

Taxpayer responsiveness to taxation

Evidence from bunching at kink points of the South African income tax schedule

Neryvia Pillay Bell

SA-TIED Working Paper #119 | June 2020



About the project

Southern Africa –Towards Inclusive Economic Development (SA-TIED)

SA-TIED is a unique collaboration between local and international research institutes and the government of South Africa. Its primary goal is to improve the interface between research and policy by producing cutting-edge research for inclusive growth and economic transformation in the southern African region. It is hoped that the SA-TIED programme will lead to greater institutional and individual capacities, improve database management and data analysis, and provide research outputs that assist in the formulation of evidence-based economic policy.

The collaboration is between the United Nations University World Institute for Development Economics Research (UNU-WIDER), the National Treasury of South Africa, the International Food Policy Research Institute (IFPRI), the Department of Monitoring, Planning, and Evaluation, the Department of Trade and Industry, South African Revenue Services, Trade and Industrial Policy Strategies, and other universities and institutes. It is funded by the National Treasury of South Africa, the Department of Trade and Industry of South Africa, the Delegation of the European Union to South Africa, IFPRI, and UNU-WIDER through the Institute's contributions from Finland, Sweden, and the United Kingdom to its research programme.

Copyright © UNU-WIDER 2020

Corresponding author: neryvia.pillay@uct.ac.za

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the of the SA-TIED programme partners or its donors.



WIDER Working Paper 2020/68

Taxpayer responsiveness to taxation

Evidence from bunching at kink points of the South African income tax schedule

Neryvia Pillay Bell*

June 2020

Abstract: The author applies the bunching methodology to South African administrative tax data over the period from 2011 to 2017 to investigate the responsiveness of individual taxpayers to changes in marginal personal income tax rates. She finds significant evidence of bunching among the self-employed but no evidence of bunching among wage earners. Among the self-employed, bunching is greatest at the highest kink in the income tax schedule and smallest at the lowest kink. Female self-employed exhibit greater bunching behaviour than male self-employed, and responsiveness appears to decrease with age. The responsiveness of the self-employed appears to be due to tax avoidance by shifting income into future periods through retirement fund deductions, as well as a real labour supply response. Despite the significant excess bunching observed, the implied elasticities of taxable income—under the assumption of a uniform heterogeneity distribution around the kink—are not very large.

Key words: bunching, elasticity of taxable income, personal income taxation, South Africa

JEL classification: H24, H31, O12

Acknowledgements: I thank Chris Axelson, Amina Ebrahim, and Jukka Pirttilä for useful comments and suggestions. The views expressed in this paper are those of the author, and do not necessarily reflect the views of the South African Reserve Bank (SARB) or Economic Research South Africa (ERSA).

*University of Cape Town, Economic Research Southern Africa, and South African Reserve Bank, Cape Town, South Africa, email: nerylvia.pillay@uct.ac.za

This study has been prepared within the UNU-WIDER project [Southern Africa – Towards Inclusive Economic Development \(SA-TIED\)](#).

Copyright © UNU-WIDER 2020

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9256-825-2

<https://doi.org/10.35188/UNU-WIDER/2020/825-2>

Typescript prepared by Lesley Ellen.

The United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland, Sweden, and the United Kingdom as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the Institute or the United Nations University, nor the programme/project donors.

1 Introduction

In order to evaluate the efficiency and welfare impacts of tax policy, it is crucial to understand taxpayers' responsiveness to taxation. These behavioural responses take many forms, including reduced labour supply and tax avoidance through income shifting. The elasticity of taxable income captures all these responses and can, under certain conditions, be used as a sufficient statistic for optimal tax analysis (Saez et al. 2012). Accordingly, empirical estimation of the elasticity of taxable income has received considerable attention in the international public finance literature. In this paper, I use South African administrative tax data over the period from 2011 to 2017 to estimate the elasticity of taxable income at kink points in the South African income distribution.

I use a bunching approach developed by Saez (2010) and Chetty et al. (2011) that seeks to identify the taxable income elasticity by studying bunching behaviour at kink points in the income tax schedule where the marginal tax rate changes. While this approach has been widely applied to data from developed countries, this paper is one of the few applications to a developing country and thus contributes to our understanding of the effects of taxation in developing countries. Kleven and Waseem (2013) study bunching in the Pakistani income tax system and Bachas and Soto (2018) study Costa Rican corporate tax, but these two studies analyse notches where the average tax rate changes and so their results are less informative about typical income tax schedules where the marginal tax rate changes. Boonzaaier et al. (2019) apply the bunching technique to the South African corporate tax system and find significant bunching with large implied elasticities. My paper is closely related to theirs as I apply the same technique to the South African personal income tax system.

The other main empirical approach to estimating the elasticity of taxable income is to use changes in marginal tax rates that are generated by tax reforms. Estimates of the elasticity of taxable income generated using a tax reform approach that exploits large changes to the tax system are usually larger than those estimated using bunching.¹ The obvious drawback of the tax reform approach is that it requires a tax reform, but there are also difficulties in finding control groups that are comparable to the treated groups experiencing the tax change and in finding an instrument for the explanatory variable of interest, the net-of-tax rate changes.² Bunching estimates only require that the tax rate changes at points in the income distribution and there is no endogeneity problem.

The absence of a large tax reform in South Africa over the period for which administrative tax data is available renders the bunching approach attractive.³ Another alternative that also does not require a tax reform is the bracket creep approach, developed by Saez (2003), which exploits the changes in marginal tax rates caused by inflation pushing taxpayers into a higher tax bracket when nominal brackets are not inflation adjusted. This approach was applied by Kemp (2019) to the South African personal income tax system. Both the bracket creep and bunching approaches depend on taxpayer awareness of the details of the tax code, and so the extent of the behavioural responses observed

¹ He et al. (2020) show that this is due to differences in the time and scope of the estimates—bunching estimates have an unknown time horizon and generate local elasticities for each kink, whereas tax reform approaches have a defined time horizon (typically 1, 2, or 3 years) and produce a global elasticity.

² Indeed, Weber (2014) shows that most instruments used in the prior literature are not exogenous.

³ At the time of writing, data was only available from 2011 to 2017. Over this period, marginal tax rates were constant until 2016 when there was a one percentage point increase across the board, which did not actually result in much change in the net-of-tax-rates (see Section 4.1 for details). In 2018, a new highest tax bracket was introduced but this falls outside the period of data availability.

when using these methods depends on informational considerations. In this sense, the elasticities yielded by these estimates are less relevant for more salient tax rate changes. In estimating the elasticity of taxable income using an approach that does not require a tax reform, my paper is similar to Kemp (2019) but also extends his analysis in important ways. First, I conduct a differential analysis by gender and age in order to understand how the tax system differentially impacts the income distributions of men and women, and the young and old, in South Africa. Second, I investigate the anatomy of taxpayer responsiveness in order to shed some light on the extent to which the observed responses reflect real labour supply changes or tax avoidance behaviour. Finally, the bunching approach also provides an additional robustness check to the bracket creep approach. The accuracy of the bracket creep estimates hinges on there being no bunching behaviour since, if taxpayers adjust their incomes to fall just below the kink point, they are less likely to fall in the treatment group creating downward bias in these estimates.

The remainder of the paper proceeds as follows: Section 2 briefly discusses the related literature, Section 3 outlines the bunching methodology, Section 4 provides information on the South African personal income tax system and the administrative tax data used in the analysis, Section 5 presents the empirical results, Section 6 conducts some robustness checks, and Section 7 concludes.

2 Related literature

The use of the bunching methodology in applied work has surged in recent years, due in large part to the increased availability of large administrative datasets. While bunching follows directly from the predictions of the standard taxable income labour supply model, it is a very local effect around the kink points in an income tax schedule and hence is difficult to observe in survey data.

The seminal paper in this area is Saez (2010), who finds evidence of bunching at the first kink point of the US Earned Income Tax Credit (EITC), although the effect is concentrated among the self-employed and is thus more likely to reflect changes in reporting behaviour rather than changes in labour supply. Saez (2010) also studies the kink points of the US federal income tax schedule but finds evidence of bunching at only the first kink point but not at any of the others located in the middle and upper parts of the income distribution.

Chetty et al. (2011) provide an important extension to the model by allowing for adjustment costs and hours constraints that attenuate econometric estimates of labour supply elasticities using individual level data. The predictions of the model are supported by an analysis of Danish tax data. In particular, larger taxable income elasticities are generated by larger kinks and by kinks that apply to larger groups of workers.

Most of the bunching papers in the area of tax policy that build on the work of Saez (2010) and Chetty et al. (2011) have studied developed countries (see Kleven 2016) for a summary of this work). Three notable exceptions are Kleven and Waseem (2013) who study bunching in the Pakistani income tax system, Bachas and Soto (2018) who study the Costa Rican corporate tax, and Boonzaaier et al. (2019) who study the South African corporate income tax. The first two of these papers are based on analysing notches in the tax system (points where the average tax rate changes) rather than kinks (points where the marginal tax rate changes). Since notches are less common than kinks in tax schedules, these papers are of limited relevance to policy makers in most developing countries.

Boonzaaier et al. (2019) find substantial bunching of small businesses at both kinks in the South African corporate tax system with large implied elasticities of between 0.7 and 1.6. This analysis is

obviously highly relevant to South African policy makers but is also more likely to be generalizable to other countries than the two analyses based on notches. My paper is closely related to Boonzaaier et al. (2019) since I apply a similar technique but investigate the South African personal income tax system instead of the corporate tax system. Kemp (2019) is the only existing study to estimate the elasticity of taxable income using South African administrative data and finds an estimated elasticity of around 0.3. Kemp (2019) follows Saez (2003) to identify taxable income responses off ‘bracket creep’, i.e. the lack of full adjustment of the nominal tax brackets to inflation. Both the bracket creep and bunching approaches depend on taxpayer awareness of the details of the tax code, and so the extent of the behavioural responses observed when using these methods depends on informational considerations.

3 Methodology

The bunching methodology was developed by Saez (2010) and Chetty et al. (2011). The technique relies on the predictions of the standard taxable income labour supply model that I briefly outline here. Individual preferences are defined over after-tax income (the value of consumption) and before-tax income (the cost of effort), and individual optimization generates a smooth earnings distribution. At baseline there is a proportional tax schedule so that all individuals face the same marginal tax rate τ_1 .

Consider the introduction of a kink at earnings level k so that all individuals earning above k face a marginal tax rate τ_2 with $\tau_2 > \tau_1$. All individuals earning below k continue to face a marginal tax rate of τ_1 . The introduction of this convex kink generates the following responses:

1. The earnings distribution to the left of the kink k is unaffected as there is no change in the incentives of individuals earning less than k .
2. Individuals initially, i.e. before the kink is introduced, earning above k will reduce their taxable income in response to the higher marginal tax rate that they now face.
3. There will be a spike in the income distribution. All individuals initially earning in the interval $[k, k + \Delta k]$ move to the kink point, while those initially above this interval reduce their earnings to an interior point of the upper bracket and do not move all the way to the kink point.

These responses produce excess bunching in the earnings distribution at the kink point. The intuitive idea behind the estimation procedure is to compute a measure of the excess bunching at the kink point by comparing the observed mass of individuals at the kink point with the mass of individuals that would be observed at the same earnings level in the absence of any kinks.

Central to the estimation approach then is the calculation of the counterfactual distribution, i.e. the earnings distribution that would have been observed in the absence of any kinks. The approach developed by Chetty et al. (2011),⁴ is to fit a flexible polynomial to the observed income distribution omitting observations located in a range around the kink. More concretely, individuals are grouped into earnings bins indexed by j and the following regression is estimated:

⁴ I closely follow the estimation procedure developed by Chetty et al. (2011) and use the Stata command `bunch_count` developed by the authors.

$$c_j = \sum_{i=0}^p \beta_i (z_j)^i + \sum_{i=-R}^R \gamma_i \cdot 1[z_j = i] + \eta_j \quad (1)$$

where c_j is the number of individuals in earnings bin j , z_j is the earnings level relative to the kink in bin j , $[-R, R]$ is the narrow, symmetric range of excluded data around the kink k , and p is the order of the polynomial. Both z_j and R are measured in units of the bin width, d . As in Chetty et al. (2011), I use a seventh order polynomial in all the main estimations but demonstrate that the results are robust to different polynomial orders in Section 6. The counterfactual distribution is estimated by a recursive procedure that starts with an initial estimate given by the predicted values of equation (1) omitting the effects of the dummies in the excluded range and then shifts the counterfactual distribution to the right of the kink upward until it satisfies the constraint that the area under the counterfactual must equal the area under the observed distribution.

The extent of excess bunching can then be computed by comparing the observed bin counts to the predicted bin counts in the excluded range $[-R, R]$, normalized by the average height of the counterfactual distribution in the excluded range:

$$\hat{b} = \frac{\sum_{j=-R}^R c_j - \hat{c}_j}{\sum_{j=-R}^R \hat{c}_j / (2R+1)} \quad (2)$$

where $(2R + 1)$ gives the number of bins in the excluded range $[-R, R]$. I use a bin width, d , of ZAR2,500 and a small window of excluded data of ZAR5,000 on either side of the kink point so that $R = 2$. The counterfactual income distribution is estimated over a large window of ZAR75,000 on either side of the kink point, i.e. the range $[-30, 30]$ expressed in units of d . Standard errors are estimated using a bootstrap procedure that randomly resamples the residuals from (1) to generate a large number of earnings distributions.

Under the assumption that the heterogeneity distribution of individuals is uniform around the kink,⁵ the compensated elasticity of taxable labour income locally at the kink point k , $\tilde{e}(k)$, can be related to the estimate of excess bunching using the following formula:

$$\tilde{e}(k) = \frac{b}{k \times \log\left(\frac{1-\tau_1}{1-\tau_2}\right)} \quad (3)$$

While the estimate of excess bunching, \hat{b} , from equation (2) depends on the size chosen for the earnings bins, the estimate of the taxable income elasticity calculated using equation (3) is invariant to the bin width d , provided that k is expressed in units of d (Bastani and Selin 2014).

It should also be noted that the elasticity in equation (3) is not a structural elasticity. In fact, $\tilde{e}(k)$ is likely to be smaller than the structural elasticity since optimization frictions like search costs and hours constraints may limit the observed responses.

⁵ Recent work by Blomquist et al. (2019) and Bertanha et al. (2019) provide methods that exploit kinks in budget sets to estimate the taxable income elasticity under less-restrictive assumptions, but I do not apply those techniques in this paper.


4 Background and data

4.1 The South African personal income tax system

Personal income tax is the largest source of tax revenue for the South African national government, contributing 47 per cent of gross tax revenue in 2016/17 (National Treasury 2018). The tax is imposed at the individual level so that spousal income does not affect an individual's tax liability. Most forms of personal income, including from wage and self-employment as well as fringe benefits, are taxable. The most prominent deductions available to individuals are pension fund and retirement annuity contributions. Medical aid contributions and expenses were deductible but these were converted to tax credits from the 2013 tax year onwards.

The South African tax year runs from 1 March to the end of February and I refer to a tax year by the calendar year in which it ends, e.g. the 2016 tax year refers to the period 1 March 2015 to 29 February 2016. Over my sample period, there were six taxable income brackets with marginal tax rates unchanged until 2016, when there was a one percentage point increase in all marginal tax rates except in the lowest bracket. The brackets are adjusted on an annual basis but do not track inflation in a consistent manner, leading to bracket creep over the years, a phenomenon that has been exploited by Kemp (2019) to estimate the elasticity of taxable income. Figure 1 shows the location of the tax bracket thresholds over the years 2011–17. There has been a steady increase in the tax bracket thresholds over the years with the exception of 2017, when the two highest brackets were left unchanged.

Figure 1: Kink points in the South African personal income tax schedule, 2011–17



Source: author's illustration based on SARS data.

Table 1 shows the changes in the marginal tax rates at five kink points in the taxable income distribution for the years before (2011–15) and after (2016–17) the marginal tax rate changes. It is evident that the net-of-tax rate changes, i.e. $\ln\left(\frac{1-\tau_1}{1-\tau_2}\right)$, are larger at the lower end of the distribution

than at the upper end. The largest change in the net-of-tax rate occurs at the first kink and the next largest change is in the middle of the distribution at the third kink. Moreover, the change in the net-of-tax rate at the first kink became much larger after the changes to the tax rate schedule while tax changes at the other kink points did not increase by as much. Larger tax changes should generate greater bunching both because the incentives are greater and because workers are more willing to pay the search costs to find new jobs (Chetty et al. 2011).

I analyse the extent of bunching at these five kink points in the taxable income distribution. There is an additional kink point at the lower end due to the primary rebate that creates a minimum tax threshold below which the tax rate is effectively zero. I do not examine bunching around this tax threshold point since the exact location of this threshold point also varies at the individual level after 2012, when the medical tax deduction was converted to a credit. The data does not contain sufficient information on the individual's medical tax credit to accurately determine the location of this kink point.

Table 1: Marginal tax rates, 2011–17

	2011–15			2016–17		
	τ_1	τ_2	$\ln\left(\frac{1-\tau_1}{1-\tau_2}\right)$	τ_1	τ_2	$\ln\left(\frac{1-\tau_1}{1-\tau_2}\right)$
First kink	0.18	0.25	0.089	0.18	0.26	0.103
Second kink	0.25	0.30	0.069	0.26	0.31	0.070
Third kink	0.30	0.35	0.074	0.31	0.36	0.075
Fourth kink	0.35	0.38	0.047	0.36	0.39	0.048
Fifth kink	0.38	0.40	0.033	0.39	0.41	0.033

Source: author's compilation based on SARS data.

4.2 Data

I use the individual level panel created from South African Revenue Service (SARS) administrative data (Ebrahim and Axelson 2019). This dataset contains detailed information on income and deductions for the universe of South African taxpayers for the years 2011 to 2017 but is relatively thin on demographic information. I have data on gender and age but not on education or marital status, for example.

I restrict the sample to those under 65 years old, as this is the typical retirement age in South Africa,⁶ but who are at least 15 years old. The gender of some individuals cannot be accurately determined and so these individuals are dropped from the gender differentiated results but remain in the full sample.⁷

I redefine the taxable income variable to reflect the distance from the bracket cut-off point in that year so that it takes on a value of zero at the kink point. I pool this data across years and examine the extent of bunching across the entire sample period of 2011 to 2017. Since the marginal tax rates were increased in 2016, I also present results separately for 2011 to 2015 and 2016 to 2017.

⁶ Individuals over the age of 65 also receive a higher primary income tax rebate.

⁷ Gender in the SARS administrative dataset is derived from an individual's national identity number and so is missing for individuals with no identity number, for example foreigners (Ebrahim and Lilenstein 2019). There are also individuals for whom multiple genders are reported across tax certificates and/or years and I omit these individuals with indeterminate genders from the gender differentiated analysis.

Earlier work has demonstrated important differences in bunching across wage earners and the self-employed (Bastani and Selin 2014; Chetty et al. 2011; Saez 2010). I define the self-employed as those taxpayers who report any business income, and this group comprises two per cent of all observations in the dataset. In all analyses in Section 5 the group of ‘wage earners’ is comprised of all taxpayers who do not report any business income, i.e. it excludes the ‘self-employed’.

5 Empirical analysis

5.1 Main results



I begin with a visual presentation of the excess bunching around the first kink in the South African income tax schedule in Figure 2 for wage earners in the left panel (a) and the self-employed in the right panel (b). The top row corresponds to the full sample period, 2011 to 2017, the middle row to the years before the changes to the marginal tax rates, 2011 to 2015, and the bottom row to the years after the marginal tax rate changes, 2016 to 2017. The blue dotted line plots the empirical distribution of taxable income around the kink point, while the smooth red line represents the seventh order fitted polynomial, which excludes the observations in the small window of $[-\text{ZAR}5,000:\text{ZAR}5,000]$ around the kink point.

Figure 2: Bunching at the first kink in the income tax schedule



(a) Wage earners

(b) Self-employed



All years, 2011–17



Before-tax rate changes, 2011–15



After-tax rate changes, 2016–17



Source: author's illustration based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

Over the full sample period, it is evident that wage-earning workers (top left graph in Figure 2) do not exhibit any bunching behaviour at the first kink point and the estimate of excess bunching is actually negative. On the other hand, the self-employed group (top right graph in Figure 2) exhibit significant excess bunching around this kink point over the full sample period. The bunching behaviour for the years before the tax rate changes (middle row in Figure 2) is very similar to that of the full sample period. The last row of Figure 2 shows that bunching behaviour at the first kink is greater in the years after the tax rate changes, 2016 to 2017. The self-employed exhibit much greater bunching after the tax rate changes than before, and the estimate is statistically significant at the ten per cent level. There is a small positive estimate of bunching for wage earners, although it is not statistically significant and is smaller than any of the bunching estimates for the self-employed. Despite the significant observed bunching, the implied elasticities—under the assumption that the heterogeneity distribution is uniform around the kink—are not large. The elasticity for the self-employed at the first kink over 2016 to 2017, which is where the greatest bunching is observed in Figure 2, is only 0.02.

These differences between the self-employed and other taxpayers, and across the years before and after the tax rate changes, are also evident at the other kink points in the income tax schedule. Table 2 provides estimates of the excess bunching at all five kink points for the sample of wage earners in columns (1) and (2) and the self-employed in columns (3) and (4). Panel A of Table 2 provides estimates of excess bunching for the full sample period, 2011-2017, Panel B for the years before the tax rate changes, 2011–15, and Panel C for the years after the tax rate changes, 2016–2017.

Table 2: Estimates of excess bunching at five kink points in the income tax distribution

	Self-employed		Others	
	b (1)	se (2)	b (3)	se (4)
Panel A: All years, 2011–17				
First kink	-0.116***	(0.040)	0.108***	(0.038)
Second kink	0.072***	(0.025)	0.155***	(0.043)
Third kink	0.041	(0.026)	0.150***	(0.043)
Fourth kink	0.029	(0.028)	0.102	(0.071)
Fifth kink	0.048**	(0.022)	0.224***	(0.074)
Panel B: Before-tax rate changes, 2011–15				
First kink	-0.170***	(0.050)	0.100**	(0.039)
Second kink	0.060**	(0.028)	0.132***	(0.045)
Third kink	0.055**	(0.025)	0.136***	(0.047)
Fourth kink	0.055**	(0.024)	0.099*	(0.058)
Fifth kink	0.066**	(0.029)	0.122	(0.082)
Panel C: After-tax rate changes, 2016–17				
First kink	0.032	(0.054)	0.146*	(0.088)
Second kink	0.102***	(0.036)	0.262***	(0.081)
Third kink	0.007	(0.042)	0.218**	(0.104)
Fourth kink	-0.028	(0.063)	0.112	(0.149)
Fifth kink	0.008	(0.051)	0.732***	(0.217)

Notes: sample includes all taxpayers aged 15–64 years old. Columns (3) and (4) are the sample of self-employed, defined as those taxpayers who report any business income, while columns (1) and (2) are the wage earners, defined as those taxpayers who do not report any business income. Columns (1) and (3) contain the estimates of excess bunching, and columns (2) and (4) give the bootstrapped standard errors of those estimates. *** significant at the 1% level, ** at the 5% level, * at the 10% level.

Source: author's calculations based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

It is clear from Table 2 that excess bunching is much greater among the self-employed than wage earners at all kink points in the income distribution in all time periods. The estimates of excess bunching for wage earners (column (1) of Table 2) are often statistically significant or negative, and any statistical significance is due to the high precision of the seventh order polynomial fit rather than significant bunching behaviour. By contrast, the estimates of excess bunching for the self-employed (column (3) of Table 2) are almost all positive and statistically significant and reflective of visible bunching. This can clearly be seen in Figure 2 where the positive and statistically significant estimates of bunching for the self-employed at the first kink in all three time periods is matched by visible bunching, while the statistically significant negative bunching for wage earners at the first kink over the full sample period and the earlier years, 2011–15, are not matched by obvious bunching in the graphs. Figure 3 further illustrates this point by providing a visual presentation of bunching at the second kink, where the largest bunching estimates for wage earners are found in Table 2. There is a clear lack of visible bunching among wage earners at the second kink across all years and in the years before the tax rate changes (first column, top and middle rows of Figure 3) despite the statistically significant positive estimates of bunching in Table 2 (column (1), Panels A and B). There is some evidence of slight bunching among wage earners after the tax rate changes (bottom left of Figure 3). On the other other hand, distinct bunching at the second kink among the self-employed can be observed in the second column of Figure 3, validating the large positive and significant estimates for the self-employed in column (3) of Table 2.

The results in Table 2 therefore demonstrate significant excess bunching among the self-employed with almost no evidence of bunching among wage earners. This is indicative of greater responsiveness to changes in marginal tax rates by the self-employed than wage earners and could be due to a greater ability among the self-employed to adjust hours of work and/or shift income in order to reduce taxable income. The sources of this responsiveness are examined more closely in Section 5.4.

Among the self-employed, the estimates of excess bunching are much greater in the years after the tax rate changes than in the years before (column (3) in Panels B and C of Table 2). Although the results for the years after the tax rate changes are less precisely estimated because of the fewer number of observations, they do indicate that taxpayer responsiveness is greater after the increases in the marginal tax rates. Recall from Table 1 that after the changes to the marginal tax rates, the net-of-tax rate changes became much larger only at the first kink and were relatively constant at the other kinks. The observed increases in bunching behaviour among the self-employed over 2016–17 relative to 2011–15 at kink points where incentives were effectively unchanged across the two time periods suggests that the increased responsiveness might be due to greater salience of the kink points after the marginal tax rate changes.

For self-employed taxpayers, the magnitude of the bunching estimates are typically largest at the top of the income distribution at the fifth kink, followed by the middle of the income distribution at the second and third kinks, and smallest at the first kink (with minimal evidence of significant bunching at the fourth kink). These differences across the income distribution are not matched by the relative changes in tax incentives—Table 1 shows that the largest tax change occurs at the first kink and the smallest at the fifth kink. Thus, the finding of greater responsiveness at the fifth kink than at the first kink also suggests that responsiveness might be due to information differences rather than greater incentives.

Indeed, the largest estimate of excess bunching, 0.732, in Table 2 is found among the self-employed at the fifth kink in 2016–17. This could be because the extent of taxpayer responsiveness depends on both the ability to adjust taxable income and awareness of the location

of kinks. Self-employed individuals are likely to be more able to adjust their labour supply and shift their incomes, and higher-income individuals might be more aware of the location of kink points particularly after the changes to the marginal tax rates. The importance of information has been demonstrated in the US where providing US taxpayers with information about the existing tax schedule has been shown to affect reported income in the next tax year among the self-employed but not wage earners (Chetty and Saez 2013).


Although there is evidence of significant bunching among the self-employed in Table 2, the implied elasticities under the assumption of a uniform heterogeneity distribution around the kink are not very large. Even at the largest bunching estimate, 0.732, the elasticity of taxable income is only 0.08. It is possible that these relatively small estimates are due to the restrictive assumption on the heterogeneity distribution, but estimation of the elasticity under less restrictive assumptions is beyond the scope of this paper.

Figure 3: Bunching at the second kink in the income tax schedule



(a) Wage earners

(b) Self-employed

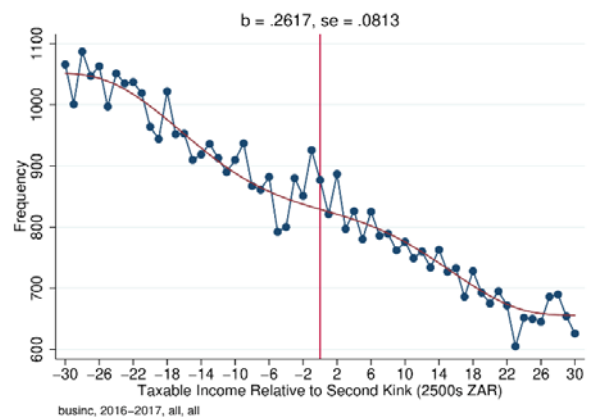

All years, 2011–17



Before-tax rate changes, 2011–15



After-tax rate changes, 2016–17



Source: author's illustration based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

5.2 Comparison of main results with other studies

My finding of greater bunching among the self-employed than among wage earners is similar to the results obtained for Sweden by Bastani and Selin (2014) and for Denmark by Chetty et al. (2011). Despite the significant bunching observed for the self-employed, the implied elasticities in both papers are small, similar to my results for South Africa. Under the assumption that the heterogeneity distribution is uniform around the kink, the implied elasticity in Bastani and Selin (2014) is 0.02 and 0.24 in Chetty et al. (2011). The largest bunching estimate I observe generates an elasticity of 0.08 and so is in line with the estimates from these two papers.

On the other hand, my estimated elasticities are much smaller than those obtained by Kemp (2019), who exploits ‘bracket creep’ in the South African income tax schedule to estimate an elasticity of taxable income of around 0.3. One reason for the difference might be the time period studied—Kemp (2019) studies the 2009–13 tax years whereas the results in this paper are for a later period over the 2011–17 tax years. Another reason for the difference could be the time horizon of the estimates—the elasticities in Kemp (2019) are estimated over a three-year period and so capture a longer-run response whereas bunching estimates pool observations in multiple time points together and so have an unclear time property.

Although our estimates are quantitatively different, they are qualitatively similar in that they both reflect relatively low responsiveness compared to the estimates for South African companies obtained by Boonzaaier et al. (2019). In addition, both Kemp (2019) and I find evidence of greater taxpayer responsiveness at higher income levels. I find general evidence of greater bunching at the top end of the income distribution than at lower points. This is similar to the results obtained by Kemp (2019) who finds a slightly higher elasticity of around 0.4 for high-income taxpayers.

5.3 Differential results by gender and age

The results in Table 2 indicate that significant bunching behaviour is demonstrated by the self-employed but not wage earners. In Table 3, I conduct separate analyses by gender and age for self-employed taxpayers over the full sample period, 2011–17. Analysis of behavioural responses by gender and age provides important information on the effects of tax policy and how the tax system differentially impacts the income distributions of men and women, and the young and old, in South Africa.

The results in Table 3 indicate that, among the self-employed, bunching behaviour is generally greater among females than among males (Panels A and B). This finding is consistent with the results of other studies (for example Chetty et al. 2011) which demonstrate that married women have much higher taxable income elasticities than single men. Although I cannot determine an individual’s marital status in my data, it is likely that many women are secondary earners and so would exhibit greater responsiveness.

Panels C through G of Table 3 differentiate bunching behaviour across five age groups and reveal a general trend of lower responsiveness as taxpayers get older. Among those approaching retirement age (Panel G), only one of the five bunching estimates are significant and all the estimates are smaller than those found for younger taxpayers. At the lower and middle points of the income distribution (i.e. the first, second, and third kinks), the bunching estimates predominantly decrease as the age group increases. At the top of the distribution (the fifth kink) the pattern is less stable, with the largest bunching estimated for the youngest group, 15–24 years old, and the second oldest group, 45–54 years old. The high-income self-employed who are aged 15–24 years have a very large bunching estimate of 4.542 with an implied elasticity of 0.54.

5.4 Anatomy of the response

The observed responsiveness of taxable income to the tax rate could be due either to changes in reporting behaviour that reflect tax avoidance or changes in labour supply behaviour that reflect changes in real economic activity. The bunching approach allows me to examine the anatomy of the taxable income response by studying the filing behaviour of individuals. Significant bunching is found for the self-employed and this could reflect a real labour supply response and/or the use of deductions to reduce taxable income.

Table 3: Estimates of excess bunching among the self-employed by gender and age, 2011–17

	b (1)	se (2)
Panel A: Females		
First kink	0.122***	(0.041)
Second kink	0.183**	(0.072)
Third kink	0.128	(0.082)
Fourth kink	-0.013	(0.113)
Fifth kink	0.291**	(0.132)
Panel B: Males		
First kink	0.065	(0.057)
Second kink	0.052	(0.069)
Third kink	0.086	(0.083)
Fourth kink	0.055	(0.089)
Fifth kink	0.120	(0.105)
Panel C: 15–24 years old		
First kink	0.199	(0.244)
Second kink	0.852*	(0.436)
Third kink	0.774	(0.569)
Fourth kink	-0.403	(0.823)
Fifth kink	4.542**	(2.090)
Panel D: 25–34 years old		
First kink	0.254***	(0.068)
Second kink	0.200**	(0.092)
Third kink	0.043	(0.082)
Fourth kink	0.209	(0.127)
Fifth kink	-0.137	(0.161)
Panel E: 35–44 years old		
First kink	0.125**	(0.056)
Second kink	0.081	(0.064)
Third kink	0.204**	(0.088)
Fourth kink	0.290**	(0.117)
Fifth kink	0.197	(0.143)
Panel F: 45–54 years old		
First kink	0.081	(0.059)
Second kink	0.200***	(0.076)
Third kink	0.151*	(0.078)
Fourth kink	-0.045	(0.110)
Fifth kink	0.465***	(0.151)

Panel G: 55–64 years old		
First kink	-0.011	(0.061)
Second kink	0.121	(0.083)
Third kink	0.143*	(0.079)
Fourth kink	-0.022	(0.108)
Fifth kink	0.157	(0.135)

Notes: sample includes all taxpayers who are self-employed, defined as those taxpayers who report any business income, for the years 2011–17. Column (1) contains the estimates of excess bunching and column (2) gives the bootstrapped standard errors of those estimates. *** significant at the 1% level, ** at the 5% level, * at the 10% level.

Source: author's calculations based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

Deductions are more prevalent among the self-employed—55 per cent of self-employed individuals take any income deductions, compared to only 37 per cent of wage earners—and this might explain the finding of significant bunching among the self-employed but not among wage earners. To investigate this, two categories of deductions are separately added back to the taxable income of self-employed individuals and the extent of excess bunching is re-estimated. In this way, the extent to which the behavioural response is income shifting rather than labour supply, is revealed. The two categories of deductions are: (1) retirement fund and income protection contributions; and (2) deductions related to income generation, such as travel expenses, subsistence allowances, and home office expenses. These represent the largest categories of deductions, in terms of the average over the sample period, for all taxpayers as well as the self-employed.⁸

Since the results in Table 2 show that significant bunching behaviour is concentrated among the self-employed, I focus on this group to determine the anatomy of the response. Table 4 examines the anatomy of the responses for self-employed individuals at four of the five kinks in the income tax schedule over three time periods. (I omit the fourth kink since there is only weak evidence of bunching at this kink in Table 2.)

Table 4: Anatomy of excess bunching, self-employed only

	2011–17		2017 only	
	b (1)	se (2)	b (3)	se (4)
Panel A: First kink				
Taxable income	0.108***	(0.038)	0.0690	(0.123)
plus retirement and income protection contributions	0.074*	(0.038)	-0.118	(0.115)
plus income-related deductions	0.097***	(0.036)	0.087	(0.127)
plus all deductions	0.074*	(0.040)	-0.091	(0.129)
Panel B: Second kink				
Taxable income	0.155***	(0.043)	0.201	(0.135)
plus retirement and income protection contributions	0.133***	(0.044)	0.210*	(0.124)
plus income-related deductions	0.167***	(0.044)	0.099	(0.124)
plus all deductions	0.218***	(0.050)	0.258**	(0.118)
Panel C: Third kink				
Taxable income	0.150***	(0.043)	0.108	(0.174)
plus retirement and income protection contributions	0.110**	(0.046)	0.176	(0.159)
plus income-related deductions	0.145***	(0.038)	0.172	(0.175)
plus all deductions	0.103**	(0.046)	0.015	(0.189)

⁸ Medical deductions are the next largest category, but this was changed to a tax credit during my sample period.

Panel D: Fifth kink				
Taxable income	0.224***	(0.074)	1.038***	(0.305)
plus retirement and income protection contributions	0.178*	(0.098)	0.373	(0.272)
plus income-related deductions	0.192***	(0.074)	0.780***	(0.301)
plus all deductions	0.114*	(0.065)	0.322	(0.251)

Notes: sample includes all taxpayers aged 15–64 years old who are self-employed, defined as those taxpayers who report any business income, for the full sample period (2011–17) in columns (1) and (2) and for 2017 only in columns (3) and (4). Columns (1) and (3) contain the estimates of excess bunching and columns (2) and (4) give the bootstrapped standard errors of those estimates. *** significant at the 1% level, ** at the 5% level, * at the 10% level.

Source: author's calculations based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

First, I study the anatomy of excess bunching over the entire sample period in columns (1) and (2) of Table 4. Over these years, adding back retirement fund and income protection contributions reduces the estimate of excess bunching at all four kinks, indicating that the self-employed are using retirement fund deductions as a way to reduce their taxable income in response to marginal tax rate changes. On the other hand, income-related deductions do not appear to be an important way in which the self-employed reduce their taxable income—the excess bunching estimates do not decrease much (and increase at the second kink) when income-related deductions are added back. Further, at all four kinks, the estimates of excess bunching are significant even after all deductions have been added back to taxable income. This suggests that there is also a real response to changes in the marginal tax rate as the self-employed have also adjusted their earnings (likely through hours worked).

Beginning in 2017, the rules for retirement fund deductions were greatly simplified so that all types of retirement funds (pension funds, provident funds, and retirement annuities) were treated the same for tax purposes, whereas previously each had their own specific deduction rules and limits. Since I find that retirement fund deductions are an important mechanism through which the self-employed manipulate their taxable income over the full sample period (columns (1) and (2) of Table 4), I separately examine the anatomy of bunching in the year after the changes to retirement fund deduction rules in columns (3) and (4) of Table 4. Since there is only one year of data under the new rules (2017), there are few observations and many of the estimates are no longer statistically significant. However, one clear result is apparent: the estimate of excess bunching of taxable income at the fifth kink in 2017 is much larger than over the full sample period and this excess bunching is almost entirely due to the use of retirement fund deductions. At the fifth kink, when retirement fund deductions are added back to taxable income, the estimate of excess bunching is no longer significant and when all deductions are added back there is no significant bunching either. Thus, these results suggest that the changes to the retirement fund deduction rules may have particularly benefited self-employed taxpayers with high earnings who utilized the changes to manipulate their taxable incomes without really affecting their labour supply.

6 Robustness checks

In this section, I test the sensitivity of the results to the bin width, size of the estimation window, and the polynomial order.

Table 5 presents results for three bin widths—500, 1,000, and 2,500—for wage earners (in Panel A) and the self-employed (in Panel B) over the full sample period, 2011–17. Column (3) of Table 5 uses a bin width of ZAR2,500 and so replicates the main results of Panel A in Table 2.

Table 5: Sensitivity of excess bunching estimates to bin width

	Bin width		
	500 (1)	1,000 (2)	2,500 (3)
Panel A: Wage earners			
First kink	-0.438*** (0.157)	-0.248*** (0.072)	-0.116*** (0.036)
Second kink	0.282*** (0.107)	0.145*** (0.048)	0.072*** (0.027)
Third kink	0.197* (0.103)	0.090* (0.049)	0.041 (0.025)
Fourth kink	0.174 (0.126)	0.080 (0.079)	0.029 (0.029)
Fifth kink	0.166 (0.125)	0.081 (0.072)	0.048** (0.022)
Panel B: Self-employed			
First kink	0.597*** (0.209)	0.290*** (0.101)	0.108*** (0.037)
Second kink	0.646*** (0.178)	0.324*** (0.088)	0.155*** (0.042)
Third kink	0.694*** (0.198)	0.362*** (0.105)	0.150*** (0.043)
Fourth kink	0.388 (0.286)	0.253* (0.148)	0.102 (0.068)
Fifth kink	1.013*** (0.366)	0.554*** (0.178)	0.224*** (0.068)

Notes: sample includes all taxpayers aged 15–64 years old over the full sample period, 2011–2017. Panel B is the sample of self-employed, defined as those taxpayers who report any business income, while Panel A is the wage earners, defined as those taxpayers who do not report any business income. Bootstrapped standard errors in parenthesis. *** significant at the 1% level, ** at the 5% level, * at the 10% level.

Source: author's calculations based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

The value of the estimates of excess bunching depend on the bin width (see the discussion in Section 3), so the magnitude of the estimates in columns (1) and (2) will differ from those in column (3) of Table 5, but they should be qualitatively similar. The results using bin widths of ZAR500 and ZAR1,000 in columns (1) and (2), respectively, are indeed similar to those obtained using a bin width of ZAR2,500 in column (3). There is significant excess bunching among the self-employed but not wage earners (any significant estimates among wage earners are due to the precision of the polynomial fit as with the main results), with similar patterns across the earnings distribution as in the main results. The one exception for the self-employed is at the fourth kink, where the estimate becomes marginally significant with a bin width of ZAR1,000 but the results still generally support the finding of weak evidence of bunching at the fourth kink. In all, the results are robust to varying the size of the bin width.

Table 6 examines the sensitivity of the results to the size of the estimation window over which the counterfactual income distribution is estimated while keeping the bin width fixed at ZAR2,500. Panel A presents the results for wage earners and Panel B for the self-employed, both estimated over the full sample period, 2011–17. Column (2) replicates the main results of Panel A in Table 2, which use an estimation window of [-75,000:75,000]. The results in column (1) use a smaller estimation window of [-50,000:50,000], which is the one used by Chetty et al. (2011), and those in column (3) use a wider estimation window of [-100,000:100,000].

Table 6: Sensitivity of excess bunching estimates to estimation window

	Estimation window		
	[-50,000:50,000] (1)	[-75,000:75,000] (2)	[-100,000:100,000] (3)
Panel A: Wage earners			
First kink	-0.073* (0.039)	-0.116*** (0.036)	-0.014 (0.056)
Second kink	0.033* (0.020)	0.072*** (0.027)	0.026 (0.032)
Third kink	0.026 (0.033)	0.041 (0.025)	0.044 (0.028)
Fourth kink	0.011 (0.028)	0.029 (0.029)	0.020 (0.027)
Fifth kink	0.030 (0.023)	0.048** (0.022)	0.039* (0.022)
Panel B: Self-employed			
First kink	0.075** (0.034)	0.108*** (0.037)	-0.003 (0.199)
Second kink	0.168*** (0.047)	0.155*** (0.042)	0.140*** (0.041)
Third kink	0.169*** (0.048)	0.150*** (0.043)	0.155*** (0.042)
Fourth kink	0.069 (0.060)	0.102 (0.068)	0.076 (0.057)
Fifth Kink	0.201*** (0.075)	0.224*** (0.068)	0.174** (0.072)

Notes: sample includes all taxpayers aged 15–64 years old over the full sample period, 2011–2017. Panel B is the sample of self-employed, defined as those taxpayers who report any business income, while Panel A is the wage earners, defined as those taxpayers who do not report any business income. Bootstrapped standard errors in parenthesis. *** significant at the 1% level, ** at the 5% level, * at the 10% level.

Source: author's calculations based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

The estimates in Table 6 indicate that the results are generally robust to the size of the estimation window. Across all three windows, there is almost no evidence of bunching among wage earners, with any statistical significance largely due to the precision of the polynomial fit rather than visible bunching. The evidence of bunching among the self-employed is still found across all three estimation windows, with the exception of the first kink over the widest window. Aside from this single anomaly, for which there is no obvious explanation, the bunching estimates for the self-employed have similar values across all three estimation windows, indicating that the results are largely robust to varying the size of the estimation window.

Finally, Table 7 demonstrates that the results are robust to using different polynomial orders to estimate the counterfactual distribution. All the estimates in the paper use a polynomial of order 7 as in Chetty et al. (2011), but the results are not dependent on this particular value. For the self-employed, the results do not differ much across different polynomial orders (with the exception of a linear fit in column (1), but this is likely to be a very poor choice for the counterfactual distribution). For wage earners, the significance of the estimates varies across the polynomial orders, confirming that any statistically significant bunching estimates are due to the polynomial fit rather than to actual bunching behaviour.

Table 7: Sensitivity of excess bunching estimates to polynomial order

	Polynomial order									
	1 (1)	2 (2)	3 (3)	4 (4)	5 (5)	6 (6)	7 (7)	8 (8)	9 (9)	10 (10)
Panel A: Wage earners										
First kink	-0.532*** (0.194)	-0.081 (0.101)	-0.048 (0.045)	-0.063 (0.050)	-0.057 (0.041)	-0.114*** (0.044)	-0.116*** (0.036)	-0.072* (0.037)	-0.072* (0.041)	-0.078* (0.042)
Second kink	-0.060 (0.050)	0.029 (0.041)	0.015 (0.027)	0.044* (0.027)	0.043 (0.028)	0.071*** (0.027)	0.072*** (0.027)	0.046 (0.028)	0.045 (0.028)	0.040 (0.029)
Third kink	-0.432** (0.172)	0.017 (0.042)	0.031 (0.027)	0.049* (0.027)	0.052** (0.024)	0.040* (0.024)	0.041 (0.025)	0.041 (0.027)	0.041 (0.026)	0.029 (0.027)
Fourth kink	-0.143** (0.063)	0.027 (0.030)	0.032 (0.028)	0.017 (0.029)	0.016 (0.028)	0.029 (0.031)	0.029 (0.029)	0.026 (0.032)	0.026 (0.030)	-0.005 (0.031)
Fifth kink	-0.031 (0.036)	0.051** (0.024)	0.055** (0.021)	0.031 (0.021)	0.031 (0.021)	0.048** (0.023)	0.048** (0.022)	0.030 (0.024)	0.030 (0.022)	0.039 (0.026)
Panel B: Self-employed										
First kink	-0.164 (0.141)	0.186** (0.093)	0.217*** (0.050)	0.117*** (0.037)	0.113*** (0.036)	0.107*** (0.036)	0.108*** (0.037)	0.097** (0.041)	0.097** (0.041)	0.065* (0.038)
Second kink	-0.012 (0.064)	0.147*** (0.043)	0.154*** (0.040)	0.150*** (0.041)	0.152*** (0.041)	0.154*** (0.043)	0.155*** (0.042)	0.175*** (0.049)	0.174*** (0.044)	0.185*** (0.049)
Third kink	-0.010 (0.063)	0.127*** (0.041)	0.132*** (0.041)	0.164*** (0.041)	0.166*** (0.039)	0.151*** (0.046)	0.150*** (0.043)	0.163*** (0.049)	0.163*** (0.046)	0.165*** (0.050)
Fourth kink	-0.007 (0.062)	0.074 (0.057)	0.084 (0.059)	0.086 (0.058)	0.086 (0.055)	0.102 (0.064)	0.102 (0.068)	0.109 (0.072)	0.109 (0.070)	0.066 (0.069)
Fifth kink	0.113* (0.065)	0.167** (0.066)	0.170*** (0.063)	0.155** (0.071)	0.157** (0.067)	0.221*** (0.067)	0.224*** (0.068)	0.172** (0.075)	0.173** (0.074)	0.176** (0.079)

Notes: sample includes all taxpayers aged 15–64 years old over the full sample period, 2011–17. Panel B is the sample of self-employed, defined as those taxpayers who report any business income, while Panel A is the wage earners, defined as those taxpayers who do not report any business income. Bootstrapped standard errors in parenthesis. *** significant at the 1% level, ** at the 5% level, * at the 10% level.

Source: author's calculations based on SARS data in Individual Panel V2018_2 (National Treasury and UNU-WIDER 2019).

7 Conclusion

I find significant evidence of bunching among the self-employed, defined as those taxpayers who report any business income, with almost no evidence of bunching among other taxpayers. This finding of greater responsiveness of the self-employed than wage earners has also been demonstrated in other countries (Bastani and Selin 2014; Chetty et al. 2011; Saez 2010) and could be due to their greater ability to adjust hours of work and shift their income in order to reduce taxable income.

Over my sample period, 2011–17, marginal tax rates were constant through 2015 and were increased in 2016 by one percentage point in all brackets except the lowest one. I find that excess bunching increases in the years after the marginal tax rate changes even though the net-of-tax rate changes remained roughly the same at most of the kink points in the income tax schedule. The observed increases in bunching behaviour among the self-employed over 2016–17 relative to 2011–15 at kink points where incentives were effectively unchanged across the two time periods suggests that the increased responsiveness might be due to greater salience of the kink points after the marginal tax rate changes.

The importance of informational considerations is further highlighted when looking at differential responses across the income distribution. Excess bunching is greatest at the top kink point and lowest at the bottom kink point even though the change in the net-of-tax rate is largest at the bottom kink and smallest at the bottom kink. These patterns suggest that differences in taxpayer responsiveness may be attributable to information rather than incentives.

I find that bunching behaviour is greater among females than among males and decreases as taxpayers get older. This suggests that the tax system in South Africa may differentially impact the income distribution of men and women, and the old and young.

Looking at the anatomy of excess bunching, I find that retirement fund deductions are particularly important for adjusting taxable income. The self-employed use this deduction, which shifts income into future periods, in order to manipulate their taxable income. There is also significant bunching in broad income (taxable income plus all deductions), suggesting that the response also involves a real labour supply response. The retirement fund deduction rules were changed in 2017 and the evidence indicates that these changes enabled the high-income self-employed to adjust their taxable incomes to a greater extent than before.

Despite the significant excess bunching observed, the implied elasticities—under the assumption of a uniform heterogeneity distribution around the kink—are much smaller than those estimated by Kemp (2019) using the bracket creep methodology on similar data. One reason for the difference could be the time horizon of the estimated elasticities since Kemp (2019) estimates an elasticity over a three-year time horizon whereas the bunching estimates are over an unclear time horizon. Further, the accuracy of the bracket creep estimates hinges on there being no bunching behaviour since, if taxpayers adjust their incomes to fall just below the kink point, they are less likely to fall in the treatment group, creating downward bias in these estimates. In this way, my results provide an additional robustness check on the results in Kemp (2019). Since I find significant bunching only among the self-employed, and the self-employed are a small proportion of all taxpayers, my results provide support for the validity of the estimates in Kemp (2019), although those results would be enhanced by excluding the self-employed.

Future extensions of this work could use the kink points in the income tax schedule to estimate the elasticity of taxable income under less-restrictive assumptions than those required by the bunching method, as suggested in recent work by Blomquist et al. (2019) and Bertanha et al. (2019). Further, the introduction of a new highest tax bracket in 2018 provides a large reform that could be studied to estimate the elasticity of taxable income for high-income taxpayers once sufficient data is available. The results in this paper suggest that high earners are likely to be more responsive than lower earners.

References

- Bachas, P., and M. Soto (2018). 'Not(ch) Your Average Tax System: Corporate Taxation Under Weak Enforcement'. Policy Research Working Paper No. 8524. Washington, DC: World Bank.
- Bastani, S., and H. Selin (2014). 'Bunching and Non-bunching at Kink Points of the Swedish Tax Schedule'. *Journal of Public Economics*, 109: 36–49. <https://doi.org/10.1016/j.jpubeco.2013.09.010>.
- Bertanha, M., A.H. McCallum, and N. Seegert (2019). *Better Bunching, Nicer Notching*. <http://dx.doi.org/10.2139/ssrn.3144539>
- Blomquist, S., W. Newey, A. Kumar, and C.-Y. Liang (2019). 'On Bunching and Identification of the Taxable Income Elasticity'. NBER Working Paper 24136. Cambridge, MA: National Bureau of Economic Research.
- Boonzaaier, W., J. Harju, T. Matikka, and J. Pirtilä (2019). 'How Do Small Firms Respond to Tax Schedule Discontinuities? Evidence from South African Tax Registers'. *International Tax and Public Finance*, 26(5): 1104–36. <https://doi.org/10.1007/s10797-019-09550-z>.
- Chetty, R., J.N. Friedman, T. Olsen, and L. Pistaferri (2011). 'Adjustment Costs, Firm Responses, and Micro Vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records'. *The Quarterly Journal of Economics*, 126(2): 749–804. <https://doi.org/10.1093/qje/qjr013>.
- Chetty, R., and E. Saez. (2013). 'Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients'. *American Economic Journal: Applied Economics*, 5(1): 1–31. <http://dx.doi.org/10.1257/app.5.1.1>.
- Ebrahim, A., and C. Axelson (2019). 'The Creation of an Individual Panel Using Administrative Tax Microdata in South Africa'. WIDER Working Paper 2019/27. Helsinki: UNU-WIDER. <https://doi.org/10.35188/UNU-WIDER/2019/661-6>
- Ebrahim, A., and K. Lilenstein. (2019). 'Gender and the South African Labour Market: Policy Relevant Research Possibilities Using South African Tax Data'. WIDER Working Paper 2019/31. Helsinki: UNU-WIDER. <https://doi.org/10.35188/UNU-WIDER/2019/665-4>
- He, D., L. Peng, and X. Wang (2020). 'Understanding the Elasticity of Taxable Income: A Tale of Two Approaches'. Working paper. Available at: <https://www.dropbox.com/s/ms1bdh5z409fssc/ETI.pdf?dl=0> (accessed 26 May 2020).
- Kemp, J.H. (2019). 'The Elasticity of Taxable Income: The Case of South Africa'. *South African Journal of Economics*, 87(4): 417–49. <https://doi.org/10.1111/saje.12232>.
- Kleven, H.J. (2016). 'Bunching'. *Annual Review of Economics*, 8: 435–64. <https://doi.org/10.1146/annurev-economics-080315-015234>.

- Kleven, H.J., and M. Waseem (2013). ‘Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan’. *The Quarterly Journal of Economics*, 128(2): 669–723. <https://doi.org/10.1093/qje/qjt004>.
- National Treasury (2018). *Budget Review*. Pretoria, National Treasury.
- National Treasury and UNU-WIDER (2019). Individual-panel 2011–2018 [dataset]. Version 2018_2. Pretoria: South African Revenue Service [producer of the original data], 2018. Pretoria: National Treasury and UNU-WIDER [producer and distributor of the harmonized dataset], 2019.
- Saez, E. (2003). ‘The Effect of Marginal Tax Rates on Income: A Panel Study of “Bracket Creep”’. *Journal of Public Economics*, 87(5-6): 1231–58.
- Saez, E. (2010). ‘Do Taxpayers Bunch at Kink Points?’. *American Economic Journal: Economic Policy*, 2(3): 180–212. <https://doi.org/10.1257/pol.2.3.180>.
- Saez, E., J. Slemrod, and S.H. Giertz (2012). ‘The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review’. *Journal of Economic Literature*, 50(1): 3–50.
- Weber, C.E. (2014). ‘Toward Obtaining a Consistent Estimate of the Elasticity of Taxable Income Using Difference-in-differences’. *Journal of Public Economics*, 117: 90–103. <https://doi.org/10.1016/j.jpubeco.2014.05.004>.

Data appendix

All analyses in this paper use the Individual Panel V2018_2 data (National Treasury and UNU-WIDER 2019) and were run in November 2019. Specifically, I make use of the following variables from each of the following datasets:

Dataset	Variable name
Income_panel	id_d
	tax_year
	age_d
	taxable_income_d
	deductions_d
Employment_panel	id_d
	tax_year
	gender
Source_code_panel	id_d
	tax_year
	amount
	category_d
	source_code
	final_d

Sample restrictions

The `id_d` and `tax_year` variables are used to merge across the datasets. In the `Income_panel` dataset, there are duplicate pairs of `id_d` and `tax_year` associated with different dates of birth. Since this creates problems with merging with other datasets and they only comprise a very small percentage of observations (0.6 per cent), I drop these duplicates from my sample.

I restrict the sample to those aged 15–64 years old. The typical retirement age in South Africa is 65 years old and individuals over the age of 65 also receive a higher primary income tax rebate, so I exclude this group of taxpayers.

I exclude the data from 2018 since this version of the data did not contain ITR12 certificates in this year.

Variable definitions

I use many variables as they are defined in the dataset, and this section describes only the variables that I created.

Information on gender is only available in the `Employment_panel` dataset, but there are some `id_d` that are associated with different genders across different employee certificates (even within the same year). I only assign a gender (male or female) to those `id_d` that are always reported as either male or female in all employee certificates, and leave gender missing for others. Because this affects a non-trivial number of observations, and because gender is derived from an individual's national identity number and so is missing for individuals with no identity number, for example foreigners (Ebrahim and Lilenstein 2019), I do not omit individuals with missing gender from the main

analysis. Of course, these individuals with indeterminate gender are omitted from the gender differentiated analysis.

I use the Source_code_panel dataset to determine the sample of self-employed taxpayers as well as the values of two specific deduction categories. I categorize a taxpayer as 'self-employed' if they have any business income and use the category_d variable to determine this. Specifically, a taxpayer is classified as self-employed if they have at least one instance of 'Business_income' under category_d in that tax year. I use the source_code variable to group deductions into two main categories—retirement fund deductions and income-related deductions. Retirement fund deductions are those deductions with source_code values of 4001, 4002, 4003 (2017 tax year only), 4004, 4006, 4007, 4018, 4026, and 4029. Income-related deductions are those deductions with source_code values of 4014, 4015, 4016, 4017, 4019, 4027, 4028, 4043, 4045, 4046, 4047, 4048, and 4050.