven, H.J. (2016). 'Bunching', *Annual Review of Economics*, vol. 8(1), pp. 435–64.
- Kleven, H.J., Knudsen, M.B., Kreiner, C.T., Pedersen, S. and Saez, E. (2011). 'Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark', *Econometrica*, vol. 79(3), pp. 651–92.
- 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*, vol. 83(330), pp. 219–46.
- Kleven, H.J. and Waseem, M. (2013). 'Using notches to uncover optimization frictions and structural elasticities: theory and evidence from Pakistan', *The Quarterly Journal of Economics*, vol. 128(2), pp. 669–723.
- Kopczuk, W. (2005). 'Tax bases, tax rates and the elasticity of reported income', *Journal of Public Economics*, vol. 89(11), pp. 2093–119.
- Kumler, T., Verhoogen, E. and Frías, J.A. (2020). 'Enlisting employees in improving payroll-tax compliance: evidence from Mexico', *Review of Economics and Statistics*, vol. 102(5), pp. 881–96.
- LaLumia, S., Sallee, J.M. and Turner, N. (2015). 'New evidence on taxes and the timing of birth', *American Economic Journal: Economic Policy*, vol. 7(2), pp. 258–93.
- le Maire, D. and Schjerning, B. (2013). 'Tax bunching, income shifting and self-employment', *Journal of Public Economics*, vol. 107, pp. 1–18.
- Mortenson, J.A. and Whitten, A. (2020). 'Bunching to maximize tax credits: evidence from kinks in the US tax schedule', *American Economic Journal: Economic Policy*, vol. 12(3), pp. 402–32.
- Paetzold, J. (2019). 'How do taxpayers respond to a large kink? Evidence on earnings and deduction behavior from Austria', *International Tax and Public Finance*, vol. 26, pp. 167–97.
- Rogerson, R., Shimer, R. and Wright, R. (2005). 'Search-theoretic models of the labor market: a survey', *Journal of Economic Literature*, vol. 43(4), pp. 959–88.
- Rosen, H.S. (1976). 'Taxes in a labor supply model with joint wage-hours determination', *Econometrica*, vol. 44(3), pp. 485–507.
- Saez, E. (2010). 'Do taxpayers bunch at kink points?', *American Economic Journal: Economic Policy*, vol. 2(3), pp. 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*, vol. 50(1), pp. 3–50.

- Schachtele, S. (2016). 'The importance of deductions in response to the personal income tax: bunching evidence from Germany', VfS Annual Conference 2016 (Augsburg): Demographic Change, Verein für Socialpolitik/German Economic Association.
- Slemrod, J. (1998). 'Methodological issues in measuring and interpreting taxable income elasticities', *National Tax Journal*, vol. 51(4), pp. 773–88.
- Slemrod, J. (2008). 'Does it matter who writes the check to the government? The economics of tax remittance', *National Tax Journal*, vol. 61(2), pp. 251–75.
- Slemrod, J. and Kopczuk, W. (2002). 'The optimal elasticity of taxable income', *Journal of Public Economics*, vol. 84(1), pp. 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*, vol. 88(2), pp. 205–16.
- Zenou, Y. (2008). 'Job search and mobility in developing countries. Theory and policy implications', *Journal of Development Economics*, vol. 86(2), pp. 336–55.