

Overclaimed Refunds, Undeclared Sales, and Invoice Mills: Nature and Extent of Noncompliance in a Value-Added Tax

Mazhar Waseem*

February 2022

Abstract

Value-added tax has seen phenomenal expansion in recent decades. Its appeal in part lies in its robustness to tax evasion relative to other tax instruments. Exploiting a tax reform from Pakistan that cuts the tax rate on five major industries of the country substantially, I estimate the size and nature of VAT evasion in the treated industries, finding that it ranges from 31–46% of the potential revenue. One important channel through which the evasion occurs is the overclaim of refunds, which constitute 11–23% of the potential revenue. Roughly two-fifths of the overclaimed refund is based on spurious invoices issued by invoice mills. Qualitatively, noncompliance is stronger in the latter stages of the supply chain, but it runs deep inside the chain, suggesting that on their own the self-enforcement mechanisms built into a VAT do not deter tax evasion fully.

Keywords: VAT, Tax Evasion, Firm Behavior

JEL Classification: H25, H26, H32

*I thank Miguel Almunia, Zareh Asatryan, Pierre Bachas, Anne Brockmeyer, Anders Jensen, Florian Neumeier, Juan Carlos Suarez Serrato, Joel Slemrod, Christian Traxler, David E. Wildasin for helpful insights. I also thank the seminar participants at the University of Manchester, the National Tax Association, the University of Freiberg, the Institute of Fiscal Studies, the University of Mannheim, Economic Research South Africa, CESifo Public Economics, RIDGE Annual Conference, and International Online Public Finance Seminar Series for helpful suggestions and comments. Financial support from the Economic and Social Research Council (ESRC), UK's Future Research Leaders Program is gratefully acknowledged. Email Address: Mazhar Waseem mazhar.waseem@manchester.ac.uk.

I Introduction

In one of the most influential results in public finance, [Atkinson & Stiglitz \(1976\)](#) show that under a fairly broad set of conditions employing a differential commodity tax when the government has access to nonlinear income taxation is not optimal. Notwithstanding this canonical result, a broad-based VAT based often on nonuniform rates continues to be applied together with income tax in both rich and emerging economies. In fact, its share in government revenue has been rising steadily in recent decades ([OECD, 2017](#); [International Tax Dialogue, 2013](#)). This apparent discord between theory and practice in large part reflects enforcement concerns because among the class of production-efficient tax instruments the VAT is perceived as the most robust to tax evasion ([Slemrod & Velayudhan, 2020](#)). This belief has underpinned VAT's steady expansion over the last half century.

Recently, however, a few cracks have begun to appear in this consensus. For example, Malaysia has replaced its VAT (called Goods and Services Tax) with a turnover tax (called Sales and Services Tax) from September 2018 and Zambia nearly did so in 2019-2020.¹ The discontent with the VAT in these and similar other settings stems from its two well-known vulnerabilities ([Keen, 2007](#)). First, the self-enforcement forces built into a VAT work only on firm-to-firm transactions and break down at the final production stage, where sales to consumers take place ([Pomeranz, 2015](#); [Naritomi, 2019](#); [Waseem, 2020](#)). This last-mile problem is particularly severe in developing economies where the final production stage is often fragmented, being composed of small, informal firms ([Almunia et al., 2022](#)). Second, the destination-based design of VAT necessitates that any tax collected on intermediates be refunded to exporters. Refund payment provides another channel through which government revenue can be misappropriated, a channel not present in alternative tax instruments. While theoretical mechanisms underlying VAT noncompliance are known, we still do not understand fully how important they are empirically. The key focus of this paper is to provide empirical evidence on the size and nature of VAT noncompliance in a representative emerging economy and to examine what mechanisms underlie it.

For this purpose, I leverage a novel tax reform from Pakistan that reduced the tax rate on five major industries of the country substantially. Before the reform, in accordance with the destination principle, only exports of these industries were zero-rated

¹Please see Richard Asquith's blog at [avalara.com](#) for the policy changes in Malaysia and Zambia. Specifically, the Malaysian change is documented [here](#) and the Zambian [here](#).

and their imports and domestic supplies were subject to the standard rate. The reform reduced the rate applicable to final goods sold and important intermediates purchased by the treated industries to zero, thereby reducing their tax liability sharply. With the reduction in tax liability, the incentive to misreport weakens so that the VAT reports of firms must converge toward their true values given that evasion remains costly but offers little benefit after the reform. This sharp change in incentives allows me to uncover the extent of evasion in the treated industries as it existed at the baseline. I do so using a simple difference-in-differences framework, comparing outcomes across firms in the treated and untreated industries. Identification in this setup presumes that the trajectory of outcomes would have been similar across the compared groups absent the tax reform. Using the standard event study plots spanning 84 pre-reform and 72 post-reform periods, I show that this indeed is highly likely. I run many robustness checks to rule out other identification concerns, including any spillovers between the two groups.

In a first-best world where firms report truthfully, reducing the VAT rate to zero should stimulate activity in the treated industries. I, however, find that the activity reported by the treated industries contracts sharply after the tax cut. Specifically, reported purchases fall by 42 log-points, sales by 22 log-points, exports by 11 log-points, and non-export sales by 8 log-points. Not only are the elasticities implied by these responses of opposite sign, they are large in size, ranging from -4.3 to -0.8. Using a simple conceptual framework, I show that the size and direction of these changes is consistent with, and can only be reconciled by, the presence of large misreporting at the baseline. I back out the level of tax evasion implied by this misreporting, finding that it ranges between 31% and 46% of the potential revenue in the four baseline years. A large proportion of this evasion occurs through the channel of overclaimed refunds, which amount to 11–23% of the potential revenue. Qualitatively, the non-compliance is stronger at the later stages of the supply chain. But it runs deep inside the chain, where transactions are largely subject to third-party information, suggesting that the self-enforcement mechanisms built into a VAT on their own do not deter tax evasion fully.

In developing economies, VAT supply chains are rarely complete. When a VAT chain breaks, the tax charged at the pre-break stages cannot be claimed back in the post-break stage. This creates arbitrage opportunities, which are exploited by firms called invoice mills (Keen & Smith, 2006). These firms engage in no real business activity and exist solely to trade in VAT invoices. Invoice mills are a poorly understood

phenomenon. There is little micro-based evidence in the existing literature on how they operate and how much revenue loss they cause.² I show that invoice mills are a key mechanism through which the overclaim of VAT refund takes place. On average, around 83% of their transactions are with exporters, who use these spurious invoices to claim money back from the government. I report above that in the baseline year the amount overclaimed by exporters was 23% of the potential revenue. Roughly 37% of this amount was appropriated using invoices issued by invoice mills.

Stiglitz (2010) argues that devising evasion-resistant tax structures must be the principal aim of the tax policy in developing economies. Understanding the nature and size of tax evasion is the first step toward devising such tax structures. The primary contribution of this paper is to do so in the context of an emerging-economy VAT. In this effort, the paper contributes to a developing literature that studies the enforcement properties of VAT in weak enforcement settings (see for example Pomeranz, 2015; Agrawal & Zimmermann, 2019; Almunia *et al.*, 2019; Shah, 2019; Waseem, 2020). This paper contributes to two other literatures as well. First, the traditional public finance theory has been developed keeping in view the compliance and information environments of high-income economies. There is a growing realization now that the policy predictions of these models do not generalize to low-compliance, high-informality settings of emerging economies (Brockmeyer *et al.*, 2019b; Best *et al.*, 2015; Basri *et al.*, 2019). To adapt the optimal tax theory to these settings, one needs to take into account how the behavior of economic agents there differs from the developed world. This paper does so in the context of VAT compliance, introducing in the process new economic agents—invoice mills—whose behavior has not been studied so far. Finally, this paper adds to a literature that uses administrative tax return data to study tax compliance as a key constraint on the development of fiscal capacity in developing and emerging economies (see for example Bachas & Soto, 2019; Brockmeyer *et al.*, 2019a; Slemrod *et al.*, 2022).

²Mittal *et al.* (2018) develop a machine learning algorithm to identify bogus VAT firms. They implement their algorithm using VAT returns from the National Capital Territory of Delhi, India. The focus of this paper goes beyond identifying to documenting the behavior of invoice mills, in particular their role in facilitating noncompliance.

II Conceptual Framework

This section develops a simple framework to highlight key mechanisms through which VAT evasion occurs and to derive a formula for its magnitude.

II.A Setup

Consider a supply chain consisting of J production stages indexed by $j \in 1, 2, \dots, J$. Figure I shows three consecutive production stages of this supply chain, illustrating inputs and outputs of firms along with their tax liability. For simplicity, I ignore firm interactions within a production stage, assuming it comprises a representative firm only. The firm in a given stage j uses intermediates c_j acquired from the preceding stage $j - 1$ to produce output valuing s_j . Out of this output, an amount $s_{E,j}$ is exported out of the country, $s_{I,j}$ is supplied to the next production stage as intermediate, and $s_{F,j}$ is sold to the final consumer. The government implements the standard destination-based VAT, whereby exports are zero-rated, imports and domestic supplies are taxed at the standard rate τ , and a full credit of VAT paid on intermediates is given. Note that purchases or some types of sales of firms in the middle stages could be zero, which means that the structure I use is general, incorporating all types of firm interactions into it. Throughout this section, I focus solely on firms' reporting behavior. I show later that the reform does not affect firms' real production decisions significantly so that abstracting away from them is not too restrictive.

II.B First-Best Benchmark

Under the first-best benchmark, firms report their sales and purchases truthfully. Panel A of the figure illustrates the input-output linkages between firms under this scenario. In this case, the government receives a total revenue of

$$(1) \quad R \equiv \sum_j \tau \cdot (s_j - s_{E,j} - c_j)$$

from the supply chain. The VAT remitted by each firm is the difference between its output tax $\tau \cdot (s_j - s_{E,j})$ and input tax $\tau \cdot c_j$, and the total government revenue R is the sum of VAT revenue across all firms.

II.C VAT with Weak Enforcement

I now consider a case closer to weak enforcement settings of developing economies, allowing firms to misreport their sales and purchases (see Figure I, Panel B). The sales and purchases reported by firm j are now denoted by $\hat{s}_{K,j}$; $K \in \{E, I, F\}$ and \hat{c}_j , distinguishing them from their true values $s_{K,j}$ and c_j . Total government revenue from the supply chain in this case is

$$(2) \quad \hat{R} \equiv \sum_j \tau \cdot (\hat{s}_j - \hat{s}_{E,j} - \hat{c}_j).$$

This revenue could be lower than the first-best for two reasons. First, the firm may not report the true value of its sales and purchases, reducing its tax liability. Second, the purchases of intermediates reported at a given production stage may not match the B2B sales reported at the preceding stage i.e. $\hat{c}_j \neq \hat{s}_{I,j-1}$. Such misreporting is referred to as one-sided evasion in the literature, signifying that the firm does not collude with its suppliers or buyers to achieve this evasion (see for example [Pomeranz, 2015](#); [Waseem, 2020](#)).³ Comparing the government revenue in this setup with the first-best, we can write the total VAT evasion in the supply chain as

$$(3) \quad \Delta R \equiv R - \hat{R} \equiv \sum_j \tau \cdot [(s_j - \hat{s}_j) - (s_{E,j} - \hat{s}_{E,j}) - (c_j - \hat{c}_j)].$$

The expression within the sum here represents the VAT evaded by firm j . Summing the firm-level evasion across all firms, we obtain the aggregate VAT evasion in the supply chain.

In my empirical application, I exploit a large tax cut that reduces the standard VAT rate on five major supply chains of the country from 15% to 0%. The reduction of the rate to zero weakens the incentives to misreport sales and purchases, pushing them toward their true values

$$(4) \quad \begin{aligned} \hat{s}_{K,j}(\tau = 0) &\rightarrow s_{K,j} \\ \hat{c}_j(\tau = 0) &\rightarrow c_j. \end{aligned}$$

It means we can use the reform driven changes in sales (\hat{s}_j), export ($\hat{s}_{E,j}$), and purchases (\hat{c}_j) to estimate the three terms in Formula (3). For example, indexing the

³It is a relatively crude form of evasion because it is not robust to cross-matching of sales-purchases records by the government.

post-reform outcomes by t' and pre-reform ones by t , the reform-driven change in sales can be used to estimate the first term in the formula

$$(5) \quad \Delta \hat{s}_j = (\hat{s}_{jt'} - \hat{s}_{jt}) \approx (s_j - \hat{s}_j).$$

The VAT evasion at the baseline, thus, can be identified at both the micro and the aggregate level.

How will firms react to a tax change that reduces their rate toward zero? Equation (3) shows that one channel through which tax evasion may occur is the overreporting of purchases. We should therefore expect a decrease in reported purchases after the reform as the tax rate goes toward zero and incentive for tax evasion weakens; $\Delta \hat{c}_j(\tau \rightarrow 0) < 0$. Exports enter equation (3) in the same way as purchases, and we should therefore expect the two variables to move in the same direction. Note, however, that firms gain no direct tax benefit from overstating exports, as they receive refunds only on inputs purchased. There are, however, two indirect ways in which they may benefit from overstating exports. First, some firms may have both domestic and export sales ($s_j > s_{E,j}; s_{E,j} > 0$) and it is less costly for them to overreport the share of exports than to underreport total sales. Second, firms can only overreport purchases up to a plausible range of value-added, and therefore to overreport purchases they must also overreport exports. There is some evidence in the VAT literature that overreporting of exports does take place.⁴ We should therefore expect reported exports to go down after the reform; $\Delta \hat{s}_{E,j}(\tau \rightarrow 0) < 0$.

In distinction to purchases and exports, we should expect an increase in domestic B2C sales reported after the reform; $\Delta \hat{s}_{F,j}(\tau \rightarrow 0) > 0$. Before the reform, these sales are subject to tax at the standard rate and recipients of these sales are not entitled for tax credit. Given that there was no benefit of overreporting them, they would be underreported before the reform and would rise once the incentive to evade VAT goes down. In comparison, the direction of movement of B2B sales is not clear. As I note above, before the reform firms have an incentive to overreport purchases. With two-sided evasion, overreporting of purchases necessarily involves overreporting of B2B transactions by the firm's suppliers (they are likely diverting their B2C transactions into B2B transactions otherwise there is no incentive for them to do so). After the reform, as reported purchases decline there would be a corresponding decline in B2B

⁴For example, [Ebrill et al. \(2001\)](#) write: "Overstatement of exports, sometimes with the support of falsified customs documents, can be a particularly attractive way to do this [overclaim refunds]".

sales to the extent of two-sided evasion. But through other mechanisms B2B sales may rise after the reform. Consider, for example, a case whereby firms collude to underreport activity through out the supply chain. After the reform, B2B sales reported by such supply chains may rise. The overall direction of movement of B2B sales is therefore not clear; $\Delta \hat{s}_{I,j}(\tau \rightarrow 0) \leq 0$. Depending upon the magnitude of the change in exports, B2C sales, and B2B sales, total sales reported by firms ($\hat{s}_j \equiv \hat{s}_{E,j} + \hat{s}_{F,j} + \hat{s}_{I,j}$) may rise or fall; $\Delta \hat{s}_j(\tau \rightarrow 0) \leq 0$. In my data, I do not observe B2B and B2C sales directly. Instead, I examine responses of four outcomes reported by firms on their tax returns—purchases (\hat{c}_j); sales (\hat{s}_j); exports ($\hat{s}_{E,j}$); and non-export sales ($\hat{s}_j - \hat{s}_{E,j}$). Together, the direction of these responses allows me to confirm that observed changes in VAT outcomes are consistent with the incentive scheme outlined above.

II.D Overclaimed Refunds

The destination-based design of VAT requires that any tax collected on intermediates used for export be refunded to firms. This creates incentives for firms to exaggerate purchases used for exports, thereby increasing the value of refunds they obtain from the government. To quantify the value of overclaimed input tax on account of exports, I rewrite formula (3) in the following format,⁵ separating the VAT noncompliance in the export and domestic sectors of the economy

$$(6) \quad \Delta R \equiv \sum_j \tau \cdot [-(c_{E,j} - \hat{c}_{E,j}) + (s_{NE,j} - \hat{s}_{NE,j}) - (c_{NE,j} - \hat{c}_{NE,j})].$$

Subscripts E and NE here denote the export and non-export values of a variable. The first term in the above expression is the overclaimed input tax on exports: it is the amount by which firms overreport the value of intermediates used to produce exports. In my empirical setting, I do not observe $\hat{c}_{E,j}$ directly. Firms in Pakistan report the total value of their purchases without breaking it down by its use (domestic vs. exports). I therefore estimate the amount of excess input tax claim on account of exports indirectly using the following formula

$$(7) \quad \Delta R_E = \sum_j -\tau \cdot f^{-1}(s_{E,j} - \hat{s}_{E,j}),$$

⁵To derive this formula, I simply substitute the two definitions $s - s_{E,j} \equiv s_{NE,j}$ and $c_j \equiv c_{E,j} + c_{NE,j}$ into (3) and rearrange terms.

where $f(\cdot)$ is the production function of exports i.e. $s_{E,j} = f(c_{E,j})$.

Note that the overclaim of export-related input tax identified by (6) is already accounted for in formula (3). Intuitively, the overclaimed input tax consists of either diverted input tax from domestic consumption toward exports or one-sided evasion. In both cases, it is already covered in the third term of formula (3). The formula therefore continues to identify the aggregate amount of VAT noncompliance in the economy.

II.E Invoice Mills

VAT chains are rarely complete, especially in developing countries. When a VAT chain breaks, the tax remitted at the pre-break stages cannot be claimed at the post-break stage, becoming a part of the price. This creates arbitrage opportunities that are sometimes exploited by firms called invoice mills (Keen & Smith, 2006). These firms do not carry out any real business activity and exist solely to trade in VAT invoices. Figure A.I shows how this process works. In this example, I focus only on the last three production stages of a supply chain. The final stage is composed of two firms: the top firm deals only in exports and the bottom only in domestic sales. The firm in the preceding stage is not registered and supplies its output to the two firms, a fraction α to the exporter and the rest to the retailer. Because the firm is not in the VAT chain, the tax remitted at the $J - 2$ and preceding stages, amounting to $\tau \cdot \hat{s}_{I,J-2}$, cannot be claimed at the final production stage. Exploiting this gap, the invoice mill places itself in the $J - 1$ stage, transferring the credit of VAT remitted at the preceding stages to the final stage. Doing so creates a gross surplus of $\tau \cdot \hat{s}_{I,J-2}$, which can be shared by the colluding firms depending upon their bargaining weights.

Note that the bottom firm in the final stage has a distinctly lower incentive to deal with the invoice mill. All of its sales are to final consumers and hence do not generate any information trails. Dealing with the invoice mills brings the transaction on to books, thus creating an information trail and forcing the firm to pay tax to the extent of value-added on it. In contrast, the top firm in the final stage has a greater incentive to deal with the invoice mill. Given that all of its output is exported, its tax liability is negative and by dealing with the invoice mill it can exaggerate its purchases, increasing the refund it obtains from the government. This coincidence of incentives means that the invoice mill would divert a proportion of input tax from domestic consumption toward the export sector ($\hat{\alpha} > \alpha$), helping exporters claim exaggerated refunds.

This phenomenon is known as the diversion fraud in the literature (Keen & Smith, 2006). It results from the interaction of the two design features of the VAT—the last mile problem and the destination-principle.

It is important to emphasize that the presence of invoice mills does not invalidate formula (3). Given the flexible structure used to derive the formula, it already accounts for the diversion fraud. Invoice mills are one mechanism through which the VAT noncompliance identified by the formula takes place. Note further that the incentive to operate as an invoice mill is strengthened if the government implements a policy commonly observed in developing countries whereby the sales of intermediates to unregistered firms are taxed at a higher rate. In our example, the gross surplus generated by the invoice mill would be $(\tau + \tau_a) \cdot \hat{s}_{I,J-2}$ if the government operates such a policy, where τ_a denotes the *additional* rate applied to the sales of intermediates to unregistered firms.

III Contextual Background

III.A Pakistani VAT System

Pakistan introduced its VAT in the 1990s. The legislation for this purpose was passed in July 1990, and although it envisaged the VAT to be a broad-based tax with standard features, the scope of the new tax was kept limited in the initial few years through large-scale exemptions. These exemptions were withdrawn rapidly from 1996 so that by the year 1998 the tax had been extended to all notable industries of the country other than energy and services sectors (Waseem, 2020).⁶ In my empirical analysis, I focus primarily on the post-1998 period during which the VAT remained applicable to almost the entire goods sector of the country.

The Pakistani VAT largely follows the standard design. Firms with annual turnover above the exemption threshold are required to register with the tax administration.⁷ Firms not required to register can do so voluntarily. While registered, firms are required to charge VAT on their sales and are allowed to adjust the tax paid on inputs.

⁶The energy sector was brought into the tax net in July 1999 and services in July 2000. Please see Waseem (2020) for more details on the introduction and growth of VAT in Pakistan

⁷Exemption threshold is applicable to manufacturers and retailers only. For manufacturers, it was PKR 1 million in 1998, and was increased to 2.5 million in 1999 and to 5 million in 2004. For retailers, it remained at PKR 5 million throughout the sample period.

In case the adjustment exceeds the output tax, they can carry forward or obtain the refund of the balance amount. A seller is required to issue a tax invoice for each sale transaction, and the buyer can claim the tax credit only if it possesses the invoice issued in its name. Firms are required to file a return and remit the tax due every month. The filing is based on the principle of self-assessment and there is no pre-audit contact between taxpayers and tax collectors. Filed returns are considered final unless selected for audit (Best *et al.*, 2021).

The tax is destination-based: imports into the country are taxed at the standard rate and exports are zero-rated. Any tax charged on inputs used for exports is therefore refunded. To obtain refund, the exporter needs to file supporting documents in addition to the VAT return, which is treated as the refund claim. The supporting documents can be filed within reasonable time after the return has been filed, and include the customs and shipping documents showing the export of goods and VAT invoices showing the purchase of intermediates. No refund is sanctioned before an audit of the claim has been completed,⁸ and hence there is a natural delay between the claim and the payment of refund.

Panel A of Figure A.II plots the standard VAT rate in the country. The rate generally remained at 15% until 2008 when it was increased to 16%. Pakistan introduced a policy in 1998 whereby supplies made to unregistered *firms* were taxed at a higher rate. Of course, the higher rate was not applicable on supplies made to end consumers. Panel B of the figure plots the *additional* rate—called Further Tax—imposed by the policy. The rate remained between 1 and 3 percentage points before it was eliminated in 2004. The policy, as I explained above, strengthens the incentive of a firm to operate as an invoice mill.

III.B Zero-rating Reform

In July 2005, Pakistan introduced a novel tax reform through which the VAT rate applicable to supply chains of five major industries of the country—textile, leather, carpets, sports goods, and surgical goods—was reduced from 15% to 0%. Before the reform, in accordance with the standard destination-based design, only exports of these industries were zero-rated and their domestic supplies were taxed at the standard rate of 15%. The reform zero-rated not only the supplies of *final* goods produced by these industries but also of their major inputs. For example, in addition to the

⁸This condition applies to refunds accruing to both exporters and nonexporters.

finished goods produced by the textile industry (fabric, garments, etc.) all its major inputs including ginned cotton, polyester, yarn, undyed fabric, and important dyes and chemicals were zero-rated. The purpose of the change was two-fold (FBR, 2005).⁹ First, a large proportion of the output of these industries was exported out of the country in one form or another and hence was zero-rated anyway. Long within-country supply chains of these industries, however, meant that VAT was to be remitted and claimed back whenever the goods moved from one production stage to the next. This created cash-flow problems for exporters who had to wait for the refund of VAT paid on their inputs.

Second, the VAT chains of these industries were rarely complete. The breaks had given rise to the phenomenon of spurious invoices. Over time, the volume of such invoices was growing, making it increasingly costly for the tax administration to distinguish between genuine and fraudulent refund claims. Ultimately the problems created by these two related issues became so severe that the government forsook the VAT revenue from the domestic consumption of these industries and zero-rated their entire supply chains.¹⁰

The reform was announced on June 06, 2005 and became applicable from July 01, 2005. Although there was no firm commitment from the government prior to its announcement, the reform to some extent was anticipated in the sense that the government had been in negotiation with the stakeholders few months leading up to the reform to finalize details, such as the list of inputs to be zero-rated. In its final form, the reform zero-rated 152 items, which included both finished goods and major inputs of the five industries. Where an input was included in the list, its supply became zero-rated regardless of whether it was used in the production of the five industries or otherwise. For this reason, only inputs *predominantly* used by the five

⁹The two purposes of the reform were described by the Federal Board of Revenue in the following words: "Delays in refunds payments has been a source of anxiety for the taxpayers. ... The measure was also necessary due to the rampant use of fake and flying invoices by unscrupulous agents to claim illegitimate refunds". The FBR hoped that the reform will lead to two benefits: "Firstly, the refund payments would be reduced considerably, and secondly there would be an improvement in the liquidity position of textile sector leading to investment and boosting export and growth further". Please see page 25-26 of FBR (2005) for details.

¹⁰Limited forms of such schemes have been implemented in Ireland and South Korea as well. In Ireland, for example, firms that export more than 75% of their output can obtain an authorization that allows their suppliers not to charge VAT. In South Korea, those who supply exporters are zero-rated in respect of selected transactions (see Ebrill *et al.*, 2001 for details of these two schemes). The Pakistani tax reform is novel in the sense that it zero-rates the entire supply chain rather than just the pre-export production stage.

industries were zero-rated.¹¹ The list of 152 items did not include electricity and gas, two important inputs of these industries. These two inputs were also zero rated, but their zero-rating—unlike that of others—was made conditional on their use in the production of the five industries. Legislative instruments zero-rating electricity and gas were therefore issued at the firm level after verification that the firm was indeed operating in one of the five treated industries. This exercise caused some delay in the zero-rating of these two inputs. The first set of orders granting such zero-rating were issued in August 2005 but the exercise was not completed till April 2007.

The reform moved the treated firms from a standard VAT regime to a novel regime, where both their output tax liability and input tax entitlement reduced toward zero.¹² It thus weakened the incentives to misreport sales or purchases. It is important to emphasize that the reform did not make other changes besides cutting the tax rate. Importantly, registration, record keeping, and return filing requirements remained unchanged for the treated firms, which as usual were required to file their tax returns and comply with other obligations.

The Pakistani tax administration began obtaining transaction-level data from firms from July 2008. This requirement was introduced by adding an annex to the VAT return wherein firms were required to provide the details of their sales and purchases during the month, aggregating them up to the level of each supplier and buyer. Simultaneously, electronic filing of both the VAT return and the annex was made mandatory for all firms. Together, these changes reduce the costs of cross-matching sales and purchase records, enabling the tax administration to detect one-sided misreporting at a relatively low cost. This change in enforcement technology in 2008 affects both treated and untreated industries similarly and therefore should not matter in my empirical setting unless there is some interaction between it and the zero-rating reform, a point I come back to in section V of the paper.

¹¹Otherwise, the loss of revenue from zero-rating would have been unsustainable for the government.

¹²Note that the reform would not reduce the output tax and input tax of the treated firms to zero, although it would reduce both these variables sharply. The output tax charged would not go down to zero if the firm sells a byproduct not included in the list of zero-rated items. And the input tax would not go down to zero because not all inputs used by the treated firms were zero-rated.

III.C Invoice Mills

One key focus of the paper is to document the role of invoice mills in VAT noncompliance. Like all fraudulent enterprises, invoice mills take great care in disguising themselves as legitimate businesses so that distinguishing them from the others is not easy. In the Pakistani setting, however, I am able to exploit a legal mechanism used by the tax authority to identify these firms.

In the initial few years after the adoption of VAT in Pakistan, the tax authority did not have any specific mechanism to deal with invoice mills. Fraudulent input tax claims based on invoices issued by such firms were dealt with generically, like any other form of noncompliance. Over time, however, the volume of such transactions grew, making it necessary to have a tailored mechanism to deal with the issue. The new mechanism, which came into force in July 2003, empowered the tax authority to suspend the registration of a firm it suspected of being involved in the issuance of fake or flying invoices. The suspension was meant to be a temporary measure aimed at protecting revenue while an inquiry against the firm could proceed. This inquiry was of a quasi-judicial nature, where the firm was confronted with the evidence against it and was afforded an opportunity to present its case. On completion of the inquiry, either the registration of the firm was restored or it was blacklisted permanently. Once a firm was blacklisted, its invoices no longer remained valid and could not be used to claim input tax credit. I observe both the suspension and blacklisting of a firm and use these to proxy if the firm is an invoice mill.

The tax authority used a simple procedure to identify invoice mills.¹³ The procedure relied on three distinguishing features of invoice mills. First, given the small margin between the price at which a VAT invoice is purchased and the price at which it is sold operating as an invoice mill is feasible only at large turnover. Second, such firms report very low value-addition to avoid making significant tax payments. Third, because they are essentially hit and run firms, they do not remain active for long duration.¹⁴ The tax authority accordingly used to desk audit newly registered firms showing large turnover and low value-addition. In case of any suspicion, more detailed audits were carried out before suspending the registration of the firm. This

¹³This description of the procedure to identify invoice mills is based on the author's discussions with Pakistan's tax authority (FBR).

¹⁴Purchase and sales of VAT invoices amounts to "tax fraud" under the Pakistani tax code. This implies that in addition to fines and penalties owners of invoices mills face the threat of prosecution. To avoid these consequences, these firms leave their registered business premises soon after registration and are difficult to trace afterward.

procedure to identify invoice mills remained in use through out the sample period. Importantly, it did not undergo any sharp changes contemporaneous with the reform, although like all business process it underwent slow incremental changes over time in accordance with the authority's experience of dealing with such firms.

III.D Data

I use administrative data from Pakistan, which include the universe of VAT returns filed in the country. The VAT return consists of three main sections. In the first section, firms report the aggregate value of their sales, breaking it down into three—domestic taxable, domestic exempt, and exports—components. In the second section, the aggregate value of inputs purchased are reported, divided likewise into the three components. In the final section, firms calculate their tax liability, indicating the tax charged on sales, the tax credited on inputs, and the final tax payable. They select one of the two options—carry forward or refund—in case the tax payable is negative. Each firm in the VAT net is assigned a unique registration number and is expected to file every tax period (month). The data, therefore, have a panel structure. In addition to the return data, I use information on firm characteristics from the tax register. This information includes the 4-digit industry, date of registration, and current registration status (suspended, blacklisted, or otherwise) of the firm. The 4-digit industry coding corresponds to the Harmonized Commodity Description and Coding System (HS Code) and classifies firms on the basis of goods or services they supply.¹⁵ The industry coding allows me to determine if a firm belongs to one of the zero-rated industries.

As I note above, exporters have to file additional documents in support of their refund claims. These documents include supplier-wise details of purchases of intermediates acquired by them. These transaction-level data are available from the tax year 2002 onward, and I use them to construct linkages between exporters and invoice mills to see what proportion of a refund claim is based on spurious invoices.

¹⁵This system is commonly used by customs administrations around the world to classify traded goods and services.

IV Empirical Methodology

Given that the reform I exploit affects a subset of firms only, the natural research design in this setting is the difference-in-differences framework. I exploit the fact that the incentives to misreport reduce sharply at the time of the reform if the firm belongs to a zero-rated industry and remain unchanged otherwise, estimating the following model

$$(8) \quad y_{it} = \alpha_i + \lambda_t + \beta \mathbf{X}_{it} + \gamma \cdot \text{zero-rated}_i \times \text{after}_t + \varepsilon_{it},$$

where α_i and λ_t are the firm and time fixed effect, \mathbf{X}_{it} is a vector of control variables, zero-rated_i denotes that firm i belongs to an industry whose rate was cut to zero by the reform, and after_t indicates a post-reform tax period (July 2005 or after). I use the model to estimate the impact of the reform on six outcomes (y_{it}) introduced as logs. In parallel specifications, I replace the double-difference term with $\log(1 - \tau_{it})$, where τ_{it} is the effective VAT rate faced by firm i on outcome y in period t . These specifications deliver the elasticity of the corresponding outcome with respect to the net-of-tax rate. The baseline specification clusters standard errors at the firm level, but in one of robustness checks I report results where the standard errors are clustered at the industry level—the most aggregate level feasible in this setting (Abadie *et al.*, 2017).

Identification in this setup requires that a given outcome would have evolved similarly in the treatment and control groups in the absence of the tax reform. I exploit the long panel of VAT records to show that this assumption is indeed reasonable in this setting. Specifically, I plot results from the following event-study model

$$(9) \quad y_{it} = \alpha_i + \sum_{j=2}^N \delta_j \cdot \text{zero-rated}_i \times 1.(\text{tax period}=j)_t + \lambda_t + u_{it},$$

where j indexes the set of tax periods (months) included in the sample. I estimate the equation on a sample from July 1998 to June 2011, dropping the dummy for the first tax period. I then plot the coefficients on the interaction terms from these regressions for all six outcomes. Using these regressions, I show that all six outcomes evolve fairly similarly across the two groups in the 84 pre-reform periods, validating the key identification assumption.

Notwithstanding parallel trends, identification in this setting may fail if the zero-rating reform creates significant spillovers in the non-zero-rated industries, violating the SUTVA assumption (see for example [Imbens & Rubin, 2015](#)). These spillovers can take two broad forms. At the extensive margin, the reform could distort the entry decision of a firm.¹⁶ I address this concern by reporting parallel results from balanced panel samples, where the composition of the sample is held fixed. Shutting down the entry and exit, however, does not rule out spillovers along the intensive margin. These spillovers may arise from general equilibrium considerations (the zero-rating reform affecting prices or the cost of compliance generally) or from demand and supply linkages of control firms with the zero-rated industries ([Waseem, 2020](#)). I present two sets of evidence to rule out this class of concerns (please see section [V.C](#)).

I estimate the impact of the reform on the following six outcomes: (i) input tax; (ii) output tax; (iii) purchases; (iv) sales; (v) exports; and (vi) non-export sales. The definitions of these variables are provided in [Appendix A.1](#). [Table A.I](#) presents summary statistics of the data. The treatment group here contains firms belonging to the five zero-rated industries; all other firms are included in the control group. The first row of the table reports firm-month observations of the two groups in the sample for the two baseline years, 2003 and 2004. The next rows present mean of the six VAT outcomes and other firm characteristics for the two baseline years. On average, the treated firms are larger and are more likely to be engaged in exports. But they are not much different from the control sample in terms of other characteristics such as location and age.

V Firm Responses to the Reform

V.A Nonparametric Evidence

I first present visual evidence on how firms respond to the reduction of the VAT rate to zero. The analysis is then formalized using the regression based framework.

¹⁶The differential tax treatment may force a firm which would otherwise have entered into a non-zero-rated industry to switch to a zero-rated industry or vice versa

V.A.1 First Stage

Before documenting firm responses to zero-rating, it is important to show that the reform creates large tax variation between the treated and untreated firms. In this section, I present eight pieces of evidence to show that it is indeed the case.

Panel A of Figure A.III plots the amount of VAT refund paid in Pakistan as a proportion of gross VAT collected in the country.¹⁷ The refund-to-gross-collection ratio was roughly 20% at the baseline. It fell by nearly 10 percentage points in the first year after the reform. It fell even further in later years, settling at around one-fourth of the baseline level within three years of the reform. Panels B–D plot the average effective tax rate faced by treated firms. Panel B shows the overall effective tax rate, whereas the latter two panels show the effective tax rate on output and input separately.¹⁸ All three effective tax rates fall sharply exactly at the time of the reform, capturing a large drop in the output tax liability, input tax credit, and tax burden of treated firms.

Continuing to showcase the strong first stage in this setup, I next use the difference-in-differences framework to highlight the reform's effects on output tax, input tax, and tax liability reported by treated firms. Figure II plots the coefficient $\hat{\eta}_j$'s from the following version of the event study equation (9)

$$(10) \quad y_{it} = \alpha_i + \sum_{j=2}^N \eta_j \cdot 1.(tax\ period=j)_t + u_{it},$$

where j indexes the tax periods (months) included in my estimation frame. I estimate the equation separately for the treatment and control groups, omitting the dummy for the first tax period (July 1998). Figure III presents the DD version of these plots, where I display the coefficients $\hat{\delta}_j$'s from equation (9) along with the 95% confidence intervals around them.

Panels A-B of the two figures show the responses of output tax and input tax. Consistent with the earlier evidence, a sharp fall occurs in both the treated outcomes exactly at the time of the reform, while the two control outcomes continue to evolve

¹⁷The figure is based on annual aggregate statistics reported by the FBR on its website, which include both treated and untreated industries. For the data used in this figure, see FBR Year Books from 1999 to 2009 available [here](#). Year books prior to 1999 are unfortunately not available on the FBR's website.

¹⁸The effective tax rate shown in Panel B is defined as the VAT remitted by firm divided by total sales reported by it. The effective tax rate on output (Panel C) is defined as the output tax liability of the firm divided by its total sales, and the effective tax rate on input (Panel D) is defined as the input tax liability of the firm divided by its total purchases.

on the preexisting trend. The dynamics of the two responses is also consistent with our expectations. Both output tax and input tax decline sharply as the reform comes into effect, but unlike the output tax the input tax continues to drop, stabilizing only around the beginning of the tax year 2008. The continuing drop of input tax, as noted in section III.B above, is due likely to the time taken in zero-rating two important inputs—electricity and gas—of the treated firms.

Panels E-F of Figure A.III repeat the above difference-in-differences analysis for two other outcomes. The outcome in Panel E is the VAT remitted by firms for each tax period, whereas the outcome in Panel F is a binary variable indicating that the firm’s reported liability is positive. The evolution of these outcomes is also similar to earlier ones. Together, these eight pieces of evidence show that the incentive to misreport falls sharply in the treatment group at the time of the reform as the output tax liability, input tax claim, and tax burden of treated firms fall to a near-zero level. Such a sharp change in incentives is likely to induce sharp behavioral responses to which I turn next.

V.A.2 Behavioral Responses

Panels C-F of Figures II and III illustrate these responses. None of the four outcomes shown in these panels is directly influenced by the tax rate, and their responses therefore isolate *pure* behavioral effects induced by the rate cut. These behavioral effects are negative for all four outcomes. Sales, purchases, exports, and non-export sales fall sharply in the treatment group after the reform. The fall is the sharpest for purchases, which decline by 30-35 log points immediately after the reform. Compared to purchases, the decline in other outcomes is slow and gradual, materializing fully in the next two years.

The direction of movement of these outcomes is consistent with the incentive scheme outlined in section II.C. Both purchases and exports fall. Non-export sales, by definition, equal the sum of domestic B2B and B2C sales $\hat{s}_{NE} \equiv \hat{s}_I + \hat{s}_F$. The model has a clear prediction on the latter (B2C sales will rise after the tax cut) but not on the former. The fall of non-export sales therefore shows that B2B sales fall after the tax cut and that their response outweighs the response of B2C sales. Note that B2B sales are mirror image of purchases: absent one-sided evasion B2B sales reported by a seller are reported as purchases by the buyer. Being so, their larger fall relative to the rise in B2C sales is plausible given the sharp fall in purchases seen in Figures II–III.

Finally, the overall reduction in reported sales (\hat{s}) is a reflection of the fact that both its components—exports and non-export sales—fall.

One other noticeable feature of the event-study results is that the sales response is slower than the purchases response. Again, such a dynamics is expected. Purchases are expected to fall immediately after the incentives to overreport them end. Past overreporting of purchases, however, would have left firms with large inventories even if only in books, which would not let the volume of reported sales and exports drop to the new equilibrium in periods immediately following the reform.

The event study results also validate my empirical strategy. The preexisting trends are fairly parallel in the two groups for all six outcomes I explore (see Figures II and III). The reform causes sharp changes in the treatment outcomes, while the control outcomes continue to evolve on the preexisting trend with no appreciable break at the time of the reform.

V.B Regression Results

Table I presents results from the difference-in-differences model (8) for the six outcomes. Table A.II report results from parallel specifications that estimate the elasticities underlying these responses, and Table A.III repeats the baseline analysis, clustering the standard errors at the industry rather than firm level. Unsurprisingly, the results are consistent with the visual evidence presented above. Both output tax and input tax drop sharply as a result of the reform, confirming that a strong first-stage exists in this setup. Columns 3–6 report pure behavioral responses induced by the rate cut. If firms report truthfully, the reduction of the rate to zero should induce no change in behavior other than some positive, real effects on activity. But the observed responses reveal a different story. Not only are all elasticities underlying the responses of opposite sign, they are also very large, ranging from -4.3 to -0.8 .¹⁹ Such a large fall in reported activity is consistent with, and can only be reconciled by, a large misreporting at the baseline. In section V.D, I use formula (3) to estimate the extent of misreporting implied by these responses and study the mechanisms underlying it.

The event study results presented above show that the key identification assumption in my empirical setting—parallel trends—is plausible. Panels B of Tables I and

¹⁹It is important to emphasize that these elasticities are own-price elasticities estimated under the assumption that the reform does not produce significant cross-price effects. I show later (see section V.C) that these cross-price effects (spillovers) are indeed insignificant.

A.II test this formally. I estimate a parallel placebo specification, where I estimate equation (8) on the pre-reform panel only, pretending that the zero-rating reform occurred in July 2002. The results support the identification assumption. There is little to distinguish the treated and untreated groups in terms of the baseline evolution of the VAT outcomes.

Table II explores the dynamics of the responses. I estimate a flexible version of model (8) by replacing the single DD term with six $\text{treat} \times \text{year}$ interactions, one each for every post-reform tax year. The results confirm the time pattern of response seen visually in Figures II and III. The first year response as a proportion of the average post-reform response is 79% for input tax, 94% for output tax, 80% for purchases, 32% for sales, 18% for exports, and 56% for non-export sales. I have already noted the likely reasons for this pattern. The first two outcomes capture the mechanical impact of the reform, which unsurprisingly is immediate.²⁰ Of the other four items, purchases were likely to be impacted first given that once input tax credit available on purchases drove down to zero there was no incentive to overreport them. In contrast, sales, exports, and non-export sales would return to the new equilibrium only after inventories built up in the books through past overreporting of purchases have been cleared.

V.C Robustness

One concern in this setting is that the zero-rating reform may create spillovers in the nontreated industries, violating the SUTVA assumption. Table A.IV shows that I get similar results from the balanced panel specification. This mitigates the spillover concern along the extensive margin, operating through the entry or exit.²¹ Figure II further illustrates that the outcomes of nontreated industries do not undergo a structural break at the time of the reform. Building on this evidence, Table A.V and A.VI explore spillovers more formally. If the reform creates significant spillovers, they would be stronger in industries producing substitutes or complements of the treated goods. The tables test this by looking at the evolution of outcomes in such indus-

²⁰The slightly lower first-year response of input tax, as I note above, was in large part due to the delay in the zero-rating of electricity and gas used by the treated firms.

²¹The VAT creates incentives for suppliers and buyers of registered firms to register (*de Paula & Scheinkman, 2010*). Spillovers along the extensive margin capture the idea that some trading partners of treated firms from nontreated industries may exit the VAT regime after the reform. Alternatively, the reform may slow down the entry of trading partners.

tries. Pakistani tax administration, as I note above, uses the HS Code to classify firms into industries. The first two digits of this eight-digit code divide firms into broad categories with industries producing similar goods getting adjoining codes.²² I exploit this coding scheme to explore the tax cut's spillovers on adjoining industries. Table A.V drops all treated industries from the sample and compares the evolution of outcomes in their adjoining industries with that of others. Here I estimate specifications parallel to ones in Table I, the only difference being that the variable *treat* now indicates firms belonging to the adjoining industries. I experiment with three definitions of adjoining. Table A.VI is structured similarly with the difference that it retains the treated industries in the sample, including the adjoining industries as an additional interaction term. Reassuringly, all specifications return trivial or insignificant spillover coefficient, putting to rest the concern that the reform might have affected outcomes of the nontreated industries as well.

Two important events occur in 2008 that may influence the interpretation of my results. First, as I note in section III.B, Pakistan introduced new filing requirements from July 2008, which mandated firms to file transaction-level data along with their returns. Second, the financial crisis hit the world markets in September 2008. Some of the negative responses documented above may reflect that these events differentially affect the treated industries. Figure A.IV addresses this class of concerns. It is a truncated version of Figure III, where I show the post-reform periods only. Zooming in on these periods illustrates that the reform started a slow, downward trend in the outcomes of the treated industries. This downward trend did not accelerate during 2008. In fact, the outcomes of the treated industries stabilized or started rising slowly from the mid of the financial year 2008-2009. The evolution of the responses thus rules out any significant *differential* impact of the two events on the treated industries.

I have noted above that the reform was anticipated in the sense that the government had been discussing it with the stakeholders before its formal announcement on June 6, 2005. To the extent that the reform was anticipated, it may trigger some intertemporal shifting of activity by forward looking firms to calibrate their response to the new incentives. I show later (section VI) that invoice mills indeed engaged in such intertemporal shifting. Here, I focus on firms other than invoice mills, which comprise 93% of the treatment sample, to see if they also engaged in such shifting.

²²For example the code 08 is assigned to edible fruit and nuts; 09 to coffee, tea, mate and spices; and 10 to cereals. The adjoining codes thus contain similar industries producing close substitute or complement goods.

I begin with Figure A.V, which plots the aggregate values of the VAT outcomes of these firms. There are no discernible signs of intertemporal shifting in either directions. Tables A.VII–A.VIII rule out anticipation formally. The former shows that dropping invoice mills—the subsample that did engage in the intertemporal shifting of activity—does not change the results. The latter defines the *after* period from the date of announcement of the reform rather than the date of its coming into force. This specification also returns similar results.

Tables A.IX and A.X rule out one alternative explanation of the results. It can be argued that the incentives to maintain records (receipts, invoices, etc.) and to report them correctly go down once the tax rate goes down to zero. In this world, the negative responses documented in Table I are explained by lazy reporting rather than a reduction in misreporting. To rule out this concern, Table A.IX compares the filing probability across treated and untreated industries using a model similar to (8). The zero-rating reform does not result in treated firms filing their tax returns less regularly than others. In fact, the filing rate goes up slightly among the treated firm. Table A.X examines the intensive margin, exploring heterogeneity in response across treated corporate and non-corporate firms. The idea behind the exercise is that lazy reporting is expected to be worse and the observed responses hence more negative among noncorporate firms whose quality of record maintenance is in general poorer. Contrary to this, the responses are on average less negative for noncorporate firms. Lazy reporting is therefore unlikely to be an explanation of the observed responses.

V.D Quantifying VAT Evasion

Section II.C shows that a lower bound on the amount of VAT evaded in the treated industries is provided by formula (3). I now use the results in Table I to estimate this lower bound. The calculations are in Table IV. Rows 1–3 compute the three terms in the formula for the four baseline years (2001–2004). To estimate each term, I multiply the percent change in the variable as implied by its response reported in Table I with its aggregate value in the baseline year. Row 4 adds the three terms to compute the total evaded amount in PKR billions using formula (3). Finally, the last row of the top panel expresses the evaded amount as a fraction of the potential revenue in the year.

Panel B of the table uses formula (7) to calculate the amount of overclaimed refund by treated firms. One important ingredient in this formula is the inverse mapping from exports to purchases of intermediates used in exports. Figure IV estimates this

mapping nonparametrically. To construct the plot, I group firms into small bins on the basis of log purchases reported by them and plot the average log exports in each bin. The sample for this binned scatter plot consists of all firm-month observations of the treated industries, excluding those with the log purchases less than the 5th or more than the 95th percentile of the corresponding distribution. The relationship between the two variables in the log-log space is approximately linear with a slope parameter of 0.5, which means that the production function of exports $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ $s_E = f(c_E)$ is Cobb-Douglas with $s_E = A \cdot c_E^{0.5}$. Note that this relationship does not need to be causal as I only use it to predict $\Sigma_j (c_{E,j} - \hat{c}_{E,j})$ implied by $\Sigma_j (s_{E,j} - \hat{s}_{E,j})$, the term identified through the DD estimator. Rows 6–8 combine the estimated production function and exports response to compute the amount overclaimed as refund. The final row expresses it as a percentage of the potential revenue in the year.

The results reveal that a large amount of VAT is evaded in the country. The evaded amount varies from 31% of potential revenue in 2001 to 46% in 2004, with an increasing trend over time (Row 5 of Table IV). A large proportion of this evasion occurs through the channel of overclaimed refunds. The overclaimed amount, reported in the last row of the table, ranges between 11–23% of the potential revenue, again with an increasing trend over time. Note that these estimates are based on the assumption that firm responses to the zero-rating reforms are homogeneous so that the average effect is broadly applicable to all firms. Practically, however, firm responses may vary based on firm characteristics, especially firm size.²³ Importantly, if the average effect I estimate is driven primarily by responses of small firms, the method I use above would overestimate aggregate tax evasion in the treated industries. To rule out this concern, Table A.XI explores heterogeneity in firm responses by firm size. I divide firms into four quartiles based on their turnover in the pre-reform period and examine if firms in different quartiles respond differently to the tax reform. They indeed do. On average, the responses of larger firms are systematically larger than those of smaller firms. The results thus suggest that the above method if anything underestimates the magnitude of aggregate tax evasion.

How do these evasion estimates compare with similar estimates from other countries? Gómez Sabaini & Jiménez (2012) compute the VAT evasion rate for many Latin American economies, finding similar levels of evasion. For example, the evasion rate is 37.5% in Guatemala; 38.1% in Nicaragua; 33.8% in Panama; and 37.7% in Peru.

²³Recent empirical evidence shows that firm size is the most important determinant of tax evasion (see Best *et al.* (2021))

Such large levels of evasion suggest that VAT's robustness to tax evasion is somewhat overstated. Theoretically, the tax indeed contains mechanisms that must make evasion harder than under its alternatives. But it appears that firms in low tax capacity setting are able to get around these mechanisms through collusion and other strategies. One manifestation of such collusion is that the evasion I find runs deep into the initial stages of the supply chain, where most transactions are subject to third-party reporting (Table III). Evasion at these stages suggests that the third-party information on its own does not deter evasion meaningfully.

V.E Real Response

The zero-rating reform, being a tax cut, was likely to stimulate activity in the treated industries. The negative behavioral responses I report in Table I are net of these positive real responses. To the extent real responses produced by the reform are significant, the evasion I compute above is underestimated. This section examines how large real responses produced by the reform are. I begin this analysis by presenting aggregate-level evidence. Figure A.VI plots the evolution of aggregate exports of Pakistan. The top panel compares exports of treated industries with the country's total exports, whereas the bottom panel compares the exports of treated and nontreated industries directly.²⁴ Exports of treated industries continue to evolve on the preexisting trend with no discernible structural break at the time of the reform. In fact, if anything the exports of nontreated industries grow faster after the reform. The zero-rating reform therefore does not seem to have stimulated activity in the treated industries significantly.

I complement the aggregate-level analysis with micro-level evidence, exploiting heterogeneity in responses over the supply chain. The idea behind the exercise is that upstream stages mostly produce intermediate goods (B2B transactions) in contrast to downstream stages whose output is mainly consumer goods (B2C transactions). To the extent that the incentive to misreport B2B transactions are distinctly lower, the heterogeneity analysis can reveal the nature of the observed responses. Concretely, real responses produced by the reform would be roughly symmetric over the supply chain, whereas reporting responses would be stronger in later stages.

²⁴I use data from Pakistan's central bank for this exercise. These data are not at the same level of granularity the micro data are. But importantly three main treated industries (textile, leather, and footwear) are separately identified so that I can compare their exports with the rest. These data are available [here](#).

For this exercise, I focus on the textile industry only, which is the largest and the most important of the five zero-rated industries. Table A.XII shows this formally, illustrating that the baseline estimates in Table I are largely driven by the textile industry. Focusing on the textile industry is useful because it has a well-defined supply chain comprising five distinct production stages: ginning, spinning, weaving, processing, and the made-up stage. In the first production stage—ginning—cotton fiber is separated from the seed and is compressed into bales. The next stage—spinning—converts these bales into cotton yarn, which then is converted into gray fabric by the next stage (weaving). The final two stages convert gray fabric into printed fabric (processing) and garments (made-up stage). Importantly, the output of the first production stage has no consumer use and is utilized entirely as raw material in the next production stage (spinning). The transactions of this stage are therefore mostly B2B transactions. I observe the production stage a textile firm belongs to and therefore can examine behavioral responses by production stage.

Table III carries out this exercise. I estimate a triple-difference version of model (8) by including interactions of the DD term with dummies indicating the production stage. I include dummies for the three upstream stages only, leaving the final two as the omitted category. The results show that responses are considerably weaker in the upstream stages, becoming progressively stronger as one moves down the supply chain. For example, both sales and purchases of the first production stage show no significant change relative to the control group after the reform. This suggests that the reform did not produce significant real responses: behavioral responses at the first production stage, as I note above, largely capture real effects as the stage predominantly produces intermediates used exclusively in the next production stage, making misreporting less likely. Lack of significant real responses allays the concern that evasion estimates I report above are seriously underestimated.²⁵

In contrast to the first production stage, treated outcomes decline significantly in later stages. Increasing response over the supply chain from the second stage onward suggests that the tax evasion is not limited to the later production stages but to some extent persists throughout the supply chain.

²⁵One alternative explanation of the above results is that the first stage—fall in the tax burden—is weaker in the upstream production stages. To rule out this explanation, Figure A.VII looks at the average effective tax rate (AETR) faced by textile firms on their outputs and input. I report the AETR separately for the three key production stages. The results show that in fact the first stage is stronger in the upstream stages.

VI Invoice Mills and VAT Evasion

One important focus of this paper is to understand the role of invoice mills in facilitating VAT's noncompliance. As I note in section II.E, invoice mills arise naturally in the low-enforcement, high-informality setting of developing countries, bridging the gaps created by broken VAT chains. Figure A.IX illustrates this empirically, plotting the stock of invoice mills in Pakistan's VAT over time. Invoice mills appeared soon after the destination-based VAT was implemented in the country. Their number grew steadily in the next few years, reaching a peak in 2003. The rising trend reversed then for a short period, as the new mechanism to blacklist firms came into effect. But a much sharper decline occurred after 2005, when the zero-rating reform became applicable.²⁶ The initial evidence thus suggests that invoice mills primarily exist to serve the export refund market (their numbers fall sharply as the tax rate on five major export-oriented industries falls to zero). To explore this point further, I next look at the *activity* reported by them.

Figure A.X examines the trajectory of VAT outcomes of invoice mills, replicating the event-study analysis of Figure II. Since invoice mills comprise only 7% of the treatment sample, the results are noisier than the baseline results. Yet the pattern of responses is similar. All four outcomes fall sharply at the time of the reform, with purchases falling more than any other outcome. Table A.XIII formalizes these results. I estimate a triple-difference version of model (8), exploring any differential response of invoice mills relative to other treated firms. Unsurprisingly, mills' responses are stronger, but given their smaller number their stronger responses do not affect the overall average significantly: the double-difference coefficient is always within the 95% confidence interval around the baseline coefficient.

Importantly, the event study plots in Figure A.X show that the outcomes spike just before the tax cut. This spike is particularly prominent in the top four panels of the figure. To probe this result further, Figure V examines the aggregate values of these six outcomes in each month. These plots show the aggregate *level* of each outcome, highlighting the size of spurious input tax credit injected into the system by invoice mills. Strikingly, all outcomes feature a sharp spike just before the tax cut. For example, aggregate purchases jump from the prereform average of around PKR 5 billion

²⁶Note that the procedure used by the tax authority to identify invoice mills did not change in 2005 (see the detail in section III.C). The sharp decline in the stock of invoice mills therefore likely arises from the change in monetary incentives created by the tax reform.

to 30 billion one month before the tax cut (May 2005). This six-fold rise is followed by even sharper fall as purchases plummet to PKR 3.6 billion in July 2005 and to 1.8 billion in November 2005.²⁷ Such a large concentration of activity on the wrong side of the reform is puzzling. Under any standard model of tax behavior a large tax cut would induce some inter-temporal shifting of activity toward the low-tax period.²⁸ Shifting in the opposite direction strengthens the conclusion that invoice mills are primarily refund generating devices: they are injecting spurious input tax credit into the system while the goods are still standard-rated. Doing so will maximize the refunds their buyers can obtain in accordance with the mechanism outlined in section II.E.

How much excess refund did invoice mills inject into the system? I examine this question using the transaction-level data exporters file to claim refund, which are available from 2002 onward. Figure VI reports the results of this exercise. The blue curve plots the aggregate value of purchases exporters report each month from invoice mills. The red curve, on the other hand, plots the aggregate value of sales invoice mills report each month to all buyers including exporters. Although the two curves use data from two different sources—the blue curve from refund claims filed by exporters and the red from VAT returns filed by invoice mills, they line up tightly. This shows that most of the sales of invoice mills are to exporters, who use their invoices to claim refunds. A regression of blue series on the red estimated for the baseline years returns a coefficient of 0.83, suggesting that on average 83% of the output of invoice mills is claimed as refund by exporters.²⁹ This ratio shoots up to 97% (PKR 8.6 billion) in 2004, the last year such a refund could be claimed in. Table IV shows that exporters overclaimed refund by nearly PKR 23 billion in 2004 (see Row 5 of the table). Roughly 37% of the overclaimed amount was based on invoices issued by mills.³⁰

²⁷One other important feature of the plots is a very large spike in exports just before the beginning of the tax year 2008. This is very likely driven by an effort by invoice mills to clear their inventories before the requirement of filing transaction-level data comes into effect from the tax year 2008.

²⁸For instance, booking a transaction that occurs just before the reform to a date just after the reform could save the seller the costs of remitting the tax, the buyer the costs of claiming the input tax credit, and any associated cash-flow costs.

²⁹This finding, however, is subject to an important caveat. I identify mills using the blacklisting procedure employed by the Pakistani tax administration. It is possible that this procedure is more effective in identifying mills connected to exporters either because the government is more concerned about the overclaim of refund than the underpayment of VAT on domestic transactions or because of any data limitation.

³⁰It is important to emphasize that this revenue loss provides only a lower bound to the real revenue

Figure A.XI shows the evolution of outcomes of invoice mills around the event of their blacklisting. It illustrates that by the time an invoice mill is blacklisted its activity has already passed its peak. Being so, most of their sales occur while their invoices are legitimate. To see why invoice mills are detected so late in their life cycle, Figure A.XII and Table A.XIV compare their characteristics with other firms. Invoice mills are essentially hit-and-run enterprises, and consistent with this notion their most distinct marker is the level of activity they report soon after registration.³¹ To catch them, the revenue authority therefore needs to move early when they are still distinguishable from other firms and when they have still not caused significant revenue damage. But doing so requires developed administrative capacity, which revenue authorities usually lack in developing economies.

VII Conclusion

The value-added tax has seen remarkable expansion in recent decades. Its popularity in large part stems from the belief that among the class of production-efficient tax instruments it perhaps has the best enforcement properties. By creating (1) third-party information on firm-to-firm transactions; (2) tax withholding at the upstream production stages; and (3) asymmetric cheating incentives between sellers and buyers, a VAT facilitates tax compliance (Waseem, 2019). But the tax also contains built-in mechanisms that may worsen tax compliance: (1) the destination principle necessitates the payment of refund on exports; (2) the last-mile problem eliminates paper trail at the end of the supply chain; and (3) broken VAT chains give rise to invoice mills. These mechanisms are emphasized a great deal in the policy literature, but due mainly to a lack of empirical evidence are largely absent from the economic literature. This paper fills the gap by casting light on the nature and strength of these mechanisms in a representative emerging economy.

I leverage a novel tax reform that weakens firms' incentives to misreport outcomes. This allows me to infer the level of misreporting in the treated industries at the baseline and study mechanisms underlying it. I document four key facts. First,

revenue loss arising through the mechanism of spurious invoices. In my setup, a firm is categorized as an invoice mill only if it has been declared so by the revenue authority. Potentially some invoice mills, especially those serving the non-export sector, escape detection by the revenue authority and as a result evasion resulting from them cannot be captured in Figure VI.

³¹Early in Figure A.XII denotes the level of activity reported in the first six months after registration

the evaded VAT constitutes a significant part of the tax base, amounting to between 31–46% of the potential revenue. Second, the evasion is not limited to the latter stages of the supply chain—where the last-mile problem is worse—but runs deep into the supply chain. This suggests that on their own the self-enforcement mechanisms built into a VAT do not deter tax evasion fully. Third, overclaiming of export-related refunds is a key channel through which VAT evasion occurs: the overclaimed amount constitutes 11–23% of the potential revenue. Finally, invoice mills are an important conduit through which the overclaim of refund takes place. More than 80% of their transactions are with exporters, who use these spurious invoices to obtain excessive refunds. Of the amount overclaimed as refund, nearly 37% is based on invoices issued by invoice mills. The results have important policy implications, highlighting the wide gulf between theory and practice of a modern tax instrument in weak enforcement setting. Optimal tax theory and policy advice to developing economies must take this difference into account.

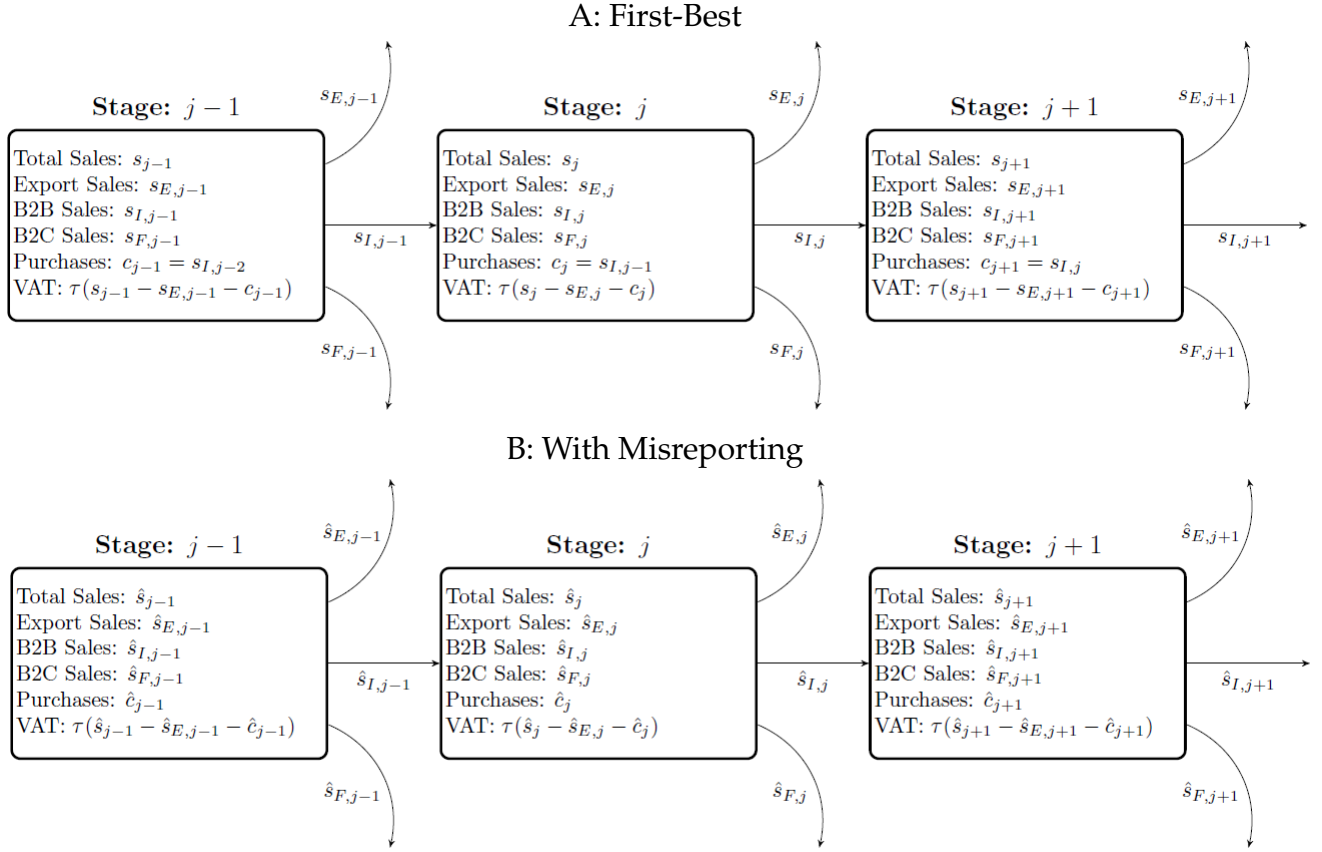
References

- ABADIE, ALBERTO, ATHEY, SUSAN, IMBENS, GUIDO W, & WOOLDRIDGE, JEFFREY. 2017. *When should you adjust standard errors for clustering?* Tech. rept. National Bureau of Economic Research.
- AGRAWAL, DAVID R, & ZIMMERMANN, LAURA. 2019. Production and evasion responses with limited state capacity. International Growth Centre Working Paper.
- ALMUNIA, MIGUEL, HJORT, JONAS, KNEBELMANN, JUSTINE, & TIAN, LIN. 2019. *Strategic or Confused Firms? Evidence from "Missing" Transactions in Uganda*. Mimeo.
- ALMUNIA, MIGUEL, BROCKMEYER, ANNE, MASCAGNI, GIULIA, NAIR, VEDANTH, & WASEEM, MAZHAR. 2022. Value-Added Tax: Theory and Practice. Unpublished.
- ATKINSON, ANTHONY BARNES, & STIGLITZ, JOSEPH E. 1976. The design of tax structure: direct versus indirect taxation. *Journal of public Economics*, 6(1-2), 55–75.
- BACHAS, PIERRE, & SOTO, MAURICIO. 2019. Not (ch) your average tax system: corporate taxation under weak enforcement. World Bank Policy Research Working Paper.

- BASRI, M CHATIB, FELIX, MAYARA, HANNA, REMA, & OLKEN, BENJAMIN A. 2019. *Tax Administration vs. Tax Rates: Evidence from Corporate Taxation in Indonesia*. Tech. rept. National Bureau of Economic Research.
- BEST, MICHAEL, SHAH, JAWAD, & WASEEM, MAZHAR. 2021. Detection Without Deterrence: Tax Audits with Limited Fiscal Capacity. Working Paper, University of Manchester.
- BEST, MICHAEL CARLOS, BROCKMEYER, ANNE, KLEVEN, HENRIK JACOBSEN, SPINNEWIJN, JOHANNES, & WASEEM, MAZHAR. 2015. Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan. *Journal of Political Economy*, **123**(6), 1311–1355.
- BROCKMEYER, ANNE, SMITH, SPENCER, HERNANDEZ, MARCO, & KETTLE, STEWART. 2019a. Casting a Wider Tax Net: Experimental Evidence from Costa Rica. *American Economic Journal: Economic Policy*, **11**(3), 55–87.
- BROCKMEYER, ANNE, ESTEFAN, ALEJANDRO MANUEL, ARRAS, KARINA RAMIREZ, & SERRATO, JUAN CARLOS SUÁREZ. 2019b. Taxing Property in Developing Countries: Theory and Evidence from Mexico. Working Paper.
- DE PAULA, AUREO, & SCHEINKMAN, JOSE A. 2010. Value-Added Taxes, Chain Effects, and Informality. *American Economic Journal: Macroeconomics*, **2**(4), 195–221.
- EBRILL, LIAM P, KEEN, MICHAEL, BODIN, JEAN-PAUL, SUMMERS, VICTORIA P, *et al.* 2001. *The modern VAT*. International Monetary Fund.
- FBR. 2005. CBR Quarterly Review. 5(1). Federal Board of Revenue, Pakistan, <https://tinyurl.com/y4tuxp3q>.
- GÓMEZ SABAINI, JUAN CARLOS, & JIMÉNEZ, JUAN PABLO. 2012. Tax structure and tax evasion in Latin America. *Macroeconomics of Development Series 118*.
- IMBENS, GUIDO W, & RUBIN, DONALD B. 2015. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- INTERNATIONAL TAX DIALOGUE. 2013. International Tax Dialogue, Key Issues and Debates in VAT, SME Taxation and the Tax Treatment of the Financial Sector. *International Tax Dialogue*.

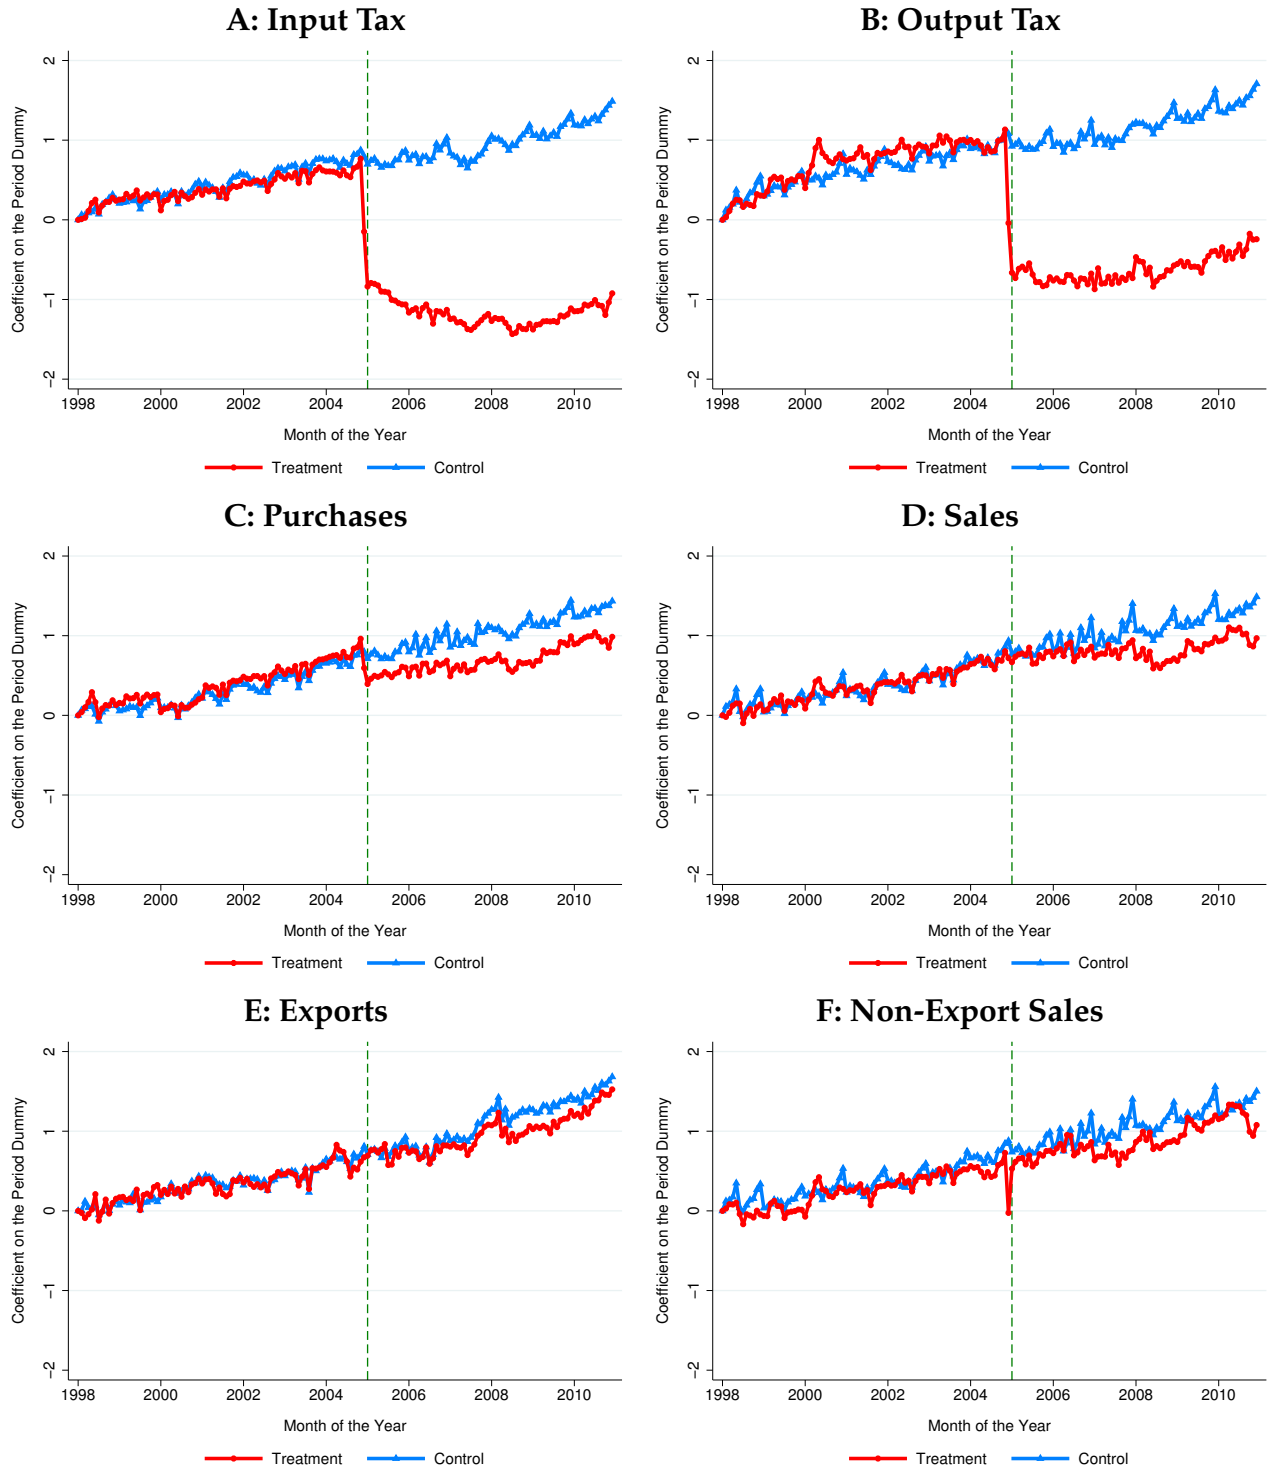
- KEEN, MICHAEL. 2007. VAT attacks! *International Tax and Public Finance*, **14**(4), 365–381.
- KEEN, MICHAEL, & SMITH, STEPHEN. 2006. VAT fraud and evasion: What do we know and what can be done? *National Tax Journal*, 861–887.
- MITTAL, SHEKHAR, REICH, OFIR, & MAHAJAN, APRAJIT. 2018. Who is bogus? using one-sided labels to identify fraudulent firms from tax returns. *Pages 1–11 of: Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*.
- NARITOMI, JOANA. 2019. Consumers as tax auditors. *American Economic Review*, **109**(9), 3031–72.
- OECD. 2017. Revenue Statistics, Organization for Economic Cooperation and Development. URL, <https://goo.gl/WC1M1L>.
- POMERANZ, DINA. 2015. No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax. *American Economic Review*, **105**(8), 2539–2569.
- SHAH, JAWAD A. 2019. *Using Computerized Information to Enforce VAT: Evidence from Pakistan*. Mimeo.
- SLEMROD, JOEL, & VELAYUDHAN, TEJASWI. 2020 (February). *The VAT at 100: A Comprehensive Health Assessment, With a Plan of Care and Suggestions for Diagnostic Tests*. University of Michigan Working Paper.
- SLEMROD, JOEL, REHMAN, OBEID UR, & WASEEM, MAZHAR. 2022. How Do Taxpayers Respond to Public Disclosure and Social Recognition Programs? Evidence from Pakistan. *The Review of Economics and Statistics*, **104**(1), 116–132.
- STIGLITZ, JOSEPH E. 2010. Development-oriented tax policy. *Taxation in developing countries: Six case studies and policy implications*, 11–36.
- WASEEM, MAZHAR. 2019. Information, Asymmetric Incentives, or Withholding: Understanding the Self-Enforcement of Value-added Tax. Working Paper, University of Manchester.
- WASEEM, MAZHAR. 2020. The Role of Withholding in the Self-Enforcement of a Value-Added Tax: Evidence from Pakistan. *The Review of Economics and Statistics*, 1–44. https://doi.org/10.1162/rest_a_00959.

FIGURE I: CONCEPTUAL FRAMEWORK



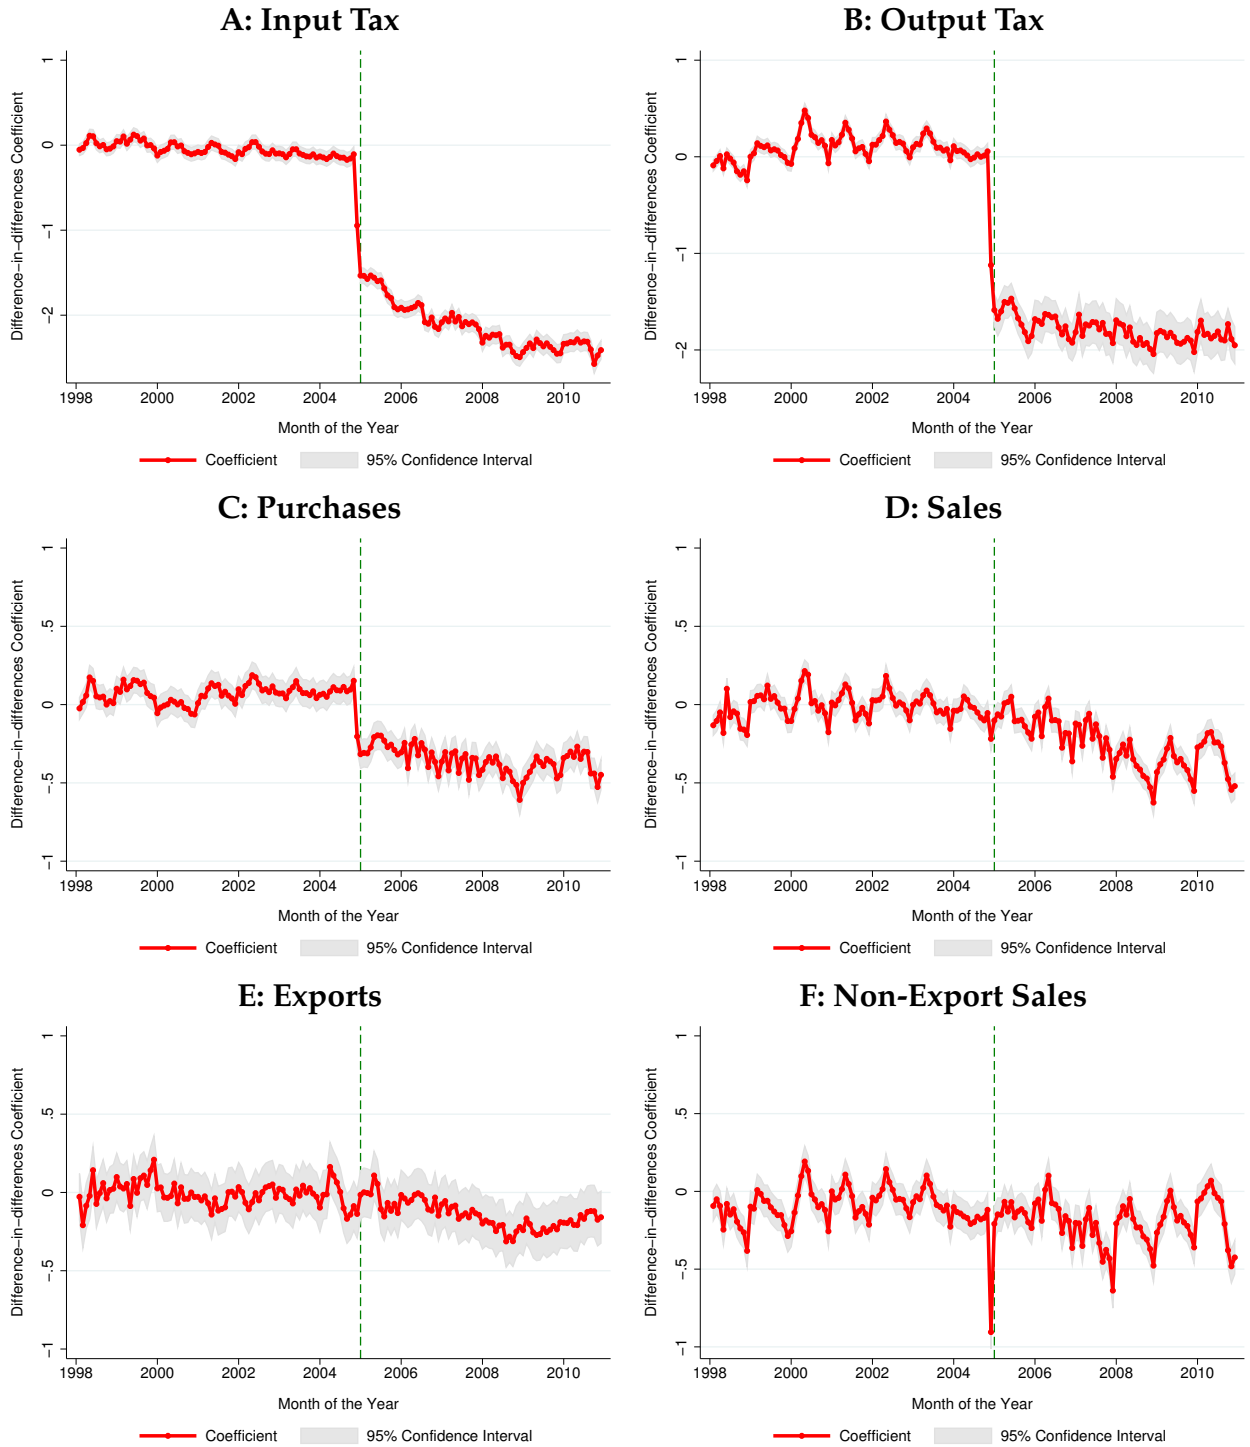
Notes: This figure displays three consecutive production stages of a supply chain. The top panel displays behavior under the first-best, where all firms report their sales and purchases truthfully. Total sales s_j of a firm consists of three components: (1) sales exported out of the country $s_{E,j}$; (2) sales made to other domestic firms $s_{I,j}$; and (3) sales made to final, domestic consumer $s_{F,j}$. Purchases here refer to the value of *taxable* intermediates, which here are all acquired from the previous production stage $c_j = s_{I,j-1}; \forall j \in 1, \dots, J$. The VAT liability of a firm equals its output tax $\tau(s_j - s_{E,j})$ minus the input tax τc_j . Panel B displays the same supply chain, but firms are now allowed to misreport. Firms' sales and purchases in this second-best world are denoted by $\hat{s}_{K,j}; K \in \{E, I, F\}$ and \hat{c}_j to distinguish them from their true values $s_{K,j}$ and c_j .

FIGURE II: FIRM RESPONSE TO THE TAX CUT




Notes: The figure compares the evolution of six VAT outcomes from the tax year 1998 to 2010 across the treatment and control groups. Treatment and control groups consist respectively of firms in the zero-rated and non-zero-rated industries. To construct these charts, I regress the log of the outcome variable shown in the title of each panel on the full set of firm and month fixed effects, dropping the dummy for July 1998. I then plot the coefficients on the time dummies of these regressions. The regressions are run separately for the two groups of firms. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines demarcate the time from which the zero-rating reform became applicable.

FIGURE III: FIRM RESPONSE TO THE TAX CUT



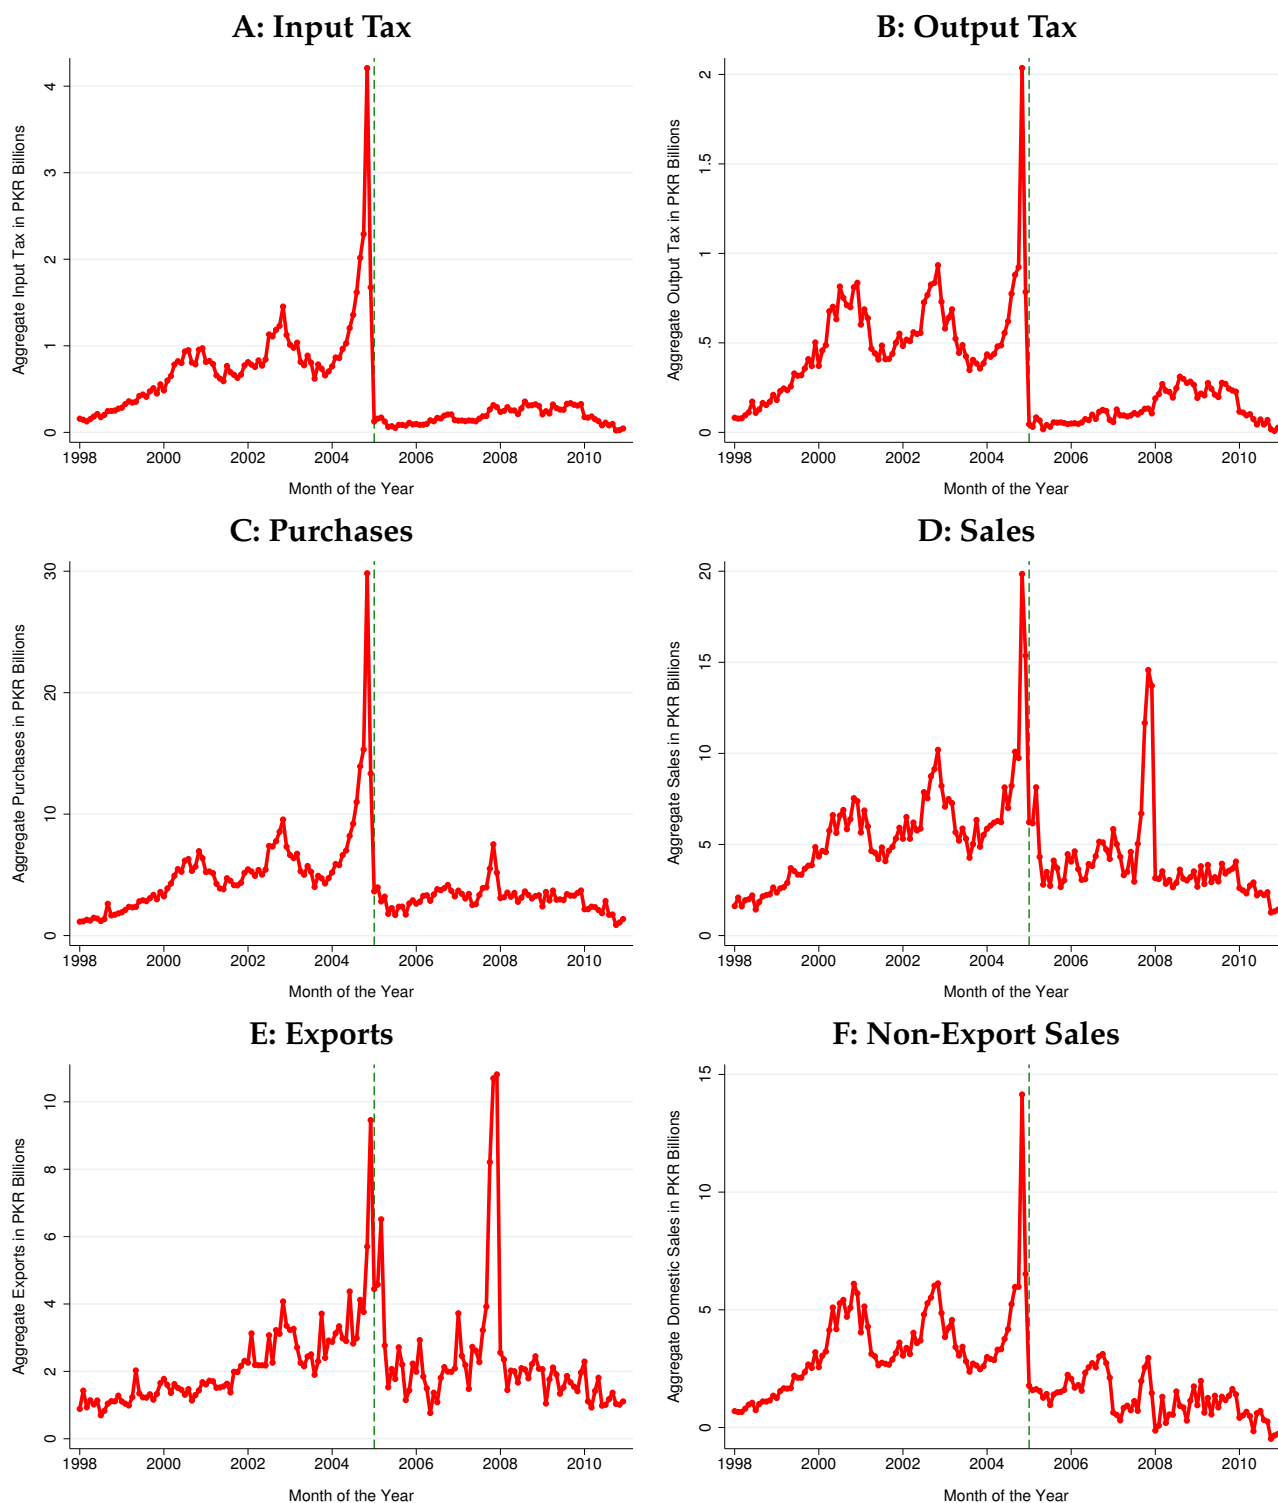
Notes: The figure shows the difference-in-differences version of the plots in Figure II. To construct these charts, I regress the log of the outcome variable shown in the title of each panel on the full set of firm, month, and month \times treat dummies, dropping the dummies for July 1998. I then plot the coefficients on the month \times treat dummies from these regressions, where $treat_i$ denotes that firm i belongs to a zero-rated industry. The gray surface plot shows the 95% confidence interval around the coefficient. I cluster standard errors at the firm level. Year t on the horizontal axis indicates July of the corresponding year. The vertical, dashed lines demarcate the time from which the zero-rating reform became applicable.

FIGURE IV: EXPORTS AS A FUNCTION OF PURCHASES



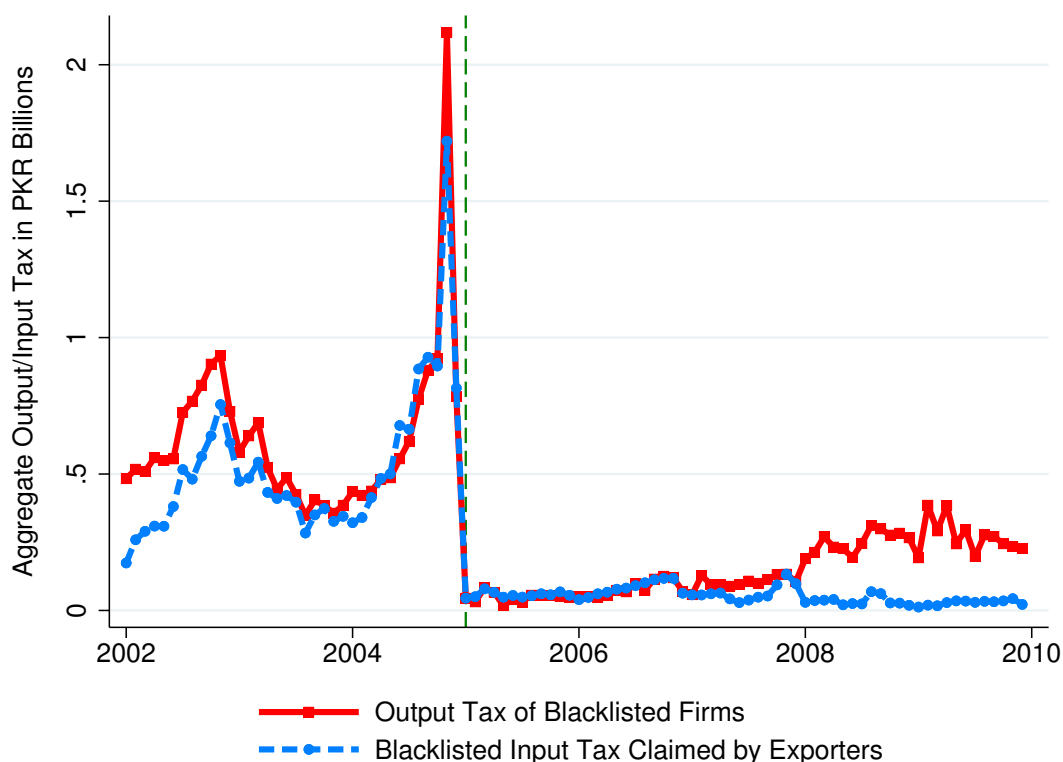
Notes: The figure explores the relationship between purchases of intermediates and exports reported by firms of the treated industries. The blue curve shows a binned scatter plot for the data. I group firms into bins on the basis of log of purchases reported by them. Each blue marker represents the average log export of firms within the bin. I restrict the sample to firms with log purchases within the fifth and 95th percentile of the aggregate log purchase distribution. The bin width is 0.05. The red curve is the straight line fitted onto the data using ordinary least squares. The slope of the fitted line and R^2 from the regression are reported in the panel. Panel A relates to the pre-reform years and Panel B to the post-reform years.

FIGURE V: AGGREGATE VALUES OF VAT OUTCOMES – INVOICE MILLS



Notes: The figure illustrates how VAT outcomes reported by invoice mills evolve around the time of the zero-rating reform. The sample contains blacklisted and suspended firms operating in the zero-rated industries. Each panel of the figure shows the aggregate value of the outcome indicated in the title of the panel for the given month. Year t indicated in the horizontal axis denotes the month July of the corresponding year. To deal with outliers, I drop ten observations with the highest values of the given outcome in the entire sample. For example, for constructing Panel A, I sort all firm-month observations on the basis of Output Tax in a descending order and drop the top-ten observations. The dashed, vertical lines in the plots demarcate the time from which the zero-rated reform takes effect.

FIGURE VI: INVOICE MILLS AND VAT REFUND



Notes: The figure explores the linkages between invoice mills and exporters. The solid, red curve in the figure plots the aggregate value of output tax involved in VAT returns filed by all blacklisted and suspended firms in the given tax period (month). The dashed, blue curve, on the other hand, shows the aggregate value of input tax claimed by exporters on the invoices of blacklisted firm for the given tax period (month). This curve has been plotted using transaction-level data filed by exporters in support of their refund claims, which provides the supplier-wise details of all purchases of intermediates made by them in the corresponding tax period. That these two curves almost lie above each other shows that the primary purpose of the existence of invoice mills is to help exporters claim exaggerated refunds. Year t indicated in the horizontal axis denotes the month July of the year. Dashed vertical lines in the plots demarcate the time from which the zero-rated reform takes effect.

TABLE I: FIRM RESPONSE TO THE TAX CUT

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Baseline Specification</u>						
treat × after	-1.961 (0.026)	-1.842 (0.071)	-0.419 (0.017)	-0.223 (0.012)	-0.106 (0.025)	-0.082 (0.014)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
<u>B: Placebo Specification</u>						
treat × after	-0.004 (0.012)	0.001 (0.011)	0.052 (0.012)	0.019 (0.010)	-0.022 (0.026)	-0.014 (0.011)
Observations	1,999,987	2,423,815	1,979,624	2,583,424	306,931	2,381,448
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores how firms respond to the reduction of the rate applicable to their supply chain to zero. The coefficient $\text{treat} \times \text{after}$ shows $\hat{\gamma}$ from model (8), where the dummy variable treat_i denotes that firm i belongs to a zero-rated industry and the dummy variable after_t that month t falls in the tax year 2005 and later. Panel B shows results from parallel placebo regressions, where I restrict the sample to pre-reform years only, defining the period beginning from July 2002 as the *after* period. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE II: FIRM RESPONSE TO THE TAX CUT – DYNAMICS

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
treat × 2005	-1.568 (0.022)	-1.732 (0.062)	-0.335 (0.015)	-0.072 (0.010)	-0.019 (0.022)	-0.046 (0.013)
treat × 2006	-1.891 (0.027)	-1.816 (0.070)	-0.391 (0.017)	-0.126 (0.012)	-0.035 (0.027)	-0.028 (0.014)
treat × 2007	-1.984 (0.029)	-1.855 (0.088)	-0.440 (0.021)	-0.203 (0.016)	-0.103 (0.030)	-0.220 (0.022)
treat × 2008	-2.240 (0.033)	-1.949 (0.090)	-0.498 (0.021)	-0.383 (0.018)	-0.216 (0.033)	-0.135 (0.017)
treat × 2009	-2.283 (0.035)	-1.956 (0.090)	-0.475 (0.023)	-0.365 (0.019)	-0.212 (0.036)	-0.083 (0.018)
treat × 2010	-2.270 (0.036)	-1.911 (0.084)	-0.432 (0.024)	-0.295 (0.020)	-0.148 (0.038)	-0.021 (0.019)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores how firm response to the zero-rating reform evolves over time. I estimate an augmented version of the difference-in-differences model (8), including interactions of the treatment variable with all post-reform years. The dummy variable $treat_i$ denotes that firm i belongs to a zero-rated industry. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE III: FIRM RESPONSE TO THE TAX CUT – ACROSS THE SUPPLY CHAIN

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
treat × after	-2.040 (0.053)	-0.468 (0.074)	-0.429 (0.034)	-0.266 (0.027)	-0.108 (0.035)	0.092 (0.033)
treat × after × ginning	2.053 (0.075)	-2.184 (0.096)	0.389 (0.067)	0.305 (0.036)	0.193 (0.057)	-0.001 (0.042)
treat × after × spinning	-0.504 (0.074)	-3.613 (0.207)	0.110 (0.045)	0.030 (0.035)	0.210 (0.071)	-0.351 (0.042)
treat × after × weaving	0.018 (0.061)	-1.831 (0.147)	0.001 (0.039)	0.026 (0.031)	0.025 (0.044)	-0.204 (0.037)
Baseline Coefficient	-1.978 (0.026)	-1.925 (0.076)	-0.394 (0.017)	-0.221 (0.013)	-0.065 (0.027)	-0.082 (0.014)
Observations	3,685,909	4,158,258	3,934,914	4,987,477	563,822	4,597,087
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores heterogeneity in firm-response to the zero-rating reform within the textile industry. I restrict the treatment sample to firms in the textile industry only and estimate a triple-difference version of model (8). The dummy variables $ginning_i$, $spinning_i$ and $weaving_i$ denote that the firm i belongs to the corresponding production stage within the textile industry. Baseline coefficient reports the $treat \times after$ coefficient I obtain from estimating the model without the triple-interaction terms. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE IV: CALCULATING VAT EVASION

	<u>2004</u>	<u>2003</u>	<u>2002</u>	<u>2001</u>
	(1)	(2)	(3)	(4)
<u>A: Underpaid VAT</u>				
1. $\sum_j \tau.(s_j - \hat{s}_j)$	-47.6	-38.9	-31.6	-27.1
	(3.4)	(2.8)	(2.3)	(2.0)
2. $\sum_j \tau.(s_{E,j} - \hat{s}_{E,j})$	-11.5	-7.6	-6.2	-5.2
	(3.9)	(2.6)	(2.1)	(1.8)
3. $\sum_j \tau.(c_j - \hat{c}_j)$	-78.5	-61.9	-49.3	-40.8
	(3.9)	(3.1)	(2.4)	(2.0)
4. $\Delta R = R - \hat{R}$	42.4	30.5	23.9	18.9
	(6.5)	(4.9)	(3.9)	(3.3)
5. $\Delta R/R$	0.459	0.386	0.359	0.312
<u>B: Overclaimed Refund</u>				
6. $\sum_j \tau.(s_{E,j} - \hat{s}_{E,j})$	-11.5	-7.6	-6.2	-5.2
	(3.9)	(2.6)	(2.1)	(1.8)
7. $\sum_j \tau.(c_{E,j} - \hat{c}_{E,j})$	-23.0	-15.1	-12.3	-10.5
	(15.7)	(10.3)	(8.4)	(7.1)
8. ΔR_E	23.0	15.1	12.3	10.5
	(15.7)	(10.3)	(8.4)	(7.1)
9. $\Delta R_E/R$	0.249	0.191	0.186	0.172

Notes: The table uses formulas (3) and (6) to quantify VAT noncompliance implied by the estimates in Table I. Panel A computes the amount of VAT evaded by the treated industries in the four baseline years using formula (3). Rows 1-3 compute the three terms on the RHS of the formula. To compute each term, I proceed in two steps. First, I convert the response of the corresponding item in Table I into percentage terms (for example, sales response of -0.223 implies a percent change of $-[\exp(0.223) - 1] = -25\%$). I multiply the percent change with the aggregate value of the variable reported in the corresponding year to estimate the amount by which it was overreported in the year. Row 4 displays the sum of Rows 1-3, and Row 5 expresses the evaded amount as a percent of the potential revenue in the year. I calculate the potential revenue as the tax due on the value added reported in the year plus the evaded amount given in row 4. Panel B estimates the first term on the RHS of formula (6), illustrating the level of refund overclaimed in the treated industries. Row 6 reproduces the exports response computed in Row 2. Using the production function of exports estimated in Figure IV, Row 7 converts the exports response into the amount by which purchases meant for export were overreported. Row 8 shows the PKR value of the excess claimed refund (negative of Row 7). Row 9 expresses the excess refund as a percent of the potential revenue in the year. All amounts are in PKR billions. To deal with outliers, I exclude ten observations with the highest values of exports and non-export sales in the entire sample (1998-2010).

A Online Appendix


A.1 Definition of Variables

- (i) **Input Tax.** The value of VAT credit claimed on purchases of intermediates made by a registered firm in a given tax period (month). It equals $\tau \cdot \hat{c}_{it}$, where τ is the applicable VAT rate and \hat{c}_{it} is the value of purchases of intermediates claimed by firm i in period t .
- (ii) **Output Tax.** The value of VAT charged on non-export sales made by a registered firm in a given tax period (month). It equals $\tau \cdot (\hat{s}_{it} - \hat{s}_{E,it})$, where τ is the applicable VAT rate and $(\hat{s}_{it} - \hat{s}_{E,it})$ is the value of non-export sales reported by firm i in period t .
- (iii) **Purchases.** The value of all taxable intermediates acquired by a firm in a given tax period (month).
- (iv) **Sales.** The value of all goods and services supplied by a firm in a given tax period (month).
- (v) **Exports.** The value of all goods and services exported by a firm in a given tax period (month).
- (vi) **Non-Export Sales.** The value of all goods and services supplied by a firm minus the value of all goods and services exported by a firm in a given tax period (month).
- (vii) **Manufacturer.** A firm whose principal business activity is the manufacture of goods. Manufacturing is the process whereby a firm converts inputs into a distinct article capable of being put to use differently than inputs and includes any process incidental or ancillary to it.
- (viii) **Wholesaler.** Wholesaler' includes a dealer and means any person who carries on, whether regularly or otherwise, the business of buying and selling goods by wholesale or of supplying or distributing goods, directly or indirectly, by wholesale for cash or deferred payment or for commission or other valuable consideration or stores such goods belonging to others as an agent for the purpose of sale; and includes a person supplying taxable goods to a person who deducts income tax at source under the Income Tax Ordinance, 2001.

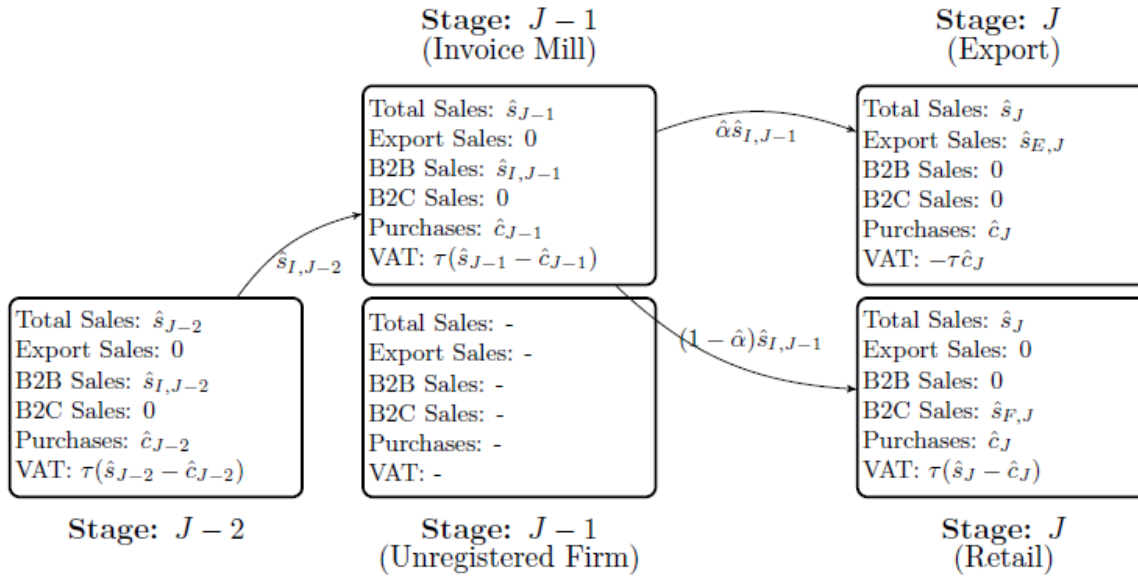
- (ix) **Retailer.** A person, supplying goods to general public for the purpose of consumption.
- (x) **Industry.** The Pakistani tax administration uses 4-digit Harmonized Commodity Description and Coding System (HS code) to classify firms into industry. The code, used by customs administrations throughout the world, divides all goods and services into 99 chapters (the first two digits in the code) and 21 sections. The sections broadly correspond to major industries in the country. I take the section a firm falls in as its industry. Table shows the sections, HS code, and description of these industries.
- (xi) **Major City** A firm registered either in Karachi or Lahore, the two largest cities in Pakistan on the basis of both population and GDP.
- (xii) **Initial Capital.** The value of initial capital of the firm, as reported by it at the time of registration for VAT.

FIGURE A.I: HOW INVOICE MILLS WORK

A: Real Input-Output Linkages




B: Reported Input-Output Linkages



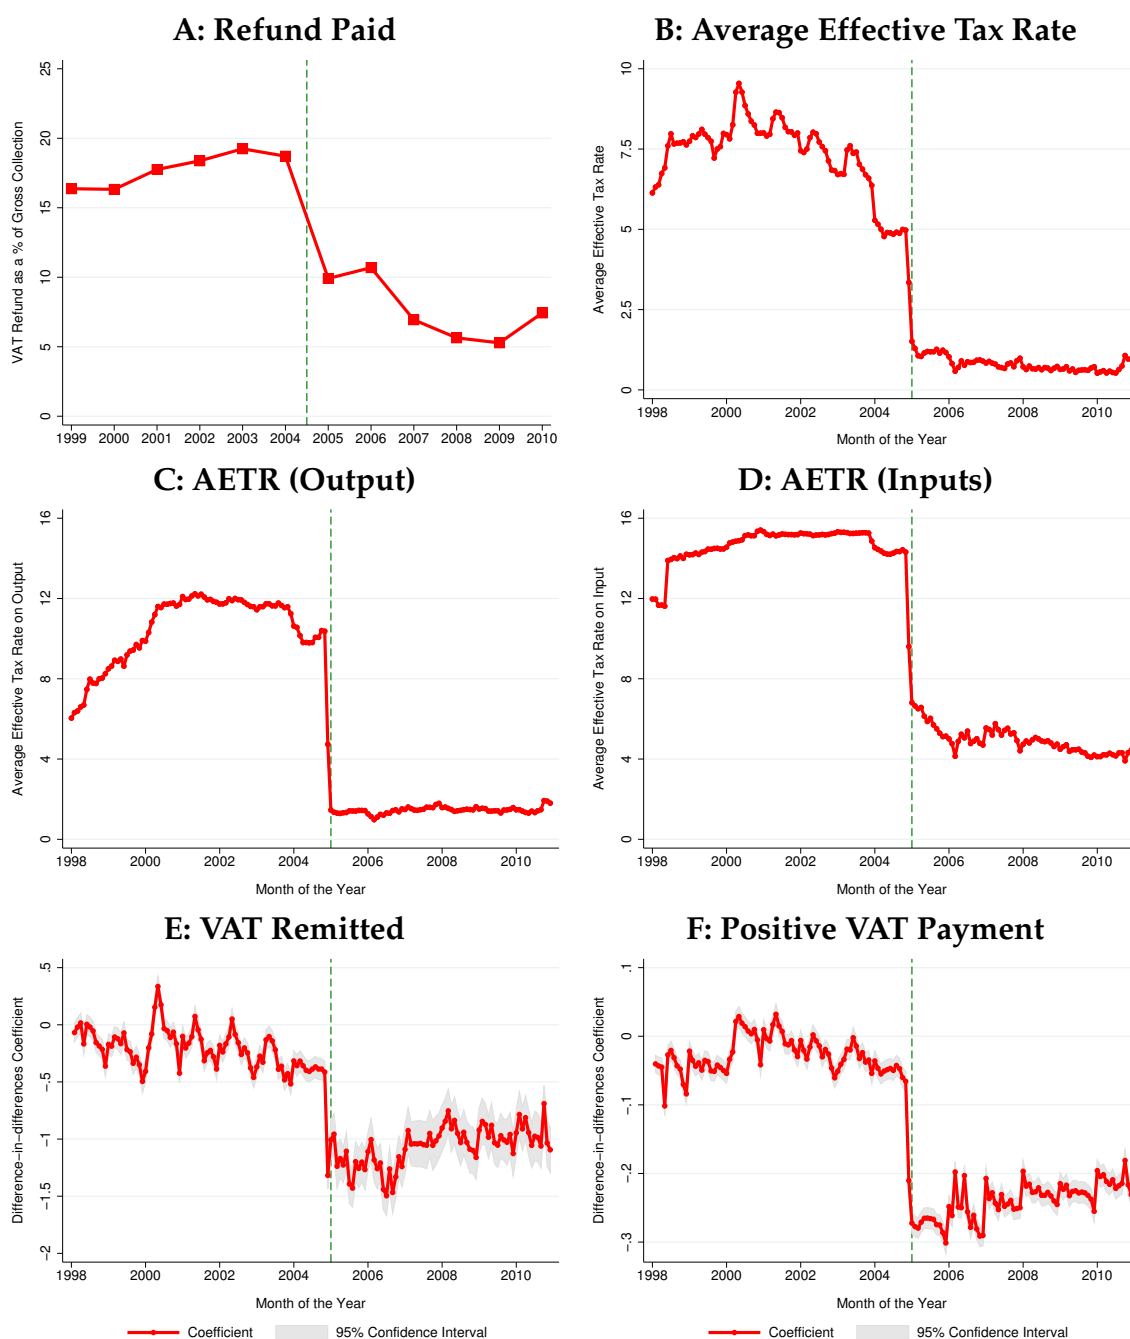
Notes: This figure shows how an invoice mills work. The top panel displays real input-output linkages in the last three stages of a supply chain. The final stage of the chain consists of two firms: the bottom firm deals exclusively in domestic sales and the top in exports. The middle firm sells a proportion α of its output to the exporter and the rest to the retailer. The middle firm is not in the VAT net. The two firms in the final stage thus cannot adjust input tax paid at the $J - 2$ and previous stages. To exploit the gap, an invoice mills places itself in the stage $J - 1$. It buys invoices from the bottom firm and sells them to the firms in the final stage. The bottom panel displays the reported input-output linkages in this setup. Given that the exporting firm has a greater incentive to exaggerate the value of its purchases, the invoice mills shifts a greater proportion of tax remitted at the bottom stages toward the exporting firms i.e. $\hat{\alpha} > \alpha$.

FIGURE A.II: VAT RATES



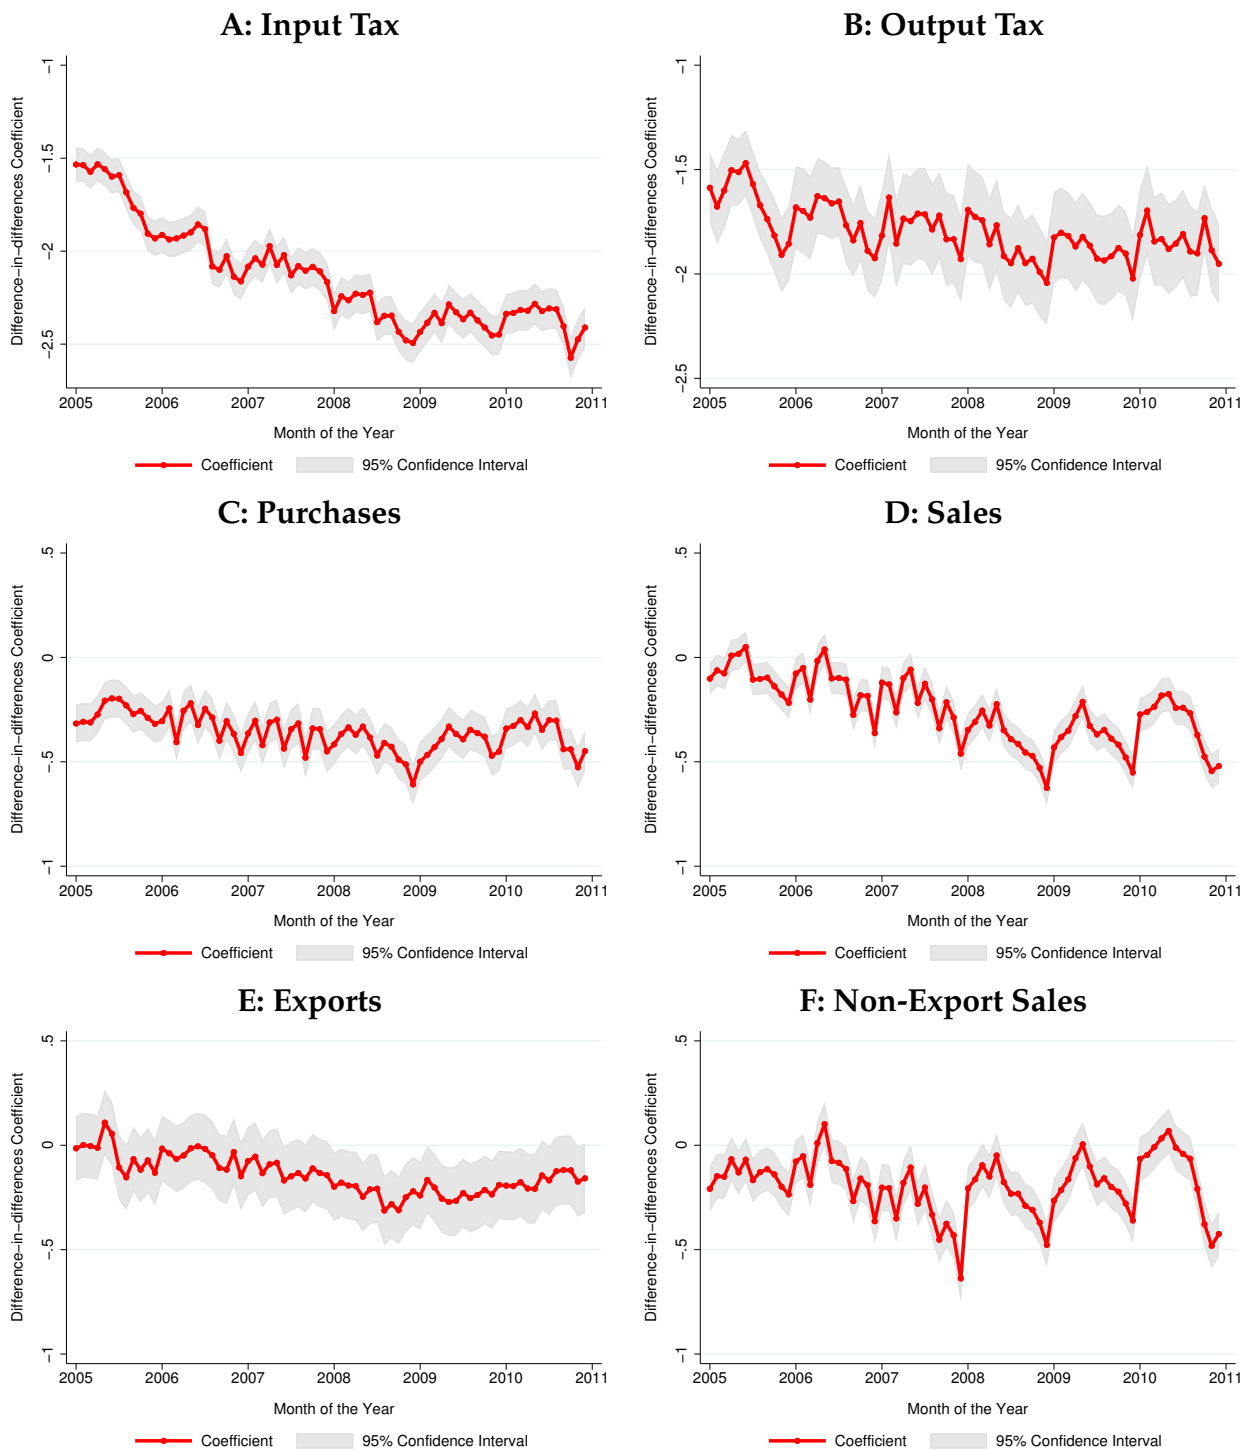
Notes: The figure shows the VAT rates applicable in Pakistan from July 1998 to June 2011. Panel A shows the standard VAT rate, which is applied to all sales made by a registered firm regardless of whether the recipient is registered or not. The rate largely stayed at 15%, but was increased to 16% from July 2008. Panel B reports the Further Tax Rate. This rate is added to the standard rate whenever the sale is made to an unregistered firm. For example, supplies made by a registered firm in July 1998 were subject to a rate of 15% if the recipient was a registered firm or an end-consumer and 16% if the recipient was an unregistered firm. To claim that a sale has been made to end-consumer, the selling firm must be registered as a retailer.

FIGURE A.III: FIRST STAGE



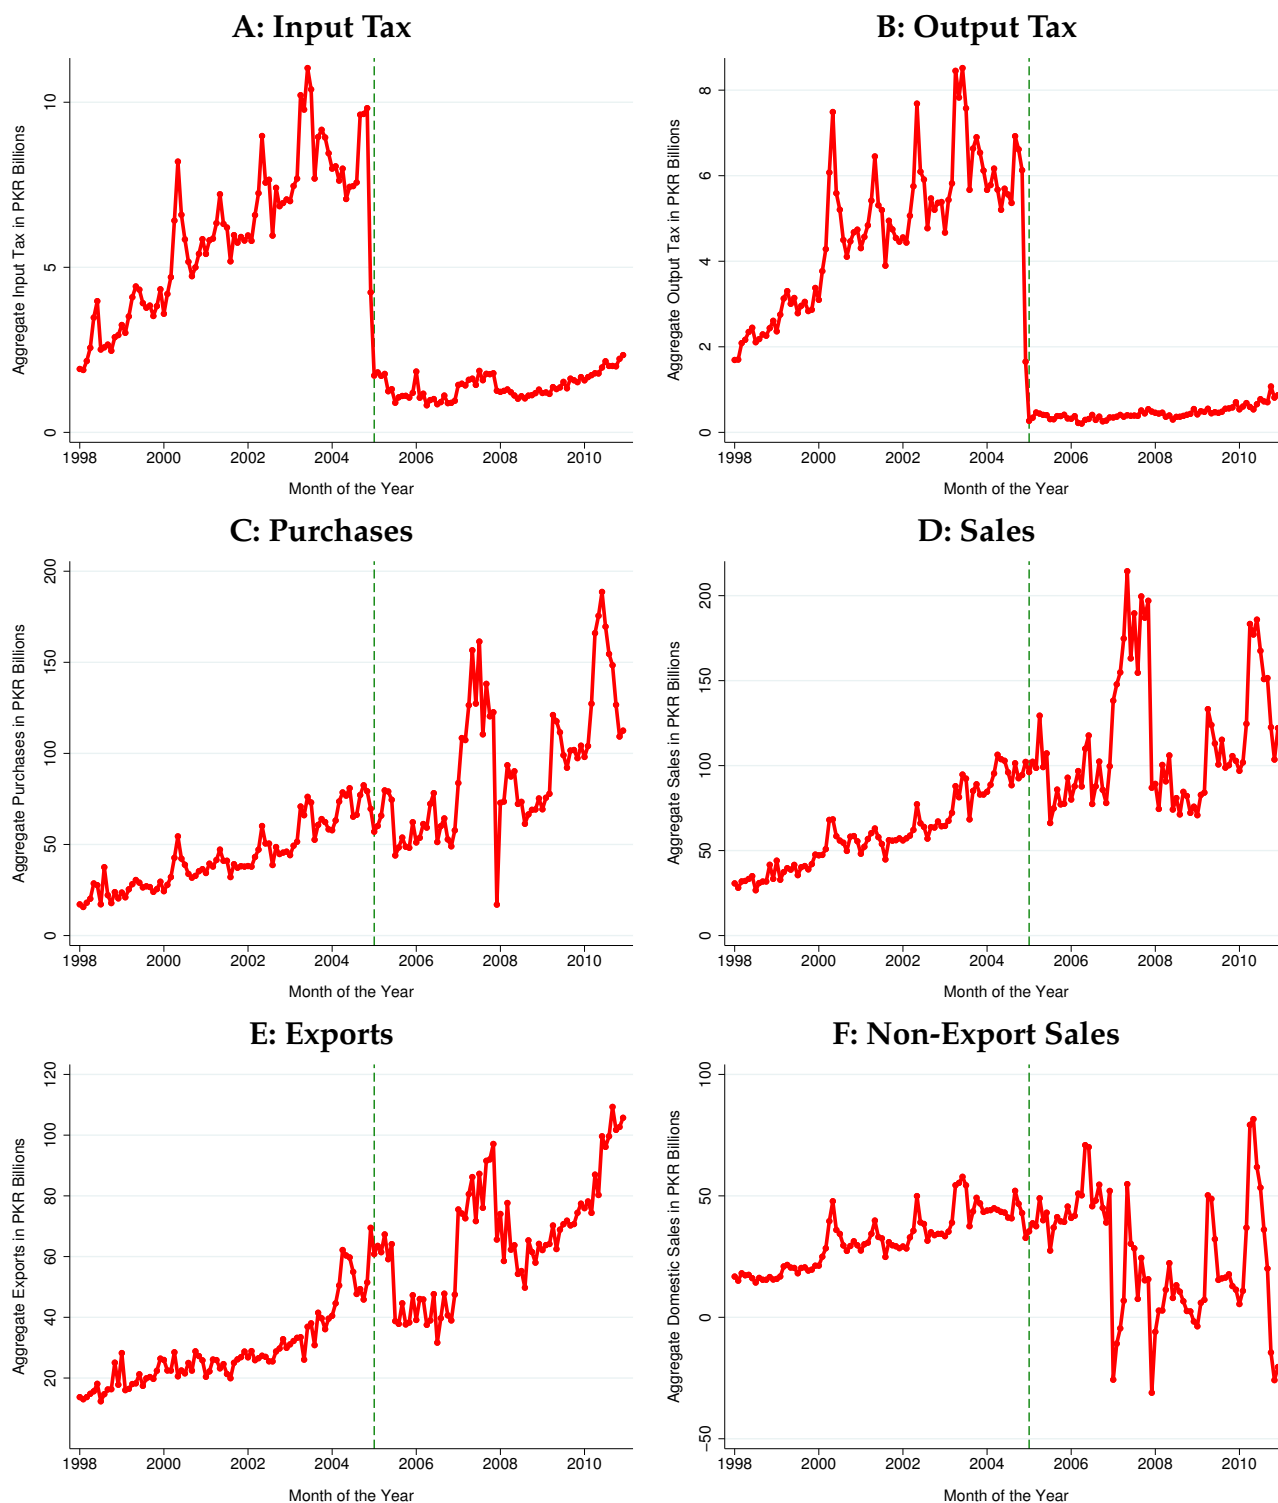
Notes: The figure shows that the zero-rating reform led to a sharp decrease in the output tax liability, input tax credit, and tax burden of treated firms. Panel A shows the evolution of the VAT refund paid in Pakistan. Each marker in the plot denotes the aggregate VAT refund paid to firms in *all* industries as a percentage of the gross VAT collection in that year. The data for this plot are compiled from the FBR yearbooks available [here](#). Panel B plots the average effective tax rate of treated firms over the sample period. The effective tax rate (ETR) is defined as VAT liability of the firm divided by its sales. Panels C-D separately show the AETR on output and input. The output ETR is defined as the output tax liability divided by total sales, whereas the input ETR is defined as the input tax credit of the firm divided by total purchases. Each marker in these plots shows the average ETR of treated firms for the given month. Panel E-F plot the estimates from the DD version of model (10). Panel E uses VAT liability as the outcome variable, whereas panel F uses an indicator that the firm’s tax liability for the given month is positive as the outcome variable. The year t in the horizontal axis denotes the month July of year t . The dashed, vertical line represents the time from which the zero-rating reform became applicable.

FIGURE A.IV: FIRM RESPONSE TO THE TAX CUT



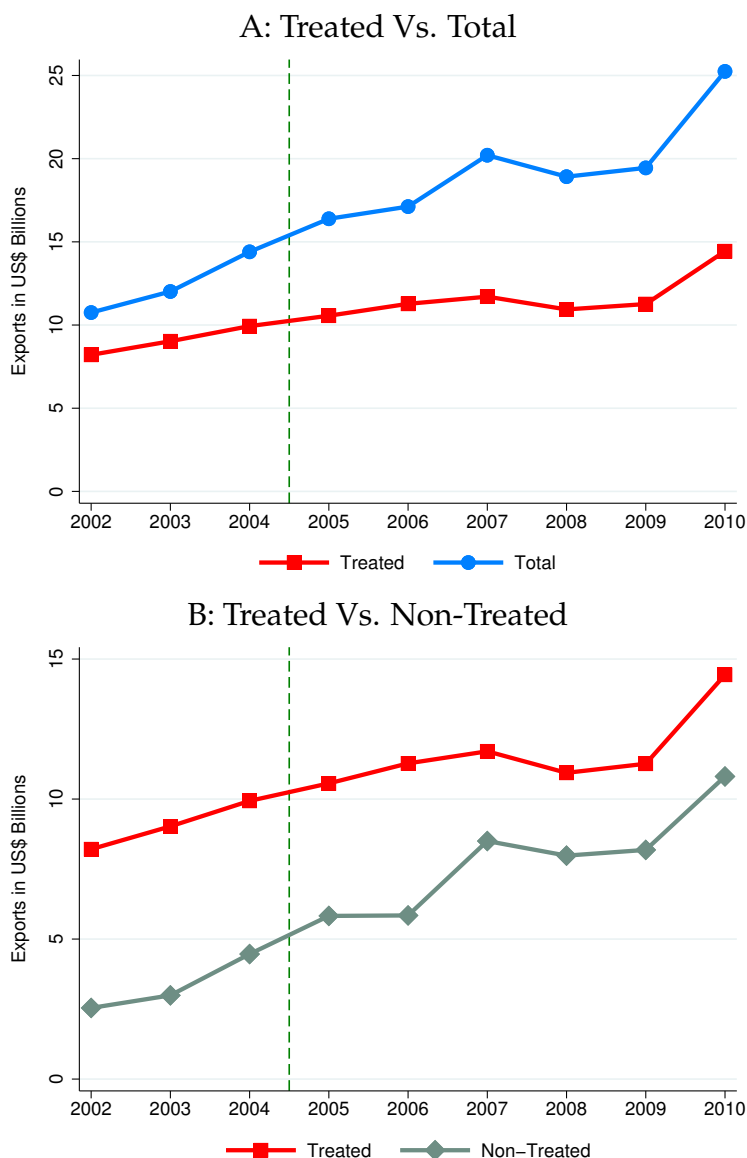
Notes: The figure reproduces the event study results reported in Figure II. The only difference between the two figure is that I show only the post-reform periods here. To construct these charts, I regress the log of the outcome variable shown in the title of each panel on the full set of firm, month, and month \times treat dummies, dropping the dummies for July 1998. I then plot the coefficients on the month \times treat dummies from these regressions, where treat $_i$ denotes that firm i belongs to a zero-rated industry. The gray surface plot shows the 95% confidence interval around the coefficient. I cluster standard errors at the firm level. Year t on the horizontal axis indicates July of the corresponding year.

FIGURE A.V: AGGREGATE VALUES OF VAT OUTCOMES – EXCLUDING INVOICE MILLS



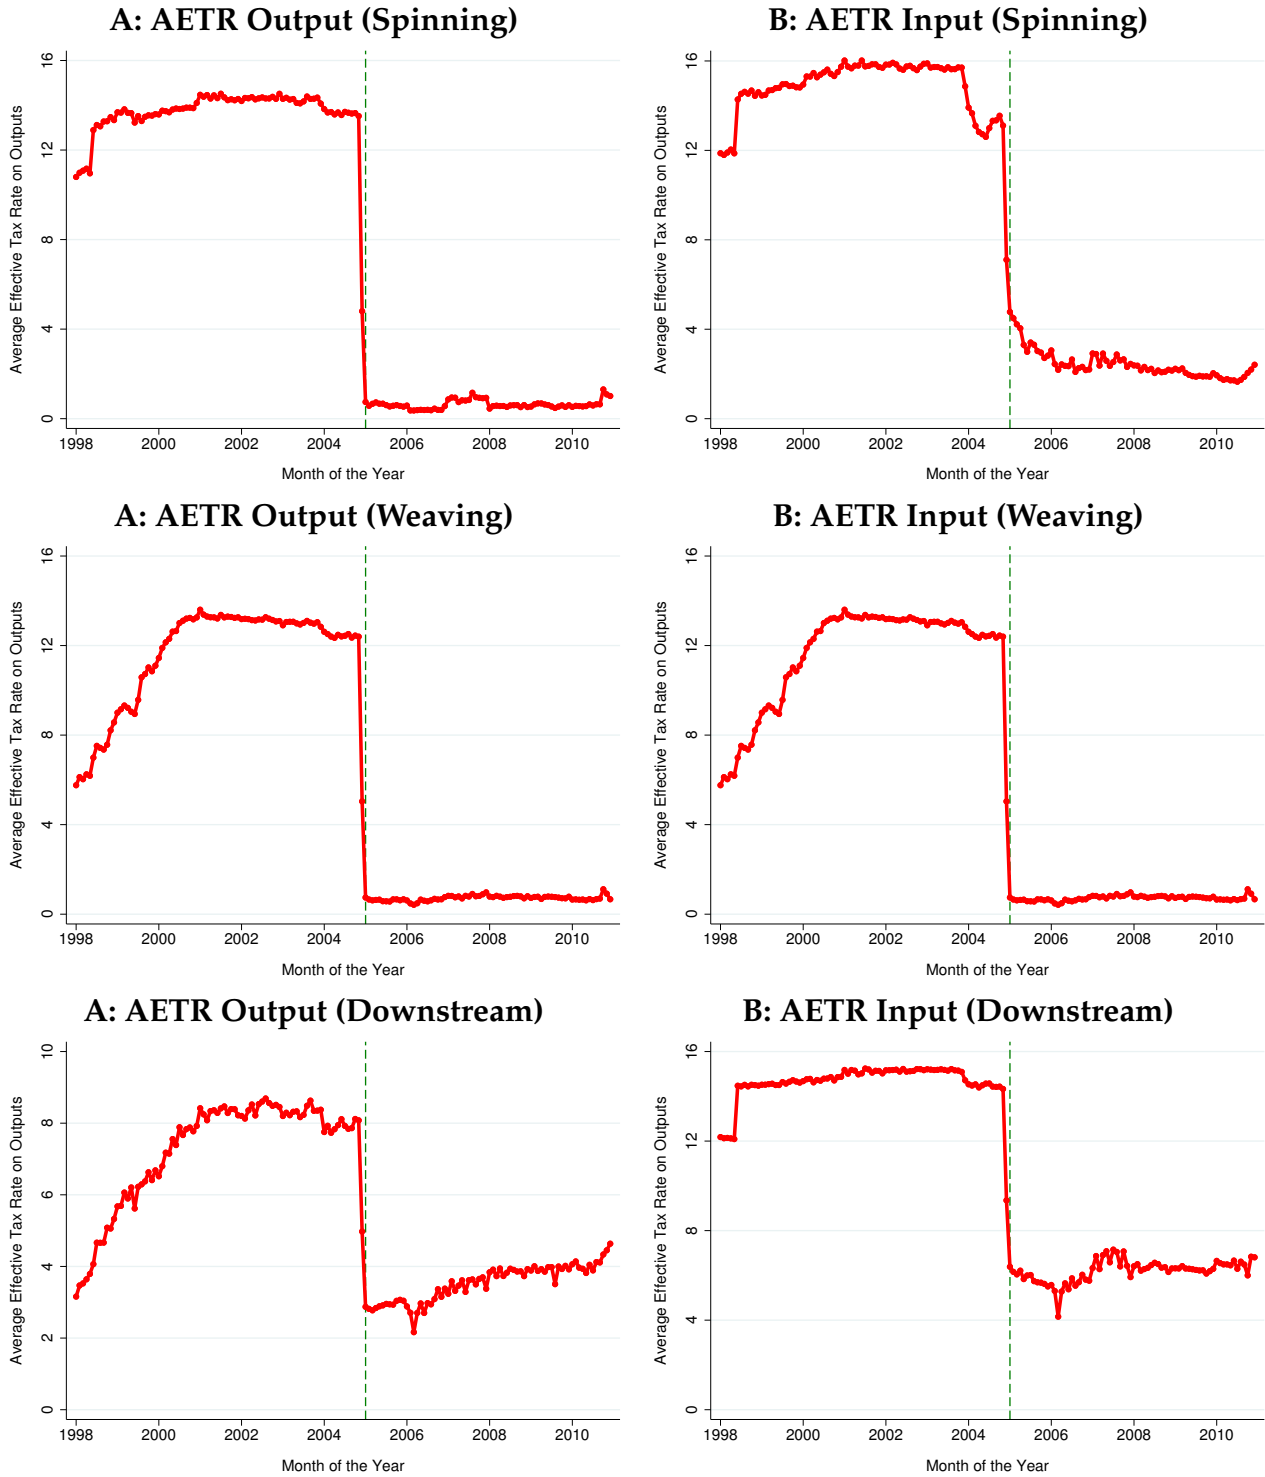
Notes: The figure illustrates how VAT outcomes evolve around the time of the zero-rating reform. The sample contains all firms operating in the zero-rated industries other than the blacklisted and suspended firms. Each panel of the figure shows the aggregate value of the outcome indicated in the title of the panel for the given month. Year t indicated in the horizontal axis denotes the month July of the corresponding year. To deal with outliers, I drop ten observations with the highest values of the given outcome in the entire sample. For example, for constructing Panel A, I sort all firm-month observations on the basis of Output Tax in a descending order and drop the top-ten observations. The dashed, vertical lines in the plots demarcate the time from which the zero-rated reform takes effect.

FIGURE A.VI: PAKISTAN'S EXPORTS OVER TIME




Notes: The figure shows the evolution of aggregate exports in Pakistan. Panel A compares the exports of treated industries in a given year with total exports of the country in that year, whereas Panel B compares exports of treated and nontreated industries. The data are from the central bank of Pakistan available [here](#). Treated industries include textile, leather, and footwear, whereas the nontreated industries include all other industries. The data are available from 2002 onward only. The timing of the zero-rating reform is indicated by the vertical dashed line.

FIGURE A.VII: AETR IN THE TEXTILE SUPPLY CHAIN




Notes: The figure shows the average effective tax rate on output and input of the three most important production stages of the textile industry. The output ETR is defined as the output tax liability divided by total sales, whereas the input ETR is defined as the input tax credit of the firm divided by total purchases. Each marker in these plots shows the average ETR of firms in the corresponding production stage for the given month. The year t in the horizontal axis denotes the month July of year t . The dashed, vertical line represents the time from which the zero-rating reform became applicable.

FIGURE A.VIII: HETEROGENEITY BY FIRM SIZE



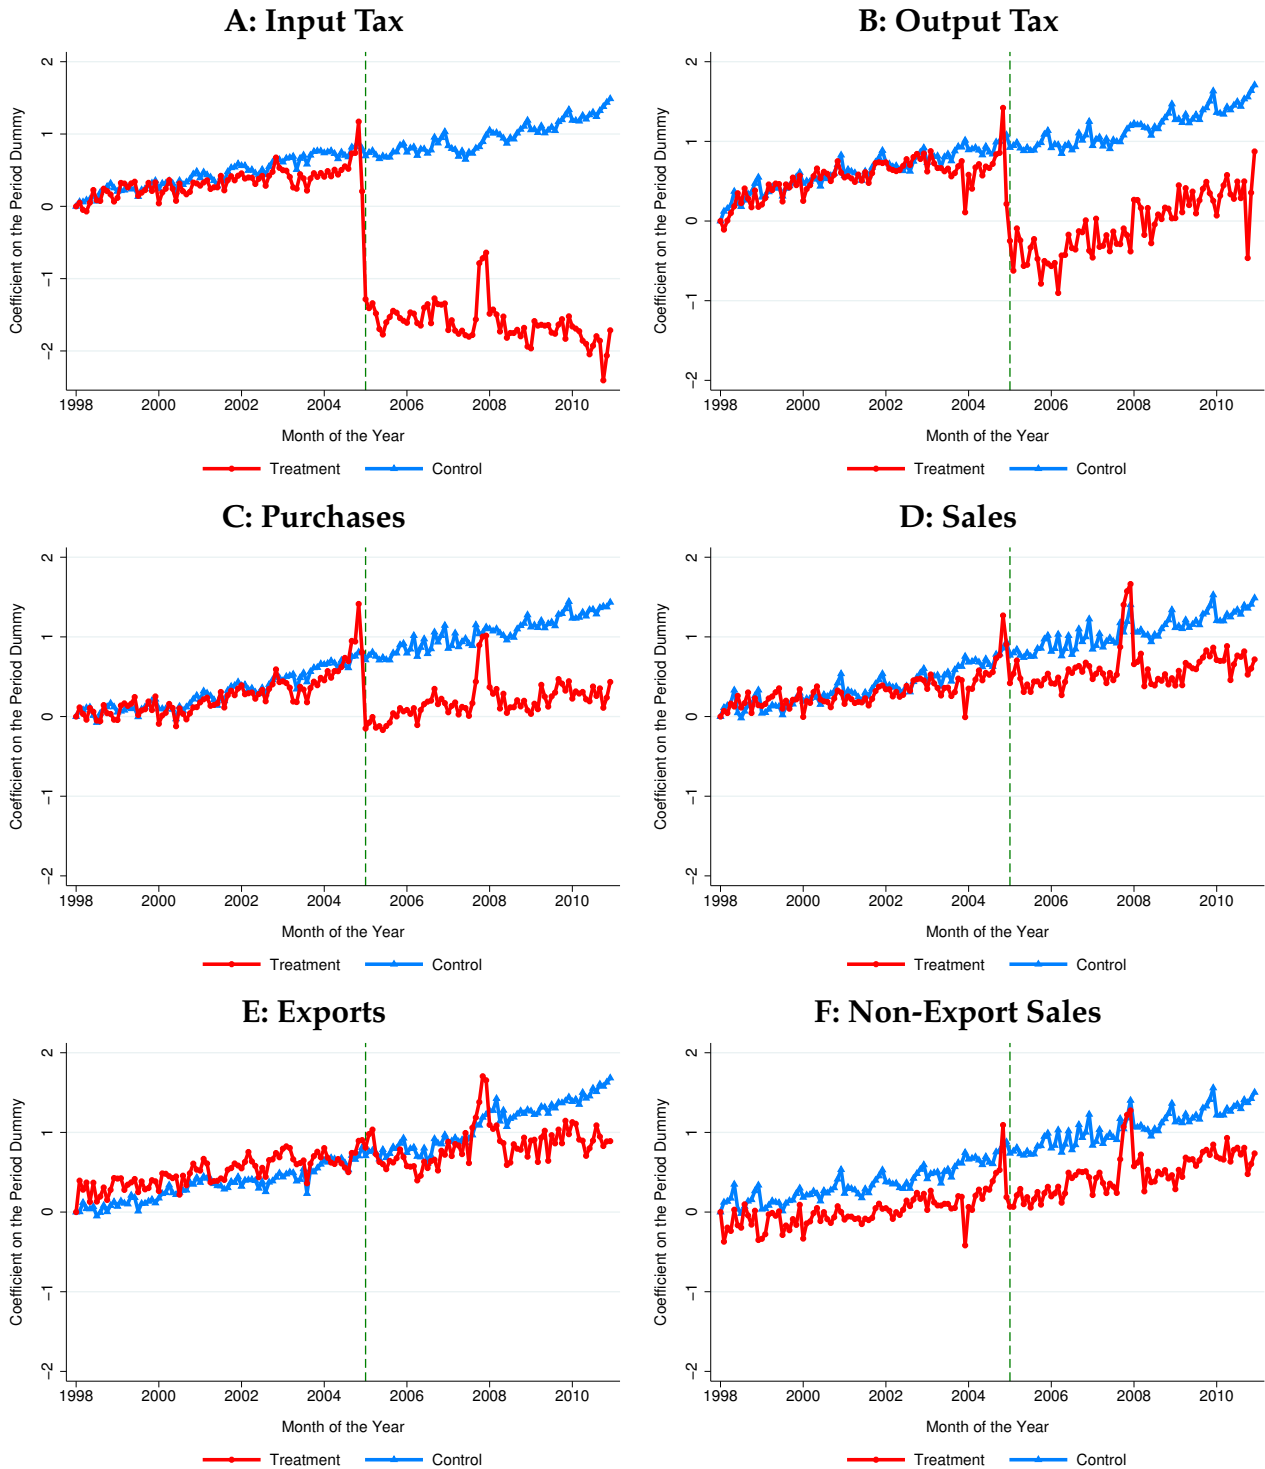
Notes: The figure explores heterogeneity in response of export sales. The baseline estimates of this response are reported in Column (5) of Table I (top panel). I estimate a triple-difference model that includes interactions of the DD term with five dummy variables indicating the firm's baseline share of exports in its total sales. I divide this share into five categories. The coefficients on these interaction terms along with 95 percent confidence interval are plotted in the above figure. The baseline estimate of the export response is indicated by the vertical line.

FIGURE A.IX: EMERGENCE, GROWTH, AND DECLINE OF INVOICE MILLS



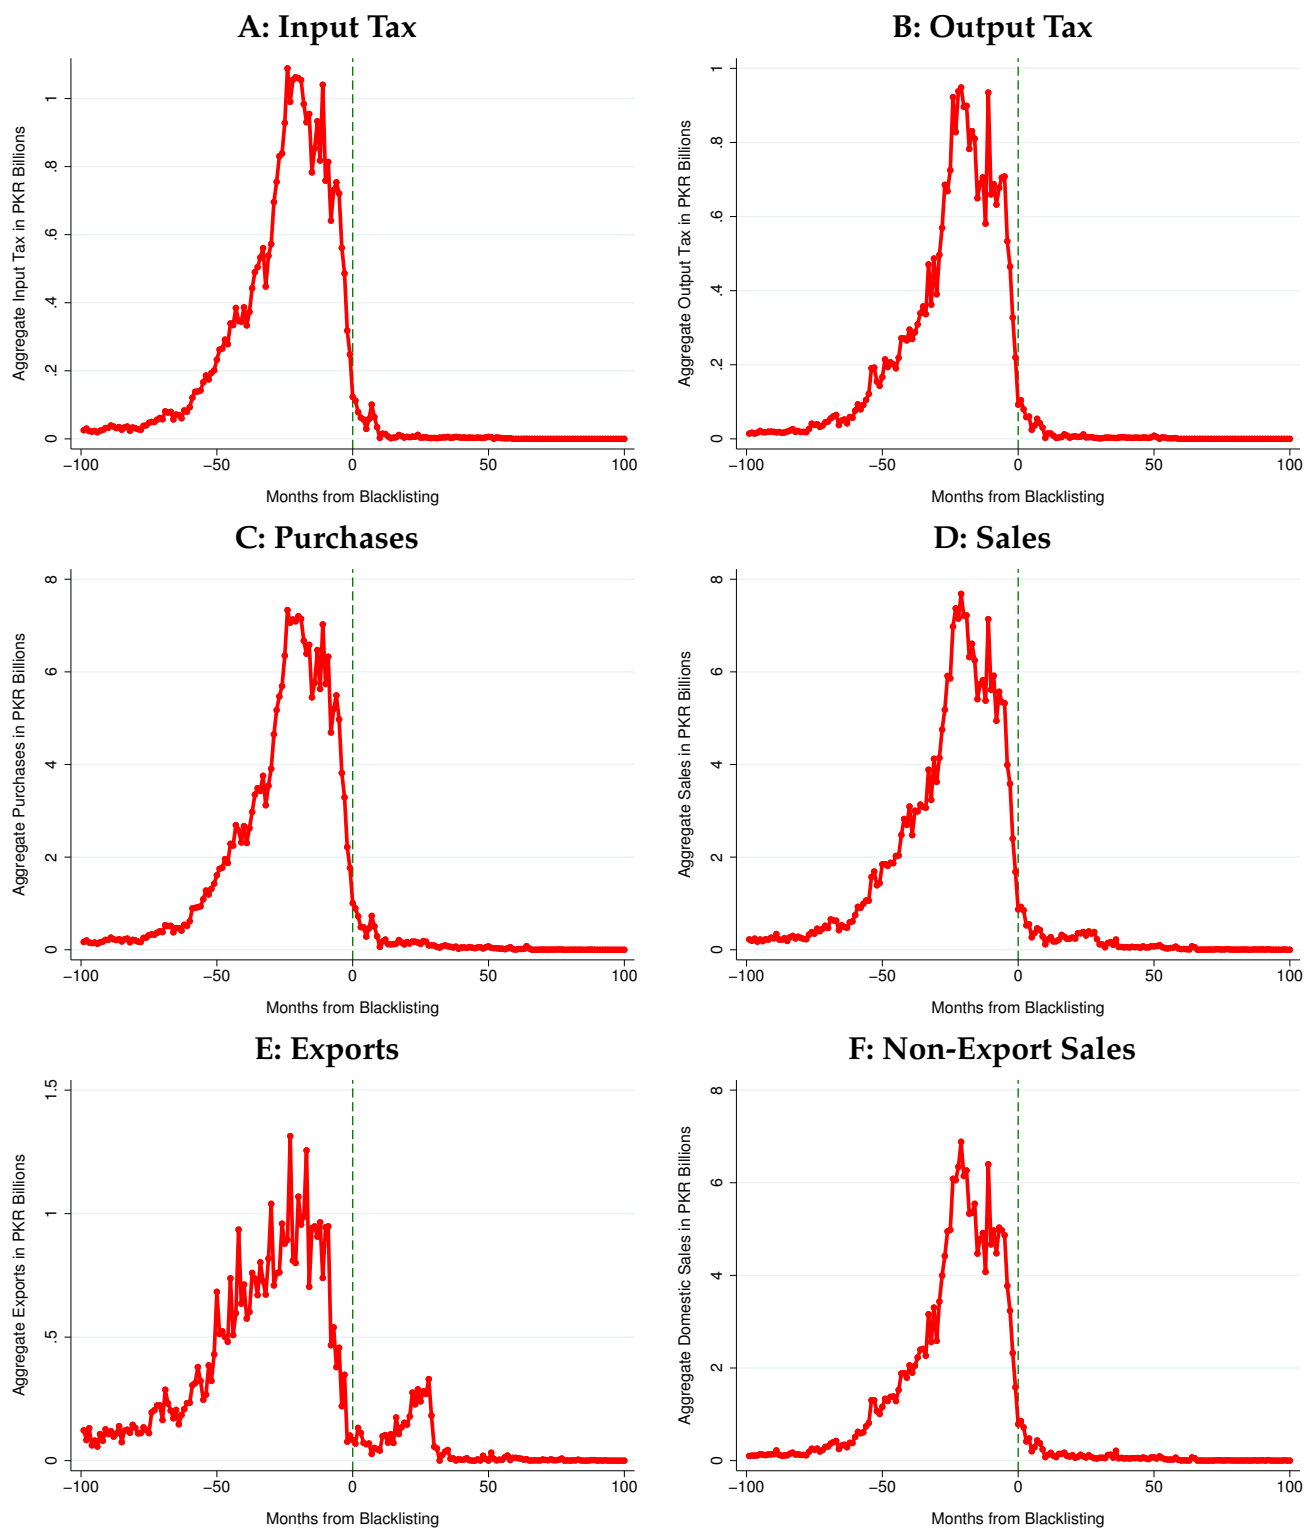
Notes: The figure shows the emergence, growth, and subsequent decline of invoice mills in the Pakistani setting. The sample begins from 1996, when a *broad-based* VAT with coverage extended to almost entire manufacturing and exports stages begun in the country. Each marker in the curve denotes the number of blacklisted firms that file a return in the given month. The year t in the horizontal axis denotes the month July of year t . The dashed, vertical line represents the time from which the zero-rating reform became applicable.

FIGURE A.X: FIRM RESPONSE TO THE TAX CUT – INVOICE MILLS



Notes: The figure compares the evolution of six VAT outcomes from the tax year 1998 to 2010 across the treatment and control groups. The treatment group here consists of blacklisted and suspended firms in the zero-rated industries only. The control group, as earlier, comprises all firms of the non-zero-rated industries. To construct these charts, I regress the log of the outcome variable shown in the title of each panel on the full set of firm and month fixed effects, dropping the dummy for July 1998. I then plot the coefficients on the time dummies of these regressions. The regressions are run separately for the two groups of firms. Year t on the horizontal axis indicates July of the year. The vertical, dashed lines demarcate the time from which the zero-rating reform became applicable.

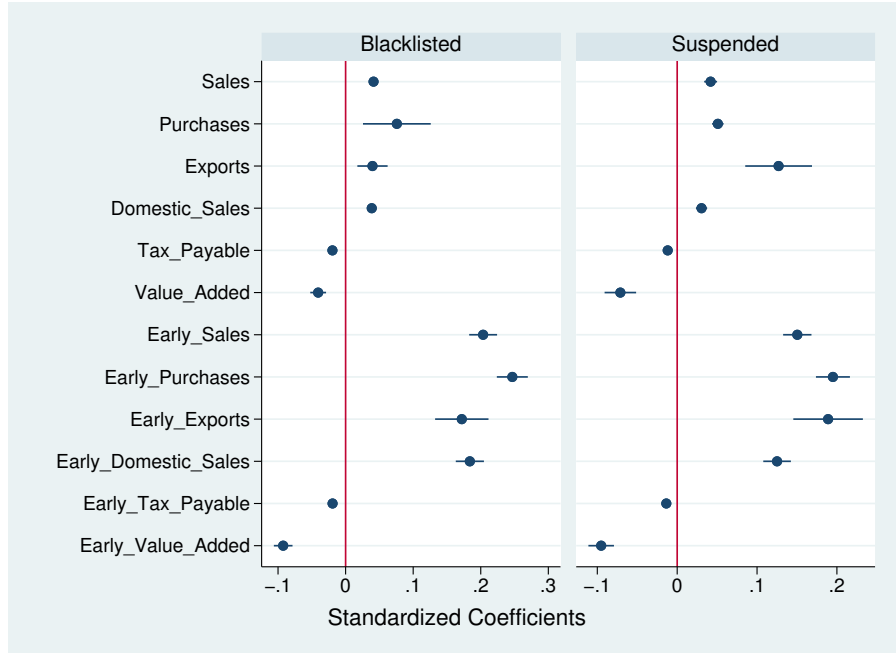
FIGURE A.XI: AGGREGATE VALUES OF VAT OUTCOMES – INVOICE MILLS



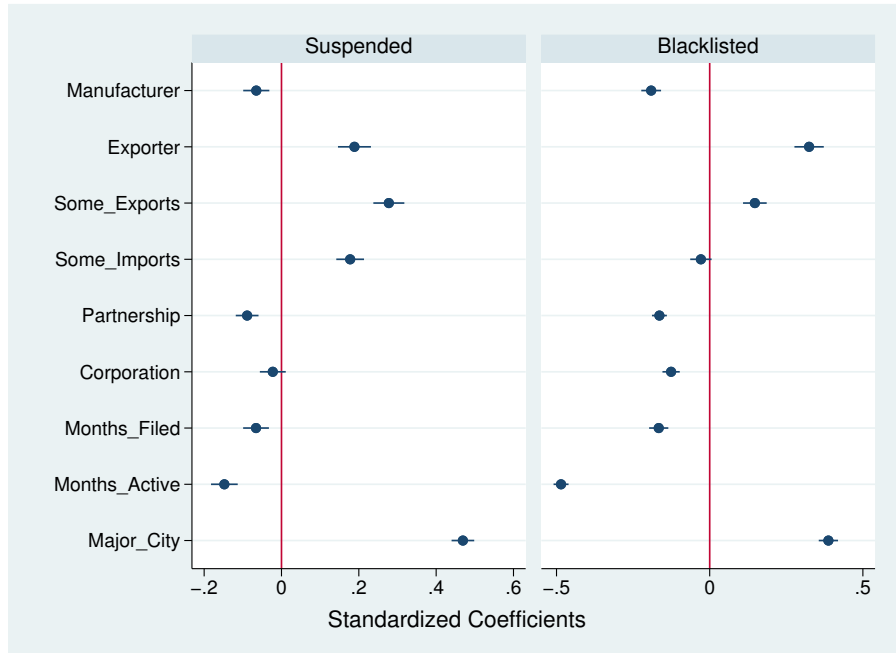
Notes: The figure explores the change in activity carried out by blacklisted firms around the event of blacklisting. The sample contains all blacklisted and suspended firms of the treated industries. Time 0 in the horizontal axis (marked by the vertical, dashed line) denotes the month in which the firm was declared blacklist or its registration was suspended. Each panel of the figure shows the aggregate value of the outcome in 100 months prior to and 100 months after the event of blacklisting. To deal with outliers, I exclude ten observations with the highest values of the given outcome in the entire sample. For example, for constructing Panel A, I sort all firm-month observations on the basis of Output Tax in a descending order and exclude the top-ten observations.

FIGURE A.XII: DISTINGUISHING MARKERS OF INVOICE MILLS

A: VAT Outcomes



B: Firm Characteristics



Notes: The figures explores if invoice mills have any distinguishing markers that could help identify them. I compare the VAT outcomes reported by (Panel A) and and firm characteristics of (Panel B) blacklisted and suspended firms with other firms. The sample here consists of pre-reform years only (1998-2004). To construct these plots, I standardize each variable by subtracting the mean and dividing by the standard deviation of each variable. I then regress the standardized variable on two dummies, indicating if the firm is a blacklisted or a suspended firm. The coefficients on these dummy variables along with the standard errors are plotted with the heading of each row indicating the LHS variable in the regression. Early here denotes the first six months after the firm’s registration For example Early_sales denotes the average sales reported by the firm in its first six returns after registration. The definitions of the compared variables are provided in Appendix A.1.

TABLE A.I: SUMMARY STATISTICS

	2003		2004	
	Treatment	Control	Treatment	Control
	(1)	(2)	(3)	(4)
1. # Observations	172,321	743,281	163,327	709,560
2. Input Tax	0.686 (0.014)	0.431 (0.017)	0.707 (0.013)	0.639 (0.031)
3. Output Tax	0.514 (0.013)	0.602 (0.028)	0.479 (0.012)	0.850 (0.045)
4. Purchases	4.600 (0.084)	2.678 (0.096)	6.158 (0.116)	4.223 (0.175)
5. Sales	6.919 (0.895)	3.782 (0.122)	7.776 (0.116)	5.623 (0.202)
6. Exports	2.614 (0.054)	0.311 (0.010)	4.207 (0.065)	0.503 (0.016)
7. Domestic Sales	4.306 (0.893)	3.471 (0.118)	3.569 (0.082)	5.120 (0.194)
8. Major City	0.486 (0.001)	0.563 (0.001)	0.507 (0.001)	0.576 (0.001)
9. # Years Registered	3.695 (0.009)	3.592 (0.004)	3.238 (0.010)	3.108 (0.004)
10. # Years Active	8.083 (0.009)	8.608 (0.004)	8.401 (0.008)	8.883 (0.004)
11. Some Export	0.472 (0.001)	0.135 (0.000)	0.545 (0.001)	0.152 (0.000)
12. Some Import	0.416 (0.001)	0.369 (0.001)	0.476 (0.001)	0.415 (0.001)

Notes: The table presents summary statistics for the treatment and control groups. Treatment group comprises firms whose supplies were zero-rated by the zero-rating reform from 2005. The control group comprises all other firms. The first row of the table compares the number of firm-month observations for the two groups in the two prereform years. Subsequent rows compare the mean of eleven VAT outcomes and firm characteristics across the two groups. Major City denotes that the firm is registered in Karachi or Lahore, the two major cities of Pakistan. The variable # Years Registered reports the number of years up to 2003 since the firm's registration; # Years Active reports the number of years the firm remained active, filing its VAT return. Standard errors of the mean are in parenthesis.

TABLE A.II: FIRM RESPONSE TO THE TAX CUT – ELASTICITIES

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Complete Panel</u>						
Elasticity	-19.973 (0.262)	-25.125 (0.972)	-4.268 (0.170)	-3.035 (0.168)	-1.078 (0.256)	-0.751 (0.124)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
<u>B: Balanced Panel</u>						
Elasticity	-0.273 (0.756)	-0.204 (2.575)	3.299 (0.792)	-4.402 (2.322)	-1.389 (1.678)	-11.753 (9.456)
Observations	1,999,987	2,423,815	1,979,624	2,583,424	306,931	2,381,448
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the elasticities underlying the responses documented in Table I. I estimate the elasticity version of the difference-in-differences model (8), where I regress the VAT outcome in each column on $\log(1 - \tau_{it})$, including in the full set of firm and period fixed effects. The tax rate τ_{it} is the effective tax rate faced by the firm on each outcome. The exogenous variation in the effective tax rate created by the reform is shown in Figure A.III. Panel B shows results from parallel placebo regressions, where I restrict the sample to pre-reform years only, defining the period beginning from July 2002 as the *after* period. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.III: FIRM RESPONSE TO THE TAX CUT – STANDARD ERRORS CLUSTERED AT THE INDUSTRY LEVEL

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Baseline Specification</u>						
treat × after	-1.961 (0.176)	-1.842 (0.380)	-0.419 (0.036)	-0.223 (0.034)	-0.106 (0.041)	-0.082 (0.046)
Observations	3,728,658	4,179,560	3,983,211	5,058,578	612,993	4,623,906
<u>B: Placebo Specification</u>						
treat × after	-0.004 (0.050)	0.001 (0.053)	0.052 (0.053)	0.019 (0.042)	-0.022 (0.054)	-0.014 (0.044)
Observations	1,999,985	2,423,814	1,979,622	2,583,423	306,931	2,381,447
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores how firms respond to the reduction of the rate applicable to their supply chain to zero. The coefficient $\text{treat} \times \text{after}$ shows $\hat{\gamma}$ from model (8), where the dummy variable treat_i denotes that firm i belongs to a zero-rated industry and the dummy variable after_t that month t falls in the tax year 2005 and later. Panel B shows results from parallel placebo regressions, where I restrict the sample to pre-reform years only, defining the period beginning from July 2002 as the *after* period. Standard errors are in parenthesis, which have been clustered at the industry level.

TABLE A.IV: FIRM RESPONSE TO THE TAX CUT – BALANCED PANEL

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Baseline Specification</u>						
treat × after	-2.337 (0.043)	-2.536 (0.139)	-0.484 (0.031)	-0.405 (0.024)	-0.158 (0.037)	-0.101 (0.029)
Observations	948,385	877,354	981,954	1,126,539	264,719	960,697
<u>B: Placebo Specification</u>						
treat × after	-0.076 (0.021)	-0.009 (0.024)	-0.008 (0.021)	0.008 (0.019)	-0.031 (0.034)	-0.051 (0.024)
Observations	560,180	562,860	564,044	646,637	142,345	564,505
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores how firms respond to the reduction of the rate applicable to their supply chain to zero. The sample here is restricted to a balanced panel of firms that file at least once in every quarter included in the sample. The coefficient $\hat{\gamma}$ from model (8), where the dummy variable $treat_i$ denotes that firm i belongs to a zero-rated industry and the dummy variable $after_t$ denotes that month t falls in the tax year 2005 and later. Panel B shows results from parallel placebo regressions, where I restrict the sample to pre-reform years only, defining the period beginning from July 2002 as the *after* period. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.V: SPILLOVER EFFECTS ON THE NON-TREATED INDUSTRIES

Industries Within:	Next Two Digits			Next Five Digits			Next Ten Digits		
	Purchases	Sales	Exports	Purchases	Sales	Exports	Purchases	Sales	Exports
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<u>A: Complete Panel</u>									
treat × after	-0.234 (0.121)	0.002 (0.083)	0.067 (0.172)	-0.059 (0.082)	0.038 (0.062)	0.012 (0.128)	-0.066 (0.022)	0.012 (0.019)	0.044 (0.077)
Observations	3,123,769	4,075,853	312,320	3,123,769	4,075,853	312,320	3,123,769	4,075,853	312,320
<u>B: Balanced Panel</u>									
treat × after	-0.152 (0.168)	-0.018 (0.130)	0.430 (0.319)	0.048 (0.121)	0.041 (0.108)	0.187 (0.228)	-0.100 (0.037)	0.014 (0.031)	0.072 (0.119)
Observations	764,271	892,064	133,421	764,271	892,064	133,421	764,271	892,064	133,421
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores if the reform produces any spillovers on the nontreated industries. I drop all industries zero-rated by the reform from the sample. I then estimate the difference-in-differences model (8). The dummy variable $treat_i$ here denotes that firm i belongs to an industry indicated in the title of each column. For example, the first three columns regard the two 2-digit industries immediately succeeding the zero-rated ones as treated. Panel B restricts the sample to a balance panel, including only the firms who file at least once in every quarter included in the sample. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.VI: SPILLOVER EFFECTS ON THE NON-TREATED INDUSTRIES – ALTERNATIVE SPECIFICATION

Industries Within:	Next Two Digits			Next Five Digits			Next Ten Digits		
	Purchases	Sales	Exports	Purchases	Sales	Exports	Purchases	Sales	Exports
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<u>A: Complete Panel</u>									
treat × after	-0.420 (0.017)	-0.222 (0.012)	-0.105 (0.025)	-0.420 (0.017)	-0.222 (0.012)	-0.105 (0.025)	-0.426 (0.017)	-0.222 (0.012)	-0.103 (0.025)
adjoining × after	-0.242 (0.119)	0.015 (0.082)	0.068 (0.165)	-0.065 (0.082)	0.044 (0.062)	0.015 (0.125)	-0.067 (0.022)	0.012 (0.019)	0.046 (0.075)
Observations	3,983,213	5,058,579	612,993	3,983,213	5,058,579	612,993	3,983,213	5,058,579	612,993
<u>B: Balanced Panel</u>									
treat × after	-0.485 (0.031)	-0.405 (0.024)	-0.154 (0.037)	-0.483 (0.031)	-0.405 (0.024)	-0.155 (0.037)	-0.496 (0.032)	-0.404 (0.025)	-0.155 (0.037)
adjoining × after	-0.154 (0.168)	-0.020 (0.131)	0.428 (0.318)	0.047 (0.121)	0.039 (0.108)	0.186 (0.226)	-0.100 (0.037)	0.012 (0.031)	0.070 (0.119)
Observations	981,954	1,126,539	264,719	981,954	1,126,539	264,719	981,954	1,126,539	264,719
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores if the reform produces any spillovers on the nontreated industries. I estimate an augmented version of the difference-in-differences model (8), including the interaction term $adjoining \times after$ into the model. The dummy variable $adjoining_i$ here denotes that firm i belongs to an industry indicated in the title of each column. For example, the first three columns regard the two 2-digit industries immediately succeeding the zero-rated ones as *adjoining*. Panel B restricts the sample to a balance panel, including only the firms who file at least once in every quarter included in the sample. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.VII: FIRM RESPONSE TO THE TAX CUT – WITHOUT BLACKLISTED FIRMS

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Complete Panel</u>						
treat × after	-1.953 (0.026)	-1.871 (0.074)	-0.406 (0.017)	-0.221 (0.013)	-0.098 (0.026)	-0.079 (0.014)
Observations	3,562,585	4,020,573	3,808,066	4,865,461	584,747	4,451,965
<u>B: Balanced Panel</u>						
treat × after	-2.316 (0.044)	-2.530 (0.141)	-0.469 (0.032)	-0.402 (0.025)	-0.152 (0.038)	-0.101 (0.029)
Observations	923,512	857,195	955,603	1,098,880	256,765	938,355
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores how firms respond to the reduction of the rate applicable to their supply chain to zero. The sample here excludes blacklisted and suspended firms. The coefficient $\text{treat} \times \text{after}$ shows $\hat{\gamma}$ from model (8), where the dummy variable treat_i denotes that firm i belongs to a zero-rated industry and the dummy variable after_t that month t falls in the tax year 2005 and later. Panel B shows results from parallel placebo regressions, where I restrict the sample to pre-reform years only, defining the period beginning from July 2002 as the *after* period. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.VIII: FIRM RESPONSE TO THE TAX CUT – FROM DATE OF ANNOUNCEMENT

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Complete Panel</u>						
treat × after	-1.944 (0.025)	-1.793 (0.064)	-0.422 (0.017)	-0.227 (0.012)	-0.114 (0.026)	-0.113 (0.013)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
<u>B: Balanced Panel</u>						
treat × after	-2.325 (0.043)	-2.476 (0.131)	-0.485 (0.031)	-0.404 (0.024)	-0.156 (0.037)	-0.121 (0.029)
Observations	948,385	877,354	981,954	1,126,539	264,719	960,697
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores how firms respond to the reduction of the rate applicable to their supply chain to zero. The only difference from the baseline specification (results reported in Table I) is that the *after* period here begins from the date of announcement of the reform i.e. from June 2005. The coefficient $\text{treat} \times \text{after}$ shows $\hat{\gamma}$ from model (8), where the dummy variable treat_i denotes that firm i belongs to a zero-rated industry and the dummy variable after_t that month t is June 2005 or a later period. Panel B shows results from parallel regressions, where I restrict the sample to a balanced panel of firms that report in every quarter included in the sample. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.IX: EXTENSIVE MARGIN RESPONSE TO THE TAX CUT

	$1(Filer_{it} = 1)$	
	Baseline	Placebo
	(1)	(2)
treat \times after	0.032 (0.001)	0.008 (0.001)
Observations	11,405,728	5,433,725
Mean of the Dependent Variable	0.927	0.932
Firm Fixed Effect	Yes	Yes
Month Fixed Effect	Yes	Yes

Notes: The table explores if treated firms began filing their tax returns less regularly after the zero-rating reform. I estimate the difference-in-differences model (8) using an indicator that the firm files its tax return for period t as the outcome variable. The binary variable $Filer_{it}$ is coded as 1 if firm i files its return for the month t . It is coded as zero if the return is not filed and the month t falls in any period after the entry and before the exit of the firm. Column (2) reports results from a parallel placebo model where I use data from pre-reform years only (1998–2004), pretending that the reform occurred in July 2002.

TABLE A.X: FIRM RESPONSE TO THE TAX CUT – CORPORATIONS VS. OTHERS

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
treat × after	-1.745 (0.030)	-1.135 (0.059)	-0.407 (0.018)	-0.148 (0.013)	-0.099 (0.028)	-0.056 (0.014)
treat × after × corporation	-0.599 (0.050)	-1.973 (0.154)	-0.040 (0.037)	-0.271 (0.027)	-0.017 (0.037)	-0.099 (0.033)
Baseline Coefficient	-1.961 (0.026)	-1.842 (0.071)	-0.419 (0.017)	-0.223 (0.012)	-0.106 (0.025)	-0.082 (0.014)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table rules out lazy reporting as an alternative explanation of the responses documented in Table I. I divide the treatment sample into corporate and non-corporate firms. I then estimate a triple-differences version of model (8), including all double interactions of the binary variables $treat_i$, $corporation_i$, and $after_t$ and their triple-interaction. Baseline coefficient reports the $treat \times after$ coefficient I obtain from estimating the baseline model (8). Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.XI: FIRM RESPONSE TO THE TAX CUT – FIRM SIZE

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
<i>treat</i> × <i>after</i>	-1.043 (0.108)	-0.369 (0.106)	-0.315 (0.080)	0.153 (0.054)	0.116 (0.083)	0.231 (0.072)
× second quartile	-0.078 (0.125)	-0.336 (0.146)	0.003 (0.093)	-0.116 (0.062)	-0.174 (0.098)	-0.061 (0.080)
× third quartile	-0.182 (0.118)	-0.686 (0.140)	0.001 (0.085)	-0.186 (0.057)	-0.129 (0.087)	-0.240 (0.076)
× fourth quartile	-1.145 (0.111)	-1.884 (0.143)	-0.138 (0.081)	-0.459 (0.056)	-0.247 (0.084)	-0.360 (0.074)
Baseline Coefficient	-1.961 (0.026)	-1.842 (0.071)	-0.419 (0.017)	-0.223 (0.012)	-0.106 (0.025)	-0.082 (0.014)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores heterogeneity in response by baseline firm size. I divide firms into four quartiles based on their total turnover reported in the baseline period. I then estimate a triple-differences version of model (8), including the interactions of the DD term with the size dummies. I omit the interaction term for the first quartile dummy. Baseline coefficient reports the *treat* × *after* coefficient I obtain from estimating the baseline model (8). Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.XII: FIRM RESPONSE TO THE TAX CUT – TEXTILE VS. OTHERS

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
treat \times after	-1.978 (0.026)	-1.925 (0.076)	-0.394 (0.017)	-0.221 (0.013)	-0.065 (0.027)	-0.082 (0.014)
treat \times after \times non-textile	0.260 (0.102)	0.958 (0.196)	-0.403 (0.078)	-0.023 (0.040)	-0.224 (0.042)	-0.001 (0.081)
Baseline Coefficient	-1.961 (0.026)	-1.842 (0.071)	-0.419 (0.017)	-0.223 (0.012)	-0.106 (0.025)	-0.082 (0.014)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table decomposes the average response to the zero-rating reform estimated in Table I into its constituent textile and non-textile components. I estimate a triple-differences version of model (8), including all double interactions of the binary variables $treat_i$, $Non-textile_i$, and $after_t$ and their triple-interaction. The dummy variable $treat_i$ denotes that firm i belongs to a zero-rated industry; the dummy variable $Non-textile_i$ denotes that the treated firm i does not belong to the textile industry; and the dummy variable $after_t$ denotes that month t falls in the tax year 2005 and later. Baseline coefficient reports the $treat \times after$ or $\log(1 - \tau)$ coefficient I obtain from estimating the baseline model (8). Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.XIII: FIRM RESPONSE TO THE TAX CUT – BLACKLISTED VS. OTHERS

	Input Tax	Output Tax	Purchases	Sales	Exports	Non-Export Sales
	(1)	(2)	(3)	(4)	(5)	(6)
treat × after	-1.948 (0.026)	-1.870 (0.074)	-0.401 (0.017)	-0.218 (0.013)	-0.095 (0.026)	-0.077 (0.014)
treat × after × blacklisted	0.208 (0.210)	0.697 (0.292)	-0.475 (0.125)	-0.306 (0.097)	-0.312 (0.229)	-0.445 (0.103)
treat × after × suspended	-0.435 (0.142)	0.694 (0.334)	-0.374 (0.090)	-0.069 (0.069)	-0.218 (0.066)	-0.043 (0.092)
Baseline Coefficient	-1.961 (0.026)	-1.842 (0.071)	-0.419 (0.017)	-0.223 (0.012)	-0.106 (0.025)	-0.082 (0.014)
Observations	3,728,660	4,179,561	3,983,213	5,058,579	612,993	4,623,907
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table compares the responses of blacklisted and suspended firms with other treated firms. I estimate a triple-difference version of model (8), partitioning the treatment dummy into three dummies $treat_i$, $blacklisted_i$, and $suspended_i$. The dummy variables $blacklisted_i$ and $suspended_i$ denote that a treated firm i is blacklisted or its registration has been suspended. Baseline coefficient reports the $treat \times after$ coefficient I obtain from estimating the model without the triple-interaction terms. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.XIV: SUMMARY STATISTICS – BLACKLISTED FIRMS

	Blacklisted (1)	Suspended (2)	Others (3)
<u>A: VAT Outcomes (PKR Millions)</u>			
1. Purchases	7.043	5.781	1.875
2. Sales	4.609	5.465	2.183
3. Exports	0.649	1.187	0.388
4. Domestic Sales	3.960	4.277	1.795
5. Tax Payable	0.010	0.038	0.082
6. First Year Sales	6.105	5.636	1.484
7. First Year Purchases	6.246	6.019	1.071
8. First Year Exports	0.899	1.069	0.309
9. First Year Domestic Sales	5.206	4.566	1.175
10. First Year Tax Payable	0.013	0.024	0.053
11. Sales Minus Purchases	-3.422	0.633	0.861
12. Output Tax Minus Input Tax	-0.121	-0.152	0.105
<u>B: Firm Characteristics</u>			
13. Manufacturer	0.267	0.286	0.373
14. Wholesaler	0.338	0.339	0.253
15. Exporter	0.175	0.146	0.080
16. Some Export	0.236	0.262	0.160
17. Some Import	0.303	0.381	0.307
18. Company	0.062	0.092	0.092
19. Partnership	0.036	0.054	0.069
20. # Months Filed	38.329	40.580	41.450
21. # Months Active	15.572	23.662	27.763
22. Major City	0.718	0.752	0.503
23. Initial Capital	0.713	1.111	7.749

Notes: The table compares VAT outcomes and firm characteristics of blacklisted and suspended firms with other firms. Each row reports the mean value of the variable for the corresponding group of firms. The definitions of the compared variables are provided in Appendix A.1.