

Does Cutting the Tax Rate to Zero Induce Behavior Different from Other Tax Cuts? Evidence from Pakistan*

Mazhar Waseem[†]

March 2019

Abstract

Using a series of Pakistani tax reforms and administrative records, I document that taxable income responses induced by *to-zero* tax cuts are orders of magnitude larger than ones induced by similar-sized other cuts. This finding is remarkably robust to alternative specifications and holds for both self-employed and wage-earners. I explore salience, selective enforcement, and discontinuous evasion costs as explanations of the observed behavior. I find that the data favor the last explanation. The difference between the two sets of responses is primarily driven by large, discrete tax evasion response, which is included in the former but not in the latter behavior. I estimate the difference as a lower bound on tax evasion, showing that at least 70% of income of low- and middle-income self-employed and 1% of low-income wage-earners goes unreported.

Keywords: Efficiency, Income Tax, Tax evasion

JEL Classification: H21, H24, H26

*I thank the Editor and three anonymous referees for helpful comments. This research has greatly benefited from suggestions and comments by seminar participants at the University of Manchester, Institute of Fiscal Studies, Oxford University Centre for Business Taxation, VATT Institute for Economic Research, the IMF, the World Bank, the Centre for European Economic Research, and the National Tax Association Annual Conference.

[†]Department of Economics, University of Manchester, Arthur Lewis Building, Oxford Road, Manchester. (mazhar.waseem@manchester.ac.uk)

I Introduction

Important policy questions such as how high the tax rate can be and how wide the tax base needs to be depend critically upon how agents react to tax changes (Feldstein, 1999; Saez, 2004). A rich body of literature leverages changes in the income tax schedule to estimate these reactions (Saez *et al.*, 2012). The changes exploited in this literature, however, are exclusively of the type where the tax rate moves within the positive region. A common feature of income tax systems around the world is that incomes below a given cutoff are not taxed. Upward revisions of the exemption cutoff create tax reforms where the rate moves from a positive value to zero. A priori, agents may not react to these *to-zero* reforms the same way they do to others. Tax evasion offers no tangible benefit when the rate is zero. *To-zero* reforms may be more salient than others. And the authorities may audit zero-rated incomes lightly. If behavior differs substantially across *to-zero* and *not-to-zero* reforms, it would have important policy implications. Yet, there is little work in the existing literature that examines the question either theoretically or empirically.

In this paper, I exploit a series of sharp changes in the Pakistani income tax system to study this question. Pakistan has two income tax schedules, one for the self-employed and one for wage earners. The schedules are not indexed to inflation, and bracket boundaries, in particular the exemption cutoff, need to be moved every few years to avoid bracket creep. During the period considered in this paper (2006–2011),¹ the schedule for self-employed was revised once, in 2010; but the exemption cutoff was moved twice, in 2010 and 2011. Similarly, the wage-earners' schedule was revised once, in 2008; but the exemption cutoff was moved four times, in 2008, 2009, 2010, and 2011. These movements create plausibly exogenous *to-zero* and *not-to-zero* rate changes, which are particularly suited to the requirements of this paper both because they are similar-sized and because they are applied to a similar area of the income distribution.

I use data from the Federal Board of Revenue (FBR) that comprise the universe of income tax returns filed between 2006 and 2011. Using the data, I present nonparametric evidence establishing that the behavioral responses produced by *to-zero* rate changes are orders of magnitude larger than ones produced by similar-sized *not-to-zero* changes. The elasticities underlying the former responses are larger than 15, while those underlying the latter are close to zero. I formalize this result using the difference-in-differences framework, comparing the outcomes across taxpayers affected and not affected by the tax changes. Identification requires that reported earnings of the compared groups would

¹Pakistani tax year runs from July to June. A year t in this paper refers to the tax year from July t to June $t + 1$.

have followed a common trend in the absence of tax changes. I confirm this using both visual and regression-based analysis. I also demonstrate that the result is robust to a series of specification checks.

Why do to-zero reforms produce much larger earnings responses than others? I consider three potential explanations. It could be that the costs of evading certain categories of income are small and of others large. Income entailing little evasion cost would be reported at a zero rate but not otherwise. On the other hand, income entailing large evasion cost would always be reported. Response to a to-zero tax cut would include both categories of income and hence would be larger. This evasion-costs-based explanation generates three testable predictions. First, to-zero responses would be substantially larger than not-to-zero responses. Second, the difference between the two would represent tax evasion. And finally, tax evasion would be nontrivial even at very low rates. The first of these predictions, as I note above, is validated by the data. To test the second prediction, I compare the evolution of easy-to-evade line items on the tax return form with hard-to-evade items. The easy-to-evade items respond much more aggressively than others, demonstrating that the to-zero responses in large part comprise changes in tax evasion and not effort. The evidence validates the third prediction too. I show that tax evasion is large (more than 70% of reported earnings) even when the tax rate is very low (just half a percent). The evasion-costs-based explanation thus fits the observed evidence quite well.

I next explore a salience-based explanation. There is growing evidence that agents do not optimize fully to taxes. Tax schedules are complex, many decision-relevant attributes of taxes are shrouded, and attention is a depletable resource. Together, this implies that agents may not pay full attention to less-salient taxes, underreacting to them (see for example [Chetty *et al.*, 2009](#); [Finkelstein, 2009](#); [Taubinsky & Rees-Jones, 2017](#)). This salience-based explanation can reconcile the observed behavior, meaning that the to-zero responses reflect true and not-to-zero responses attenuated behavior so that the difference between the two represents optimization errors. The evidence, however, does not support this explanation. The to-zero responses are too large to be taken as true responses to a *typical* tax change and the not-to-zero responses are comparable to salience-adjusted, structural responses estimated for the same set of taxpayers using other source of variation ([Kleven & Waseem, 2013](#)).

The final explanation I consider is that the enforcement function may not be neutral across incomes in various brackets. It, for example, may treat zero-rated incomes favorably considering that no tax is payable. But this mechanism is completely absent in the Pakistani setting. The Pakistani tax administration audits around 2-5% of tax returns annually. These returns are selected at random through a publicly held ballot. The audit

and enforcement functions in my empirical setting are therefore independent of income brackets or any other taxpayer trait and cannot explain the observed responses.

Having concluded that the empirical evidence favors the evasion-costs based explanation, I proceed to show that the large difference observed between the observed to-zero and not-to-zero responses identifies a lower bound on tax evasion. The intuition for this result is simple. At a zero tax rate, it is optimal for a taxpayer to report her true income, as evasion offers no pecuniary benefit but still entails costs. As the rate increases marginally above zero, evading the component of income that entails trivial evasion cost becomes optimal. Reported income thus jumps as the rate moves to or away from zero. Because this jump represents the component of income that will not be reported at *any* positive tax rate, it identifies a lower bound on tax evasion. Comparing the to-zero and not-to-zero responses, I estimate this lower bound to be 70% for self-employment income and 1% for wage income, meaning that at least 70% of reported self-employment and 1% of reported wage income is evaded by zero-rated taxpayers in Pakistan. In the most parsimonious formulation of the model, these lower bounds reflect actual evasion rates tightly. But in richer settings the baseline result—true incomes are reported at the zero tax rate—may not hold and the lower bounds may not be tight. I explore three such settings: (1) possibility of downward revision of the exemption cutoff; (2) cross-checks in other tax bases; and (3) threat of future audits. In each case, the evidence suggests that incorporating the richer element of behavior is unlikely to take us too far away from the baseline results.

It is important to emphasize that this tight lower-bound interpretation is primarily relevant to the Pakistani setting only. In order to recover the level of evasion by comparing to-zero and not-to-zero responses, it is crucial that both actual and perceived enforcement functions do not change discontinuously at the point the tax rate rises from zero to a positive value. This requirement, as shown above, is satisfied in the Pakistani setting but may not be satisfied in other settings for one or more of the reasons mentioned above.² One other factor limiting the generalizability of the result is that while the exemption cutoff in most of the developing countries—similar to Pakistan—is located high up in the income distribution (near the 80th percentile), it is located quite low in rich countries (below the 20th percentile). In both these cases (discontinuous enforcement function and low exemption cutoff), the difference between to-zero and not-to-zero responses would continue to recover a lower bound on tax evasion, but this lower bound would not be informative on the actual evasion level in the economy.

²Indeed, there is some evidence from another developing country context that taxpayers are worried about reporting true income even when facing a zero tax rate due both to changes in the audit function and to dynamic enforcement considerations (see [Tourek, 2019](#)).

This paper contributes to a growing literature that uses quasi-experimental variation created by tax reforms to estimate behavioral responses to taxation (see [Saez et al., 2012](#) for a survey). In particular, it adds to a recent strand of this literature that uses administrative microdata to study tax compliance in low-enforcement-capacity environments, emphasizing the role of information in compliance (see for example [Pomeranz, 2015](#); [Best et al., 2015](#); [Naritomi, 2018](#); [Carrillo et al., 2017](#); [Waseem, 2018b,c](#)). Non-incremental, sizable to-zero reforms are frequent in both rich and developing countries,³ and their policy implications are potential serious. Yet, there is very little work in the existing literature that examines the distinction between to-zero and not-to-zero reforms either theoretically or empirically.⁴ This paper fills the gap, documenting how behavior differs substantially depending upon whether the taxpayer faces a zero or positive tax rate.

II Context, Data and Research Design

This section describes important features of the Pakistani income tax system and the research design I use for the empirical analysis.

II.A Context

Like other developing countries, personal income tax is an important and growing source of revenue for Pakistan. Its share in federal tax receipts has been rising steadily in recent years, accounting for roughly 13% of the receipts in 2013 ([FBR, 2014](#)). The tax is collected through two distinct schedules, one each for the self-employed and wage earners. A taxpayer is classified as self-employed (wage earner) if her wage income does not exceed (exceeds) 50% of the taxable income and is then taxed according to the assigned schedule on the entire taxable income. The two schedules, shown in [Figure A.I](#), specify *average* tax rate as a function of taxable income. The Pakistani tax system is quite simple. To calculate tax liability, a taxpayer simply multiplies her taxable income with the rate applicable in the corresponding bracket. The schedules are individual-based, there is no universal deduction other than that earnings below the exemption cutoff are not taxed, itemized deductions such as charitable donations are applied only after the tax liability has been calculated, and there is no system of tax credits or transfers interacting with the

³I have mentioned the Pakistani case above. See [Piketty & Qian \(2009\)](#) for China and India. Since 2010, the UK's exemption cutoff has moved every year, almost doubling from £6,475 in 2010 to £11,850 in 2018.

⁴Besides this paper, one other works that I am aware of which examines taxpayer behavior at a zero tax rate is [Tourek \(2019\)](#).

schedules.⁵

The most important feature of the tax system from the perspective of this paper, however, is that the two schedules are not indexed to inflation and need revision every few years to avoid bracket creep.⁶ During the period considered in this study (2006–2011), the schedule for self-employed was comprehensively revised once, in 2010; but the exemption cutoff was moved twice, in 2010 and 2011. Similarly, the wage-earners' schedule was comprehensively revised once, in 2008; but the exemption cutoff was moved four times, in 2008, 2009, 2010, and 2011.⁷ These reforms create plausibly exogenous tax variation, which for at least two reasons is particularly suited to the requirements of this paper. First, the to-zero and not-to-zero rate changes resulting from the reforms, illustrated in Figure I, are almost of the same size and are applied roughly to a similar area of the income distribution.⁸ Second, as the main motivation behind these reforms was to avoid bracket creep, they are essentially narrow in focus and do not make significant changes to the tax code other than adjusting the bracket boundaries.

One additional advantage of the Pakistani context is that earnings reported at zero tax rate are also observed. Two provisions in the tax code make it possible. First, a provision introduced in 2009 mandates all *registered* taxpayers to file a return even if no tax is payable. Before 2009, another provision in the code required taxpayers to file for period t if income in any of the two previous periods, $t - 1$ and $t - 2$, was above the exemption cutoff. Table A.I assesses compliance with these filing requirements (also see Figure A.II for a nonparametric counterpart of this exercise.). I regress an indicator that a tax filer in period t also files in period $t + 1$ on a dummy indicating if the tax filer experiences a to-zero rate change. The regression is run separately for self-employed and wage-earners, and I also report results from placebo regressions where the rate changes are predated by one year. Overall, around 80% of the self-employed continue to file in the next period, but more importantly this probability does not drop for tax filers whose rate gets reduced to zero. In fact, the placebo exercise shows that such taxpayers are slightly more likely to file a return. This should not be surprising as filing is an easily-verifiable discrete variable, and most tax administrations including the FBR use automated processes to identify and

⁵Pakistan has a small, means-tested income transfer program targeted to extremely poor households. Given, however, that the income tax exemption cutoff is set around the 80th percentile of the income distribution, the sets of taxpayers and transfer recipients do not overlap.

⁶Inflation is generally high in Pakistan and hovered around 10% during the periods considered in this study 2006-11.

⁷All these movements were in the upward direction. In fact, the exemption cutoff has never been revised downward in the history of the country. This creates strong, legitimate expectation that once reduced to zero the tax rate would not be raised back to the positive territory.

⁸For example, the 2008 not-to-zero change and 2011 to-zero change for wage earners are exactly similar other than that the latter reduces the rate to zero whereas the former does not.

penalize nonfiling.

The Pakistani income tax system is based on the principle of self assessment. Returns filed in a tax year are considered final unless they are selected for audit. Audit therefore is the only mechanism through which compliance can be secured. The FBR, like their counterparts in other countries, have limited resources for audit, which means they can audit only a small fraction of returns filed every year.⁹ The Pakistani tax code provides that the selection of returns for audit can only be based on objective criteria. Over the years, the superior courts of the country have narrowed down the definition of *objective* considerably. In fact, the FBR have been finding it difficult to defend any parametric selection criterion as objective. To avoid further litigation on the issue, they have adopted the practice of selecting audit cases randomly through a computer ballot. These ballots are carried out publicly, and the results are displayed on the FBR website. The audit function faced by taxpayers in my sample is, therefore, quite simple: every tax filer faces a small (around 2-5%), exogenous probability of audit. The probability does not change discontinuously at the exemption cutoff, nor does it increase or decrease upon declaring income in the zero-rated region.

II.B Data

I use administrative data from the FBR that include income tax returns filed by the self-employed and wage earners in 2006-2011 and a set of taxpayer characteristics. The tax-return dataset contains variables corresponding to line items on the return form, including a brief profit and loss account, the decomposition of taxable income by source, and tax computations. The taxpayer characteristics dataset contains information captured at the time of registration such as the date of registration, gender, and location of a taxpayer. Appendix A.1, provides a detailed description of the variables used in the empirical analysis.

Table A.II reports descriptive statistics of the data. The analysis sample (columns 3-4) differs from the full sample (columns 1-2) on three dimensions. First, the research design used in this paper is based on panel analysis, comparing within-taxpayer changes in earnings ($\log \frac{z_{it+1}}{z_{it}}$) over time. Consequently, the analysis sample for period t gets restricted to taxpayers for whom $\log \frac{z_{it+1}}{z_{it}}$ is defined. Second, because the main focus of the paper are earnings responses to the movement of the exemption cutoff, I do not include taxpayers who have base period earnings (z_{it}) too far away from the exemption

⁹For example, only 10,271 (2.1%) of the 485,420 nonsalaried returns were audited in the tax year 2010.

cutoff.¹⁰ For the self-employed, the analysis sample, accordingly, includes taxpayers with $z_{it} \in (80k \text{ } 500k]$, which constitutes around 94% of the population. The wage income distribution is more dispersed, and the analysis sample, therefore, includes all taxpayers with $z_{it} \in (140k \text{ } 700k]$, which constitutes around 62% of the population. In one of the robustness checks, I show that the results are not affected if this sample restriction is relaxed. Third, I drop taxpayers from the analysis sample for whom the log change in earnings ($\log \frac{z_{it+1}}{z_{it}}$) is less than the 1st percentile or is in excess of the 99th percentile of the corresponding pooled distribution. Such winsorizing is common in the literature to deal with the extreme outliers (see for example [Gruber & Saez, 2002](#)).

All empirical results in this paper, unless otherwise specified, are based on the analysis sample with the following three categories of taxpayers dropped: (1) female taxpayers because the exemption cutoffs for them are slightly higher than male taxpayers in 2006–09,¹¹ (2) partners in partnership firms as their earnings are subject to a different tax regime ([Waseem, 2018c](#)), (3) taxpayers who switch from self-employed to wage earner and vice versa from concerns that such switching may be endogenous to tax changes.¹² These taxpayers are only a small fraction of the population (rows 4, 5, and 11 of the table), and the empirical results, therefore, are based on more than 96% of the potential analysis sample.

II.C Research Design

I use a simple difference-in-differences research design to estimate earnings responses generated by the to-zero and not-to-zero rate changes. The research design—based on the workhorse empirical model in the tax responsiveness literature (see for example [Saez et al. 2012](#))—leverages the fact that taxpayers in different brackets of the two tax schedules experience differential rate changes over time. It is particularly suited to the Pakistani setting as taxpayers in few brackets undergo no tax change at all and can therefore serve

¹⁰Of course, I do not impose any restriction on z_{it+1} .

¹¹Doing the analysis separately for the two genders is difficult because female taxpayers are less than 3% of the analysis sample (see row 11 of the table).

¹²Table [A.III](#) assesses if the switching probability differs across years. Switching between the two bases is rare, and there are no meaningful differences in switching across years. I also investigate if switchers are concentrated disproportionately around the exemption cutoff. The regression of an indicator variable that a self-employed in period t becomes a wage-earner in period $t + 1$ on a dummy indicating that the taxpayer is located within 50k of the exemption threshold (300k) returns a coefficient of 0.0064 with a standard error of 0.0012. This coefficient is in fact smaller than the one I obtain from the placebo regression, which is run on the prereform periods only (when 300k is not the exemption cutoff).

as a clean control group. I estimate the following model

$$(1) \quad \Delta \log z_{it}^k = \alpha + \mathbf{treat}_i \beta + \mathbf{year}_t \gamma + \mathbf{treat}_i \times \mathbf{post}_t \delta + \mathbf{X}_{it} \mu + u_{it},$$

where $\Delta \log z_{it}^k$ is the log change in income of type k from period t to $t + 1$ for taxpayer i , \mathbf{treat}_i is a vector of two dummies [$to\text{-}zero_i$ $not\text{-}to\text{-}zero_i$] which turn on whenever the corresponding tax change is experienced, \mathbf{year}_t is a vector of year fixed effects, \mathbf{post}_t are dummies indicating the year in which the particular change takes place, and \mathbf{X}_{it} are a set of controls. Given that the identification here comes from the comparison of taxpayers in different areas of the income distribution, the major threat to identification is mean-reversion. I take three steps to rule out this and related concern. First, I provide nonparametric evidence showing that the earnings growth rate ($\Delta \log z_{it}^k$) remains remarkably uniform throughout the income distribution during the periods of no tax change. Second, I also estimate augmented specifications corresponding to (1), where controls for mean-reversion suggested in the literature (see [Saez et al., 2012](#))—the log base period income and a ten-piece spline of log base period income—are included into the model. The results with and without these controls are very similar. Third, I conduct placebo analysis pretending that each reform took place one year prior to its actual implementation. If areas of the income distribution affected by the reforms experience significantly different earnings growth for nontax reasons, it would show up in the placebo regressions.

III Does behavior differ across to-zero and other tax cuts?

In this section, I estimate behavioral responses produced by the Pakistani to-zero and not-to-zero reforms to see if they differ substantially from each other. I begin by presenting nonparametric evidence. The results are then formalized through the regression-based framework. I finally show that the results are robust to a series of specification checks.

III.A Self-employment Income

III.A.1 Nonparametric Evidence

Figure [IIA](#) plots the evolution of self-employment income from 2006 to 2011 for taxpayers classified as self-employed by the tax code. To construct the diagram, I group taxpayers into PKR 20,000 bins on the basis of their base period income (z_{it}^S) and plot the mean log change in income from year t to $t + 1$, $\mathbb{E} [\log \frac{z_{it+1}^S}{z_{it}^S} | z_{it}^S \in b]$, in each bin b . These plots show how self-employment income growth in various areas of the income distribution

responds to the 2010-11 rate changes. Two features of the evidence are noteworthy. First, the growth rate is remarkably stable over time and homogeneous across the income distribution in periods of no tax change. Second, the responses produced by the two types of rate changes are strikingly different from each other: while reported income of taxpayers experiencing the reduction of the rate to zero jumps dramatically, that of taxpayers experiencing a similar-sized not-to-zero rate reduction does not change at all.

Was the dramatic income growth at the bottom of the distribution in 2010-11 *caused* by the reduction of the rate to zero? That it is (1) concentrated precisely in the region between the old and new cutoffs, (2) very strong at the bottom and then tapers off monotonically as we move towards the new cutoffs, and (3) indistinguishable from the prereform level just above the new cutoffs strongly suggests that it was. To further reinforce the causal link, Panel B of the figure looks at the evolution of self-employment income of taxpayers classified as wage earners by the tax code. These taxpayers do not experience the 2010-11 rate changes. Therefore, to the extent that self-employment income is subject to common macro shocks, any nontax factors affecting it in 2010–11 could be detected here. However, in sharp contrast to Panel A all curves in Panel B are tight to each other, establishing that the 2010–11 responses are indeed driven by the tax changes. Panels C-F, which are the difference-in-differences versions of Panels A-B, formalizes this conclusion by demonstrating that (1) there are no significant preexisting differences across the two groups (Panel C); (2) the reduction of the rate to zero causes a sharp jump in income reported by the treated group (Panels D and E); and (3) the additional reported income is as large as 70% of the base period income at the bottom (Panel F). Figure A.III show that the results are indistinguishable if I restrict the sample to a balanced panel of taxpayers who file in all six years 2006-2011.

III.A.2 Regression-based Estimates

Table I reports the results from equation (1). The first column contains the estimates from the baseline specification, while the rest of the columns add additional control variables or experiment the replacement of year fixed effects with a parametric time trend. To address any concerns from a change in the composition of the sample, I restrict it to a balanced panel of taxpayers who file every year included in the estimation. Table A.IV drops this restriction, replicating the results for the complete panel. The two sets of results are indistinguishable. Two findings emerge from the analysis. First, consistent with the visual evidence the to-zero changes generate extremely strong responses. The corresponding estimates are always large, statistically significant, and remarkable robust to alternative specifications. Column (1) of the table, for example, shows that the reduction of the rate

to zero causes a 27 log-points additional income growth in the treatment group in the first year after the reform. This is around 5 times larger than the prereform average of 5 log-points per year. Considering that the average net-of-tax rate change behind the response is only 1.7 log-points, the estimate translates into an enormous elasticity of greater than 15. Second, the similar-sized, not-to-zero tax cut generates no response at all. The corresponding estimate is always of opposite sign, small, and statistically insignificant in all but one specification.

III.A.3 Robustness

Given the difference-in-differences research design, the key identification requirement in this setup is of parallel trends: reported earnings of the compared groups (tax brackets that experience and do not experience the tax rate changes) must have followed the same path in the absence of the rate changes. The nonparametric evidence in Figure II illustrates that this is indeed true. The earnings growth rate remains remarkably stable and uniform over time and across the income distribution during the periods of no-tax-change. Figure A.IV provides an aggregate counterpart to this result. It shows that the average earnings growth was very similar in the prereform years for both groups but surged sharply in the treated group exactly at the time of the reform. In fact, the preexisting earnings trends were so flat and stable, that the time-series estimates, reported in Table A.V, are indistinguishable from the corresponding difference-in-differences estimates. Tables I and A.IV-A.V further confirm that the results are insensitive to (1) including additional control variables; (2) replacing year fixed effects with the parametric time trend; (3) adding a full set of year, industry, region, industry \times year, and region \times year fixed effects; and (4) keeping the composition of the estimation sample fixed.

Table A.VI conducts an additional set of robustness checks. Column (2) drops taxpayers who *bunch* at the notches in the baseline tax schedule from concerns that their reported income might be affected by the strong, local incentives created by the notches or that these taxpayers might be special. Column (3) drops taxpayers around the income-composition notch, where the classification of a taxpayer switches from self-employed to wage-earner and vice versa. Columns (4)–(5) increase the range of the data from $z_{it} \in (80k\ 500k]$ in the baseline results to $z_{it} \in (0\ 500k]$ in column (4) and $z_{it} > 0$ in column (5). Columns (6)–(9) add additional control variables into specification (1). Reassuringly, the results from all these alternative specifications are very similar to the baseline results.

Another common concern in the tax responsiveness studies is mean-reversion. I take two steps to alleviate this concern. First, I test if mean reversion is a significant problem in this setting by estimating equation (1) on the prereform periods only. Table A.VII

shows the results. The coefficient on the two interaction dummies are extremely small and statistically indistinguishable from zero in all but one specification. This demonstrates that the relative difference in reported income from one year to the next does not change significantly across the treated and untreated groups for any nontax reason including mean-reversion. I complement this test with another strategy (see Table A.VIII). I reestimate equation (1) after including the standard controls for mean-reversion i.e. log of base period income and a ten-piece spline of log base period income (Saez *et al.*, 2012). I obtain very similar results with and without these controls, reinforcing the conclusion that mean-reversion is not a significant concern in this setting.¹³

Finally, I rule out one alternative explanation of the observed behavior. Suppose that agents do not like to pay taxes (or there are some fixed costs of actually making the payments), but they also do not like to cheat more than they have to. In this setting, agents would evade only up to the point of zero tax liability, reporting income just below the exemption cutoff. They would simply move to the new cutoff after a reform that increases the cutoff. This alternative model generates two testable predictions: bunching would be stronger at the exemption cutoff than a similar other notch, and it would shift to the new cutoff after the reform. Figure A.V tests the latter prediction. It compares the earnings growth rate needed to hit the new exemption cutoff with the actual earnings growth rate observed in the data. For example, earnings of a taxpayer bunching at the baseline exemption cutoff of PKR 100k have to increase by 200% if it, in accordance with this explanation, simply moves to the new exemption cutoff of PKR 300k after the reform. The data clearly reject this alternative explanation. The earnings growth rate around the old cutoff is roughly one-third of the rate needed to hit the new cutoff. In fact, the two curves do not coincide in any area of the income distribution. The first prediction of this alternative model is also rejected by the data: the bunching at the exemption cutoff is not significantly larger than the one at similar other notches.¹⁴

¹³In addition to these, a working paper version of this paper (Waseem, 2018a) carries out further robustness tests. These include, inter alia, reporting estimates from two variants of equation (1), where I use wage-earners with positive self-employment income as the control group. These alternative research designs compare the self-employment income of taxpayers classified as self-employed by the tax code with the self-employment income of taxpayers classified as wage-earners by the tax code, who do not undergo the tax rate changes in 2010-11. These research designs allow additional set of robustness checks, including built-in tests for the parallel trends assumption. The results from these double- and triple-difference specifications are strictly consistent with those in this version of the paper.

¹⁴In a working-paper version, Kleven & Waseem (2013) find that the elasticity implied by the bunching at the baseline exemption cutoff is 0.077. In comparison, the elasticity implied by the bunching at the next three notches is 0.097, 0.083, and 0.091 (please see Table 1 of May 2011 version of the paper). The exemption cutoff during these years was at PKR 100k and the next three notches were at PKR 110k, 125k and 150k. The notches were also of a similar size, involving a jump in average tax rate of 0.5 (for the first two) and 1 (for the others) percentage points.

III.B Wage Income

III.B.1 Nonparametric Evidence

Figure IIIA-B shows the evolution of wage income from 2006 to 2011 for taxpayers classified as wage earners by the tax code. The diagram is constructed analogously to Figure IIA-B and plots the growth of wage income from period t to $t + 1$ as a function of the base period income $\mathbb{E}[\log \frac{z_{it+1}^W}{z_{it}^W} | z_{it}^W \in b]$, where b are bins of PKR 20K. The comparison of Figures III and II shows that the growth of wage income, in distinction to that of self-employment income, is not homogeneous across years. Because of this, it is hard to differentiate between the tax-induced behavior and yearly shocks in the simple plot of raw data.

I follow a simple strategy to obtain a first-pass evidence on the tax-induced behavior, regressing the log change in wage income from period t to $t + 1$ on a full set of year fixed effects. The residuals from the regression are then regressed on four yearly dummies, one each for 2007 to 2010. These later regressions are run separately in the PKR 20k bins, and the estimated coefficients and 95% confidence intervals are plotted in Panels C-F. Clearly, once a common year effect is partialled out, the residual income growth is homogeneous across years and over the income distribution. It, however, spikes in the areas of the distribution where the rate was brought to zero. Though this spike is not as prominent as self-employment income's, the overall pattern is consistent with the earlier result that to-zero and not-to-zero reforms elicit substantially different behavior. The next section formalizes this analysis by presenting the regression-based estimates.

III.B.2 Regression-based Estimates

Table II reports the results from equation (1). I begin with the baseline specification in column (1) and then successively add more control variables, permuting among the combinations of controls for mean-reversion—log base period income and a ten-piece spline of log base period income—and other controls in the rest of the columns. To test the adequacy of the mean-reversion controls, panel B reports the estimates from placebo regressions, where I pretend that all rate changes took place one year earlier than they actually did. Table A.IX runs additional robustness checks.

The main findings are the following. First, the to-zero coefficient is always economically meaningful, statistically significant, and considerably stable across specifications. Given that the rate changes underlying the response are extremely small (always less than one percent, Figure IC-F), the estimate translates into a huge elasticity of more than three. This elasticity is orders of magnitude larger than the one, 0.04, estimated by Kleven &

Waseem (2013) for wage earners in Pakistan. Second, despite the fact that the not-to-zero changes are on average twice the size of the to-zero changes, the earnings responses generated by them are statistically and economically insignificant. Third, the placebo coefficient corresponding to the to-zero estimate is always trivial and statistically insignificant in most of the specifications. This shows that mean-reversion is not much of a concern in this setting and that the base-period income controls are able to account for it adequately.

IV Why does behavior differ across to-zero and other tax cuts?

The above evidence convincingly shows that reported income responses produced by to-zero tax cuts differ substantially from those produced by not-to-zero tax cuts. In this section, I explore three potential explanations of the finding.

IV.A Reverse-L Shaped Evasion Costs

The standard way to think about evasion costs is that they are expected fine and penalty payments which would be recovered in case the evasion is detected by the government (Allingham & Sandmo, 1972; Slemrod, 2001). Recent empirical evidence shows that the probability that evasion gets detected is quite high if reported income is covered by the third-party information the government obtains from sources such as employers and financial institutions and quite low otherwise (see Slemrod, Forthcoming for a recent survey of the evidence).¹⁵ In Appendix A.2, I propose a simple model that incorporates this government's information problem into the standard model of tax reporting under imperfect enforcement (see for example Chetty, 2009). The model assumes that \underline{e} units of income of the agent are not covered by any third-party information, while the rest are covered. Note that this assumption is without any loss of generality as \underline{e} potentially varies from 0 to the maximum. The assumption means that the agent faces a reverse-L shaped evasion costs function, as shown in Figure IV.A. The costs of evasion are low at the bottom because earning or consumption of income up to \underline{e} leaves no verifiable information trails and therefore its evasion entails little detection probability. The costs turn sharply afterward once the third-party reported units of income begin. An optimizing agent facing such evasion costs would report true income when the tax rate is zero but would evade \underline{e}

¹⁵Throughout the paper, I maintain the assumption that the third-party reports the government receives are complete and not themselves distorted by evasion. Otherwise, they would not deter tax evasion as effectively (see for example Brockmeyer & Hernandez, 2017).

as the rate increases marginally above zero, producing a discontinuous earnings supply function of the form shown in Figure IVB.

The discontinuity in the earnings supply function can explain the large, substantive difference observed between the two types of responses. The intuition for this result is provided in Figure IVB, which considers the effects of two equal-sized rate cuts on income reported by the agent. The response triggered by the to-zero cut $\Delta z_A(\tau_A \rightarrow 0)$ is considerably larger, as it consists of both the discrete change in tax evasion (movement along the horizontal axis) and the continuous change in reported income (movement along the supply curve). By contrast, the response induced by the not-to-zero tax cut $\Delta z_B(\tau_B \rightarrow \tau'_B)$ is smaller, as it consists of the latter component only. The model thus fits the observed pattern of response quite well. To probe this point further, I take the other two predictions of the model to the data.

If the difference between the two sets of response reflects evasion costs then it must be that the to-zero response largely arises from a changes in tax evasion and not effort. Figure A.VII tests this prediction of the model. I look at how individual line items on the tax return form react to the two types of rate cuts. The idea behind the exercise is to see which factor—adjustments in tax evasion or effort—drives the larger response to the to-zero tax cuts. The six items considered here form the profit and loss account of a taxpayer, and while all of them are expected to increase with effort, some can be misreported easily than others. My focus here is to identify any differential response between the easy- and hard-to-misreport items. Each panel of the figure plots the mean log change in the line item from period t to $t + 1$ as a function of the self-employment income in period t . Since the sets of taxpayers in various bins here are the same as in Figure IIA, the analysis should be seen as the decomposition of the response depicted there. Figures A.VIII and A.IX formalize this analysis, showing the difference-in-differences version of these plots. Clearly, the line items do not respond uniformly: annual sales and costs respond aggressively, profit and loss expenses (in part third-party-reported and therefore harder to misreport) respond moderately,¹⁶ and imports do not respond at all. Of all the items, imports is perhaps the hardest to misreport because such misreporting can easily be detected through the Customs and Excise records. Its non-responsiveness, therefore, provides the cleanest evidence that the large jump in reported earnings is driven by a drop in tax evasion. Panels E–F strengthen this conclusion. A surge in real activity triggered by an unanticipated decrease in taxes is likely to result in the running down of inventory. Contrary to this,

¹⁶Profit and loss expenses are input costs such as wages, rents, accounting and legal fees, electricity, and interest paid on loans. Although these costs can be over-reported, it is difficult to do so considering that these can potentially be verified at the time of audit.

inventories at the end of 2010 and 2011 rise sharply.¹⁷ Thus overall behavior of line items is consistent with a tax evasion explanation of the observed behavior. Had the growth of self-employment income been a result of an increase in effort, all line items would have responded uniformly. Instead, easy-to-misreport items respond more aggressively than others.

The third prediction of the model is that tax evasion would be high even at a very low tax rate. This prediction simply reflects that the cost of evading income up to \underline{e} is quite low and therefore not reporting this component of income becomes optimal even at a very low tax rate (this can be seen from Figure IV). The evidence presented above is consistent with this prediction of the model as well. Specifically, Figure IIA shows that the reported earnings of taxpayers with baseline income in the range (100k, 110k] on average rise by around 70% as their tax rate reduces from 0.5% to 0%. This demonstrates that consistent with the structure in the model a large component of tax evasion is *fixed* in nature: around 70% of reported income is evaded even when the tax rate is as low as 0.5%.

IV.B Salience

There is growing evidence in literature that agents do not optimize fully to taxes. Tax schedules are complex, many decision-relevant attributes of taxes are shrouded, and attention is a depletable resource. Together, this implies that agents may not pay full attention to less-salient taxes, underreacting to them (see for example Chetty *et al.*, 2009; Finkelstein, 2009; Taubinsky & Rees-Jones, 2017). Salience can also explain the behavior documented above. It would mean that the to-zero tax cuts are fully salient and hence elicit true behavior, while the not-to-zero cuts are less salient and hence elicit attenuated behavior. Representing true earnings response by $\hat{\Delta}z$ and observed response by Δz , this statement can be translated into two testable conditions

$$(2) \quad \begin{aligned} \hat{\Delta}z_A(\tau_A \rightarrow 0) &\approx \hat{\Delta}z_B(\tau_B \rightarrow \tau'_B) \\ \Delta z_B(\tau_B \rightarrow \tau'_B) &= \theta \cdot \hat{\Delta}z_B(\tau_B \rightarrow \tau'_B). \end{aligned}$$

¹⁷The Pakistani tax cuts of 2010 were announced on June 6, 2010 but took effect from the beginning of the new financial year *i.e.* July 1, 2010. This gave taxpayers a window of around three weeks to plan for the next year. Forward-looking taxpayers would have beefed up inventories had they anticipated a large increase in output in the coming year. But this is not what we observe. The start-of-year inventories in 2010 are exactly similar to those in 2009 (Panel E shows no growth in 2009). If start-of-year inventories do not increase and production increases enormously during the year, the end-of-year inventories must go down as firms have limited resources to invest. But again, this is not what we observe. The end-of-year inventories in fact go up significantly (Panel F shows significant growth in 2009).

True responses of both types are nearly equal,¹⁸ but the observed not-to-zero response is attenuated by a factor $\theta \in [0, 1]$. In the extreme case, agents completely ignore a not-to-zero rate change $\theta = 0$ so that the difference between the two captures optimization error only.

Figure A.VI assesses the first of these conditions. It replicates Figures IIA but plots the elasticity of taxable income on the vertical axis in place of the earnings response. The elasticity implied by the to-zero response is implausibly large (>60 at the bottom). It is in fact orders of magnitude larger than the corresponding estimate reported in the existing literature for both developed and developing countries (see Chetty, 2012 for the meta analysis and Kleven & Waseem, 2013 for Pakistan). It therefore cannot represent the true responsiveness to a *typical* rate change. To assess the second condition, I compare the observed not-to-zero response with the *structural* elasticity estimated in Kleven & Waseem (2013). The structural elasticity reflects true responsiveness—free of all frictions including salience—of the same set of Pakistani taxpayers. For taxpayers with income in the range 400-500k, it is estimated to be 0.06, a response not different from the one observed here. Both conditions in (2) therefore lack empirical support, suggesting that salience is not the first-order mechanism driving the large difference between the two behaviors. I establish this further in section V of the paper, where I show that the difference remains unchanged even if the not-to-zero response is corrected of the salience bias by assuming an extreme value of the attenuation factor θ , such as 0.1.

IV.C Enforcement Environment

If the audit function has a discontinuity at the exemption cutoff, it may make an agent's decision of how much to report contingent upon the tax rate (zero or positive). We have seen in section II.A that it is not the case in Pakistan. The audit probability faced by a tax filer in the country is independent of the tax rate, income, or any other characteristic. But even more generally, a discontinuity in the audit function at the exemption cutoff is unlikely in the Pakistani setting. Figure II illustrates that taxpayers just below the baseline exemption cutoff of PKR 100k respond to the to-zero tax cut exactly similar to ones just above the exemption cutoff. The first two bins of the figure contain taxpayers in the income range (80k, 100k] and (100k, 120k] respectively, and the average earnings response is virtually the same in both bins. To the extent that these responses capture tax evasion, the evidence thus shows that the evaded amount does not change discontinuously at the

¹⁸Note that I assume throughout this section that both types of rate changes are roughly equal in magnitude and are applied to very similar taxpayers.

exemption cutoff. Being so, auditing taxpayer below the exemption cutoff less (or more) aggressively would not have been optimal.¹⁹

V Lower Bound on Tax Evasion

Figure IVB shows that the difference in responses induced by to-zero and not-to-zero tax cuts identifies e . This result is formally shown in Appendix A.2 (please see equation 15). Because e represents the component of income earning or consumption of which leaves no verifiable information trail and evasion of which is therefore optimal even at a very low tax rate, it represents a lower bound on tax evasion. In this section, I use formula (15) to estimate this lower bound for the zero-rated Pakistani taxpayers. Table III shows the results for the self-employed. I divide the region below the new exemption cutoff into six segments and present estimates separately for each segment. The columns of the table correspond to terms in the formula: column (1) to the income segment; column (2) to the earnings response produced by the to-zero change; column (3) to the earnings response produced by the equal-sized not-to-zero change; column (4) to the difference between the two responses; and column (5) to the average evasion rate in the segment.

To compute the estimates in column (2), I follow three steps. I first replace the two *to-zero* \times *post* dummies in equation (1) with twelve *segment* \times *post* dummies to estimate the self employment income response separately in each segment in 2010 and 2011. I next add the two yearly estimates for each segment to compute a medium-run estimate of the response in the segment. And finally, I multiply the estimate from the second step with the average income in the segment to convert it into rupees. The estimates in column (3) are computed using the taxable income elasticities reported in Kleven & Waseem (2013).²⁰ Using these elasticities instead of the not-to-zero estimates in Table I is preferable for two reasons. First, these elasticities are estimated from the same area of the income distribution where the 2010-2011 to-zero rate changes were applied to, making the comparison between the two more like-for-like.²¹ Second, Kleven & Waseem (2013) are able to ac-

¹⁹In an optimal plan, a revenue-maximizing tax authority would allocate audit resources towards cases where the enforcement action is likely to be most productive. This means that among similar taxpayers ones evading the most would be targeted. This rule is unlikely to generate discontinuous enforcement given that there is no difference in amount evaded on both sides of the cutoff.

²⁰In computing the estimates for column (3), I keep the magnitude of the tax rate change the same as for column (2). More specific, for a segment k I convert the elasticity ε_k into earnings response Δz_k using the formula $\Delta z_k = \varepsilon_k \cdot \bar{z}_k \cdot \hat{\Delta}(1 - \tau_k)$, where \bar{z}_k is average income in the segment, and $\hat{\Delta}(1 - \tau_k)$ is the proportional net-of-tax rate change experienced by taxpayers in the segment because of the reform.

²¹Kleven & Waseem (2013) use notches in the baseline Pakistani tax system (2006-2009) to identify these elasticities. These notches were in the same area of the income distribution where the 2010 and 2011 to-zero reforms were applied. Column (2) and (3) thus compare how very similar taxpayers react to equal-sized

count for all frictions including inattention nonparametrically. Using the *structural* elasticities estimated by them mitigates any salience-related concerns in my lower bounds. More specifically, Kleven & Waseem (2013) adjust their estimates of observed elasticities to compute structural elasticities that are free from salience and other related frictions. This adjustment is based on the attenuation factor they estimate from the variation created by the tax notches. The attenuation factor estimated by them for the zero-rated self-employed ranges between 0.3-0.5. The estimates in column (3) are therefore already inflated by a factor of 2-3 to alleviate any frictions including inattention. To establish that salience plays little role in my setting, I replicate Table III assuming an extreme value of 0.1 for the attenuation factor θ . The results remain virtually unaffected (compare Tables III and A.X). This should not be surprising given that the large difference between the to-zero and not-to-zero responses documented here reflects not that the latter responses are too small but rather that the former responses are too large.

Two features of the results need emphasizing. First, tax evasion is large even at very low rates, such as 0.5% at the bottom. This behavior, as I note above, is consistent with the predictions of the model, illustrating that evasion jumps to $\underline{e} \gg 0$ whenever the rate increases marginally above zero. Second, the evasion rate is roughly constant up to the income of PKR 150k and then declines monotonically. This pattern of response is consistent with an extended version of the model where the post-reform tax schedule—similar to the Pakistani setting—takes a nonlinear form, containing a notch at the earnings level z_N (see Appendix A.2.3 for details). With the nonlinear tax schedule, earnings responses of taxpayers close to the new cutoff (notch) are constrained, as reporting true income would take them into the positive-rate region. Formally, it is not optimal for a taxpayer experiencing the reduction of the rate from $\tau > \underline{\tau}$ to $\tau = 0$ to report true income as long as $z(\tau) + \underline{e} \geq z_N$, where z_N is the new cutoff. There, thus, exists an interval $[z_N - \underline{e}, z_N]$ below the new cutoff where evasion does not approach zero even when the rate drops to zero. This interval is larger if taxpayers have dynamic considerations so that they try to keep not only their current but also future income below the cutoff. A monotonically declining response, therefore, arises naturally in this model and means that the unconstrained evasion rate is observed only at the bottom where taxpayers are too far away from the new cutoff to be influenced by it. On the basis of this consideration, I conclude that the lower-bound on the evasion rate of zero-rated self-employed, as implied by the first four rows of the table, is around 70%.

Table II shows that wage income also behaves according to the predictions of the discontinuous evasion costs based model. Of course, its response is much smaller than the

to-zero and a not-to-zero rate changes.

self-employment income’s because it is third-party reported. But to the extent that formula (15) holds for wage income as well, the difference between the to-zero and not-to-zero responses identifies the lower bound on the evasion of wage income. The extent of such evasion, however, is small at around 1% of the reported income.

V.A Applicability

The Pakistani to-zero reforms bite deep into the self-employment income distribution. Figure A.X shows this formally by superimposing the baseline CDF of self-employment income on Figure III.F. The exercise shows that the minimum evasion rate I report in the above section applies to more than 50% of the population of self-employed tax filers (Panel A). Note, however, that although broadly applicable, this estimate is essentially local in nature. To claim global applicability, I need to assume that the evasion technology available to top-income taxpayers is the same as the one available to low- and middle-income taxpayers. Or in other words top-income taxpayers face the same evasion costs as do low- and middle-income taxpayers. It is, however, a strong assumption²² and the one I cannot test, as the to-zero reforms do not extend farther to the right of the distribution. The estimates in Table III therefore apply to both low- and middle-income taxpayers but not to top-income taxpayers. The evasion rate of wage income, on the other hand, applies to the bottom 20% of taxpayers only (see Panel B).

V.B Heterogeneity

To explore heterogeneity in the self-employed evasion rate, I estimate the following triple-difference version of equation (1)

$$\begin{aligned}
 (3) \quad \Delta \log z_{it}^S &= \alpha_0 + \alpha_1 \text{to-zero}_{it} + \alpha_2 \text{trait}_i + \text{year}_t \gamma \\
 &+ \text{to-zero}_{it} \times \text{post}_t \delta + \beta_1 \text{to-zero}_{it} \times \text{trait}_i + \beta_2 \text{trait}_i \times \text{post}_t \\
 &+ \text{to-zero}_{it} \times \text{trait}_i \times \text{post}_t \eta + \mathbf{X}_{it} \boldsymbol{\mu} + u_{it},
 \end{aligned}$$

where trait_i is a taxpayer characteristic indicator. To avoid making strong functional-form assumptions, all traits are introduced into the equation nonparametrically. The coefficients on the triple-interaction dummies capture differential responses of taxpayers with

²²High-income taxpayers may have access to more sophisticated methods of evasion (they can hire accountants, move income offshore). On the other hand, they may also be subject to stricter enforcement (higher proportion of digital/third-party-reported transactions). Thus a priori it is not clear if the evasion rate amongst them would be higher or lower than the low- and middle-income taxpayers.

the given trait j . Since traits are not randomly assigned, one problem with this specification is that the estimates might simply reflect that taxpayers with different traits are located in different areas of the income distribution.²³ To make the comparison more meaningful, I also show results from an alternative approach, where I group taxpayers into bins of PKR 20k on the basis of their base period income and run the above regression separately in each bin. I then generate aggregate estimates as the weighted average of the bin-level estimates, with the weights provided by the distribution of trait j in the binned income distribution. This approach is similar to matching and compares taxpayers of different traits by matching them on the basis of their base period income. Figure A.XI and Tables A.XI-A.XIII show the results. Clearly, once the position of a taxpayer in the income distribution is controlled for nonparametrically, the response does not vary much across small vs. large taxpayers (Table A.XI), manufacturers vs. non-manufacturers, regular tax filers vs. irregular tax filers, VAT-registered taxpayers vs. other taxpayers, electronic return filers vs. manual return filers (Table A.XII), and young taxpayers vs. old taxpayers (A.XIII). The evidence thus suggests that within the population of self-employed there is little variation in the the extent of evasion.

V.C How tight is the lower bound?

In the baseline model, taxpayers report their true income as the tax rate approaches zero: evasion offers no benefit at the zero rate but entails a strictly positive cost. Although quite intuitive, the assumption might not be satisfied in richer settings. In this section, I take up three such settings, examining in each case if the extension can cause a significant departure from the baseline result.

Downward revision of the exemption cutoff.—Taxpayers may fear that future reforms could lead to a strictly positive rate in the currently zero-rated region, thereby not revealing their true income as the rate approaches zero. While such downward revision of the exemption cutoff is possible and does occur in some countries, it is not common. In the seventy years history of Pakistan, the cutoff has never been revised downward, although its upward revisions are quite frequent. Piketty & Qian (2009) study the income tax exemption cutoff and its consequences in India and China. Between 1986 and 2008, the cutoff increased from Rs. 18,000 to 150,000 in India and from Yuan 10,000 to 20,000 in China, and never was it revised downward in any of the two countries. Similarly, the exemption cutoff in the UK—known as the Personal Allowance—has increased from £1,165 in 1979-80 to £11,500

²³This is especially problematic in the current setup because the responses, as shown in Figure IIA, decline monotonically along the income distribution.

in 2017 and likewise has never been revised downward. There is thus some evidence that the downward revision of the exemption cutoff is a relatively rare event. Note, however, that even if the exemption threshold is nominally rigid, the *real* exemption threshold may move downward if the nominal threshold is under-indexed to inflation and economic growth (see the comparison of China and India in [Piketty & Qian, 2009](#)). It is one of the reasons that I estimate my lower bound on the evasion rate from the behavior of taxpayers who are reasonably away from the nominal threshold (see the discussion at the end of this section).

Cross-check in other bases.—Reporting true income in one tax base may cause taxpayers difficulties in other tax bases. Pakistan raises almost all its revenue through five tax instruments: Income Tax, VAT, Customs Duty, Excise Duty, and Petroleum Levy. Of these instruments, VAT has the greatest overlap with income tax. If the difficulty of cross-matching declared outcomes in other tax-bases is really important, the to-zero responses of VAT-registered taxpayers must be substantially smaller than taxpayers who are not registered for VAT. But the responses of the two sets of taxpayers—VAT-registered and others—are not significantly different from each other (see [Figure A.XIE](#) and [Table A.XII](#) for the evidence). That the possibility of cross-verification across multiple bases does not worry taxpayers too much should not be surprising, as the evidence from other contexts also shows that when taxpayers report a higher base in one tax, they can leave the base reported in an overlapping tax unchanged by adjusting along some other margin (see for example [Carrillo *et al.*, 2017](#); [Waseem, 2018c](#)).

Threat of future audits.—Taxpayers may not report their true income even at a zero tax rate if they fear that it would make them more likely to face an audit or that it would be used against them in future audits. Note, however, that the first of these two considerations is entirely absent in the Pakistani setting. As explained in [section II.A](#), the probability of audit faced by Pakistani tax filers is small and exogenous. The probability does not increase or decrease with a change in reported income. On the second of these considerations, note that the ability of an audit to assess the legitimacy of a change in the reported income categorized \underline{e} in this paper is limited. The earning or consumption of this component of income leaves no verifiable information trail, making it extremely difficult for an auditor to claim, much less prove, that the change in income is not legitimate.

The above discussion shows that none of the above three extension is likely to push us too far away from the baseline setting. The lower bounds reported in [Table III](#) therefore must be close to the actual evasion rate of these taxpayers. One argument reinforcing

the conclusion is that any disadvantage of revealing true income, even if it exists, must diminish as we move left within the zero-rated region. It is because all the forces mentioned above, and similar other forces, must decline as we move toward the lower end of the distribution.²⁴ Empirical evidence is consistent with this argument. As we move left of the new exemption cutoff, the to-zero responses initially grow stronger. They reach their peak at PKR 150k and remain static thereafter (see Figure II). This pattern provides perhaps the strongest evidence that the responses of taxpayers with baseline income up to PKR 150k are free from the above considerations and therefore capture the amount evaded by these taxpayers reasonably.

VI Conclusion

I leverage a series of sharp changes in the Pakistani income tax schedule to document that behavioral responses induced by to-zero tax reforms—reforms that move the rate to or away from zero—are orders of magnitude larger than those induced by others. This result is remarkably robust across specifications and is true for both self-employed and wage-earners, although as expected the responses of wage-earners are much smaller. I explore three explanations of the observed behavior. It can be argued that taxpayers pay full attention when the rate reduces to zero but not otherwise. While a salience-based model can explain the weak and insignificant not-to-zero responses, it cannot explain the large to-zero responses documented in this paper (taxable income elasticity as large as 60). I also reject discontinuous enforcement—zero-rated incomes receive lighter enforcement—as an explanation of the result. The likelihood of audit faced by Pakistani income tax filers is small and independent of the tax rate they face or any other of their trait. I conclude that the divergence in behavior is driven primarily by tax evasion. The costs of not reporting some categories of income are small and of others large. Income entailing little evasion cost is reported at the zero rate but not otherwise. On the other hand, income entailing large evasion cost is always reported. Responses to to-zero tax reforms are larger because they include both categories of reported income. In the evasion cost based model, the difference between the to-zero and not-to-zero responses provides a lower bound on tax evasion. Exploiting the Pakistani tax reforms, I estimate that at least 70% of self-employment and 1% of wage income is not reported by the low- and middle-income taxpayers of the country.

²⁴Taxpayers at the lower end of the distribution become increasingly less likely to (i) face a positive tax rate in future (the first point here); (ii) be registered for other taxes (the second point here); or (iii) face a heightened risk of audit (the third and fourth point here).

The reporting behavior I uncover has important policy implications. First, the elasticity of taxable income is an important parameter for key tax policy choices such as setting the optimal tax rates. The existing literature estimates this elasticity from not-to-zero reforms only. I show that the most important reaction to taxes takes place as the rate moves from zero to a small, positive value. Ignoring this can result in suboptimal policy choices. For example, the shape of the optimal income tax schedule—especially at the bottom of the income distribution—that takes these reactions into account would be substantially different from the one that does not. Second, I find that a substantial proportion of taxable income goes unreported. This suggests that the return from investment in the enforcement capacity could be large (Besley & Persson, 2013). Research shows that altering the tax mix (Best *et al.*, 2015), redesigning the incentives of tax collectors (Khan *et al.*, 2016), promoting documentation (Naritomi, 2018), strengthening traditional enforcement methods (Almunia & Lopez-Rodriguez, 2018; Waseem, 2018b), and priming social and psychological factors (Slemrod *et al.*, 2018) can promote tax compliance. Exploiting such measures is even more desirable in developing economies where the third-party information—which drives tax compliance in rich countries—is limited in both its scope and effectiveness (Jensen 2019; Carrillo *et al.* 2017; Waseem 2018b). Third, I show that tax evasion lowers the effective tax rate faced by the low- and middle-income self-employed to one-half the statutory rate. This creates significant horizontal equity concerns across self-employed and wage-earners, which need to be addressed to build trust in the tax system and to avoid misallocation of physical and human capital.

References

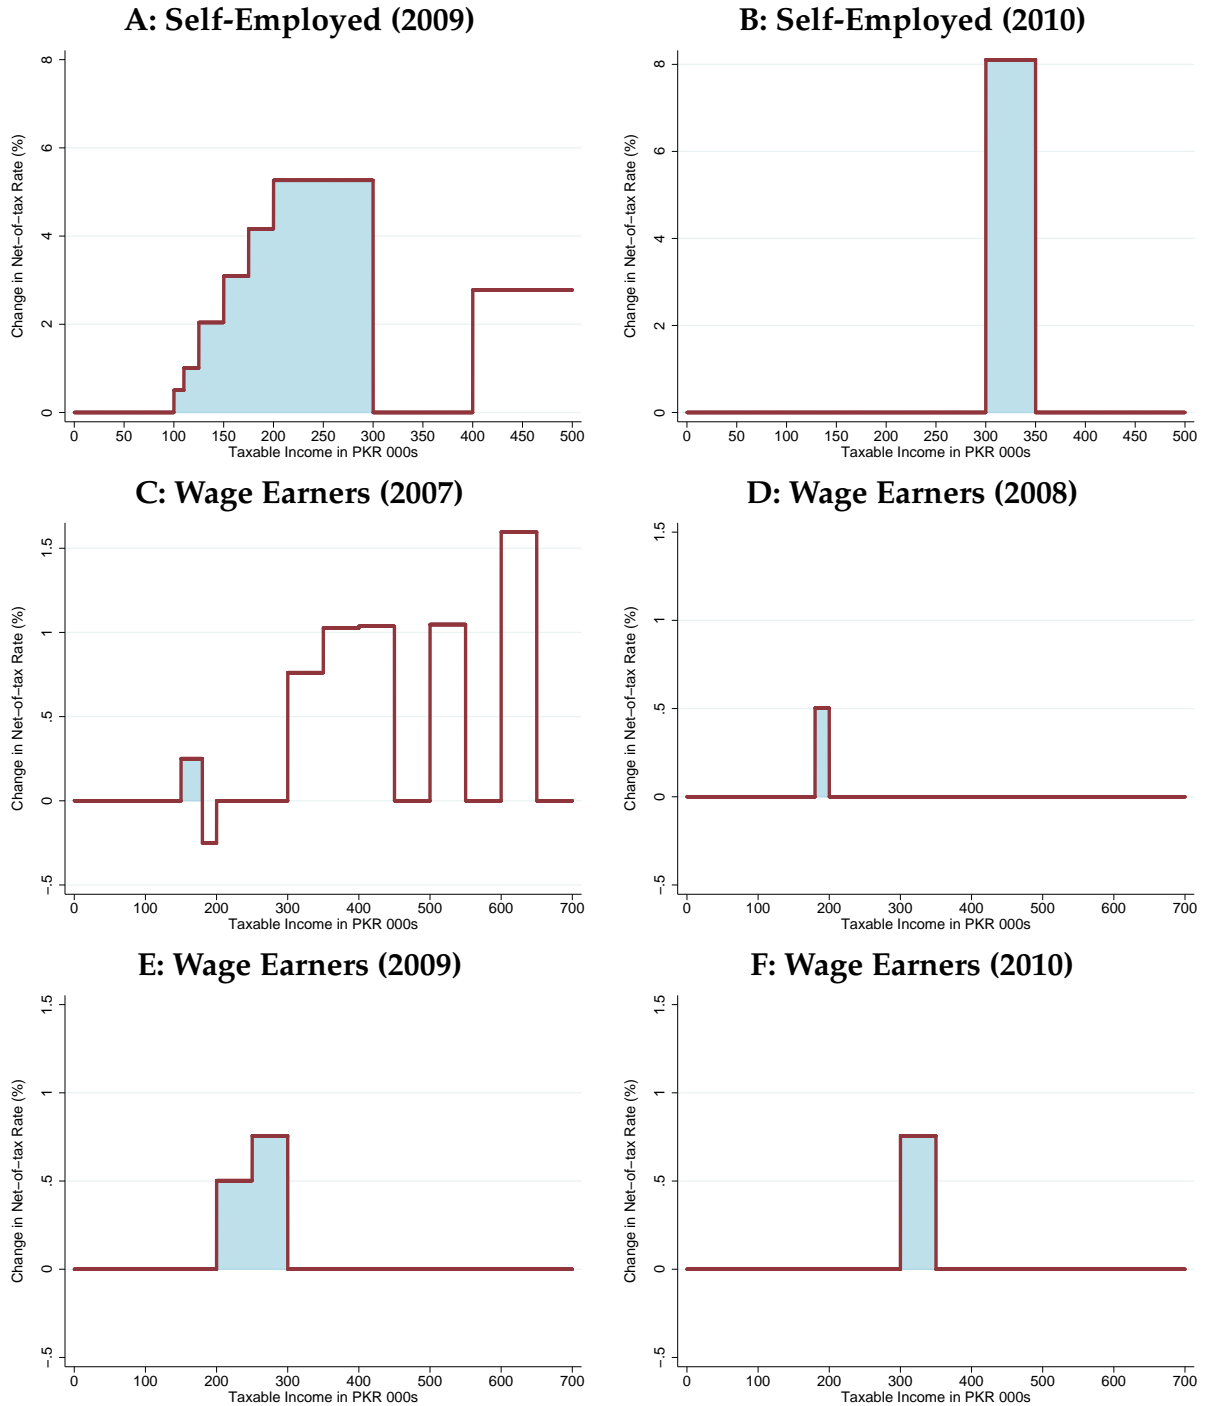
- ALLINGHAM, MICHAEL G., & SANDMO, AGNAR. 1972. Income Tax Evasion: A Theoretical Analysis. *Journal of Public Economics*, **1**, 323–338.
- ALMUNIA, MIGUEL, & LOPEZ-RODRIGUEZ, DAVID. 2018. Under the Radar: The Effects of Monitoring Firms on Tax Compliance. *American Economic Journal: Economic Policy*, **10**(1), 1–38.
- BESLEY, TIMOTHY, & PERSSON, TORSTEN. 2013. Taxation and Development. In: ALAN J. AUERBACH, RAJ CHETTY, MARTIN FELDSTEIN, & SAEZ, EMMANUEL (eds), *handbook of public economics*, vol. 5.
- BEST, MICHAEL CARLOS, BROCKMEYER, ANNE, KLEVEN, HENRIK JACOBSEN, SPINNEWIJN, JOHANNES, & WASEEM, MAZHAR. 2015. Production versus Revenue Effi-

- ciency with Limited Tax Capacity: Theory and Evidence from Pakistan. *Journal of Political Economy*, **123**(6), 1311–1355.
- BROCKMEYER, ANNE, & HERNANDEZ, MARCO. 2017. *Taxation, Information and Withholding: Evidence from Costa Rica*. Working Paper.
- CARRILLO, PAUL, POMERANZ, DINA, & SINGHAL, MONICA. 2017. Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement. *American Economic Journal: Applied Economics*, **9**(2), 144–64.
- CHETTY, RAJ. 2009. Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance. *American Economic Journal: Economic Policy*, **1**(2), 31–52.
- CHETTY, RAJ. 2012. Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica*, **80**(3), 969–1018.
- CHETTY, RAJ, LOONEY, ADAM, & KROFT, KORY. 2009. Salience and Taxation: Theory and Evidence. *American Economic Review*, **99**(4), 1145–77.
- FBR. 2014. Biannual Review, January-June 2013-14. **13**(4). Federal Board of Revenue, Pakistan, <http://goo.gl/gYHPoA>.
- FELDSTEIN, MARTIN. 1999. Tax Avoidance and the Deadweight Loss of the Income Tax. *The Review of Economics and Statistics*, **81**(4), 674–680.
- FINKELSTEIN, AMY. 2009. E-ztax: Tax salience and tax rates. *The Quarterly Journal of Economics*, **124**(3), 969–1010.
- GRUBER, JON, & SAEZ, EMMANUEL. 2002. The elasticity of taxable income: evidence and implications. *Journal of Public Economics*, **84**(1), 1 – 32.
- JENSEN, ANDERS. 2019. *Employment Structure and the Rise of the Modern Tax System*. NBER Working Paper No. 25502.
- KHAN, ADNAN Q., KHWAJA, ASIM I., & OLKEN, BENJAMIN A. 2016. Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors. *The Quarterly Journal of Economics*, **131**(1), 219–271.
- KLEVEN, HENRIK J., & WASEEM, MAZHAR. 2013. Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *Quarterly Journal of Economics*, **128**(2), 669–723.

- NARITOMI, JOANA. 2018. Consumers as Tax Auditors. London School of Economics Working Paper.
- PIKETTY, THOMAS, & QIAN, NANCY. 2009. Income Inequality and Progressive Income Taxation in China and India, 1986-2015. *American Economic Journal: Applied Economics*, **1**(2), 53–63.
- POMERANZ, DINA. 2015. No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax. *American Economic Review*, **105**(8), 2539–2569.
- SAEZ, EMMANUEL. 2004. Reported Incomes and Marginal Tax Rates, 1960-2000: Evidence and Policy Implications. In: *Tax Policy and the Economy, Volume 18*. National Bureau of Economic Research.
- SAEZ, EMMANUEL, SLEMROD, JOEL, & GIERTZ, SETH. 2012. The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, **50**, 3–50.
- SLEMROD, JOEL. 2001. A General Model of the Behavioral Response to Taxation. *International Tax and Public Finance*, **8**(2), 119–128.
- SLEMROD, JOEL. Forthcoming. Tax Compliance and Enforcement. *Journal of Economic Literature*.
- SLEMROD, JOEL, REHMAN, OBEID UR, & WASEEM, MAZHAR. 2018. Pecuniary and Non-Pecuniary Motivations in Tax Compliance: Evidence from Pakistan. Working Paper.
- TAUBINSKY, DMITRY, & REES-JONES, ALEX. 2017. Attention variation and welfare: theory and evidence from a tax salience experiment. NBER Working Paper.
- TOUREK, GABRIEL. 2019. *Targeting in Tax Behavior: Evidence from Rwandan Firms*. Working Paper.
- WASEEM, MAZHAR. 2018a. Does Cutting the Tax Rate to Zero Induce Behavior Different from Other Tax Cuts? Evidence from Pakistan. Working Paper, University of Manchester.
- WASEEM, MAZHAR. 2018b. Information, Asymmetric Incentives, or Withholding: Understanding the Self-Enforcement of Value-added Tax. Working Paper, University of Manchester.

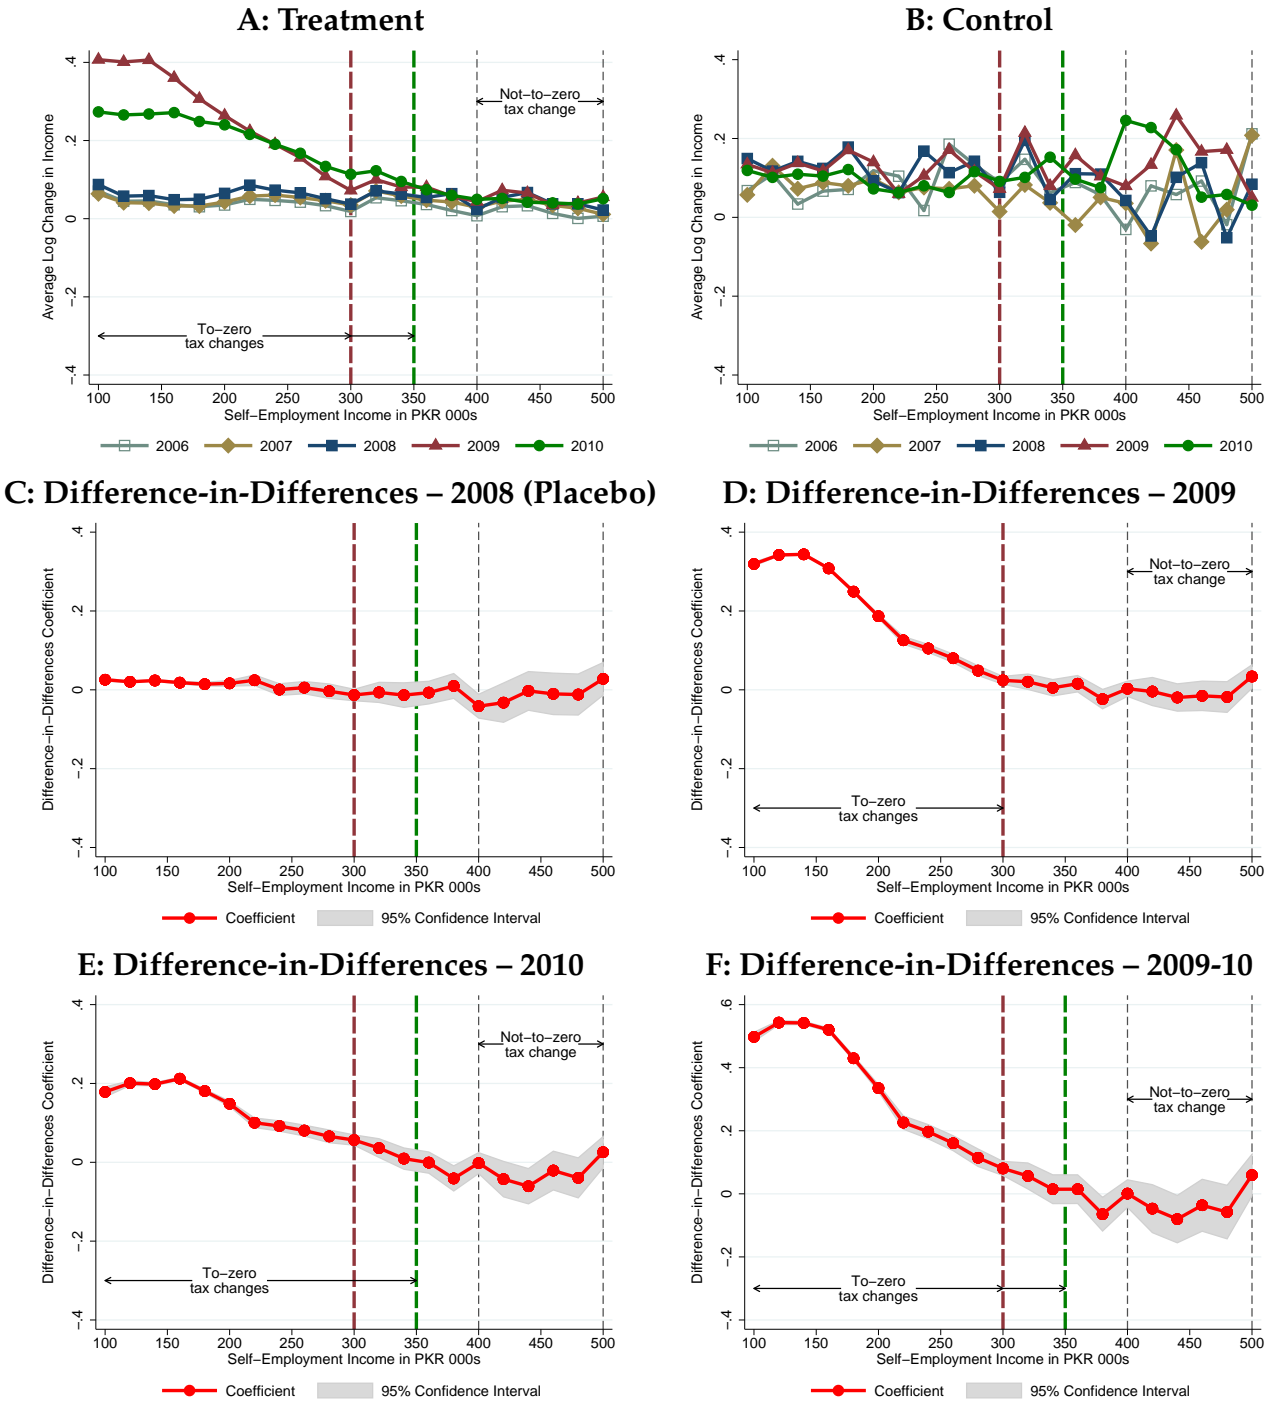
WASEEM, MAZHAR. 2018c. Taxes, Informality and Income Shifting: Evidence from a Recent Pakistani Tax Reform. *Journal of Public Economics*, **157**, 41–77.

Figure I: Tax Variation



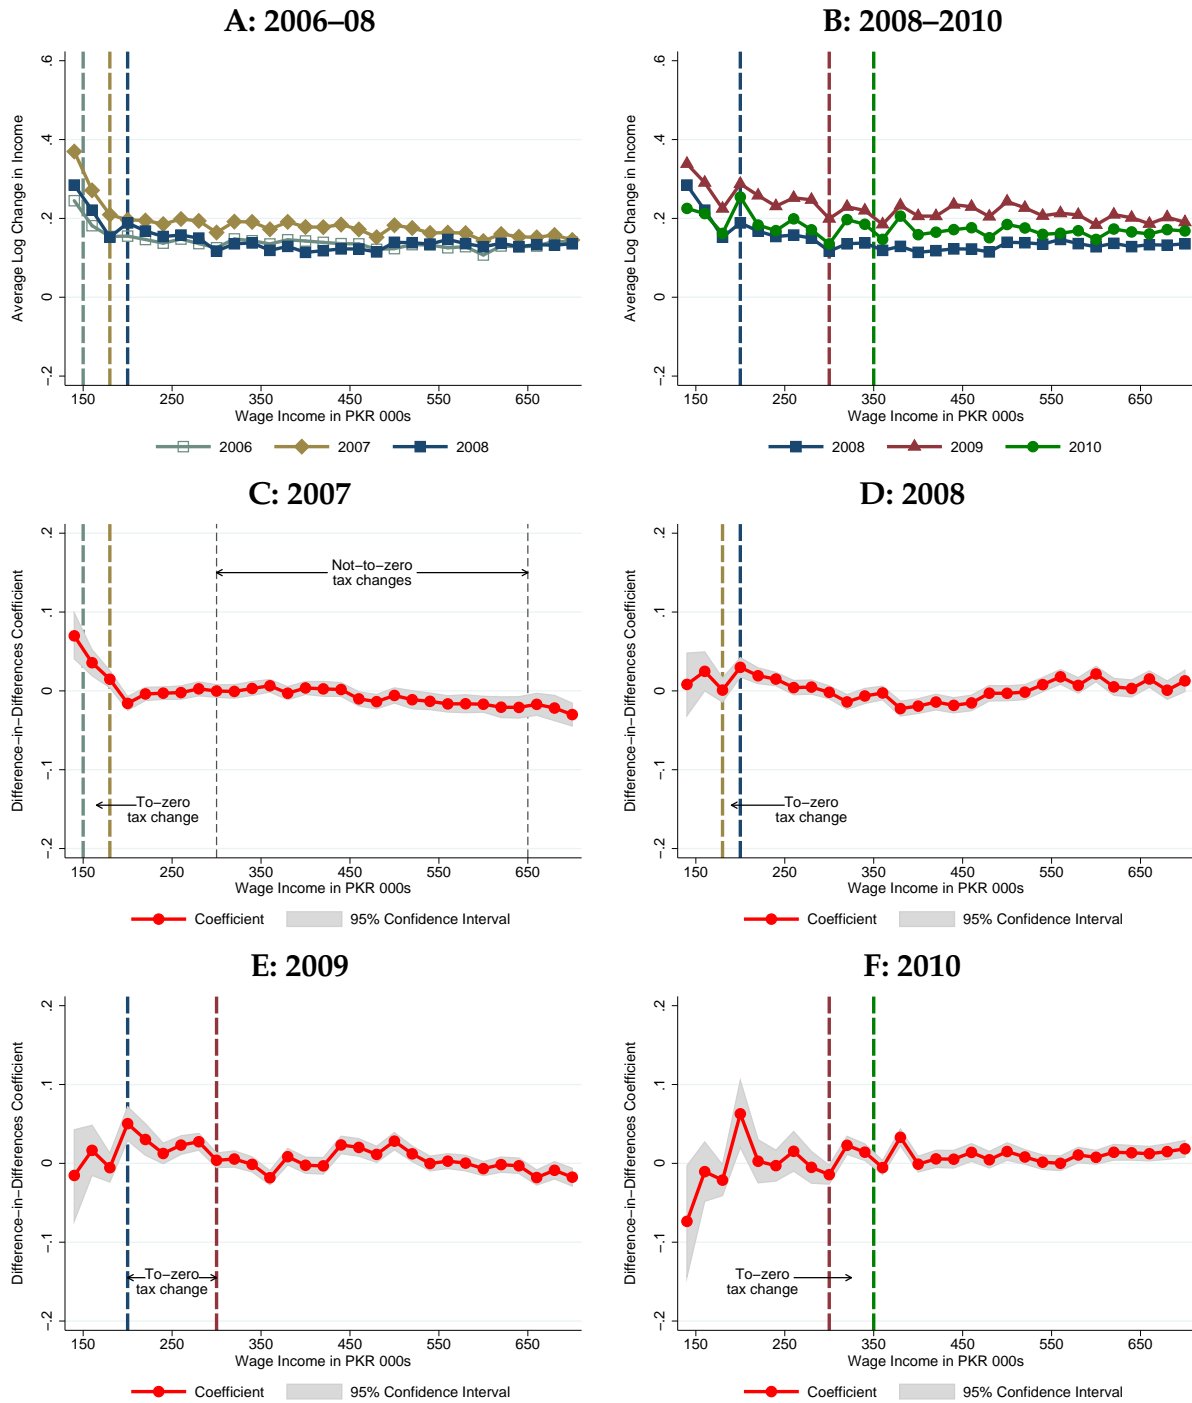
Notes: The figure displays tax variation created by the Pakistani reforms from 2006 to 2011. All curves plot percentage change in the net-of-tax rate from period t to $t + 1$ as a function of the base period income. Other than a small, narrow tax increase in 2007-08 for wage earners (Panel C), all reforms—both rate changes and movement of the bracket boundaries—result in a reduction of the tax rate, meaning percent changes in the net-of-tax rate created by them are positive. The to-zero changes, which reduce the rate to zero, have been shaded with the dark grey color.

Figure II: Self-Employment Income Response



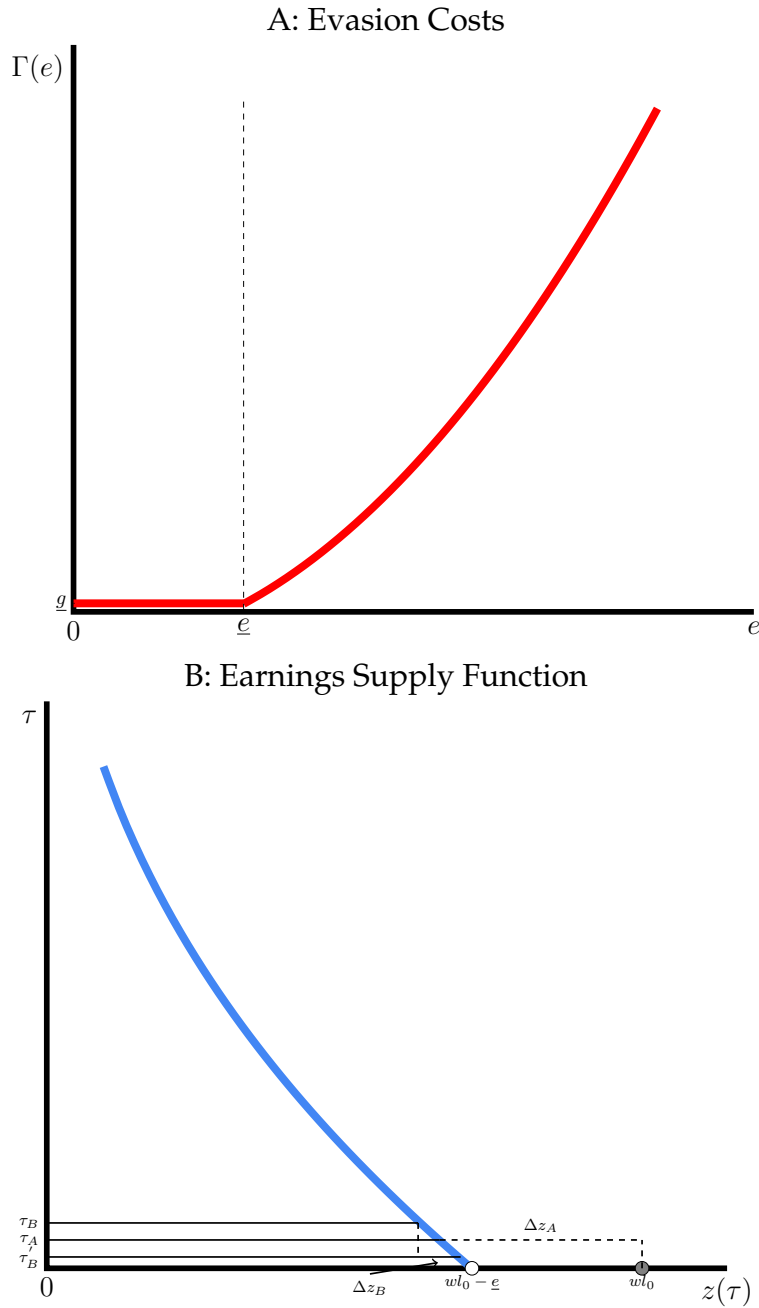
Notes: The figure compares the self-employment income response to the two types of rate changes. To construct the top-two plots, taxpayers are grouped into bins of PKR 20K on the basis of their base period income. Then, average log change in income from year t to $t + 1$ is plotted as a function of the base period income. The bottom-four panels display a difference-in-differences version of the top-two plots. Panel F illustrates the sum of 2009 and 2010 coefficients. The standard errors have been clustered at the individual level.

Figure III: Wage Income Response



Notes: The figure shows wage income growth from 2006–2011. For the top-two panels, taxpayers are grouped into bins of PKR 20K on the basis of their base period income. Then, the log change in income from period t to $t + 1$ averaged across taxpayers in the bin is plotted as a function of the base period income. For the bottom-four panels, I regress the log changes in wage income from period t to $t + 1$ on a full set of year fixed effects. The residuals from the regression are then regressed on a treatment group dummy and four yearly dummies one each for 2007 to 2010. The figure plots the estimated coefficients and 95% confidence intervals on the four yearly dummies from these regression. The standard errors have been clustered at the individual level.

Figure IV: Conceptual Framework



Notes: Panel A of the figure illustrates the evasion cost function (7). The agent can evade \underline{e} units of income on paying a small fixed cost of g . The costs turn sharply at the cutoff \underline{e} from where the third-party reported units of income begin. Panel B displays the earnings supply function (12), illustrating how an optimizing taxpayer facing such evasion costs would behave at various tax rates: the taxpayer would report true income wl_0 at the zero rate but discretely lower income $wl_0 - \underline{e}$ at a rate marginally above zero. The discontinuity means that the difference between the taxable income response to a to-zero reform $\Delta z_A(\tau_A \rightarrow 0)$ and a not-to-zero reform $\Delta z_B(\tau_B \rightarrow \tau'_B)$ identifies \underline{e} . Intuitively, any smooth change in reported income caused by a change in rate (movement along the curve) is netted out, leaving behind the discrete change in income from $wl_0 - \underline{e}$ to wl_0 (movement along the horizontal axis).

Table I: Self-Employment Income Response

	Dependent variable: Log change in self-employment income from period t to $t + 1$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$to-zero \times 2009$	0.273 (0.003)	0.287 (0.001)	0.258 (0.006)	0.236 (0.006)	0.271 (0.003)	0.263 (0.003)	0.229 (0.006)
$to-zero \times 2010$	0.126 (0.004)	0.136 (0.001)	0.109 (0.006)	0.091 (0.006)	0.127 (0.004)	0.124 (0.004)	0.089 (0.006)
$not-to-zero \times 2009$	-0.014 (0.008)	-0.007 (0.007)	-0.006 (0.012)	-0.003 (0.012)	-0.017 (0.008)	-0.015 (0.008)	-0.003 (0.012)
Fixed effects:							
Year	Yes	No	Yes	Yes	Yes	Yes	Yes
Industry	No	No	Yes	Yes	No	No	Yes
Industry \times year	No	No	No	Yes	No	No	Yes
Region	No	No	No	No	Yes	Yes	Yes
Region \times year	No	No	No	No	No	Yes	Yes
Time trend:							
Linear	No	Yes	No	No	No	No	No
Pre-reform mean of the dependent variable	0.050	0.050	0.066	0.066	0.050	0.050	0.066
Observations	526,150	526,150	93,525	93,525	525,260	525,260	93,319

Notes: The table reports the estimates from equation (1). I restrict the sample to a balanced panel of taxpayers who file in all six year 2006-2011. The first columns correspond to the baseline specification; the second column replaces the year fixed effects with a linear time trend; and the rest of the columns add year, industry, industry \times year, region and region \times year fixed effects. The details of the industry and region classifications are provided in Appendix A.1. I do not observe industry classification for all taxpayers, owing to which the numbers of observations in the last two columns are lower. The control group here comprises self-employed who do not experience any rate change. Standard errors are in parenthesis, which have been clustered at the individual level.

Table II: Wage Income Response

	Dependent variable: Log change in wage income from period t to $t + 1$								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A: Tax-Driven Response									
<i>to-zero</i>	0.016 (0.002)	0.016 (0.002)	0.013 (0.002)	0.014 (0.002)	0.014 (0.002)	0.011 (0.002)	0.014 (0.002)	0.014 (0.002)	0.011 (0.002)
<i>not-to-zero</i>	0.000 (0.002)	0.000 (0.002)	0.002 (0.002)	0.001 (0.002)	0.002 (0.002)	0.003 (0.002)	0.001 (0.002)	0.002 (0.002)	0.003 (0.002)
B: Placebo									
<i>to-zero</i>	0.006 (0.001)	0.004 (0.001)	0.003 (0.001)	0.001 (0.002)	-0.000 (0.002)	-0.002 (0.002)	0.001 (0.002)	-0.000 (0.002)	-0.002 (0.002)
<i>not-to-zero</i>	0.005 (0.002)	0.006 (0.002)	0.006 (0.002)	0.010 (0.002)	0.010 (0.002)	0.011 (0.002)	0.010 (0.002)	0.010 (0.002)	0.011 (0.002)
Fixed effects:									
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Region \times year	No	No	Yes	No	No	Yes	No	No	Yes
Base-period income controls:									
Log base-period income	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes
Spline of log base-period income	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Pre-reform mean of the dependent variable	0.162	0.162	0.162	0.162	0.162	0.162	0.162	0.162	0.162
Observations	240,649	240,643	240,643	240,649	240,643	240,643	240,649	240,643	240,643

Notes: This table reports the estimates from equation (1). Column (1) corresponds to the baseline specification, and the subsequent columns add more control variables, permuting among the combinations of controls for mean reversion—log base period income and a ten-piece spline in log base period income—and other controls. The definition of the region variable is provided in Appendix A.1. Panel B carries the results from placebo regressions corresponding to each specification, assuming that all tax changes took place one year earlier than they actually did. Standard errors are in parenthesis, which have been clustered at the individual level.

Table III: Evasion Rate of Self-Employment Income

Income	Earnings Response ($\tau \rightarrow 0$)	Earnings Response ($\tau \neq 0$)	Difference	Evasion Rate (%)
(1)	(2)	(3)	(4)	(5)
80-100K	71,754 (1,251)	333 (188)	71,421 (1,265)	74.2 (1.3)
100-150K	92,280 (784)	603 (341)	91,676 (855)	71.9 (0.7)
150-200K	92,440 (938)	573 (209)	91,867 (961)	52.2 (0.5)
200-250K	70,537 (1,116)	763 (278)	69,774 (1,150)	30.7 (0.5)
250-300K	28,868 (1,172)	692 (252)	28,176 (1,198)	10.0 (0.4)
300-350K	20,124 (918)	419 (153)	19,704 (931)	6.0 (0.3)

Notes: The table presents the estimates of the rates of evasion of self-employment income from equation (15). Column (1) shows the income segment; column (2) the earnings response produced by the to-zero change; column (3) the earnings response produced by the equal-sized, not-to-zero change; column (4) the difference between the two responses; and column (5) the average evasion rate in the segment. The estimates in column (2) are computed from the self-employment income response to the movements of the exemption cutoff in 2010–11 shown in Table I. The estimates in column (3) are computed using elasticities estimated for the same group of taxpayers in Kleven & Waseem (2013). The details on how the estimates in columns (2)–(3) have been computed are in section V. The difference between the two set of estimates represents average \underline{e} for taxpayers in the segment. I obtain the percentage evasion rate implied by the average \underline{e} by dividing it with the average income in the segment—see column (5). The standard errors are in parenthesis.

A Online Appendix

A.1 Details of Variables

- (i) **Taxable Income.** Income derived by a person in a year from all sources such as employment, business, ownership of capital, minus three minor deductions Zakat, Workers' Welfare Fund, and Workers' Participation Fund.²⁵ The deductions, however, are extremely rare for wage earners and self-employed, and practically taxable income is the aggregate income of a taxpayer in a year.
- (ii) **Wage Income.** Income received by an employee in a tax year from employment including wages, leave pay, bonuses, commission, fees, and gratuity.
- (iii) **Self-Employment Income.** Profits from any business carried out by a person other than those from partnerships or corporations.
- (iv) **Self-Employed.** The tax code classifies a person self-employed if her wage income does not exceed 50% of taxable income.
- (v) **Wage Income.** The tax code classifies a person wage earner if her wage income exceeds 50% of taxable income.
- (vi) **Sales.** The variable is from the brief profit and loss account included in the tax return form, and denotes the full revenue for the year from selling goods or services.
- (vii) **Costs.** The variable is also from the profit and loss account, and denotes what in accounting is referred to as the cost of sales. The cost is calculated by adding opening stock, net purchases, and manufacturing and trading expenses, and then taking away the closing stock.
- (viii) **Profit and Loss Expenses.** The variable is also from the profit and loss account, and denotes input costs such as wages, rents, accounting and legal fees, electricity, and interest paid on loans.
- (ix) **Imports.** Imports are costs net of import duties and taxes that are incurred on imported inputs / investment goods.

²⁵Zakat is a charitable payment made annually by Muslims as a religious obligation. Workers Welfare Fund is paid by industrial establishment having taxable income exceeding Rs. 500,000 in a year. Workers' Participation Fund is paid by industrial establishments having more than 50 workers.

- (x) **Opening Stock.** Inventory at the start of the year.
- (xi) **Closing Stock.** Inventory at the end of the year.
- (xii) **Electronic Return Filer.** I categorize a taxpayer electronic return filer if any of the six returns from 2006–2011 was filed electronically.
- (xiii) **VAT-registered.** The variable indicates if a self-employed taxpayer is registered with the FBR to remit VAT on its sales.
- (xiv) **Region.** The variable codes the tax district a taxpayer’s registered office is located in.
- (xv) **Tax office.** The variable represents the tax office that is responsible for audit and enforcement of an individual’s tax liability. Pakistan has two types of tax offices: (1) Large Taxpayer Units, which are located in Karachi, Lahore and Islamabad and cater to the top tax contributors; (2) Regional Tax Offices, which are located in twelve cities and administer the rest of taxpayers.
- (xvi) **Industry.** The variable represents the 2-digit industry code a self-employed taxpayer operates in.

A.2 Tax Reporting Under Imperfect Enforcement

In this section, I adopt the standard model of tax reporting under imperfect enforcement (see for example [Chetty, 2009](#)) to the Pakistani setting. The purpose of the model is to understand the mechanisms driving the large difference between the to-zero and not-to-zero responses, as documented in section III of the paper. Consider an agent who decides how many hours to work (l) at a fixed wage rate w . The agent is required to pay income tax at a rate $\tau \in [0, 1]$ on its labor income wl but can reduce tax liability by hiding e units of income on paying a resource cost of $\Gamma(e)$. Denoting consumption by c and disutility of labor by $\psi(l)$, the agent’s utility maximization problem can be written as

$$\begin{aligned}
 (4) \quad \max_{l,e} \quad & u(c, l, e) = c - \psi(l) - \Gamma(e) \\
 \text{s.t.} \quad & c = (1 - \tau)(wl - e) + e.
 \end{aligned}$$

The agent will work up to the point that the disutility of work equals the net-of-tax wage rate and will evade income up to the point that the marginal cost of evasion equals the tax rate.

$$(5) \quad \begin{aligned} (1 - \tau)w &= \psi'(l^*) \\ \tau &= \Gamma'(e^*). \end{aligned}$$

The combination of optimal hours and evasion determines how much earnings $z(\tau) \equiv wl(\tau) - e(\tau)$ will be reported to the authorities.

Now consider two small, similar-sized tax cuts, indexed by $j \in \{A, B\}$. The first of these reduces the rate to zero, while the other does not. The evidence in section III shows that earnings response to the former type of tax cut is always orders of magnitude larger than one to the latter type

$$(6) \quad \begin{aligned} \Delta z_A(\tau_A \rightarrow 0) &\ggg \Delta z_B(\tau_B \rightarrow \tau'_B) \\ \tau_A, \tau_B, \tau'_B &> 0 \end{aligned}$$

Potential explanations of this finding can be divided into three classes. First are the discontinuous evasion costs based explanations. It is possible that the costs of evading certain categories of income are small and of others large. Income entailing little evasion cost would be reported at a zero rate but not at any positive rate. On the other hand, income entailing large evasion cost would always be reported. Response to a to-zero tax cut will include both categories of income and hence will be larger. Second are the behavioral explanations, in particular the salience. It can be argued that to-zero tax cuts receive more attention of taxpayers than do not-to-zero cuts. Third are the enforcement environment based explanations. The audit and penalty functions chosen by the government may not be neutral across the two types of rate changes. For example, the government may audit zero-rated income more leniently, as no tax is payable. I investigate which of these explanations fits the Pakistani evidence better in sections IV.A -IV.C of the paper. Below, I describe the first explanation in more details.

A.2.1 Reverse-L Shaped Evasion Costs

The standard way to think about evasion costs is that they are expected fine and penalty payments which would be recovered in case the evasion is detected by the government (Allingham & Sandmo, 1972; Slemrod, 2001). Recent empirical evidence shows that the probability that evasion gets detected is quite high if reported income is covered by the third-party information the government obtains from sources such as employers and financial institutions and quite low otherwise (see Slemrod, *Forthcoming* for a recent survey of the evidence).²⁶ To introduce this interaction between the information environment and reporting behavior into the model, I assume that \underline{e} units of income of the agent are not covered by any third-party information, while the rest are covered. Note that this assumption is without any loss of generality as \underline{e} potentially varies from 0 to the maximum. The costs of evasion faced by the agent take the following form (see Figure IV A)

$$(7) \quad \Gamma(e) = \begin{cases} \underline{g} & \text{if } e \leq \underline{e} \\ g(e) & \text{if } e > \underline{e}. \end{cases}$$

They are extremely low at the bottom where income is not covered by any third-party information. In fact, to keep the exposition simple I assume that this component of income can be evaded on paying a small fixed cost \underline{g} (the marginal cost is zero). The costs turn sharply afterward once the third-party reported units of income begin. The agent faces regular (increasing, convex, and twice-differentiable) evasion costs $g(e)$ with $g(\underline{e}) = \underline{g}$ on this third-party-reported component of income.

The optimum in this extended model is described by the following conditions

$$(8) \quad \psi'(l^*) = (1 - \tau)w$$

$$(9) \quad [g'(e^*) - \tau]e^* = 0$$

$$(10) \quad \tau e^* \geq \underline{g} \quad \text{if } e^* > 0.$$

In addition to the standard FOCs, we now have a participation constraint providing that the agent will enter into evasion only if the benefit of doing so τe^* exceeds the

²⁶Throughout the paper, I maintain the assumption that the third-party reports the government receives are complete and not themselves distorted by evasion. Otherwise, they would not deter tax evasion as effectively (see for example Brockmeyer & Hernandez, 2017).

fixed cost \underline{g} . The constraint implicitly defines the reservation tax rate $\underline{\tau} \equiv \frac{\underline{g}}{\underline{e}}$ at which evasion jumps from 0 to \underline{e} .²⁷ I assume that regularity conditions $\psi'(\cdot) > 0$; $\psi''(\cdot) > 0$ and $\psi(0) = \psi'(0) = 0$ are satisfied so that the choice of hours is interior. I further assume that $g'^{-1}(\cdot)$ exists for all $\tau > \underline{\tau}$ such that it can be represented by a normalized inverse function $\kappa(\tau) \equiv g'^{-1}(\tau) - \underline{e}$ having the property that $\kappa(0) = \kappa'(0) = 0$. These assumptions generate an evasion function

$$(11) \quad e(\tau) = \begin{cases} 0 & \text{if } \tau < \underline{\tau} \\ \underline{e} + \kappa(\tau) & \text{if } \tau \geq \underline{\tau}. \end{cases}$$

which has three properties: (1) there is no evasion as long as the tax rate is below $\underline{\tau}$ i.e. $e(\tau < \underline{\tau}) = 0$; (2) evasion jumps to \underline{e} whenever the tax rate reaches $\underline{\tau}$; and (3) evasion increases smoothly with the rate from that point onward $\kappa'(\tau) > 0$. This pattern of evasion in turn feeds into the agent's tax report, producing the earnings supply function of the following form (see Figure IVA)

$$(12) \quad z(\tau) = \begin{cases} wl(\tau) & \text{if } \tau < \underline{\tau} \\ wl(\tau) - \underline{e} - \kappa(\tau) & \text{if } \tau \geq \underline{\tau}. \end{cases}$$

A corollary of the discontinuous earnings supply function is that the behavioral responses to the two types of rate changes differ substantively. The intuition for this result is provided in Figure IVB. The response to the to-zero change $\Delta z_A(\tau_A \rightarrow 0)$ consists of the movement along the earnings supply curve $z(\tau)$ and the movement along the horizontal axis from $wl_0 - \underline{e}$ to wl_0 . By contrast, the response to the other change $\Delta z_B(\tau_B \rightarrow \tau'_B)$ consists of the movement along the supply curve only. This can be shown formally as the following. The earnings responses induced by the two tax cuts are given by

$$(13) \quad \begin{aligned} \Delta z_A(\tau_A \rightarrow 0) &= w\Delta l(\tau_A \rightarrow 0) - \Delta e(\tau_A \rightarrow 0) \\ \Delta z_B(\tau_B \rightarrow \tau'_B) &= w\Delta l(\tau_B \rightarrow \tau'_B) - \Delta e(\tau_B \rightarrow \tau'_B) \end{aligned}$$

If these cuts, as I assume above, are small, similar-sized $\Delta\tau_A \approx \Delta\tau_B$, and fairly ad-

²⁷The jump occurs because the marginal cost of evading income up to \underline{e} is zero and once the participation constraint gets satisfied at the rate $\underline{\tau}$, evading \underline{e} units of income becomes optimal.

adjacent to each other $\tau'_A \approx \tau'_B$, then the fact that $l(\tau)$ and $\kappa(\tau)$ are smooth, continuous functions means that

$$(14) \quad \begin{aligned} \Delta l(\tau_A \rightarrow 0) &\approx \Delta l(\tau_B \rightarrow \tau'_B) \\ \Delta \kappa(\tau_A \rightarrow 0) &\approx \Delta \kappa(\tau_B \rightarrow \tau'_B) \end{aligned}$$

so that the difference between the two responses captures \underline{e}

$$(15) \quad \Delta z_A(\tau_A \rightarrow 0) - \Delta z_B(\tau_B \rightarrow \tau'_B) \approx \underline{e}.$$

Thus as long as \underline{e} is nontrivial, the model indeed reconciles the empirical distinction between the to-zero and not-to-zero behavior documented in section III.

A.2.2 Microfounding the Evasion Costs

In the baseline case, I model evasion costs in a reduced-form way through the deterministic function $g(e)$ assuming that $g'(e) > 0$ and $g''(e) > 0$. As a result of this assumption, the amount of undeclared income e increases with the marginal tax rate τ . In this section, I use the Allingham and Sandmo type setting to show that $\frac{de}{d\tau} > 0$ is a reasonable assumption in the Pakistani context.

Suppose that the government detects the undeclared amount with probability $p(e)$. Reflecting the intuition that detection becomes more likely as evasion becomes more aggressive, the probability is assumed to be increasing and convex in e . In the event that evasion is detected, the agent is required to pay the undeclared tax along with a penalty. The Pakistani tax code provides that a taxpayer found misreporting its income would pay a penalty of 25,000 rupees or 100% of the tax avoided whichever is higher, meaning that the penalty takes the functional form $\max\{\vartheta, \tau e\}$, where $\vartheta = 25,000$.²⁸ Assuming that the agent is risk-neutral and the taxation is linear, evasion costs are given by

$$(16) \quad g(e) = p(e)[\tau e + \max\{\vartheta, \tau e\}].$$

As earlier, the agent would evade up to the point that the marginal cost of evasion equals the tax rate $g'(e) = \tau$. Optimal evasion thus solves one of the following two

²⁸Having a floor on the penalty is not unique to Pakistan and is quite common across tax administrations. Its purpose is to cover the fixed costs of audit and to deter evasion even when tax liability is small.

conditions

$$(17) \quad \begin{aligned} [p'(e)e + p(e)] + \frac{p'(e)\vartheta}{\tau} &= 1 \quad \text{if } \vartheta > \tau e \\ 2 [p'(e)e + p(e)] &= 1 \quad \text{if } \vartheta < \tau e. \end{aligned}$$

Differentiating these expressions with respect to τ leads to the following comparative statics

$$(18) \quad \frac{de}{d\tau} = \begin{cases} \frac{p'(e)\vartheta}{\tau} \cdot \frac{1}{\tau[p''(e)e + 2p'(e)] + p''(e)\vartheta} > 0 & \text{if } \vartheta > \tau e \\ 0 & \text{if } \vartheta < \tau e. \end{cases}$$

Evasion increases with the tax rate as long as the evaded tax is less than the minimum penalty ϑ . The intuition of the result is simple: as long as the penalty floor is binding increasing the marginal tax rate makes evasion more attractive without causing a proportional increase in the expected penalty. When the floor is not binding the two effects cancel each other so that the tax rate has no effect on the unreported amount.

The Pakistani income tax rates are quite low, meaning that the condition $\vartheta > \tau e$ is binding for most of taxpayers. For example, assuming that the evaded tax equals the reported tax so that the evasion rate is 100% of the reported tax, the condition would be binding for around 90% of the self-employed. Thus, for the purposes of this paper $\frac{de}{d\tau} > 0$ is a reasonable assumption, especially for taxpayers close to the exemption cutoff. Note that this conclusion is likely to survive even if we discard the assumption of risk-neutrality. When the agent is risk-averse, the disutility of evasion also depends on the income effect created by a change in the marginal tax rate. For example, when the marginal tax rate increases a risk-averse agent would feel poorer and would experience greater disutility of evasion if absolute risk aversion is decreasing in income. In the Pakistani context, however, this effect is likely to be trivial. It is because the tax changes created by the reforms are extremely small and do not change after-tax income of taxpayers by much (the change in after-tax income for most cases is less than 1%).

A.2.3 Nonlinear Tax Schedule

The baseline model assumes a linear tax schedule, but the empirical application considers a setting where the tax schedule is linear within brackets only. Figure A.XII relaxes the linearity assumption. I now focus solely on the to-zero reform, assuming

it changes the tax schedule from $T_A = \tau_A$ to $T'_A = \tau_A \cdot 1(z > z_N)$. The postreform schedule now contains a notch at the earnings level z_N , where the rate jumps from 0 to τ_A . The reform, accordingly, moves the rate from τ_A to zero on baseline earnings up to z_N but has no effect beyond that. Panel A of the figure shows that the notch does not affect a taxpayer whose baseline earnings are too far away from it. The taxpayer continues to behave exactly in the manner described in the baseline model so that its declarations still identify \underline{e} in accordance with equation (15). In contrast, a taxpayer whose baseline earnings are in the interval $[z_N - \underline{e}, z_N]$ bunches at the notch after the reform (see the bottom panel of the figure). For this taxpayer, the marginal cost of evasion is strictly less than the postreform tax rate in the region just above the notch. It accordingly increases evasion, and reduces reported earnings, until it hits the notch. Earnings response to such a reform will therefore depend upon the baseline earnings of the taxpayer. The response will be unconstrained if the baseline income is less than $z_N - \underline{e}$ and constrained if it is not. This implies that the average earnings response to a tax reform that reduces the rate to zero on earnings below a cutoff z_N only will vary across the income distribution. It will be trivial close to the notch and will increase monotonically before becoming constant as we move sufficiently left of the notch. Note that this prediction of the model is consistent with what we observe in the data (see Figure II). To account for the fact that the responses of taxpayers close to the notch are constrained, I estimate my evasion rates from the declarations of taxpayers with baseline income up to PKR 150k only.

A.2.4 Separability

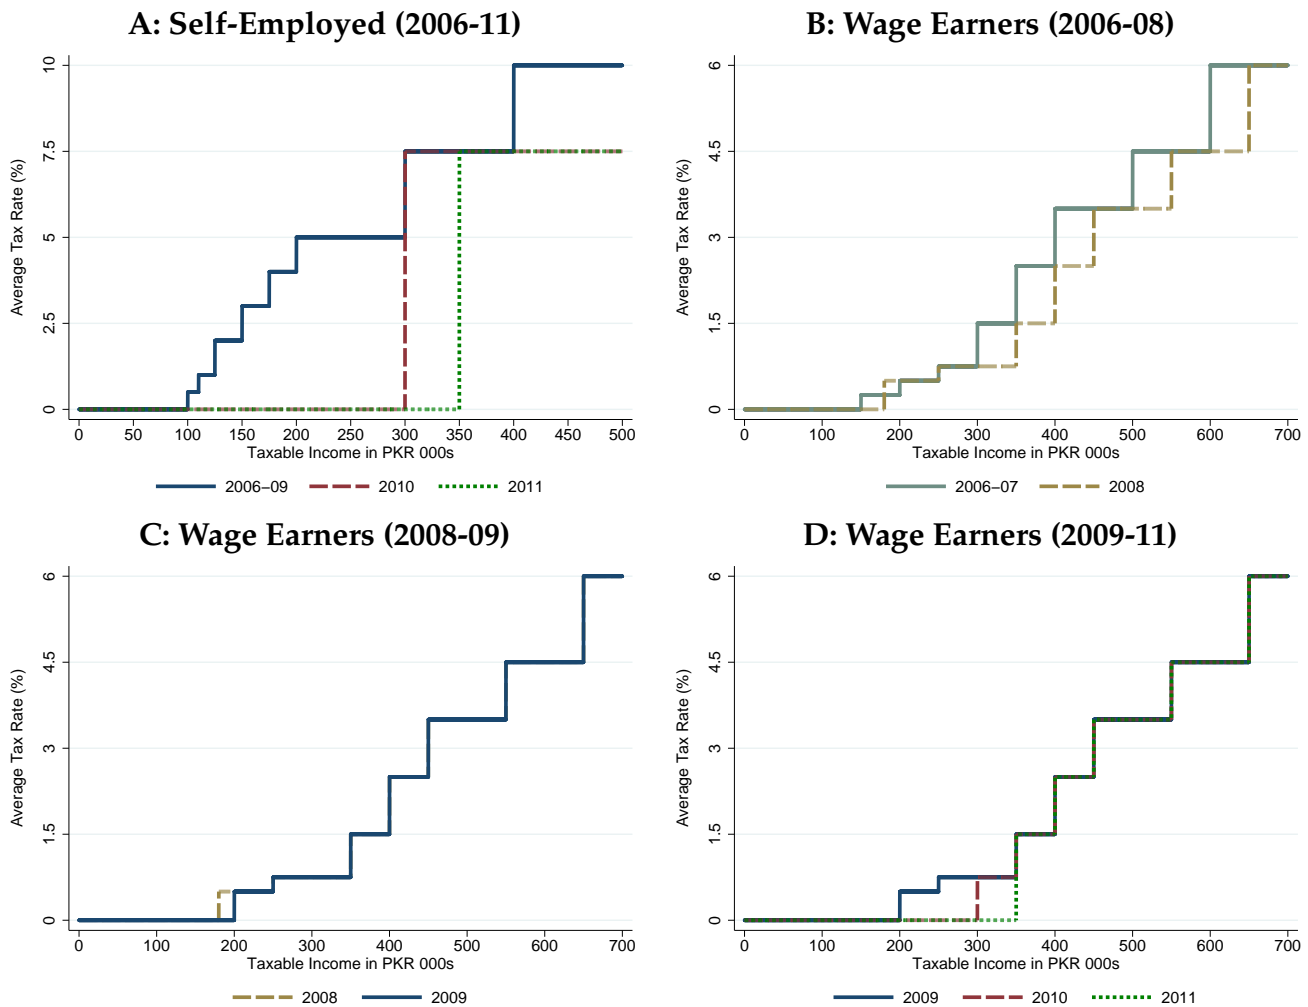
The baseline model assumes that the production and evasion costs are additively separable. Such separability is a common assumption in the tax compliance literature. It is because taxpayers report their income to the government strictly after it has been earned. This time lag (3-15 months in the case of Pakistan) together with the low costs of evading non-third-party-reported income means that two-stage optimization—in the first stage, taxpayers decide how much to produce and in the second how much to report—is the most natural decision framework in this setting. Separability is, thus, quite a natural assumption in this setup. In fact, the only plausible channel discussed in the literature that can violate it is if the fines and penalties are proportional to true, rather than reported, taxable income (Chetty, 2009). As I mention in Appendix A.2.2, the penalty for undeclared taxable income in Pakistan,

to a first-order approximation, is fixed and does not depend on either reported or true income. More generally, it is not feasible to make the penalty contingent on true income, as the government does not observe it. Note further that if separability is not satisfied in the Pakistani setting, meaning that real production has some dependence on evasion, it would have been picked up empirically by the exercise on the anatomy of response (see Figures A.VIII and A.IX). But the exercise shows little signs that the rate changes considered in the empirical application affect real production.

A.2.5 Income effects


The absence of income effects is an innocuous assumption in case of small tax rate changes. Note that while the earnings responses documented in this paper are large, the rate changes driving these responses are extremely small (ranging between 0.25-8%; see Figure I). Assuming no income effects is, therefore, a reasonable approximation in the Pakistani context. More generally, evasion rate in this method is identified through the comparison of the earnings responses to the equivalent-sized to-zero and not-to-zero tax reforms. The comparison nets out all elements of the responses that change smoothly with the tax rate, including any income effects. The strategy is therefore robust to allowing income effects.

Figure A.I: Tax Schedules



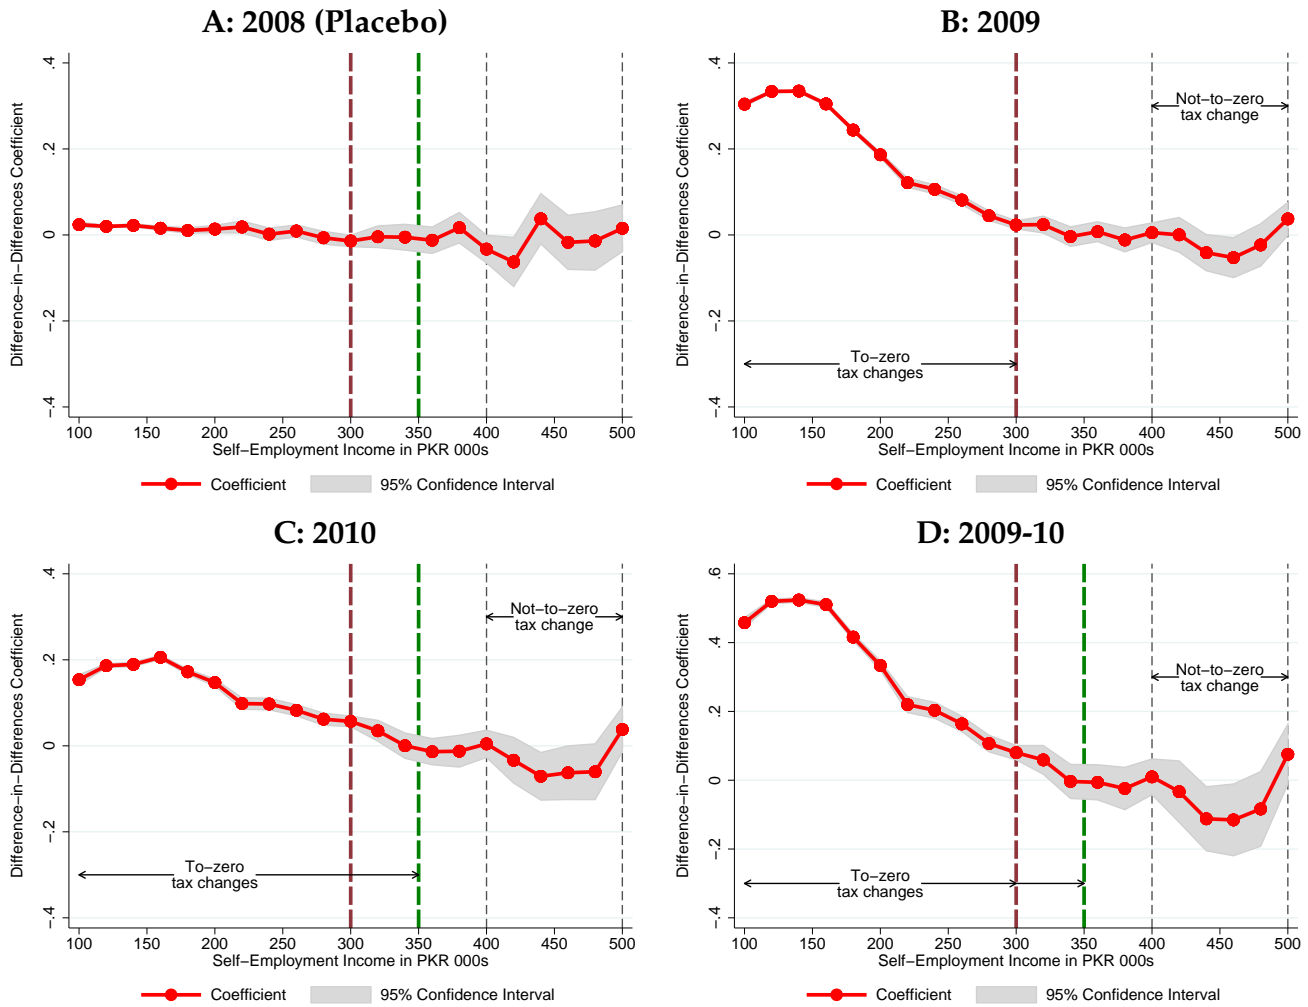
Notes: The figure shows the two personal income tax schedules of Pakistan from 2006 to 2011. A period t plot illustrates the statutory tax rate in the period as a function of taxable income. The Pakistani schedules specify average rather than marginal tax rate in a bracket. The schedule for self-employed underwent three changes during the period. First, the exemption cutoff was increased from PKR 100K to PKR 300K in 2010 and to PKR 350K in 2011. Second, the average tax rate on taxable income more than PKR 400K but not more than PKR 500K was reduced from 10% to 7.5% in 2010. There were more changes to the schedule higher up the income distribution in 2010, which I do not show here. The schedule for wage earners was comprehensively revised in 2008 (Panel B). The exemption cutoff for wage earners was increased from PKR 150K to PKR 180K in 2008, to PKR 200K in 2009, to PKR 300K in 2010, and to PKR 350K in 2011. The net-of-tax rate changes implied by these reforms are plotted in Figure I.

Figure A.II: Filing Probability



Notes: The figure plots the probability that a tax filer in period t also files in $t + 1$. Taxpayers are grouped into bins of PKR 20,000 on the basis of their base period self-employment income. The average probability in each bin for the corresponding year is then plotted as a function of the baseline income. The figure illustrates that the probability is not correlated with the to-zero or not-to-zero rate changes.

Figure A.III: Self-Employment Income Response (Balanced Panel)




Notes: The figure replicates the analysis in Figure IIC-F, holding the composition of the sample fixed. The sample now is a balanced panel, comprising only the taxpayers who file in all years from 2006 to 2011. The panel for period t plots the corresponding coefficient δ_t from the following regressions

$$\Delta \log z_{it}^S = \alpha + \beta SE_i + \gamma \text{year}_t + \mathbf{SE}_i \times \mathbf{T}_t \delta + u_{it},$$


where year_t is a linear time trend and $\mathbf{SE}_i \times \mathbf{T}_t$ is a vector of interaction dummies. For the placebo regressions (Panel A), these interactions comprise three treatment \times year dummies, one each for 2008–2010. For the regressions in Panels B and C, the 2008 interaction is dropped. Panel D illustrates the sum of 2009 and 2010 coefficients, reflecting aggregate response to the tax changes. The regressions are run separately in each bin so that the displayed coefficients reflect average additional earnings growth from period t to period $t + 1$ experienced by the treated taxpayers with the base period income in the bin relative to the control taxpayers in the same bin. The 95% confidence interval around the coefficient is represented by the gray area in the plots. The standard errors have been clustered at the individual level. Vertical lines demarcate the areas where the to-zero and not-to-zero tax changes shown in Figure IA–B were applied to.

Figure A.IV: Self-Employment Income Trend



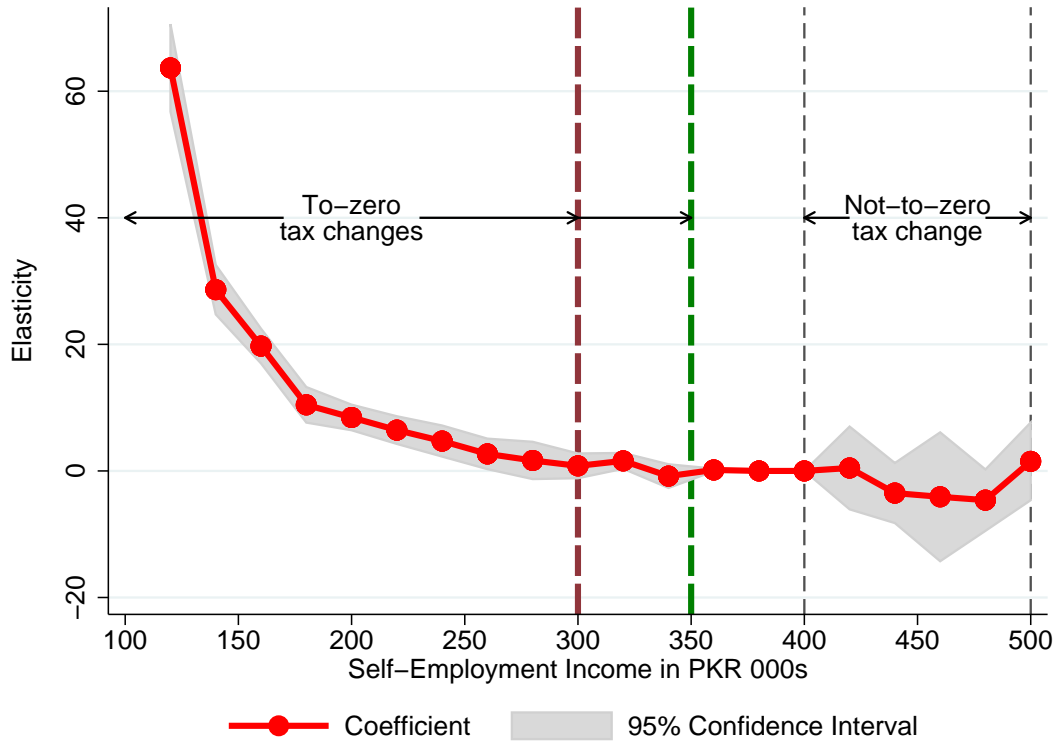
Notes: The figure plots the time path of self-employment income growth from 2006 to 2010. Each dot on the two curves illustrates log change in self-employment income from period t to $t + 1$ averaged across taxpayers in the corresponding group. The treatment group comprises the self-employed who experienced a to-zero or a not-to-zero tax change in 2010-2011, whereas the control group comprises the self-employed who did not experience any tax change during the sample period. Vertical dashed line demarcates the time from which the to-zero and not-to-zero tax changes applied to the schedule for self-employed in 2010-11, as shown in Figure IA-B, affect income growth in the treatment group.

Figure A.V: Earnings Growth Rate Needed to Hit the New Exemption Cutoff



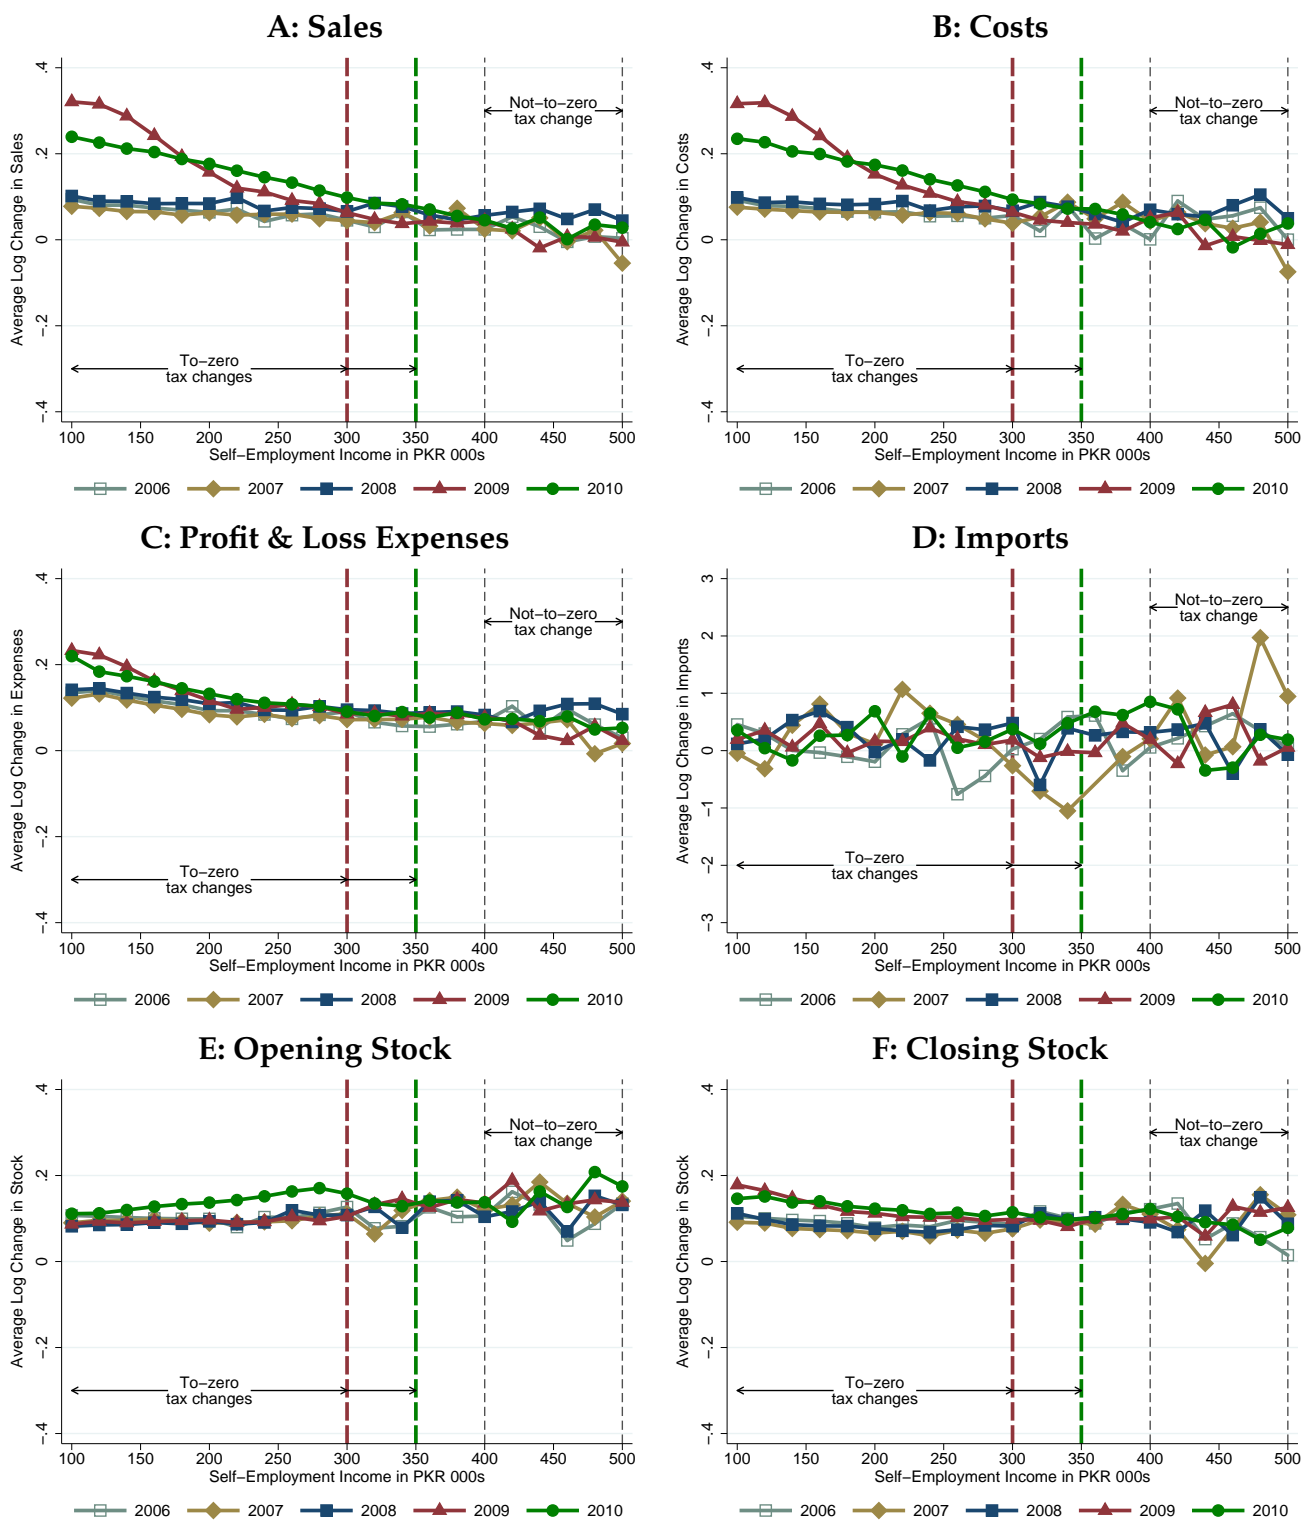
Notes: The figure rules out one alternative explanation of the large self-employment income response observed in the two post-reform years. Taxpayers are grouped into bins of PKR 20k on the basis of their self-employment income. The dashed curve in each plot depicts the average growth rate needed for taxpayers in the bin to hit the new exemption cutoff (PKR 300k in Panel A and PKR 350k in Panel B). If taxpayers simply move to the new exemption cutoff after a reform that increases the cutoff, the dashed curve should coincide with the actual earnings growth rate observed in the data for the corresponding year (the maroon curve in Panel A and the green curve in Panel B). The corresponding sets of growth-rate curves, however, do not coincide in any area of the income distribution.

Figure A.VI: Elasticity Implied by the Self-Employment Income Response



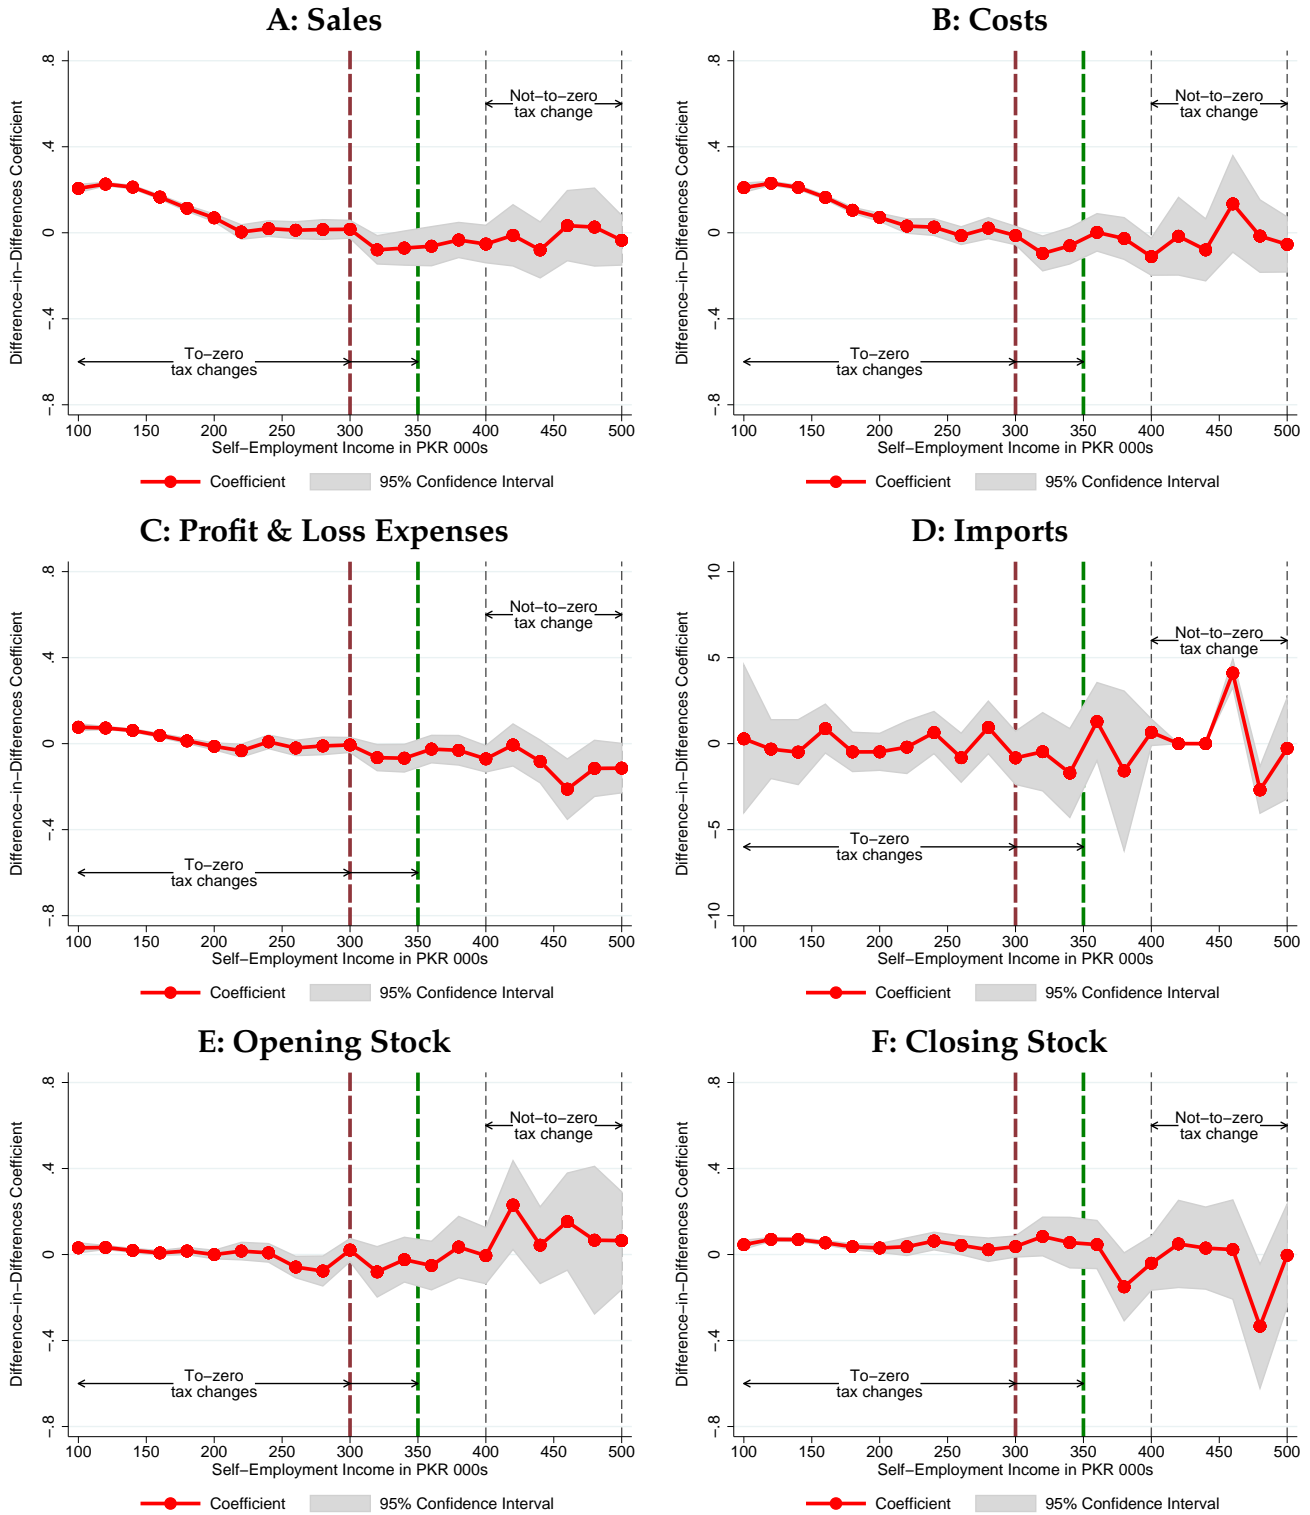
Notes: The figure illustrates the elasticity implied by the self-employment income response. The figure is created exactly similar to as Panel F of Figure II, the only difference being that the change in net-of-tax rate from period t to $t + 1$ experienced by taxpayers is instrumented by the double interaction term $SE_i \times post_t$. Each marker on the diagram, accordingly, represents the elasticity implied by the self-employment income response observed in the PKR 20k bin for the two post-reform years 2010-2011. The gray surface plot represents the 95% confidence interval around the coefficient. Note that the earnings responses exhibited in this paper are discontinuous responses that depend more on the *type* of tax change than the *size* of tax change. Elasticity is therefore not a very meaningful concept in this setup. Nevertheless the implausibly large elasticity implied by the response supports the claim that the response captures changes in tax evasion rather than effort.

Figure A.VII: Line Items Response



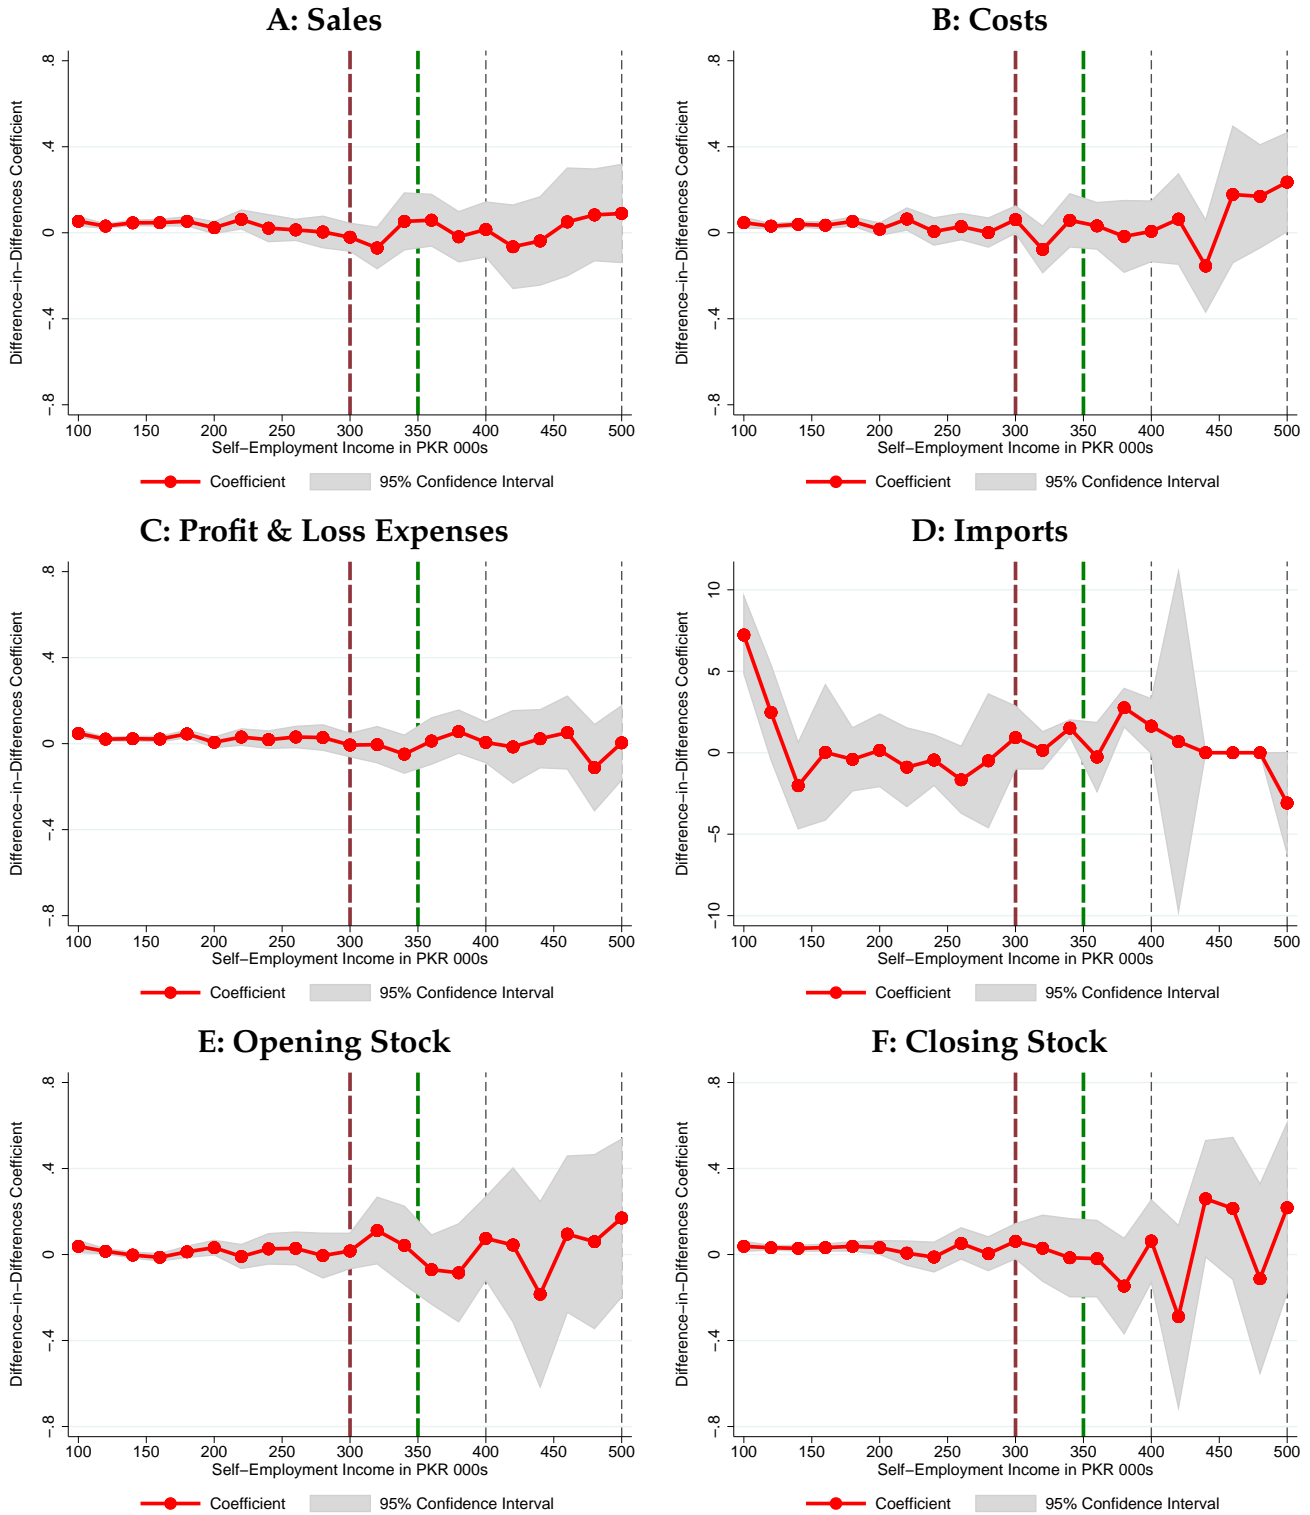
Notes: The figure looks at the anatomy of the self-employment income response depicted in Figure II, investigating if easy-to-misreport components of income—line items—respond differently from the other components. The detailed description of the line items used here are provided in Appendix A.1. To construct the plot, taxpayers are grouped into bins of PKR 20,000 on the basis of their base period income. Then, log change from period t to $t + 1$ averaged across taxpayers in the bin is plotted as a function of the base period income in the bin. Vertical lines demarcate the areas where the to-zero and not-to-zero changes shown in Figure IA–B were applied to.

Figure A.VIII: Line Items Response – 2009



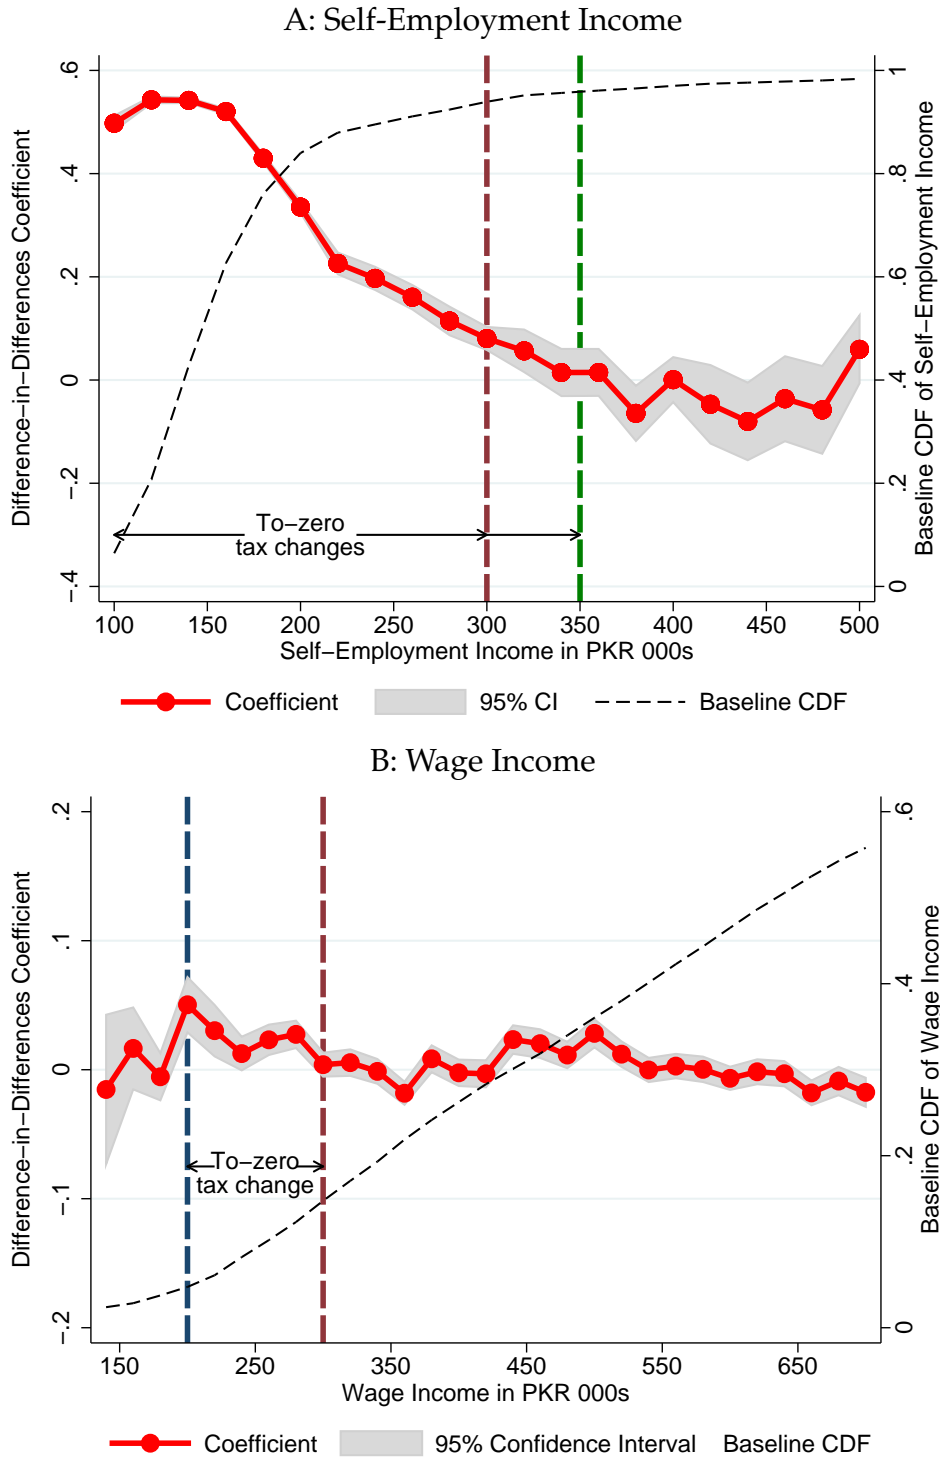
Notes: The figure plots the difference-in-difference versions of the line-items-response plots in Figure A.VII. To construct the plots, taxpayers are grouped into bins of PKR 20,000 on the basis of their base period income. Then, for each line item the analysis in Panel D of Figure II is replicated. Accordingly, a given marker in the plot represents the additional growth in the line item from year 2009 to 2010 observed for the treated taxpayers relative to the untreated taxpayers. The gray surface plots depict the 95% confidence interval around the coefficients. Vertical lines demarcate the areas where the to-zero and not-to-zero changes shown in Figure IA–B were applied to. Please see Appendix A.1 for the definitions of line items depicted here.

Figure A.IX: Line Items Response – Placebo



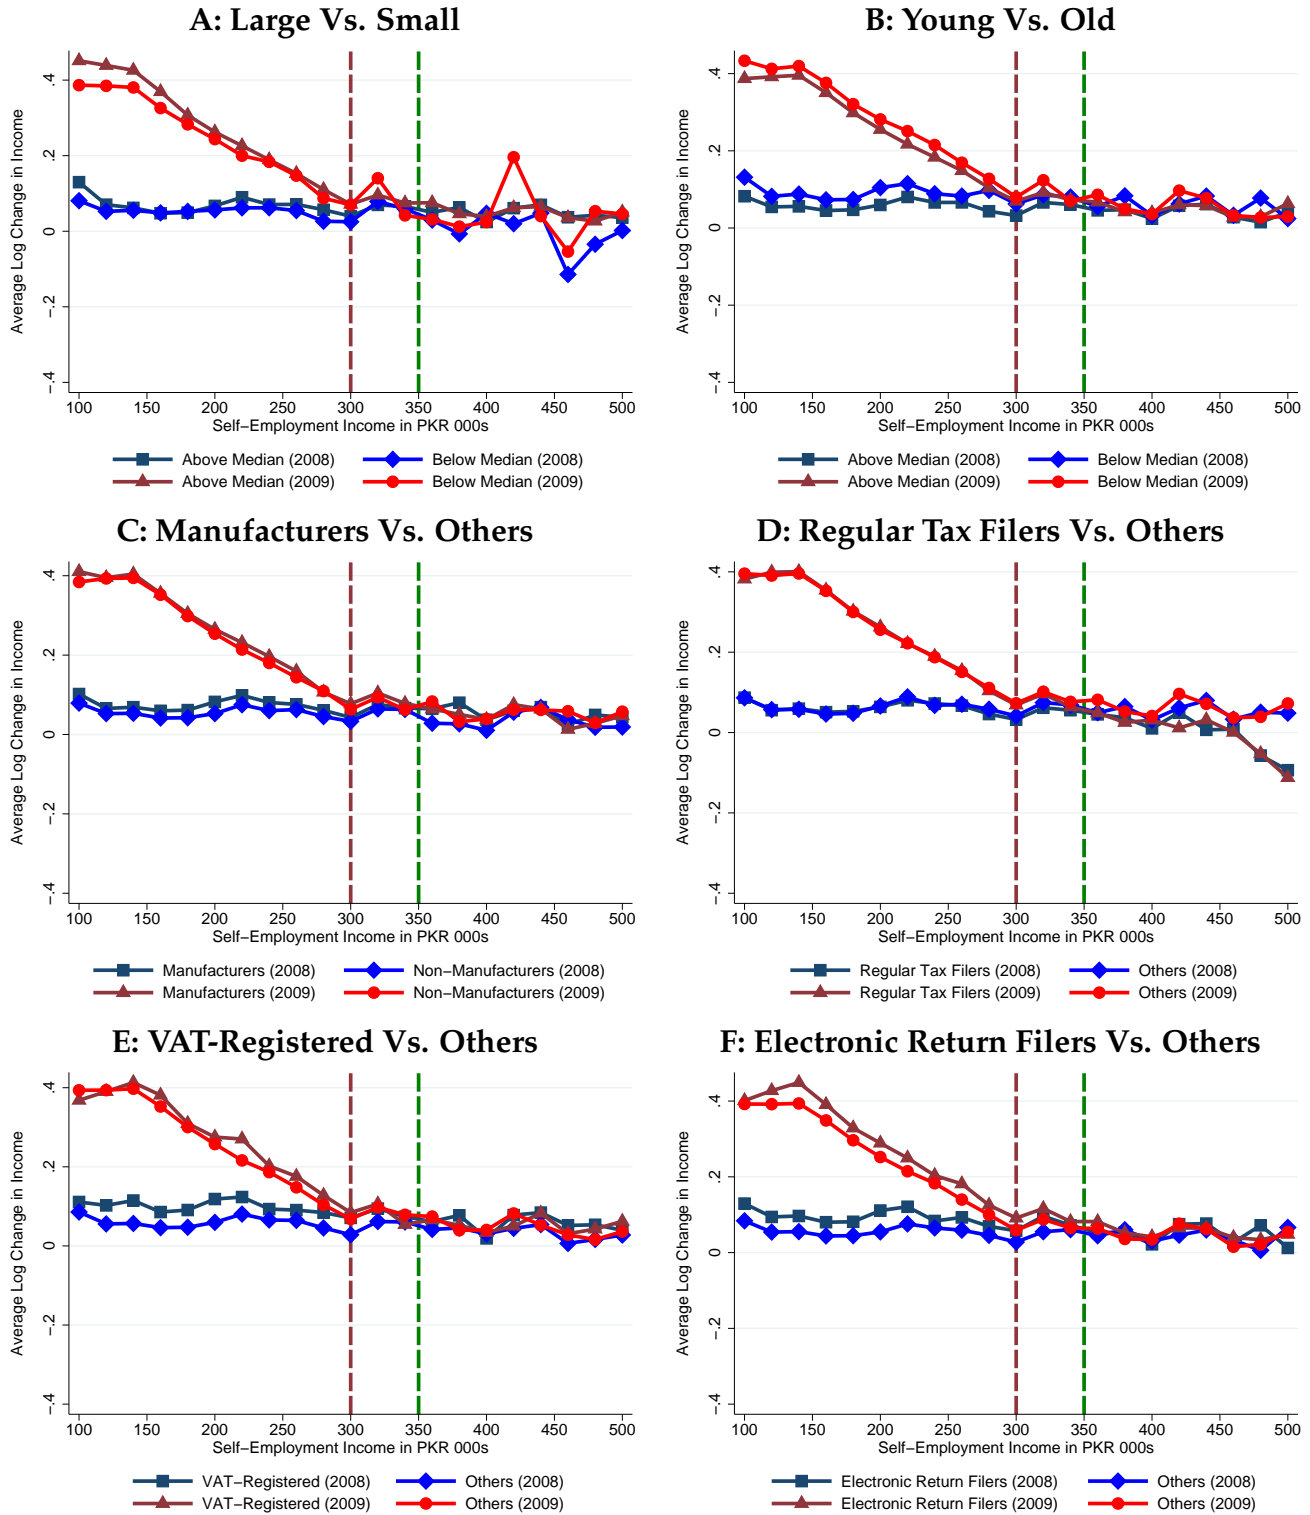
Notes: The figure plots the difference-in-difference versions of the line-items-response plots in Figure A.VII. To construct the plots, taxpayers are grouped into bins of PKR 20,000 on the basis of their base period income. Then, for each line item the analysis in Panel C of Figure II is replicated. Accordingly, a given marker in the plot represents the additional growth in the line item from year 2009 to 2010 observed for the treated taxpayers relative to the untreated taxpayers. The gray surface plots depict the 95% confidence interval around the coefficients. Vertical lines demarcate the areas where the to-zero and not-to-zero changes shown in Figure IA–B were applied to. Please see Appendix A.1 for the definitions of line items depicted here.

Figure A.X: Applicability of the Evasion Rates



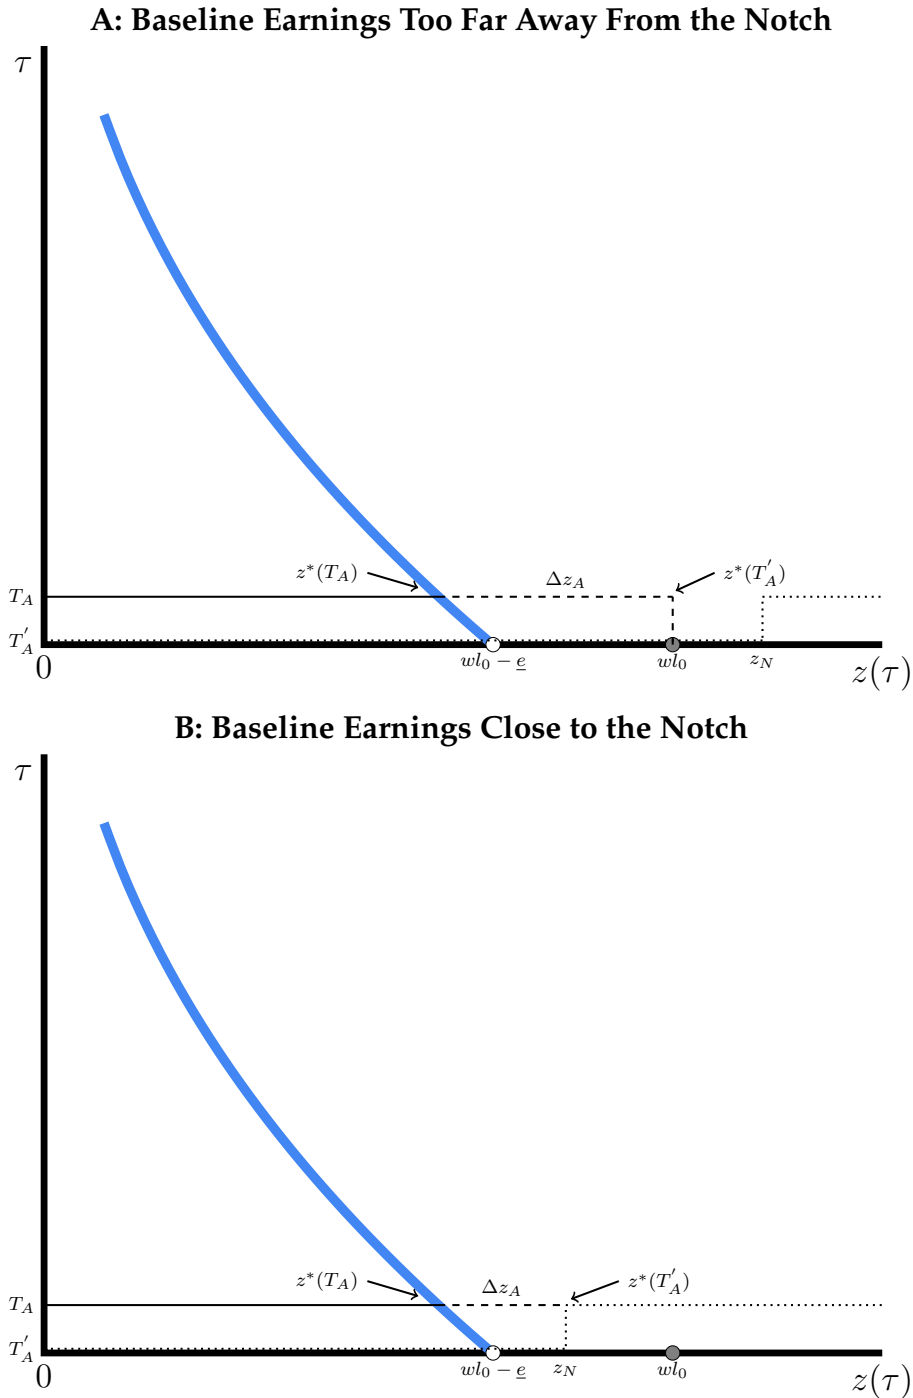
Notes: How representative are the evasion rates I report? The figure explores this question. Panel A superimposes the baseline (2009) CDF of self-employment income on Panel F of Figure II, demonstrating that PKR 150k represents more than the 60th percentile of the baseline self-employment income distribution. Since the 70% evasion rate of self-employment income I report is computed from the responses of taxpayers with baseline income up to PKR 150k only, it applies to both low- and middle-income taxpayers. Panel B superimposes the baseline (2009) CDF of wage income on Panel E of Figure III. It shows that the evasion rate of wage income applies only to the bottom-two deciles of the wage-income distribution only.

Figure A.XI: Heterogeneity in Self-Employment Income Response



Notes: The figure explores heterogeneity in self-employment income response. Taxpayers are grouped into bins of PKR 20,000 on the basis of their base period self-employment income, and then average log change in income from period t to $t + 1$ is plotted as a function of the base period income in the bin, stratifying the sample by the given trait j . The figure shows that there is hardly any difference in response across taxpayers with different characteristics but located in the same area of the income distribution. The details of taxpayer characteristics variables used here are given in Appendix A.1.

Figure A.XII: Earnings Responses with a Notched Post-Reform Tax Schedule



Notes: The figure relaxes the assumption of a linear tax schedule. The postreform tax schedule is now nonlinear, having a notch at z_N . Panel A illustrates that a taxpayer too far away from the notch would behave exactly similar to as in the baseline model (compare Panel A with Figure IVB). In contrast, a taxpayer close to the notch would not report its true income as the rate approaches zero (Panel B). Its to-zero response would, accordingly, be constrained (compare Δz_A in Panels A and B). This is strictly consistent with what we observe in the data. Average earnings response is nearly zero close to the notch in the postreform schedule, increases monotonically as we move left of the notch, reaches its peak at PKR 150k, and becomes stable thereafter (please see Figure II).

Table A.I: Filing Compliance

	Dependent Variable: Filer in period $t + 1$			
	Self-Employed		Wage Earners	
	(1)	(2)	(3)	(4)
$to-zero_{it}$	0.028 (0.001)		-0.066 (0.003)	
$placebo_{it}$		-0.024 (0.001)		-0.170 (0.002)
Constant	0.785 (0.000)	0.797 (0.000)	0.559 (0.001)	0.569 (0.001)
Observations	2,346,351	2,346,351	781,803	781,803

Notes: This table assesses if the propensity to file a return differs significantly between taxpayers experiencing the reduction of the rate to zero and the others. The dependent variable is a dummy that takes the value one if a tax filer in period t also files in the period $t + 1$. The to-zero variable is also a dummy, indicating if the taxpayer experiences a to-zero rate change. In placebo regressions, all tax changes are antedated by one year. I interpret the coefficient on the to-zero variable relative to the placebo coefficient to account for the fact that the propensity to file might vary across different areas of the income distribution. The results show that taxpayers are more likely to file a return when their tax rate is brought to zero.

Table A.II: Summary Statistics

	Full Sample		Analysis Sample	
	Self-Employed	Wage Earners	Self-Employed	Wage Earners
	(1)	(2)	(3)	(4)
<i>Observations</i>				
1. Taxable Income > 0 (#)	2,187,943	985,293	1,279,516	274,478
2. Self-Employment Income > 0 (%)	100	4.30	100	6.40
3. Wage Income > 0 (%)	0.60	100	0.40	100
4. Partnership Income > 0 (%)	3.50	1.10	1.80	1.20
5. Switchers (%)	0.90	1.30	0.90	2.70
<i>Outcomes</i>				
6. Taxable Income	3,825,812 (3,334,397,184) [147,800]	3,509,781 (1,408,862,848) [512,185]	160,154 (69,773) [135,000]	403,456 (153,240) [390,234]
7. Self-Employment Income	3,820,831 (3,362,148,352) [146,000]	154,147 (244,261) [88,889]	178,195 (16,954,248) [135,000]	86,146 (77,459) [65,500]
8. Wage Income	356,891 (1,236,341) [160,000]	3,049,276 (1,351,159,296) [512,887]	142,476 (187,928) [113,460]	397,086 (164,085) [382,493]
<i>Characteristics</i>				
9. Years Registered	7.04 (4.95)	7.80 (5.19)	7.39 (4.83)	7.79 (4.99)
10. Large City	0.39 (0.49)	0.52 (0.50)	0.38 (0.49)	0.47 (0.50)
11. Male	0.98 (0.12)	0.97 (0.18)	0.99 (0.11)	0.97 (0.18)
12. Age	39.55 (15.48)	43.41 (12.23)	40.39 (15.22)	43.27 (13.08)

Notes: This table presents the summary statistics of the data. The analysis sample for period t contains taxpayers for whom (i) $\log \frac{z_{it+1}}{z_{it}}$ is defined, (ii) $\log \frac{z_{it+1}}{z_{it}}$ is greater than the first percentile and less than the 99th percentile of the corresponding pooled distribution, and (iii) $z_{it} \in (80K \ 500K]$ if i is a self-employed and $z_{it} \in (140K \ 700K]$ if i is a wage earner. The detailed description of variables used here are provided in Appendix A.1. The first row of the table reports the number of self-employed and wage earners in the two samples who report positive taxable income. Rows 2-5 of the table shows the share of taxpayers with the given characteristic in the corresponding sample. Rows 6-12 report the mean of the outcome / characteristic variable in the corresponding sample with the standard error in parenthesis and the median in square brackets. The variable in row 10 of the table indicates that the taxpayer is located in one of the three big cities of Pakistan, Lahore, Karachi, and Islamabad.

Table A.III: Switching Probability

	Dependent Variable: Switches base in period $t + 1$	
	Self-Employed	Wage Earners
(1)	(2)	(3)
2007	-0.008 (0.000)	0.008 (0.001)
2008	-0.007 (0.000)	-0.004 (0.001)
2009	-0.007 (0.000)	-0.004 (0.001)
2010	-0.008 (0.000)	0.011 (0.001)
Constant	0.015 (0.000)	0.025 (0.001)
Observations	1,279,516	274,478

Notes: This table assesses if the propensity to switch the tax base differs significantly across years. I regress an indicator that a self-employed in period t becomes a wage-earner in period $t + 1$ on the full set of year dummies. Column (2) reports the results. The coefficient for a given year t represents the additional probability of switching in the given year relative to the omitted category (2006). Column (3) reports the results from a parallel regression of switching in the opposite direction. Two results are noteworthy. First, very few taxpayer switch from the self-employed to wage-earner and vice versa. Second, there is little meaningful difference in the switching probabilities across years.

Table A.IV: Self-Employment Income Response (Complete Panel)

		Dependent variable: Log change in self-employment income from period t to $t + 1$						
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$to\text{-}zero \times 2009$	0.271 (0.003)	0.290 (0.001)	0.262 (0.004)	0.240 (0.004)	0.269 (0.003)	0.262 (0.003)	0.233 (0.004)
	$to\text{-}zero \times 2010$	0.134 (0.003)	0.137 (0.001)	0.116 (0.004)	0.101 (0.004)	0.134 (0.003)	0.130 (0.003)	0.094 (0.004)
	$not\text{-}to\text{-}zero \times 2009$	-0.003 (0.005)	0.014 (0.005)	0.001 (0.008)	0.004 (0.008)	-0.005 (0.005)	-0.001 (0.006)	0.006 (0.008)
	Fixed effects:							
	Year	Yes	No	Yes	Yes	Yes	Yes	Yes
	Industry	No	No	Yes	Yes	No	No	Yes
	Industry \times year	No	No	No	Yes	No	No	Yes
	Region	No	No	No	No	Yes	Yes	Yes
	Region \times year	No	No	No	No	No	Yes	Yes
	Time trend:							
	Linear	No	Yes	No	No	No	No	No
	Pre-reform mean of the dependent variable	0.049	0.049	0.068	0.068	0.049	0.049	0.068
	Observations	1,214,538	1,214,538	255,153	255,153	1,212,252	1,212,252	254,264

Notes: The table reports the estimates from equation (1). The first columns correspond to the baseline specification; the second column replaces the year fixed effects with a linear time trend; and the rest of the columns add year, industry, industry \times year, region and region \times year fixed effects. The details of the industry and region classifications are provided in Appendix A.1. I do not observe industry classification for all taxpayers, owing to which the numbers of observations in the last two columns are lower. The control group here comprises self-employed who do not experience any rate change. Standard errors are in parenthesis, which have been clustered at the individual level.

Table A.V: Self-Employment Income Response – Before-After Research Design

		Dependent variable: Log change in self-employment income from period t to $t + 1$							
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$to-zero \times 2009$	0.292 (0.002)	0.295 (0.002)	0.288 (0.002)	0.288 (0.002)	0.287 (0.003)	0.283 (0.003)	0.259 (0.004)	0.255 (0.004)
	$to-zero \times 2010$	0.146 (0.002)	0.143 (0.002)	0.142 (0.002)	0.142 (0.002)	0.143 (0.003)	0.138 (0.003)	0.142 (0.003)	0.137 (0.003)
	$not-to-zero \times 2009$	-0.004 (0.005)	-0.001 (0.005)	-0.003 (0.005)	-0.004 (0.005)	-0.010 (0.006)	-0.009 (0.006)	-0.021 (0.007)	-0.020 (0.007)
	$post$	0.014 (0.002)	-0.004 (0.002)	0.017 (0.002)	0.017 (0.002)	0.006 (0.003)	0.007 (0.003)	-0.061 (0.061)	0.103 (0.074)
	Fixed effects:								
	Year	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
	Region	No	No	Yes	No	No	Yes	No	Yes
	Tax Office	No	No	No	Yes	No	Yes	No	Yes
	Industry	No	No	No	No	Yes	Yes	Yes	Yes
	Industry \times year	No	No	No	No	No	No	Yes	Yes
	Time trend:								
	Linear	No	Yes	No	No	No	No	No	No
	Pre-reform mean of the dependent variable	0.049	0.049	0.049	0.049	0.068	0.068	0.068	0.068
	Observations	1,214,538	1,214,538	1,214,524	1,152,314	255,153	255,140	255,153	255,140

Notes: This table reports the estimates from the before-after analog of equation (1). The first column in the table correspond to the baseline specification; column (2) replaces the year fixed effects with a linear time trend; and the rest of the columns add additional control variables. The definitions of the control variables used here are provided in Appendix A.1. I do not observe industry classification for all taxpayers, owing to which the numbers of observations in the last four columns are lower. Standard errors are in parenthesis, which have been clustered at the individual level.

Table A.VI: Robustness of the Self-Employment Income Response Estimates

	Dependent variable: Log change in self-employment income from period t to $t + 1$								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$to-zero \times SE \times 2009$	0.271 (0.003)	0.245 (0.003)	0.271 (0.003)	0.280 (0.003)	0.281 (0.003)	0.262 (0.004)	0.269 (0.003)	0.268 (0.003)	0.260 (0.004)
$to-zero \times SE \times 2010$	0.134 (0.003)	0.131 (0.003)	0.134 (0.003)	0.128 (0.003)	0.136 (0.003)	0.116 (0.004)	0.134 (0.003)	0.134 (0.003)	0.114 (0.004)
$not-to-zero \times SE \times 2009$	-0.003 (0.005)	-0.007 (0.006)	-0.003 (0.005)	0.003 (0.006)	0.006 (0.005)	0.001 (0.008)	-0.005 (0.005)	-0.005 (0.006)	0.001 (0.008)
Fixed effects:									
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry	No	No	No	No	No	Yes	No	No	Yes
Region	No	No	No	No	No	No	Yes	No	Yes
Tax office	No	No	No	No	No	No	No	Yes	Yes
Sample:									
$z_{it} \in (80K \ 500K]$	Yes	Yes	Yes	No	No	Yes	Yes	Yes	Yes
Bunchers dropped	No	Yes	No	No	No	No	No	No	No
$z_{it} \in (0 \ 500K]$	No	No	No	Yes	No	No	No	No	No
$z_{it} \in (0 \ \infty)$	No	No	No	No	Yes	No	No	No	No
Pre-reform mean of the dependent	0.049	0.052	0.049	0.052	0.052	0.049	0.049	0.049	0.049
Observations	1,214,538	770,989	1,213,768	1,257,607	1,267,921	255,132	1,214,524	1,152,314	255,119

Notes: This table assesses the robustness of the estimates from equation (1). Column (2) drops taxpayers who *bunch* at the notches from concerns that their base period income might be low owing to strong, local incentives created by the notches or because these taxpayers might be special. Column(3) drops taxpayers around the income-composition notch, where the classification of a taxpayer switches from self-employed to wage-earners and vice versa. Columns (4)–(5) increase the range of the data from $z_{it} \in (80K \ 500K]$ in the baseline results to $z_{it} \in (0 \ 500K]$ in column (4) and $z_{it} > 0$ in column (5). Columns (6)–(9) add additional control variables into specification (1): industry fixed effects in column (6), region fixed effects in column (7), tax office fixed effects in column (8), and all these fixed effects together in column (9). The control group here ($SE=0$) comprises taxpayers whose self-employment income is subject to the schedule for wage earners. Standard errors are in parenthesis, which have been clustered at the individual level.

Table A.VII: Self-Employment Income Response – Placebo Specification

	Dependent variable: Log change in self-employment income from period t to $t + 1$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>to-zero</i> × 2008	-0.003 (0.003)	0.019 (0.001)	0.002 (0.004)	0.001 (0.004)	-0.004 (0.003)	-0.004 (0.003)	0.001 (0.004)
<i>not-to-zero</i> × 2008	0.001 (0.006)	0.023 (0.005)	0.002 (0.008)	0.002 (0.008)	-0.000 (0.006)	-0.000 (0.006)	0.001 (0.008)
Fixed effects:							
Year	Yes	No	Yes	Yes	Yes	Yes	Yes
Industry	No	No	Yes	Yes	No	No	Yes
Industry × year	No	No	No	Yes	No	No	Yes
Region	No	No	No	No	Yes	Yes	Yes
Region × year	No	No	No	No	No	Yes	Yes
Time trend:							
Linear	No	Yes	No	No	No	No	No
Pre-reform mean of the dependent variable	0.041	0.041	0.052	0.052	0.041	0.041	0.052
Observations	754,011	754,011	136,137	136,137	754,011	754,011	136,137

Notes: This table reports the results from placebo regressions corresponding to equation (1). I restrict the sample for this analysis to prereform periods (2006-2009) only and pretend that the *to-zero* and *not-to-zero* changes took place one year earlier than they actually did. The first columns correspond to the baseline specification; the second column replaces the year fixed effects with a linear time trend; and the rest of the columns add year, industry, industry × year, region and region × year fixed effects. The details of the industry and region classifications are provided in Appendix A.1. I do not observe industry classification for all taxpayers, owing to which the numbers of observations in the last two columns are lower. The control group here comprises self-employed who do not experience any rate change. Standard errors are in parenthesis, which have been clustered at the individual level.

Table A.VIII: Self-Employment Income Response – With Controls for Mean-Reversion

	Dependent variable: Log change in taxable income from period t to $t + 1$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>to-zero</i> × 2009	0.259 (0.004)	0.291 (0.001)	0.229 (0.005)	0.210 (0.005)	0.258 (0.004)	0.249 (0.004)	0.203 (0.005)
<i>to-zero</i> × 2010	0.119 (0.004)	0.166 (0.001)	0.136 (0.005)	0.119 (0.005)	0.119 (0.004)	0.112 (0.004)	0.109 (0.005)
<i>not-to-zero</i> × 2009	0.026 (0.008)	0.036 (0.007)	-0.016 (0.009)	-0.012 (0.009)	0.025 (0.008)	0.028 (0.008)	-0.009 (0.009)
Fixed effects:							
Year	Yes	No	Yes	Yes	Yes	Yes	Yes
Industry	No	No	Yes	Yes	No	No	Yes
Industry × year	No	No	No	Yes	No	No	Yes
Region	No	No	No	No	Yes	Yes	Yes
region × year	No	No	No	No	No	Yes	Yes
Time trend:							
Linear	No	Yes	No	No	No	No	No
Pre-reform mean of the dependent variable	0.046	0.046	0.065	0.065	0.065	0.065	0.065
Observations	1,214,524	1,214,524	255,150	255,150	1,212,238	1,212,238	254,261

Notes: The table illustrates the robustness of the results against mean-reversion. I replicate Table A.IV after including the standard controls for mean reversion into equation (1). These controls include the log of base period income and a ten-piece spline of log of base period income. Inclusion of these additional controls makes very little difference to the results, showing that mean-reversion is not a significant concern in this setting. The first two columns correspond to the baseline specification; columns (3)-(4) replace year fixed effects with a linear time trend; columns (5)-(6) include year, industry and industry × year fixed effects, allowing taxpayers in each industry their own earnings growth trend. The details of the industry classification are provided in Appendix A.1. I do not observe industry classification for all taxpayers, owing to which the numbers of observations in the last two columns are lower. The control group here comprises self-employed who do not experience any rate change. All columns contain controls for mean-reversion that include the log of base period income and a ten-piece spline of the log base period income. Standard errors are in parenthesis, which have been clustered at the individual level.

Table A.IX: Robustness of the Wage Income Response Estimates

	Dependent variable: Log change in wage income from period t to $t + 1$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>A: Tax-Driven Response</u>								
<i>to-zero</i>	0.014 (0.002)	0.014 (0.002)	0.014 (0.002)	0.008 (0.004)	0.010 (0.002)	0.010 (0.002)	0.014 (0.002)	0.020 (0.002)
<i>not-to-zero</i>	0.001 (0.002)	0.002 (0.002)	0.002 (0.002)	0.007 (0.003)	0.002 (0.002)	0.002 (0.002)	-0.000 (0.002)	0.003 (0.002)
<u>B: Placebo</u>								
<i>to-zero</i>	0.001 (0.002)	-0.000 (0.002)	0.000 (0.002)	-0.013 (0.004)	0.001 (0.002)	-0.001 (0.002)	0.001 (0.002)	-0.002 (0.002)
<i>not-to-zero</i>	0.010 (0.002)	0.011 (0.002)	0.009 (0.002)	0.001 (0.004)	0.007 (0.002)	0.008 (0.002)	0.010 (0.002)	0.005 (0.002)
Fixed effects:								
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region	No	Yes	No	No	No	Yes	No	No
Tax office	No	No	Yes	No	No	Yes	No	No
Age deciles	No	No	No	Yes	No	No	No	No
Years-registered deciles	No	No	No	No	Yes	Yes	No	No
Sample:								
$z_{it} \in (140K\ 700K]$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Bunchers dropped	No	No	No	No	No	No	Yes	No
$z_{it} \in (0\ 1000K]$	No	No	No	No	No	No	No	Yes
Pre-reform mean of the dependent variable	0.162	0.162	0.162	0.162	0.162	0.162	0.162	0.166
Observations	240,649	240,643	236,482	85,537	236,490	236,468	226,889	309,921

Notes: This table reports the estimates from equation (1). The first column corresponds to the baseline specification; columns (2)–(6) add additional control variables; column (7) drops taxpayers who *bunch* at the notches from concerns that their base period income might be low owing to the strong, local incentives created by the notches, or because these taxpayers might be special; columns (8) increase the range of the data from $z_{it} \in (140K\ 700K]$ in other specifications to $z_{it} \in (0\ 1000K]$. The definitions of variables used here are in Appendix A.1. All regression include log base period income as a control to account for mean-reversion. Standard errors are in parenthesis, which have been clustered at the individual level.

Table A.X: Evasion Rate of Self-Employment Income

Income	Earnings Response ($\tau \rightarrow 0$)	Earnings Response ($\tau \not\rightarrow 0$)	Difference	Evasion Rate (%)
(1)	(2)	(3)	(4)	(5)
80-100K	71,754 (1,251)	1,765 (999)	69,989 (1,601)	72.7 (1.7)
100-150K	92,280 (784)	3,198 (1,810)	89,082 (1,972)	69.9 (1.5)
150-200K	92,440 (938)	4,070 (1,425)	88,370 (1,706)	50.2 (1.0)
200-250K	70,537 (1,116)	5,421 (1,897)	65,117 (2,201)	28.6 (1.0)
250-300K	28,868 (1,172)	4,914 (1,720)	23,954 (2,081)	8.5 (0.7)
300-350K	20,124 (918)	2,978 (1,042)	17,146 (1,389)	5.2 (0.4)

Notes: The table illustrates that the evasion rates I report are robust to any salience-related concerns. I present the estimates of the rates of evasion of self-employment income from equation (15). Column (1) shows the income segment; column (2) the earnings response produced by the to-zero change; column (3) the earnings response produced by the equal-sized, not-to-zero change; column (4) the difference between the two responses; and column (5) the average evasion rate in the segment. The estimates in column (2) are computed from the self-employment income response to the movements of the exemption cutoff in 2010–11 shown in Table I. The estimates in column (3) are computed using elasticities estimated for the same group of taxpayers in Kleven & Waseem (2013). I inflate the observed elasticities estimated by Kleven & Waseem (2013) by a factor of 10 (which implicitly assumes a value of 0.1 for the attenuation factor θ) to illustrate that applying even an extreme level of salience correction to the not-to-zero responses does not make any difference to my results. The details on how the estimates in columns (2)–(3) have been computed are in section V. The difference between the two set of estimates represents average \underline{e} for taxpayers in the segment. I obtain the percentage evasion rate implied by the average \underline{e} by dividing it with the average income in the segment—see column (5). The standard errors are in parenthesis.

Table A.XI: Heterogeneity in Self-Employment Income Response (Size)

	Dependent variable: Log change in self-employment income from period t to $t + 1$					
	(1)	(2)	(3)	(4)	(5)	(6)
$to-zero \times 2009$	0.295 (0.009)	0.267 (0.022)	0.293 (0.010)	0.252 (0.021)	0.294 (0.013)	0.241 (0.022)
$to-zero \times 2010$	0.160 (0.009)	0.165 (0.024)	0.168 (0.010)	0.142 (0.022)	0.162 (0.012)	0.115 (0.025)
$to-zero \times trait \times 2009$	-0.038 (0.010)	0.058 (0.036)	-0.077 (0.011)	0.040 (0.042)	-0.120 (0.013)	0.066 (0.048)
$to-zero \times trait \times 2010$	-0.043 (0.010)	0.037 (0.048)	-0.072 (0.011)	-0.015 (0.053)	-0.081 (0.013)	0.052 (0.059)
$to-zero \times trait$	0.012 (0.007)	-0.001 (0.019)	0.030 (0.007)	0.001 (0.021)	0.044 (0.009)	-0.059 (0.027)
$trait \times post$	-0.002 (0.010)	-0.039 (0.036)	0.008 (0.010)	-0.008 (0.042)	0.010 (0.013)	-0.073 (0.050)
Trait = Size	Above Vs. Below Median		Top Vs. Bottom Quartile		Top Vs. Bottom Decile	
Controls for Base Period Income	No	Yes	No	Yes	No	Yes
Pre-reform mean of the dependent variable	0.048	0.048	0.052	0.052	0.053	0.053
Observations	823,552	823,552	414,183	414,183	165,099	165,099

Notes: This table explores heterogeneity in self-employment income response across taxpayers with different firm size. The size thresholds are defined on the basis of observed yearly firm-size distribution. The odd-numbered columns report the results from equation (3) with no controls for the base period income. The even-numbered columns control for the base-period income in a non-parametric way: (i) taxpayers are grouped into bins of PKR 20K on the basis of their base period income, (ii) the regression (3) is run separately in each bin, and (iii) the aggregate estimates are generated as the weighted average of the bin-level estimates, with the weights provided by the distribution of firm-size in the binned income distribution. I do not observe firm-size for all taxpayers, owing to which sample size here is lower than in Table A.IV. The estimates show that there is hardly any difference in the self-employment income response across taxpayers with different firm size, once taxpayers in the same area of the income distribution are compared.

Table A.XII: Heterogeneity in Self-Employment Income Response

	Dependent variable: Log change in self-employment income from period t to $t + 1$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$to-zero \times 2009$	0.280 (0.004)	0.274 (0.010)	0.270 (0.004)	0.262 (0.011)	0.282 (0.003)	0.271 (0.010)	0.288 (0.004)	0.260 (0.015)
$to-zero \times 2010$	0.141 (0.003)	0.165 (0.012)	0.131 (0.003)	0.152 (0.013)	0.147 (0.003)	0.164 (0.012)	0.153 (0.003)	0.194 (0.027)
$to-zero \times trait \times 2009$	-0.019 (0.005)	-0.010 (0.028)	0.029 (0.005)	0.011 (0.017)	-0.100 (0.007)	-0.049 (0.039)	-0.065 (0.005)	0.014 (0.017)
$to-zero \times trait \times 2010$	-0.000 (0.005)	-0.030 (0.031)	0.029 (0.005)	0.009 (0.019)	-0.075 (0.007)	-0.057 (0.041)	-0.062 (0.005)	-0.029 (0.030)
$to-zero \times trait$	0.021 (0.004)	0.067 (0.022)	0.029 (0.004)	-0.014 (0.010)	0.049 (0.004)	0.048 (0.031)	0.026 (0.004)	-0.002 (0.013)
$trait \times post$	-0.011 (0.005)	0.007 (0.026)	-0.037 (0.005)	-0.015 (0.016)	0.006 (0.006)	0.032 (0.037)	0.009 (0.005)	0.006 (0.022)
Trait	Manufacturers		Regular Tax Filers		VAT-Registered		Electronic Filers	
Percent with Trait	24.1		41.8		4.9		11.7	
Controls for Base Period Income	No	Yes	No	Yes	No	Yes	No	Yes
Pre-reform mean of the dependent variable	0.049	0.049	0.049	0.049	0.049	0.049	0.049	0.049
Observations	1,221,165	1,221,165	1,221,165	1,221,165	1,221,165	1,221,165	1,221,165	1,221,165

Notes: This table explores heterogeneity in self-employment income response across taxpayers with different characteristics. The odd-numbered columns report the results from equation (3) with no controls for the base period income. The even-numbered columns control for the base-period income in a non-parametric way: (i) taxpayers are grouped into bins of PKR 20K on the basis of their base period income, (ii) the regression (3) is run separately in each bin, and (iii) the aggregate estimates are generated as the weighted average of the bin-level estimates, with the weights provided by the distribution of characteristic j in the binned income distribution. The estimates show that there is hardly any difference in the self-employment income response across taxpayers with different characteristics, once taxpayers in the same area of income distribution are compared.

Table A.XIII: Heterogeneity in Self-Employment Income Response (Age)

	Dependent variable: Log change in self-employment income from period t to $t + 1$					
	(1)	(2)	(3)	(4)	(5)	(6)
$to-zero \times 2009$	0.251 (0.008)	0.257 (0.026)	0.232 (0.010)	0.221 (0.031)	0.205 (0.014)	0.119 (0.052)
$to-zero \times 2010$	0.105 (0.007)	0.242 (0.033)	0.100 (0.009)	0.195 (0.040)	0.087 (0.013)	0.091 (0.049)
$to-zero \times trait \times 2009$	0.003 (0.009)	-0.007 (0.027)	-0.009 (0.012)	-0.006 (0.035)	0.002 (0.017)	0.041 (0.066)
$to-zero \times trait \times 2010$	0.011 (0.009)	-0.098 (0.037)	0.006 (0.012)	-0.032 (0.046)	0.013 (0.017)	0.013 (0.071)
$to-zero \times trait$	-0.046 (0.007)	-0.009 (0.021)	-0.043 (0.009)	-0.024 (0.029)	-0.036 (0.012)	0.064 (0.057)
$trait \times post$	-0.019 (0.009)	0.052 (0.029)	-0.016 (0.011)	0.013 (0.038)	-0.028 (0.016)	-0.044 (0.065)
Trait = Age	Above Vs. Below Median		Top Vs. Bottom Quartself-employmentile		Top Vs. Bottom Decile	
Controls for Base Period Income	No	Yes	No	Yes	No	Yes
Pre-reform mean of the dependent variable	0.065	0.065	0.064	0.064	0.062	0.062
Observations	252,750	252,750	133,063	133,063	55,004	55,004

Notes: This table explores heterogeneity in the self-employment income response across taxpayers of different ages. The age thresholds are defined on the basis of observed yearly age distribution of self-employed. The odd-numbered columns report the results from equation (3) with no controls for the base period income. The even-numbered columns control for the base-period income in a non-parametric way: (i) taxpayers are grouped into bins of PKR 20K on the basis of their base period income, (ii) the regression (3) is run separately in each bin, and (iii) the aggregate estimates are generated as the weighted average of the bin-level estimates, with the weights provided by the distribution of the taxpayer's age in the binned income distribution. I do not observe ages of all taxpayers, owing to which sample size here is lower than in Table A.IV. The estimates show that there is hardly any difference in the self-employment income response across taxpayers of different ages, once taxpayers in the same area of the income distribution are compared.