

Do Value-Added Taxes Distort Input Decisions? Evidence from Pakistan

Faraz Hayat*

February 6, 2024

Abstract

We examine how firms' input choices respond to the value-added tax (VAT) on purchases, using detailed VAT return data from Pakistan and a quasi-experimental variation in electricity VAT rates. Contrary to traditional theories that intermediate VATs don't affect production decisions, our study reveals a significant elasticity of electricity demand to VAT rates (-0.81). These findings aren't attributed to tax evasion shifts but to frictions in the VAT refund system. Specifically, many firms, not collecting VAT on output, depend on government reimbursements for their VAT credits, explaining our results.

*International Growth Center, LSE ; Department of Economics, University of Chicago. Email: shayat1@uchicago.edu. We are very grateful to Michael Greenstone, Mikhail Golosov, Jonathan Dingel, Patrick Schwarz, Neil Cholli, and Antonis Tsiflis, for their helpful comments.

The value-added tax (VAT) is the world's most common consumption tax, used in over 160 countries and accounting for over 20 percent of the world's tax revenue (Keen and Lockwood, 2007). This widespread adoption is motivated by both practical and efficiency considerations. From a practical standpoint, the collection of VATs at each stage of production helps overcome issues surrounding enforcement and evasion, as buyers and sellers separately report taxes paid or received and have competing incentives over whether to over- or under-report transaction values. Additionally, since firms are eligible to recover whatever VAT they pay on inputs, VATs are also a theoretically efficient tax system in the sense that they do not distort production decisions or the prices that producers face when buying from one another (Ebrill et al., 2001).

Real-world applications of VATs, however, rarely perfectly resemble textbook models, raising concerns over their efficiency in practice. Common departures from textbook VAT models include size thresholds for registration, incomplete VAT chains from informality, differing VAT rates by products, and delays in processing and disbursing VAT refunds. A growing empirical literature has begun to study these departures and how they affect the productive and revenue efficiency of VATs in practice. Gadenne et al. (2019), for instance, finds significant segmentation in trade between VAT and non-VAT registered firms around VAT registration thresholds in India, and Gerard et al. (2018) finds similar results among firm-to-firm trade in Brazil. In a similar vein, using data from Europe, Liu et al. (2021) and Harju et al. (2019) find that VAT registration thresholds affect firm growth and inter-firm trade. A key takeaway from these studies is that since VATs distort firm behavior, the revenue efficiency of VATs must be weighed against the loss of production efficiency.

In this paper, we add to this growing literature by testing a key principle of the theoretical efficiency of VATs: that VAT rates on intermediate transactions do not affect firms' input decisions. As discussed above, collecting VATs at intermediate stages of production is critical to the revenue efficiency of the VAT because it creates a paper trail of transactions and a system of self-enforcement among VAT registered firms. In principle, this tax is non-distortionary because firms are eligible to receive whatever tax they pay on inputs back in the form of tax credits or refunds. If there

are frictions in the VAT system, however, and firms aren't able to fully recover the VAT they pay on inputs, then the revenue efficiency of VAT made possible by taxing intermediates comes at the cost of reduced productive efficiency. This is because firms effectively pay an additional tax on their inputs in the form of unrecovered tax credits. We directly test whether input decisions are sensitive to VAT rates, quantify the magnitude of any distortion in production decisions we observe, and investigate underlying mechanisms.

The empirical application of this paper is based on a quasi-experiment that leverages policy variation in the VAT rate on electricity purchases in Pakistan combined with detailed administrative microdata on Pakistani firms between 2017-2019. Over this period, Pakistan's tax authority had set the intermediate VAT rate on the majority of inputs in five export oriented sectors to 0% (instead of the standard VAT rate of 17%), but required a separate application to regional tax offices to additionally zero-rate electricity purchases. In addition to having the power to approve or reject applications to zero-rate electricity purchases, regional tax offices also had the power to revoke, either temporarily or permanently, the zero-rating status of electricity for any firm at their own discretion. We leverage the fact that although an application or revocation of the zero-rating status is likely endogenous in firm behavior, the timing of decisions by regional tax offices is likely uncorrelated with contemporaneous factors that might affect a firm's input decisions. We therefore estimate a difference-in-differences regression at the firm-month-year level to estimate the effect of moving to and from a zero-rating status on firm demand for electricity.

Our variation in input tax rates on electricity has several key advantages relative to other analyses of VAT rates and their relation to firm decisions. First, in contrast to system-wide VAT changes, the variation we use is product- and firm-specific. This allows us to separately identify the effect of the input VAT rate on firm decisions from other factors such as upstream or downstream effects of more global tax changes, changes in final good demand, or contemporaneous shocks. Second, since the federal government of Pakistan has near total control of the distribution of electricity and self-generate electricity bills, we can rule out misreporting of electricity purchases as a driver of our results. Third, we are also able to rule out the possi-

bility that firms misreport their output in response to VAT rate changes, given that the majority of firms in our sample have their output zero-rated and therefore face no tax benefit of misreporting sales. Finally, we can rule out input price changes in response to changing VAT rates given that the timing of VAT rate changes are firm-specific and because electricity prices are fully determined by a fixed schedule issued by the government. The estimated effect we recover, therefore, can be attributed to the direct effect of the VAT rate itself on input demand.

In contrast to predictions from the standard VAT model, we find that electricity demand is indeed sensitive to the VAT rate. Moving in or out of the 17% standard VAT rate causes a 13% change in electricity demand. Our results are robust to a wide set of econometric specifications, and we find effects of similar magnitude for both moving in and out of electricity zero rating. Our events study graphs also do not exhibit any evidence of pre-trends prior to status changes, suggesting that firms were unable to precisely anticipate when their electricity zero-rating was going to be established or revoked. Importantly, we also show that firms' non-electricity VAT rates, such as the VAT they pay on other inputs or the VAT they collect on sales, do not change around the time they gained or lost their electricity zero-rating.

We argue that electricity demand responds to changes in the VAT rate because of delayed or never-paid VAT refunds. Refund non-payment is a well-known issue in Pakistan and is often cited by the local trade press (see, e.g., Rana, 2016, 2021), although the problem appears to extend to many less-developed countries. For instance, an IMF report from 2005 describes the refund system as the "Achilles heel" of VAT systems and shows that VAT refunds as a percent of total VAT collections are starkly lower in developing countries relative to other developed nations. Consistent with this explanation, we find a sharp and sustained increase in the amount of taxes owed to firms as a result of the withdrawal of an electricity zero-rating. We are also able to rule out other common potential mechanisms, such as changes in input or output evasion, broader general equilibrium effects such as price changes, changes in the final demand of firms' output, or changing dynamics between registered and unregistered firms.

We then benchmark the magnitude of this effect by comparing our estimated VAT rate elasticity to a standard price elasticity of demand for electricity that we

measure using the roll-out of reduced electricity rates among a similar set of firms in 2019. We find that a reduction in the price of electricity of 22% increased electricity consumption by approximately 13% (95% CI: [0.05, 0.21]), implying a price elasticity of 0.59, although an alternative specification including more detailed fixed effects is measured with imprecision (point estimate: 0.08; 95% CI: [-0.04, 0.20]). We, however, rely on these estimates because the existing literature on the price elasticity of electricity demand among firms in Pakistan appears exclusively based on correlations and therefore may suffer from standard endogeneity concerns (see, e.g., Chaudhry, 2010; Khan and Qayyum, 2009; Ishaque, 2018; Alter and Syed, 2011; Khan and Abbas, 2016). Comparing this estimate to our estimated VAT elasticity of 0.81 (s.e. 0.32) suggests that firms treat the VAT on electricity roughly as a standard price increase. The fact that firms appear to be fully discounting the VAT they pay on electricity suggests the existence of significant frictions in the ability of firms to recover the VAT they pay on inputs.

Finally, we explore the implications of the frictions we identify on a broader set of firm outcomes. We do not find evidence that the VAT rate change on electricity impacts firms' sales, non-electricity purchase, exports or imports, although our estimates are measured with significant noise. We view these null results as a limitation of our study; while the limited scope of the policy variation to a single input to production allows us to rule out several channels and identify key elasticities, studying only a single input restricts us from precisely and systematically analyzing the effect of the refund friction more broadly on firm outcomes and productivity that are critical to welfare analyses. We therefore leave this as an area of exploration for future research.

Our research intersects several strands of literature, focusing on VAT system frictions and the trade-offs of different tax policies. Existing studies often highlight how VAT systems induce distortions in firm behavior, particularly around VAT registration thresholds (e.g., Liu et al., 2021; Harju, 2019; Onji, 2009) and the impact on inter-firm trade due to VAT registration (e.g., Gerard et al., 2018; Gadenne et al., 2019; de Paula and Scheinkman, 2010). We extend this literature by delving into a relatively unexplored area of how VAT can create productive inefficiencies. Our study suggests these inefficiencies stem from refund delays, aligning with anecdotal

reports from Pakistani policymakers and trade media. This complements Chandra and Long (2013), who used an instrumental variable strategy with six years of annual firm-level data from China’s National Bureau of Statistics to demonstrate the impact of increased VAT rebates for exporters on China’s export volumes. In contrast to their focus on exporters and export volumes, our study utilizes monthly VAT filings from both exporters and non-exporters in Pakistan, a less export-driven economy, to examine the influence of VAT on firms’ input decisions. Our approach allows us to explore how VAT input taxes influence production decisions, contributing to the dialogue on balancing productive and revenue efficiency in tax systems in developing nations (refer to Best et al., 2015; Cage and Gadenne, 2018; Gordon and Li, 2009). Additionally, by utilizing administrative data from Pakistan, a country with one of the lowest GDP per capita rankings globally (World Bank, 2021), our study enhances the understanding of VAT systems in low-income countries. This is particularly relevant as Gerard and Naritomi (2018) note the recent VAT research primarily uses data from middle- to high-income developing countries.

The rest of the paper is organized as follows. Section 2 provides a background of VAT in Pakistan and our policy variation. Section 3 contains our conceptual framework. Section 4 discusses the administrative data we use for our analysis. Section 5 introduces our main empirical strategy and alternative specifications. Section 6 presents our results and discusses potential mechanisms. Section 7 presents our standard price elasticity of demand results that we use to benchmark the effect size of the impact of VAT rates on electricity demand. Section 8 concludes.

I Background

A. VAT in Pakistan

Since its first implementation as a Goods and Services Tax (GST) in the 1990s, Pakistan’s VAT system has been characterized by a mix of administrative and political challenges. Early on, the VAT struggled with issues of incomplete coverage and frequent changes in the tax base by the country’s tax authority– the Federal Board of Revenue (FBR)– through the use of Sales Tax Notification and Statutory Regu-

latory Orders (SROs), which had broad power to change tax exemptions across the economy and required no approval from Parliament. The government also relied on ad-hoc methods and fixed prices to determine VAT liabilities and VAT refunds. Recognizing these problems, the government with the support of the World Bank made a number of attempts to improve the administration and coverage of VAT in the late 1990s and early 2000s, although the expected gains in domestic revenue from these changes didn't fully occur (Ahmad, 2010). Additional changes meant to improve VAT coverage and administration were introduced in the 2010s, including new penalties for non-registration and an automated system deployed in 2014 that reviews VAT invoices for fraud (Shah, 2021). However the country still grapples with low tax revenue, with a tax-to-GDP ratio of 13% in 2021 versus 11.3% in 1985, placing the country in the bottom 20% of all tax-to-GDP ratios for nation states around the world (Ahmad and Stern, 1991; Heritage Foundation, 2021).

Today, VAT in Pakistan follows the standard input credit system and accounts for about 30% of the country's total tax revenue (Govt. of Pakistan: Finance Division, 2020). Firms are required to register for VAT if their annual turnover exceeds an industry-dependent threshold (FBR, 1990). Once in the system, firms collect tax credits on their purchases and tax debits on their sales, logging each transaction in the FBR's electronic filing system. If tax debits exceed tax credits at the end of the reporting period (typically each month), the firm is required to pay the FBR the difference. If credits exceed debits, firms have the choice to carry the credits forward to the next tax period, or apply for a tax refund.

The standard VAT rate for goods on both domestic purchases and sales is 17% while exports are entirely zero-rated to maintain competitiveness in the international trade market (PWC, 2021; Deloitte, 2020). As we discuss more below, many VAT exemptions still exist throughout the economy, with the FBR continuing to have broad discretion over defining the tax base and VAT exemptions through the use of SROs. Since 2015 more than 200 SROs and SRO amendments were issued by the FBR pertaining to VAT, the majority of which adjusted exemptions or zero-ratings of VAT for particular firms or products (FBR, 2021).

B. Pakistan's Zero-Rating Policy

In July 2005, Pakistan began a policy of zero-rating domestic trade in a large set of goods in its major exporting sectors via a SRO (Waseem, 2023).¹ Prior to the policy change, exporters in these sectors would pay VAT on their inputs but collect no VAT on their exports, often generating excess amounts of input credits. In turn, exporters with excess credits would apply to have their input credits refunded, but the refund process was marred by long delays and fraudulent invoice receipts that limited the government's ability and willingness to grant refunds.² The zero-rating policy eliminated the need to refund exporters by setting the VAT rate for both final and intermediate goods in exporting sectors to 0%, meaning that many domestic firms primarily engaged in the purchase or sale of zero-rated goods also effectively became exempt from the VAT.

Electricity purchases used in the production of zero-rated goods were also intended to be zero-rated, but their zero-rating was made conditional on the recommendation of regional tax offices. To obtain an electricity zero-rating, firms would apply to their regional tax office with evidence that their facilities were primarily engaged in the production of zero-rated goods. Once an application was approved, the regional tax offices would make a recommendation to the FBR, which would subsequently release an SRO to formally begin the zero-rating. Firms could also have their zero-rating status revoked based on the recommendation of regional tax offices. Our conversations with local tax experts suggest that the application process is arduous and that firms are mostly left unaware about if and when their application would be approved.

The rationale for the zero-rating policy was to support the exporting sector by freeing up capital otherwise locked up in the refund system, and to avoid paying refunds to exporters who had abused the refund system by submitting fraudulent input invoices to reduce their tax liability. This came at the cost of breaking the VAT chain for large parts of the economy and creating new potential for evasion

¹These sectors are: carpets and rugs, leather products, sports goods, surgical instruments, and textiles.

²For more details on invoice fraud and its interaction with introduction of zero-rating in 2000s, see Waseem, 2023.

by falsely claiming that sales or purchases were in zero-rated goods. Despite these drawbacks, the zero-rating policy was largely unchanged for the next 14 years, and many of the refund requests made prior to 2005 were never fully processed and released back to firms.

The most recent version of the zero-rating policy that applies during our sample period comes in the form of SRO 1125, which was first established in 2011 and then incrementally amended through 2016. SRO 1125 largely kept the original zero-rating policy unchanged, but limited the tax exemption to firms that were VAT registered as manufacturers, importers, exporters, and wholesalers. Supplies to and from retailers, final consumers, and to VAT unregistered firms were subject to the standard rate of 17%. The final form of the zero-rating policy therefore acted as a tax exemption on the trade of intermediate inputs purchased by firms in the export-oriented sectors, but not on the domestic or retail sale of final goods.

In May of 2019, Pakistan agreed to abruptly end its zero-rating policy as part of a \$6 billion loan deal with the IMF, raising the VAT rate on all products covered by SRO 1125 from 0% to the standard rate of 17% (Ahmed, 2019). The changes took effect the following month in June 2019. The policy change was met with harsh backlash from the exporting community, with some industry representatives claiming that the ensuing liquidity crunch from refunds delays would cause exports to decline by 25-30% (The Nation, 2020).

As we discuss more in our empirical framework section below, we use regional tax offices' approvals and revocations of electricity zero-ratings to identify the effect of VAT rates on input decisions. Since the zero-rating policy underwent changes till 2016 and was universally abolished in May 2019, our analysis focuses on the time between January 2017 and May 2019.

II Conceptual Framework

We build a static model of a firm that uses and manufactures intermediate goods in a VAT regime to show that frictions in the VAT credit system can distort input decisions. The model shows that an accurate measurement of the VAT induced financial friction requires knowledge of how input purchases respond to a change

in both the VAT rate and price.

We model a VAT paying firm that collects taxes on its output, pays taxes on its input, and pays the difference to the government. A friction in the system arises when the firm pays more tax on its inputs than it collects on its output and the government does not reimburse the firm for the excess tax it paid. Assume that the firm uses a single input and takes input prices as given. The firm's profit maximization problem is:

$$(1) \quad \begin{aligned} \max_{e,y} \quad & (1 + \tau_0) \cdot p(y) \cdot y - (1 + \tau_1) \cdot q \cdot e - \kappa \cdot [\tau_0 \cdot p(y) \cdot y - q \cdot \tau_1 \cdot e] \\ \text{s.t.} \quad & y = f(e) \end{aligned}$$

where p is the output price, $f(e)$ is the production function, q is the input price, e is input quantity, τ_0 is the VAT rate on output, τ_1 is the VAT rate on input, and κ captures frictions in the VAT regime, which we describe in more detail below.

The square-bracketed term in equation(1) represents the VAT liability. The firm owes the government the difference in the tax it collects on its output and the tax that it pays on its inputs. If the term is positive, then the amount is the sales tax liability of the firm. If the term is negative, then the amount is the tax credit liability of the government.

Under conventional models, the government fully pays all tax credit liabilities and VAT is non-distortionary. To see this in the model, consider the case when $\kappa = 1$. The τ_0 and τ_1 terms drop out from equation(1). Each firm collects taxes on its output and pays it to the government, while at the same time writing off or recovering any tax paid on its inputs. Consequently, the VAT does not burden the firm and the VAT does not distort intermediate production decisions.³

In practice, however, firms cannot easily write-off taxes paid on inputs, which is captured in the model as the case where $\kappa < 1$.⁴ This leads to VAT being dis-

³The only entity that cannot write-off taxes on purchases is the final consumer, who bears the full financial burden of VAT.

⁴In practice, the friction may enter non-linearly in the sense that the government readily collects money it is owed but creates delays when it owes money to firms. Since our empirical application focuses on firms whose output is mostly zero-rated (i.e. $\tau_0 \approx 0$), modeling the friction in this way would not affect the interpretation of our results.

tortionary. Solving the maximization problem and applying the implicit function theorem gives the following equation for the effect of a change in the input VAT rate on input purchases:

$$(2) \quad \frac{\partial e}{\partial \tau_1} = \frac{q \cdot (1 - \kappa)}{(1 + \tau_0 - \kappa \cdot \tau_1) \left(\frac{d^2 P(e) \cdot f(e)}{de^2} \right)}$$

Note that if $\kappa = 1$, then the derivative of input purchases w.r.t. a change in VAT on the input purchases equals 0. In other words, a firm's input purchase decisions do not respond to the VAT rate. This verifies the intuition that when κ equals one the VAT does not distort the firm's decisions.

Under mild assumptions the derivative of input purchases w.r.t. a change in VAT on input purchases is negative when $\kappa < 1$.⁵ When τ_1 increases and the government doesn't allow firms to write-off taxes paid on inputs, the de facto unit cost of inputs rises. Consequently, firms use less inputs in response to an increase in VAT and the VAT regime distorts purchase decisions.

In the extreme case where $\kappa = 0$, equation(1) implies that an increase in τ_1 is equivalent to a $\tau_1\%$ increase in the price of the input. This is seen more formally by applying the implicit function theorem, and deriving the following expression that links κ with the VAT and price elasticities,

$$(3) \quad \frac{\epsilon_{1+\tau_1}}{\epsilon_q} = \frac{(1 - \kappa) \cdot (1 + \tau_1)}{(1 + \tau_1 - \kappa \cdot \tau_1)}$$

where ϵ_k for $k \in \{1 + \tau_1, q\}$ is the elasticity of e w.r.t k and is equal to $\frac{k}{e} \cdot \frac{\partial e}{\partial k}$. Equation(3) implies that $\epsilon_{1+\tau_1} = \epsilon_q$ if and only if $\kappa = 0$. In other words, if κ equals 0 then an increase in VAT from 0 to τ_1 is equivalent to a $\tau_1\%$ increase in the input price.

Equation(3) is important because it can be used to measure κ . The econometrician needs to measure the elasticity of input purchases w.r.t. the VAT rate and the price, and use their ratio to pin down κ (τ_1 is observed). In other words the

⁵A set of sufficient conditions are concavity of $f(e)$ and the assumption that $\frac{d^2 P(e)}{de^2} \leq 0$. This includes a model with linear downward sloping demand and a model where the firm takes their output price as given.

two elasticities help put a magnitude on the distortion that results from the firms' inability to write-off taxes paid on inputs.

Due to our focus on the effect of the government's non-repayment of VAT liability to firms, it is important to highlight how the government can owe firms money. This is not obvious assuming that most firms create value, in which case taxes collected on outputs exceed taxes paid on inputs. However, there are several reasons why governments may end up owing firms. One reason is that τ_0 may be less than τ_1 , and may in fact equal zero, which implies that the VAT liability term would almost always be negative. As described in the background section, this is applicable in our case. Other reasons include purchases of capital that do not lead to an immediate increase in output, and delays in production, where output is sold several periods after the purchase of inputs. This can make firms sensitive to a change in the VAT rate because of late recovery of VAT input credits through sales.

The model discussed in this section makes several simplifying assumptions. First, it assumes that the price of the input is not impacted by a change in the VAT rate. Second, it abstracts away from tax evasion. Finally, the model condenses all frictions in the VAT into κ , which raises the question of other confounding frictions. We defend the first two assumptions in the empirical strategy section, which describes the policy variation in more detail and establishes that these channels should not impact the results. The issue of confounding frictions, such as lags between production and sales, is an important one and an example of such a model is given in the Appendix. Despite the possible confounders, we believe unpaid tax credits are the driving force behind our results for two reasons. First, the results section documents a persistent change in unpaid tax credits after a change in the VAT rate. Second, our empirical estimate of the effect of VAT rate on input purchases is too large for it to be reasonably explained by delays in production.⁶

III Data

We make use of Pakistan's complete record of firm-level monthly electronic VAT filings. These returns include rich information on gross and taxable purchases and

⁶Further discussion on the effect sizes can be found in the Appendix.

sales across a wide variety of categories, such as exports, imports, domestic transactions, electricity purchases, the stock of accumulated VAT input credits, and the amount of any VAT refund firms apply for during the reporting period. The aggregate reported sales and purchases are further broken down into transaction-level invoices. Each transaction invoice includes the tax identification number of the counterparty, the transaction value, any sales tax collected or paid, as well as the specific law exempting the transaction from the standard VAT rate if any exemption or zero-rating was applied. We use this invoice-level data to estimate firm-level exposure to the zero-rating policy by aggregating sales and purchases by whether the transaction was zero-rated under SRO 1125.

We make two primary sample restrictions for the purposes of our empirical analysis. First, we limit our analysis to between January 2017 and May 2019. We do so because several amendments were made to SRO 1125 in mid-2016 that included changes to the VAT rate charged on certain intermediaries, and since the withdrawal of the zero-rated regime was announced in June 2019 (SRO 491(1)/2016; Hanif, 2019). These global changes to the zero-rating regime plausibly had significant effects on the electricity demand of firms that confound the effects of an electricity zero-rating specifically. Second, in order to maintain a balanced panel, we require that firms report non-zero electricity purchases for at least 26 out of the 29 months in our sample period. This limits the scope of our analysis to larger, and likely more profitable, firms that did not close during the sample period, although these firms also make up relatively more of the value-added in the Pakistani economy.⁷

Table 1 contains summary statistics of the firms in our final sample. Column (1) summarizes for all the firms in the final sample, while column (2) contains a summary of only the firms that received or lost an electricity zero-rating in our sample period. As expected, treated firms are much more likely to have a zero-rating status on sales and non-electricity purchases, and by virtue of being zero-rated, also tend to be more export oriented.

⁷Out of 18,967 firms that purchase electricity in at least one month over the period January 2017 to May 2019, we retain 11,505 in the final sample due to the balanced sample restriction. These 11,505 firms constitute 90% of the sales of the total 18,967 firms over the period January 2017 to May 2019.

IV Empirical Strategy

A. Policy Variation

To measure the VAT elasticity of electricity demand, we leverage monthly changes in the zero-rating of electricity, which allow us to analyze the effects of VAT changes on electricity purchases. As previously mentioned, the zero-rating facility for electricity can be both withdrawn and reinstated each month based on the decisions of regional tax offices.

Firms seeking zero-rating status for electricity must apply to these regional tax offices, providing a valid meter identification number and various documents for review. This process, according to our discussions with local tax experts, involves a thorough and time-consuming examination by the tax officials, leaving firms uncertain about the approval timeline. Once approved, a directive is sent to the state-owned power sector to stop collecting sales tax on electricity purchases associated with the specified meter, a change reflected in the electricity purchase receipts indicating zero collected value-added tax.

Regional tax offices also have the authority to audit firms and rescind their electricity zero-rating status. Official regulations stipulate that this can occur due to errors in a firm's documentation, expansion of a firm's business beyond zero-rated sectors, electricity supply or redistribution to un-authorized recipients, or for reasons at the tax office's discretion. Our discussions with local tax experts reveal that revocations often occur without explicit reasons, leaving firms unprepared and adversely affected.

Despite the arbitrary nature of zero-rating status revocations, it remains crucial to address the issues that could potentially distort our VAT elasticity estimates. For instance, if a firm loses its zero-rating due to redistributing electricity, the resulting drop in its electricity purchases might not reflect a true reduction in its production-related electricity use. This is because a VAT change would impact not only the firm's own demand but also that of any recipients of the redistributed electricity, assuming some pass-through. To address this, we drop all firms that lost their zero-rating status for supplying electricity to others from the final sample using details on penalties issued for this activity that are recorded in our VAT data. Additionally,

we consider the influence of firms expanding their operations beyond zero-rated sectors by examining changes in their mix of zero-rated sales in the Appendix.

Using changes in firms' zero-rating status on electricity purchases as a source of variation has a couple of advantages that help alleviate concerns about whether we indeed recover the VAT elasticity discussed in section III irrespective of the empirical strategy that we may employ. These concerns are:

Evasion via misreporting: The possibility of firms altering their reported figures to reduce tax liability is a common concern when changes in VAT rates on inputs or outputs occur. Such misreporting complicates the task of distinguishing real effects from those driven by evasion responses. However, the risk of evasion in our study is considerably mitigated by the highly regulated nature of Pakistan's electricity sector. With electricity distribution predominantly under government operation or stringent regulation, it becomes challenging for firms to inaccurately report their electricity purchases. These purchases are consistently cross-verified with government-issued data, a process detailed by Shah (2021), reducing the likelihood of evasion-influenced misreporting. Additionally, our summary statistics in Table 1 reveal that firms experiencing zero-rating status changes frequently have a significant portion of their outputs and inputs already zero-rated. This situation limits the potential tax benefits of misreporting, consequently reducing the incentive for firms to falsify their sales and purchase reports.

Changes in the price of electricity: Another aspect to consider is that our data captures the monthly electricity expenditures of firms rather than the actual quantity of electricity purchased. A concern arises if the price of electricity is affected by changes in the VAT rate applied to a firm's electricity expenses, or if price changes occur simultaneously with VAT rate changes. In such cases, observed shifts in electricity expenditure could be attributable to price fluctuations, in addition to VAT rate changes. This issue is particularly pertinent in policy variations exploiting system-wide VAT changes, where price correlations could emerge due to government policies or equilibrium responses.

We think that due to institutional factors this is not a concern in our empirical setting. The electricity sector is heavily regulated by the government, which sets electricity tariffs across the country. Changes in tariffs are global and do not co-

incide with firm-specific month to month changes in the zero-rating of electricity. This allows us to interpret changes in the electricity expenditure as changes in the quantity of electricity purchased. Consequently, we maintain that our policy variation enables us to measure the correct parameter of interest i.e. the VAT elasticity of electricity.

B. Empirical Specification

In our analysis, we use a differences-in-differences (DD) methodology to isolate the causal effects of changes in the zero-rating status of electricity purchases on various outcome variables, which we will describe subsequently in this section. The DD methodology uses two distinct sources of variation. First, we compare the periods before and after a zero-rating status change. Second, we compare firms whose zero-rating status remained constant throughout our sample period with those experiencing changes in their status.

Specifically we estimate the following regression:

$$(4) \quad y_{it} = \alpha_i + \gamma_t + \beta Z_{it} + \epsilon_{it}$$

Here, y represents the outcome variable, with i and t indicating the firm and month-year, respectively. Z_{it} is an indicator variable that equals one if firm i benefits from zero-rated electricity in period t on at least one of its electricity meters, and ϵ_{it} is the error term. We perform this estimation on a pooled sample, as well as separately among firms that lose (an increase in VAT on electricity) and gain (a decrease in VAT on electricity) zero-rating status on their electricity purchases. The coefficient β is interpreted as the causal effect of the zero-rating status change on the outcome y .

Our analysis focuses on the following outcomes y :

1. **VAT rate paid on electricity purchases:** This variable is central to the first-stage of our analysis, and is directly affected by the policy. It is computed as the ratio of the total VAT paid on electricity to the total pre-tax value of these purchases. The results section will delve deeper into the importance of this

outcome.

2. **Log of electricity purchases:** This is our primary outcome variable. The coefficient β is interpreted as the percentage change in electricity purchases in response to the zero-rating policy.
3. **Log of input credits and refunds:** These variables measure the government’s financial obligations to firms, and are vital for establishing the mechanism through which VAT rate changes impact firms.
4. **Sales, Purchases, Exports, and Imports:** These variables assess the policy’s impact on a broader set of firm-level outcomes.

To examine pre-trends and the dynamics of firms’ responses, we also employ the following event study specification:

$$(5) \quad y_{it} = \alpha_i + \gamma_t + \sum_{j \in L} \beta_j 1(l = j, l = t - E_i) + \epsilon_{it}$$

Here, E_i denotes the specific period when the firm’s zero-rating status changes, and $1(l = j, l = t - E_i)$ is an indicator for l periods since this change for treated firms. The coefficients β_j measure the average difference between treatment and control firms for each period $l \in L = \{-12, \dots, 12\}$, excluding the period $l = -1$ as the reference. The absence of trends in β_j for $j < 0$ suggests firms did not anticipate the policy change.

Our regression approach leverages month-year-firm variation to estimate the coefficient β . The two-way fixed effects model compares outcomes y_{it} in periods of zero-rated electricity against non-zero-rated periods, accounting for constant differences between firms, such as sector, and time-specific factors, like monthly exchange rate fluctuations. Two key assumptions underpin this model. Firstly, we assume uniform treatment effects across different cohorts and over time. This assumption, critical for the accurate estimation of treatment effects in two way fixed effect designs as highlighted by Sun and Abraham (2021), Callaway and Sant’anna (2021), and Chaisemartin and D’Haultfoeuille (2020), is admittedly strong. However, as we demonstrate in the Appendix, our primary results maintain robustness

when employing the alternative estimators recommended by these authors.

Secondly, and more significantly, our identification strategy hinges on the assumption that changes in electricity's zero-rating status are exogenous to the trajectory of outcomes under the status-quo zero-rating regime. Essentially, this means that the identification of β might be compromised by factors influencing a firm's outcomes over time, specifically if these factors are correlated with, or cause changes in the zero-rating status. One potential violation of this assumption could occur if the zero-rating status changes in anticipation of electricity price changes, which could otherwise negatively impact firms' outcomes. However, due to institutional dynamics, we consider this a minor concern. Price changes in electricity are implemented nationally and uniformly, hence not aligning with the month-to-month, firm-specific adjustments in the zero-rating of electricity. Another potential violation could arise if zero-rating status changes are preemptive responses to anticipated future output demand shocks that necessitate higher electricity purchases, leading to misinterpretation of changes in electricity purchases as VAT-driven, rather than demand-driven. Given the control exerted by regional tax offices over the timing and status of zero-rating, we believe this risk to be minimal. We acknowledge, however, that we cannot entirely dismiss the possibility of firms influencing tax officials for zero-rating in anticipation of increased demand for their products.

We further estimate an instrumental variables regression to estimate the elasticity of electricity demand with respect to the VAT rate. The second-stage regression is:

$$(6) \quad y_{it} = \alpha_i + \gamma_t + \beta \cdot \text{VAT}_{it} + \epsilon_{it}$$

and the first stage regression is:

$$(7) \quad \text{VAT}_{it} = \alpha_i + \gamma_t + \mu Z_{it} + \epsilon_{it}$$

where the y is log of electricity purchases, VAT is the realized VAT rate on electricity purchases, and the definitions of all other terms are the same as before. In contrast to equation(4), the estimated β corresponds to a percentage change in the

purchase of electricity for a percent change in the VAT on electricity.⁸

We check the robustness of our results in several ways and report the results in Appendix Table A.1, and Appendix Table A.3. This includes adding firm-by-quarter fixed effects for idiosyncratic seasonality, allowing for treatment effect heterogeneity across cohorts and time, including firms who lose and regain the zero-rating status multiple times over the sample timeframe, and donut DiDs that drop observations in small windows around the event time to exclude anticipatory or other spurious short-run effects.

V Results

A. Average VAT Rate on Electricity

We begin with an examination of the first-stage effects of changes in zero-rating status on the effective VAT rate paid by firms. The first-stage estimation is crucial since the treatment is delivered at the electricity meter level, while our analyses are conducted at the firm level. This is because the VAT data does not contain meter identifiers. Consequently, a firm with multiple meters might experience a change in the zero-rating status of only one meter, potentially having a minimal impact on its overall VAT payments on electricity. Such a scenario could imply varying treatment intensity across firms. To address this issue, we leverage the specifications described in equation(4) and equation(5) on the realized VAT rate paid on electricity to measure the relevance of the policy variable Z_{it} in influencing the VAT rate.

The results of our regression analyses are presented in Table 2, Panel A. Specifically, Column (1) in Panel A of Table 2 illustrates the impact of zero-rating status on the VAT rate across the entire sample. We observe that the acquisition of zero-rating status leads to a reduction in the VAT rate by 15 percentage points, a result that is statistically significant at the 99% level. Columns (2) and (3) in the same panel further dissect this effect for split samples: firms acquiring and firms los-

⁸In the case of equation(4), we recover the effect of the change in the zero-rating policy on electricity demand and other outcomes

ing their zero-rating status, respectively. Given that the Z_{it} variable is assigned a value of 1 for a firm's meter under zero-rating, the negative coefficients reported in both columns reflect this operational definition.⁹ We find that the zero-rating status changes the VAT rate for the different samples by similar magnitudes and in expected directions.

Figures 1 Panels A and B visually depict the β_j estimates from equation (5), reinforcing the findings reported in Columns (2) and (3) of Table 2 Panel A. In Panel A of Figure 1, we observe a distinct and sustained reduction in the VAT rate on electricity purchases for firms acquiring zero-rating status. Conversely, Panel B of Figure 1 illustrates that firms losing their zero-rating status experience a notable and persistent increase in their VAT rate on electricity purchases. Importantly, the figure also indicates an absence of pre-trends in the average VAT rate on electricity before any changes in zero-rating status, further validating our analysis.

These findings clearly indicate that firms exhibit a substantial difference in the VAT rate they pay in reaction to policy changes, thereby affirming the importance of the policy variable Z_{it} . The observed shift in realized VAT rates, approaching the full 17%, suggests that the absence of meter identification numbers in our data does not significantly impede the variation we capture. Furthermore, this outcome implies that the variation in VAT rates primarily occurs at the firm level rather than within the firm at the individual meter level.

B. Electricity Demand

Having showed that changes in zero-rating status lead to notable alterations in VAT rates, we now turn our attention to examining its impact on firms' log electricity expenditure. In Pakistan, the government sets electricity prices, which remain unaffected by individual firm's variations in VAT rates. Therefore, analyzing the influence of zero-rating on electricity expenditure effectively equates to assessing its effect on the quantity of electricity purchased. Consequently, this measure of electricity expenditure aligns with the conceptual outcome variable we outlined in Section 3.

⁹For firms losing their zero-rating status, the Z_{it} variable shifts from 1 to 0, resulting in a negative coefficient.

Table 2, Panel B, presents the results obtained from estimating equation (4) to both the logs and levels of electricity expenditure. In Column (1) of Panel B, the impact of zero-rating status on electricity expenditure across the full sample is detailed. We observe that acquiring zero-rating status causes an average increase of approximately 0.11 million PKR per month in electricity expenditures. Comparing with the pre-treatment means this translates to an effect size of roughly 8 percent. The regressions using logs show a 13 percent rise in monthly electricity expenditure. Furthermore, Columns (2) and (3) of the same panel dissect this effect for the distinct subsets of firms—those that gain and those that lose their zero-rating status, respectively. Across these subsets, we consistently find that having zero-rating status notably elevates electricity expenditure.

Figures 1 Panels C and D visually represent the β_j estimates derived from the model outlined in equation (5), lending further support to the results detailed in Columns (2) and (3) of Table 2, Panel B. In Panel C of Figure 1, we observe an initial dip in electricity purchases by firms just before they receive zero-rating status, followed by a sustained increase in monthly electricity purchases thereafter. This pattern suggests that firms, anticipating the activation of zero-rating status in the following month, may defer their electricity purchases to benefit from the lower rates. Conversely, Panel D illustrates that firms losing their zero-rating status experience an abrupt decline in electricity purchases during the month of the change, with no prior indications, indicating that the loss of zero-rating status often comes as a surprise. In the months following the loss of zero-rating status, while electricity purchases partially rebound, they consistently remain below the levels observed during periods when zero-rating status was in effect.

As indicated in our conceptual framework, in the absence of VAT-related frictions on electricity purchases, changes in VAT rates should not affect the quantity of electricity purchased. However, our findings reveal the opposite: significant and substantial effects of VAT rate changes on electricity on the purchasing behaviour of firms. This clearly suggests the existence of a friction in the VAT system. Moreover, the observed trends in firms that gain or lose zero-rating status align with the interpretation that firms perceive a rise in VAT as an increase in the unit cost of electricity.

C. 2SLS Estimates

In this subsection, we present results from the instrumental variables strategy outlined in equations (6) and (7). The resulting 2SLS estimates can be interpreted as the elasticity $\epsilon_{(1+\tau_1)}$, as discussed in Section 3. This interpretation is valid because a VAT rate shift from τ_a to τ_b , with $\tau_a = 0$, represents a τ_b percent change in $(1 + \tau_a)$. As explained about equation (3) of Section 3, measuring this elasticity, alongside the price elasticity of demand, ϵ_q , is pivotal in quantifying the extent of the VAT friction, denoted as κ .

Our analysis reveals that electricity demand is inelastic in relation to the VAT rate. In Column (1) of Table 2, Panel C, we present the VAT elasticity of electricity expenditure for the pooled sample, illustrating that a 1% increase in VAT rate leads to a 0.81% decrease in electricity expenditure.¹⁰ Additionally, Columns (2) and (3) of the same panel assess this elasticity for the distinct groups of firms—those obtaining and those losing their zero-rating status—and report consistent findings. The magnitude of this elasticity underscores the significant friction introduced by the VAT. It suggests that firms are either reducing their operational scale in response to higher VAT rates or are shifting towards alternative, non-grid power sources. In Section 7, we will rigorously compare this elasticity with standard price elasticity of demand to quantify the extent of the friction, denoted as κ .

D. Robustness

In order to validate the reliability of our findings, we have conducted a series of robustness tests. The event study plots in Figure 1 provide reassurance, as they do not reveal any pre-trends. Furthermore, we have explored a range of alternative specifications to test the resilience of our results. This includes expanding our sample to encompass firms that experience multiple changes in their zero-rating status throughout the study period, incorporating firm-quarter fixed effects to capture unique quarterly seasonal variations, using DD estimators that allow for treatment effect heterogeneity across cohorts and time, and applying 'donut' difference-in-

¹⁰The F-Stat for the first stage is 604.32, which, while large, is expected given the direct impact of zero-rating status changes on the VAT rate on electricity.

differences (DiDs) approaches. In these donut DiDs, we omit data from up to three months surrounding the policy change events, thereby accounting for any anticipatory behaviors and other transient effects that may emerge immediately following the policy shift. Detailed outcomes of these robustness checks are provided in Appendix Table A.1 and Appendix Table A.3.

E. Mechanisms

We propose that the friction identified in our findings is primarily a result of delays in the VAT refund system. Although we lack exogenous variation in the government’s repayment of VAT credits, which limits our capacity to definitively ascertain that the delay in VAT credit repayments is the driving mechanism behind the observed friction, we offer supporting evidence for this hypothesis. We demonstrate that changes in the zero-rating status of electricity lead to significant and persistent changes in the amounts owed by the government to firms. Additionally, we successfully exclude alternative mechanisms.

We start by highlighting that the issue of unpaid credits and refunds in the developing world is well-known to policymakers. A 2005 IMF report characterizes the refund system as the ‘Achilles heel’ of VAT (Harrison and Krellove, 2005), noting that in Asia and Africa’s developing nations, VAT refunds represent merely 6-7% of total VAT collections, a stark contrast to developed countries. For instance, VAT refunds in Canada account for 50% of total VAT collections, while in the EU, the figure stands at 38%. Focusing on Pakistan, it is a widely acknowledged issue that the government frequently does not reimburse firms, often citing political or financial reasons. For example, an article from the Express Tribune, a local newspaper, highlights that withheld refunds are sometimes used to artificially boost tax collection figures, enhancing the country’s ability to secure loans from international organizations like the IMF (Rana, 2016).

In our analysis, we examine two key variables relevant to VAT credits. First, we observe the value of excess input tax credits at the conclusion of each VAT reporting period. This variable represents the VAT paid on inputs which the firm could not offset against taxes collected from sales and has chosen to carry forward, with the

intention of offsetting it against future sales. Second, we track VAT refund requests submitted by firms to the government, reflecting the amount of VAT paid on inputs that could not be offset against sales taxes and for which the firm seeks a refund. The cumulative total of these two variables essentially signifies the government's VAT liability to a firm in any specific month.

The rationale for firms to request refunds is straightforward: it is generally more advantageous for them to have the money owed to them in their own accounts rather than on the government's ledger. However, firms may opt to retain tax credits instead of immediately seeking refunds for two primary reasons. First, when a refund is requested, all related paperwork is removed from the official tax documentation, leaving the repayment at the government's discretion. If the government fails to disburse the refund, the firm loses the opportunity to offset this amount against future tax liabilities. Therefore, if there is a high risk of non-repayment by the government, firms might prefer to keep tax credits on their sales tax forms, preserving the ability to use these credits to offset future tax obligations. Second, requesting a refund involves substantial paperwork, additional tax form submissions, and further scrutiny, which can be burdensome for firms.

This discussion suggests that while the government's financial obligation to firms is at least the total of excess input tax credits and refund requests, the precise amount owed by the government remains unknown. The lack of official data on the processing status of refund requests in tax records precludes exact measurement. Consequently, in our reporting, we use the sum of excess input tax credits and refund requests as our main outcome variable. Additionally, we also present results for each of these variables individually.

Table 3 reports the estimates from the empirical specification in equation (4) for the log of input credits, refunds, and their sum. Our analysis reveals that obtaining zero-rating status results in a decrease of 2.21 log points in a firm's input credits, 0.94 log points in refunds, and 2.94 log points in the total of credits and refunds, all significant at the 99% level. While these figures are substantial, they align with expectations. As shown in Table 1, the firms undergoing zero-rating status changes predominantly have their inputs and outputs zero-rated. Therefore, any change in the zero-rating status of a vital input such as electricity leads to considerable fluc-

tuations in the amounts owed by the government to these firms, relative to what the government owed them previously.

Figure 2, Panels A and B, graphically depict the β_j estimates derived by estimating equation (5), and corroborate the findings highlighted in Column (3) of Table 3. In Panel A of Figure 2, we observe a pronounced and ongoing decrease in the total sum owed by the government to firms that acquire zero-rating status. This steady reduction over time can be attributed to the accumulating credits within the control group. Conversely, Panel B illustrates that firms losing their zero-rating status see a substantial surge in the amount owed to them by the government, which initially increases but then stabilizes over time. This trend is consistent with a scenario where both the treated and control groups see parallel increases in credits following the initial impact of the policy change.

These results offer evidence that the government's financial obligations to firms escalate following the loss of their zero-rating status. Given that a substantial portion of these firms' outputs are zero-rated, they face challenges in offsetting these tax credits against taxes collected on future sales. As a result, we argue that firms perceive an increase in VAT rate as a higher unit cost of electricity. This prompts them to modify their input choice behaviour, as previously discussed in our analysis.

One potential confounding factor in our analysis is the zero-rating status of other inputs or outputs. For instance, if the zero-rating status of electricity coincides with changes in the zero-rating status of other inputs, then the observed shifts in the government's liability to firms might not solely be attributable to the change in electricity's zero-rating status. To address this concern, we investigate whether there are corresponding changes in the fraction of zero-rated sales and zero-rated non-electricity purchases when the zero-rating status of electricity is changed. The results, detailed in Appendix Table A.2, indicate that these variables remain largely unaffected by changes in the zero-rating status of electricity.

Although our results thus far suggest that the amount the government owes firms is influenced by changes in the zero-rating status of electricity, it is important to consider other potential mechanisms that could drive the observed variations in monthly electricity expenditures. One such confounding factor could be demand shocks, particularly those arising from government contracts. It is conceivable that

firms receiving government contracts might be granted zero-rating status for electricity as part of their production arrangement. To investigate this possibility, we estimate equation (4) on additional outcomes like sales, exports, imports, and non-electricity purchases. If demand shocks were influencing our results, one would expect an increase in sales and purchases concurrent with changes in the zero-rating status of electricity. However, as outlined in Table 4, our analysis reveals that these outcomes are not significantly affected, suggesting that demand shocks are not a major driver of the changes we observe.

The absence of significant changes in sales and other purchases, and the pronounced response of electricity purchases to the zero-rating status, prompts a question: why do electricity purchases react so distinctly while other aspects like sales and purchases remain largely unaffected? Several explanations could account for this phenomenon. First, our dataset does not encompass all inputs used by firms, some of which might vary in response to changes in electricity demand. Notably, the VAT data lacks details on labor expenses.¹¹ Second, it is common for firms in developing countries to generate their own electricity due to unreliable grid supply. Consequently, our findings might reflect a shift from grid-supplied to self-generated electricity in response to changes in VAT on electricity, offering a plausible explanation for the substantial -0.81 VAT elasticity we observe.¹²

Lastly, the null results can also be viewed as a limitation of our study. The narrow scope of the policy variation, while beneficial for excluding various channels and identifying key elasticities, simultaneously limits our capacity to comprehensively analyze the broader impact of refund frictions on firm outcomes and productivity, which are vital components of welfare analyses. Consequently, we acknowledge this as a gap and propose it as a valuable avenue for future research.

¹¹This aspect is particularly relevant in the developing world, where the production function may not strictly adhere to a Leontief relationship between electricity and other inputs. For instance, a textile manufacturer with both powerlooms and handlooms could increase labor hiring to compensate for higher electricity costs.

¹²We find suggestive evidence for this phenomenon when we estimate the price elasticity for electricity and find the elasticity to be particularly large amongst firms with active gas connections. This is documented in the Appendix

VI Comparison to the Price Elasticity of Demand

In this section, we focus on measuring the price elasticity of electricity demand. This is achieved by analyzing the effects of reduced electricity tariffs implemented for firms in zero-rated sectors. Subsequently, we combine this price elasticity estimate with the VAT elasticity determined in Section 6.3. Together these two measures enable us to empirically quantify the VAT friction, denoted as κ , which is referenced in equation (3).

A. Estimating the Price Elasticity of Electricity

1. Policy Variation

We leverage changes in the price of electricity that are plausibly exogenous to quantify their impact on the quantity of electricity purchased by firms from January 2018 to September 2019. This analysis specifically utilizes the introduction of reduced tariffs for manufacturing firms within the five zero-rated sectors. This policy adjustment lowered the electricity price to 7.5 cents per kilowatt-hour, effectively amounting to a 22% reduction in the unit cost of electricity.¹³ Initiated in January 2019, the policy was implemented incrementally, with eligible firms being progressively enrolled for these reduced tariffs.

The staggered enrollment of firms into the reduced tariff scheme stemmed from the necessary application procedures. To access this benefit, firms were required to apply to the respective electricity distribution companies. Once an application was submitted, the distribution company would forward the firm's profile to the Federal Board of Revenue (FBR) for verification. This step was essential to confirm the firm's eligibility based on its classification within the zero-rated sectors. The activation of the reduced tariff occurred only after the FBR approved the application, a process communicated back to the distribution company. The inherent administrative delays at each stage resulted in a phased enrollment of firms into the program. Our discussions with local manufacturers revealed that firms were generally un-

¹³The methodology for calculating the change in unit price resulting from this policy shift is detailed in the appendix.

aware of the exact month their approval would be granted, typically discovering their enrollment only upon noticing the adjusted charges on their electricity bills

We estimate the price elasticity of demand from this specific policy change, rather than relying on pre-existing studies, for two primary reasons. First, existing literature estimating the price elasticity of electricity for firms in Pakistan, such as studies by Chaudhry (2010), Khan and Qayyum (2009), Ishaque (2018), Alter and Syed (2011), and Khan and Abbas (2016), predominantly employs Ordinary Least Squares (OLS) methodologies. These OLS-based studies may be prone to standard endogeneity issues, where price variations are correlated with unobserved changes in demand. Secondly, the reduced tariffs we examine were targeted specifically at firms in the zero-rated sector and were implemented over a time frame that closely aligns with our main analysis. This enhances the relevance and comparability of our price elasticity findings to the VAT elasticity we have estimated.

2. Data and Empirical Specification

To calculate the price elasticity of electricity, we focus on a sample of manufacturing firms operating within the zero-rated sectors.¹⁴ The dataset underpinning this analysis was sourced from Pakistan's Ministry of Water and Power. It encompasses detailed monthly records of electricity consumption for these firms, the specific periods during which they were enrolled for reduced tariffs, and the PKR value of the monthly electricity subsidies received by each firm. It is important to note, however, that this dataset does not include information on billing, and VAT charged on electricity. Additionally, due to the absence of common identifiers, it cannot be merged with the anonymized VAT data.

We use a differences-in-differences (DD) methodology to isolate the causal effects of enrolling in the reduced electricity tariff regime on the quantity of electricity units purchased. Specifically we estimate the following reduced form regression:

¹⁴Notably our sample contains firms that eventually received subsidy till December 2019, except those from LESCO and KE, which are two large electricity distribution companies in Pakistan.

$$(8) \quad y_{it} = \alpha_i + \gamma_{m(t)} + \kappa_{y(t)} + \beta T_{it} + \epsilon_{it}$$

In this equation, y_{it} represents the log of the quantity of electricity purchased, where i denotes the firm, $m(t)$ the month, $y(t)$ the year, and t the month-year. The term T_{it} is an indicator variable that takes the value of one when the price paid by firm i for electricity drops by 22%, in line with the government's policy. The error term is denoted by ϵ_{it} . The coefficient β is interpreted as the percentage change in the quantity of electricity purchased in response to a 22% price reduction.

Two assumptions underpin our identification strategy. First, we assume uniform treatment effects across different cohorts and over time. This assumption, critical for the accurate estimation of treatment effects in two way fixed effect designs as highlighted by Sun and Abraham (2021), Callaway and Sant'anna (2021), and Chaisemartin and D'Haultfoeuille (2020), is admittedly strong. However, as we demonstrate in the Appendix, our results maintain robustness when employing the alternative estimators recommended by these authors. Second, we assume that enrollment into the program is exogenous to the trajectory of outcomes had there been no program to begin with. One potential violation of this assumption could occur if firms whose electricity consumption would decline sharply without assistance from the program enrolled first.

To visualize our results and assess for any pre-trends, we also present results using the event study model specified below:

$$(9) \quad y_{it} = \alpha_i + \gamma_{m(t)} + \kappa_{y(t)} + \sum_{j \in L} \beta_j 1(l = j, l = t - E_i) + \epsilon_{it}$$

In this model, y_{it} represents the log quantity of electricity purchased. The term E_i identifies the specific time period when a firm's electricity price changes, while $1(l = j, l = t - E_i)$ serves as an indicator for the number of periods l since the price change for treated firms. The coefficients β_j quantify the average difference between treatment and control firms across each period $l \in L = -12, \dots, 8$, with the period $l = -1$ omitted and used as a reference. A flat evolution of β_j for $j < 0$

values for $j < 0$ would suggest that firms did not anticipate the policy change, supporting our assumptions.

B. Results

Table 5, column (2) reports the estimate of β obtained from the specification in equation(8). We find that a 22% reduction in electricity prices leads to a 13% increase in electricity consumption. This finding translates to a price elasticity estimate of 0.59, with a standard error of 0.18. In the next section, we compare the previously measured VAT elasticity with this estimate of the price elasticity.

In Table 5, Column (3), we report the β estimate derived from the model in equation (8), this time incorporating month-year fixed effects. The resulting estimate of an 8% increase in electricity consumption, though positive and substantial, is not statistically significant. Consequently, we opt not to use this particular estimate for the comparative analysis between VAT and price elasticity. Instead, we rely on the more substantial 13% estimate as it offers a more conservative basis for estimating κ in Section 7.3.¹⁵

Figure 3, panel A plots the β_j estimates from the specification in equation(9) where the outcome variable is the PKR value of the subsidy firms received under the policy. The figure shows a flat profile of β_j in the pre-period, which is followed by a substantial increase in the rupee value of the subsidy received by the firms. This graph indicates that firms did indeed receive a subsidy from the government.

Figure 3, Panel B, displays the β_j estimates from equation (9), reinforcing the findings from Column (2) of Table 5. The figure illustrates a somewhat erratic yet generally flat trajectory for β_j in the pre-policy period. It then depicts a noticeable increase in electricity consumption a few months subsequent to the price reduction. The relatively flat β_j trajectory prior to the policy change mitigates concerns about the exclusion of month-year fixed effects, which are typically employed to capture broader, non-linear time trends.

¹⁵A higher VAT elasticity in comparison to price elasticity suggests a more negative value of κ .

C. Discussion

In the preceding section, we observed that the estimated price elasticity, at 0.59, is lower than the VAT elasticity, which stands at 0.81.¹⁶ This outcome is intriguing, especially when considered within the context of our model. Theoretically, if the VAT paid is never reimbursed, it should have a financial impact equivalent to an unrecovered increase in electricity prices. Thus, the reaction to a VAT change should not surpass that to a price change, even when considering the VAT system's inherent frictions. This implies that the model's κ , representing the friction's magnitude in the VAT system, should not be less than 0. Utilizing our calculated VAT and price elasticities, and applying equation (3) from our conceptual framework, we derive an empirical estimate of $\hat{\kappa} = -0.43$.¹⁷ In this section, we delve into how, despite theoretical considerations, our empirical findings suggest that firms essentially treat VAT on electricity akin to a standard price increase.

Note that the 95% confidence intervals for the estimated elasticities are $[-1.43, -0.18]$ for VAT elasticity, and $[-0.95, -0.23]$ for price elasticity. Given the substantial overlap between these intervals, we cannot statistically distinguish between the two elasticities. Additionally, employing the delta method allows us to calculate the standard error and 95% confidence interval for $\hat{\kappa}$, though this requires an assumption about the correlation between the variables, which is not directly observable.¹⁸ Addressing this, we explore two scenarios: one where the correlation between elasticity estimates is 0 and another where it is 1. In the former, the standard error for $\hat{\kappa}$ is 0.54, while in the latter, it increases to 0.67. Opting for the more cautious standard error of 0.67, we determine a 95% confidence interval for $\hat{\kappa}$ to be $[-1.7, 0.8]$. If we set aside the theoretically implausible negative values within this interval, we can still refute the notion that $\hat{\kappa} = 1$. This finding, as laid out in our conceptual framework, indicates that the VAT system in Pakistan is not with-

¹⁶It is important to note that the data sets used to estimate price and VAT elasticities were distinct, with the former derived from separate records and the latter from anonymized tax data. This separation, due to an inability to merge these data sets, means that sample differences could naturally lead to variations in the estimates.

¹⁷While we use the formula to illustrate the calculation of κ , we recognize that applying it might be theoretically inconsistent, given that κ should not fall below 0.

¹⁸In our delta method application, we assume the measured standard errors represent the standard deviations of the random variables.

out frictions. Crucially, however, the possibility that $\hat{\kappa} = 0$ cannot be dismissed, suggesting a scenario where firms entirely offset the VAT paid on inputs.

Combining our findings with corroborating evidence from existing literature further strengthens our conclusion. For instance, Chaudhry (2010) found a price elasticity of -0.81 for Pakistani firms in the textile sector, which is one of the zero-rated sectors. This supports our assertion that firms in Pakistan essentially treat a VAT change as a shift in the price of their inputs, or in other words, view VAT as analogous to a sales tax on inputs. This insight is crucial, as it underscores the potential pitfalls of a VAT system that is not implemented correctly. Despite its theoretical efficiency, a VAT regime that is applied improperly can inadvertently create distortions in the market for intermediate goods. International organizations, like the IMF and the World Bank, which are instrumental in encouraging VAT adoption in developing countries (Mahon, 2004), should be mindful of these risks, and not only advocate for the adoption of VAT systems but also emphasize their proper and effective implementation.

VII Conclusion

In this paper, we explore the significance of unpaid tax credit refunds in a VAT regime, which is a relatively understudied topic. We demonstrate that policies that vary the amount of tax credit generation have real impact on the input purchases of firms. We show this in the context of electricity, which was zero-rated on a meter-by-meter basis in Pakistan. The distortionary effect of such policies is explained by a significant and persistent change in the government's liability to the firms. Our findings underscore a critical point: while VAT regimes are lauded for their theoretical efficiency, their practical implementation, especially in developing countries like Pakistan, can be a source of significant inefficiency. Future research should focus on unpacking the wider implications of these frictions, particularly their impact on welfare and productivity, which are pivotal in understanding VAT systems in developing economies.

References

- Ahmad, Ehtisham.** 2010. “Why is it so difficult to implement a GST in Pakistan?” International Growth Centre Working Paper.
- Ahmad, Ehtisham, and Nicholas Stern.** 1991. “The Theory and Practice of Tax Reform in Developing Countries.” Cambridge Books from Cambridge University Press.
- Ahmed, Khurshid.** 2019. “Pakistani exporters warn of 20-30% hit if zero rating sales tax regime abolished.” Arab News.
- Alter, Noel, and Shabib Haider Syed.** 2011. “An Empirical Analysis of Electricity Demand in Pakistan.” International Journal of Energy Economics and Policy, 1(4): 116–139.
- Best, Michael Carlos, Anne Brockmeyer, Henrik Jacobsen Kleven, Johannes Spinnewijn, and Mazhar Waseem.** 2015. “Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan.” Journal of Political Economy, 123(6): 1311–1355.
- Cagé, Julia, and Lucie Gadenne.** 2018. “Tax revenues and the fiscal cost of trade liberalization, 1792–2006.” Explorations in Economic History, 70: 1–24.
- Callaway, Brantly, and Pedro HC Sant’Anna.** 2021. “Difference-in-differences with multiple time periods.” Journal of econometrics, 225(2): 200–230.
- Chandra, Piyush, and Cheryl Long.** 2013. “VAT rebates and export performance in China: Firm-level evidence.” Journal of Public Economics, 102: 13–22.
- Chaudhry, Azam.** 2010. “A Panel Data Analysis of Electricity Demand in Pakistan.” Lahore Journal of Economics, 15: 75–106.
- De Chaisemartin, Clément, and Xavier d’Haultfoeuille.** 2020. “Two-way fixed effects estimators with heterogeneous treatment effects.” American Economic Review, 110(9): 2964–2996.
- Deloitte.** 2020. “Pakistan Highlights 2020.” International Tax Summaries.
- de Paula, Aureo, and Jose A Scheinkman.** 2010. “Value-Added Taxes, Chain Effects, and Informality.” American Economic Journal: Macroeconomics, 2(4): 195–221.

- Ebrill, Liam, Michael Keen, and Victoria Summers.** 2001. The Modern VAT. Washington, D.C.:International Monetary Fund.
- FBR.** 1990. “Sales Tax Act of 1990.”
- FBR.** 2021. “Statutory Regulatory Orders (SROs).”
- Foundation, Heritage.** 2021. “Pakistan Country Profile.” 2021 Index of Economic Freedom.
- Gadenne, Lucie, Tushar Nandi, and Roland Rathelot.** 2019. “Taxation and Supplier Networks: Evidence from India.” IFS Working Papers W19/21, Institute for Fiscal Studies.
- Gerard, Francois, and Joana Naritomi.** 2018. “Value Added Tax in developing countries: Lessons from recent research.” International Growth Centre: Growth Briefs.
- Gerard, Francois, Joana Naritomi, and Arthur Seibold.** 2018. “Tax Systems and Inter-Firm Trade: Evidence from the VAT in Brazil.” Mimeo.
- Gordon, Roger, and Wei Li.** 2009. “Tax structures in developing countries: Many puzzles and a possible explanation.” Journal of Public Economics, 93(7-8): 855–866.
- Hanif, Usman.** 2019. “PTI govt withdraws zero-rated status for major exporters.” The Express Tribune.
- Harju, Jarkko, Tuomas Matikka, and Timo Rauhanen.** 2019. “Compliance costs vs. tax incentives: Why do entrepreneurs respond to size-based regulations?” Journal of Public Economics, 173: 139–164.
- Ishaque, Hanan.** 2018. “Revisiting income and price elasticities of electricity demand in Pakistan.” Economic research-Ekonomska istraživanja, 31(1): 1137–1151.
- Khan, Muhammad Arshad, and Abdul Qayyum.** 2009. “The demand for electricity in Pakistan.” OPEC Energy Review, 33(1): 70–96.
- Khan, Muhammad Arshad, and Faisal Abbas.** 2016. “The dynamics of electricity demand in Pakistan: A panel cointegration analysis.” Renewable and Sustainable Energy Reviews, 65: 1159–1178.

- Krelove, Russell, and Graham Harrison.** 2005. VAT Refunds: A Review of Country Experience. International Monetary Fund. OCLC: 942504651.
- Liu, Li, Ben Lockwood, Miguel Almunia, and Eddy H. F. Tam.** 2021. “VAT Notches, Voluntary Registration, and Bunching: Theory and U.K. Evidence.” The Review of Economics and Statistics, 103(1): 151–164.
- Lockwood, Ben, and Michael Keen.** 2007. Value-Added Tax: Its Causes and Consequences. OCLC: 942504193.
- Mahon Jr, James E.** n.d.. “61 Taxation and State Capacity.”
- of Pakistan: Finance Division, Govt.** 2020. “Summary of Consolidated Federal and Provincial Budgetary Operations, 2019-20.”
- Onji, Kazuki.** 2009. “The response of firms to eligibility thresholds: Evidence from the Japanese value-added tax.” Journal of Public Economics, 93(5-6): 766–775.
- PWC.** 2021. “Pakistan Country Profile.” Worldwide Tax Summaries.
- Rana, Shahbaz.** 2016. “Dar accused of delaying tax refunds.” The Express Tribune.
- Rana, Shahbaz.** 2021. “FBR flouts orders to pay tax refunds.” The Express Tribune.
- Shah, Jawad.** 2021. “Essays on Evasion and Enforcement in Value Added Tax.” PhD diss. UK Libraries.
- Sun, Liyang, and Sarah Abraham.** 2021. “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects.” Journal of Econometrics, 225(2): 175–199.
- The Nation.** 2020. “Exports to suffer 30pc drop if zero-rating regime not restored.”
- Waseem, Mazhar.** 2023. “Overclaimed refunds, undeclared sales, and invoice mills: Nature and extent of noncompliance in a value-added tax.” Journal of Public Economics, 218: 104783.
- World Bank.** 2021. “World Development Indicators.”

VIII Tables and Figures

Table 1: Descriptive Statistics

	All Firms	Treated Firms
	(1)	(2)
Panel A: Output Characteristics (100,000 PKR)		
Total sales	196 (555)	275 (702)
Exports	30 (144)	70 (230)
Panel B: Input Characteristics (100,000 PKR)		
Elec. purchases	5 (8)	7 (10)
Non-elec. purchases	120 (365)	170 (467)
Imports	33 (151)	37 (158)
Panel C: VAT Characteristics		
Input credits + refunds (100,000 PKR)	17 (57)	17 (50)
Non-elec. purchases zero-rated (%)	0.33 (0.45)	0.75 (0.37)
Sales zero-rated (%)	0.43 (0.49)	0.88 (0.32)
Observations	305,755	21,228

Notes: This table contains descriptive statistics of the firms in Pakistan's tax database during our sample period. Column (1) reports summary statistics for all firms in the database, while column (2) reports summary statistics only among firms that gained or lost an electricity zero-rating. Panel (A) contains characteristics about a firm's sales and export behavior. Panel (B) contains characteristics about firms' input demand. Panel (C) reports the average % of sales and purchases that are zero-rated, and the average amount of accumulated VAT input credits and requested VAT refunds on firms' balance sheets. Observations are defined by firm and month.

Table 2: Effect of Zero-Rating on Average VAT Rate and Electricity Demand

	(1)	(2)	(3)
Panel A: VAT rate paid on electricity (first stage)			
Levels (%)	-0.15 (0.01)	-0.15 (0.01)	-0.15 (0.01)
Panel B: Electricity demand (reduced form)			
Levels (Millions of PKR)	0.11 (0.03)	0.15 (0.05)	0.07 (0.03)
Mean (Pre-Treatment)	1.42	0.98	1.76
Log demand	0.13 (0.05)	0.13 (0.06)	0.14 (0.08)
Panel C: 2SLS estimates of VAT rate change on electricity demand			
Electricity demand (log)	-0.81 (0.32)	-0.96 (0.43)	-0.82 (0.44)
Treated Sample	Pooled	Get ZR	Lose ZR

Notes: This table presents results from difference-in-differences regressions estimating the effect of zero-rated electricity on the average VAT rate paid on electricity and total electricity demand. Each cell represents the estimated treatment effect of the zero-rating from a separate regression. Panel A contains effects on the average VAT rate paid on electricity (first stage). Panel B contains effects on electricity demand in both levels (millions of PKR), the mean of electricity demand in the pre-treatment period (millions of PKR) and logs (reduced form). Panel C contains DiD-IV estimates of the effect of a 1 p.p. change in the VAT rate on log electricity demand. Column (1) estimates effects using all firms that either received or lost their electricity zero-rating over the sample period. Column (2) estimates effects only using firms that received a zero-rating. Column (3) estimates effects only using firms that lost their zero-rating. The pre-treatment means are mean values of the outcome variable in the treated group prior to the treatment, and in case of pooled sample is a weighted average. All regressions include firm and year-by-month fixed effects. The number of observations in Columns (1), (2) and (3), respectively, are 294,603, 289,245, and 286,510. Standard errors are clustered by firm.

Table 3: Effect of Zero-Rating on VAT Credits and Refund Requests

	Input credits	Refunds requests	Sum
	(1)	(2)	(3)
Electricity zero-rating (log)	-2.21 (0.26)	-0.94 (0.15)	-2.94 (0.27)
Observations	300,359	300,359	294,409
R-Squared	0.77	0.74	0.79

Notes: This table contains estimates of the effect of zero-rated electricity on a firm's accumulated VAT input credits and VAT refunds requests made to Pakistan's tax authority. All outcome variables are in log units. All regressions include firm and year-by-month fixed effects. Standard errors are clustered by firm.

Table 4: Effect of Zero-Rating on Broader Firm Outcomes


	Sales	Non-elec. purch	Exports	Imports
	(1)	(2)	(3)	(4)
Electricity zero-rating (log)	-0.14 (0.17)	0.08 (0.15)	-0.03 (0.08)	0.01 (0.08)
Observations	294,584	294,613	294,710	294,454
R-Squared	0.83	0.78	0.84	0.73

Notes: This table contains estimates of the effect of zero-rated electricity on a set of broader firm outcomes over the 12 month period proceeding the zero-rating change. All outcome variables are in log units. All regressions include firm and year-by-month fixed effects. Standard errors are clustered by firm.


Table 5: Estimates of the Price Elasticity of Demand for Electricity

	Subsidy	Elec. demand (log)	Elec. demand (log)
	(1)	(2)	(3)
Estimate	22%	0.13	0.08
	-	(0.04)	(0.06)
Firm FE	-	✓	✓
Month FE	-	✓	-
Year FE	-	✓	-
Month x Year FE	-	-	✓
Observations	-	51,736	51,736
R-Squared	-	0.82	0.82


Notes: This table contains estimates of the effect of Pakistan's electricity subsidy on prices and log electricity demand. The estimated subsidy percentage in column (1) is based on the authors' calculation described in the paper's text. Columns (2)-(3) contain the difference-in-differences estimate of the effect of the subsidy on log electricity demand using year and month, and year-by-month fixed effects, respectively. Standard errors are clustered by firm.




(a) Average VAT rate: receive zero-rating



(b) Average VAT rate: lose zero-rating




(c) Electricity demand (log): receive zero-rating



(d) Electricity demand (log): lose zero-rating

Figure 1: Effect of Zero-Rating on Average VAT Rates and Electricity Demand

Notes: This figure presents results from an event study DiD specification estimating the effect of an electricity zero-rating on the average VAT rate paid on electricity and log electricity demand. Panels (A) and (B) present effects on the average VAT rate paid by firms for receiving and losing a zero-rating, respectively. Panels (C) and (D) present effects on log electricity demand. Intervals around point estimates represent 95 percent confidence bands. Standard errors are clustered by firm. The corresponding regression estimates from a pre-/post- analysis can be found in Table 2




(a) Receive zero-rating


(b) Lose zero-rating

Figure 2: Effect of Zero-Rating on Input Credits + Refund Requests (log)

Notes: This figure presents results from an event study DiD specification estimating the effect of an electricity zero-rating on firms' accumulated VAT input credits and VAT refund requests. Intervals around point estimates represent 95 percent confidence bands. Standard errors are clustered by firm. The corresponding regression estimates from a pre-/post- analysis can be found in Table 3.



(a) First stage: subsidy (PKR)



(b) Reduced form: electricity demand

Figure 3: Effect of Electricity Subsidies on Demand

Notes: This figure presents results from an event study DiD specification of the effect of Pakistan's electricity subsidy on prices and electricity demand. Intervals around point estimates represent 95 percent confidence bands. Standard errors are clustered by firm. The corresponding regression estimates from a pre-/post- analysis can be found in Table 5

IX Appendix

A. Adding Delays in Production and Sales to Conceptual Framework

In this subsection, we extend our conceptual framework to include delays in production and sales, which introduces a friction not considered in the model discussed in the main text. Specifically, we consider a scenario where a VAT-compliant firm, which levies taxes on its output and incurs taxes on its inputs, cannot synchronize its sales with its input purchases. This discrepancy necessitates that the firm borrows to fund its input costs, with the borrowed amount increasing due to the VAT applied to inputs.¹⁹ The firm's profit maximization problem is:

(A.1)

$$\begin{aligned} \max_{e,y} \quad & (1 + \tau_0) \cdot p(y) \cdot y - (1 + r) \cdot (1 + \tau_1) \cdot q \cdot e - [\tau_0 \cdot p(y) \cdot y - q \cdot \tau_1 \cdot e] \\ \text{s.t.} \quad & y = f(e) \end{aligned}$$

Here, p represents the price of output, $f(e)$ the production function, q the price of input, e the quantity of input, τ_0 the VAT rate on output, τ_1 the VAT rate on input, and r the interest rate.

In this model, the firm's outcomes are affected by variations in τ_1 , because the government disallows the immediate offset of the additional financing cost incurred due to a positive τ_1 . An elevated τ_1 exacerbates the friction by increasing the cost of borrowing, leading to a reduction in input purchases as a direct response to an increase in VAT. This dynamic serves as a potential explanation for the findings observed in our empirical analysis. To mitigate this friction, the government could implement one of two solutions. Firstly, it could provide immediate reimbursement for the VAT paid on inputs, rendering the firm's decisions unaffected by changes in τ_1 . Alternatively, the government could permit the firm to deduct the additional financing cost in a subsequent period, potentially in the form of a reimbursement, if $[\tau P(y)y - (1 + r)\tau_1 qe] < 0$.

¹⁹This is under the government policy that processes tax credits only subsequent to sales.

We think that this mechanism is not an important confounder. First, the implied increase in the borrowing cost is low. Suppose the annual interest rate is 20%. Monthly r will equal roughly 2%.²⁰ An increase in τ_1 of 0.17, which will be refunded next month, translates to an approximate cost increase of 0.2%. This is a small increase in cost and cannot explain the large changes in electricity purchases that we observe empirically. This assumes that the time-to-build friction is 1 month long and the firm is able to claim tax credits in the next month.

We consider the mechanism outlined above to be an unlikely confounder. The implied increase in borrowing costs is small. For instance, if we assume an annual interest rate of 20%, the corresponding monthly rate, r , would be approximately 2%.²¹ An increase in τ_1 by 0.17, anticipated to be refunded in the subsequent month, would result in an estimated cost increase of just 0.2%. Such a marginal increase is insufficient to account for the significant shifts in electricity purchases we empirically observe. This calculation presumes a one-month time-to-build friction and that the firm can successfully claim tax credits the following month. If we assume that the firm faces significantly longer delays in production, in effect, we assume that the government does not pay what it owes the firm “on time”. In other words, the policy to process tax credits post sales introduces refund delays not as a willful act of the government but as a consequence of the manner in which the VAT regime has been implemented.

The friction highlights the need for the VAT reimbursements to account for increased upfront borrowing costs. In other words, as τ_1 increases, the firms are not being paid back enough. We provide evidence on firms not being paid back the base tax amount, let alone the borrowing cost from a higher purchase price.

²⁰We consider monthly rates because VAT reporting and payments happen on a monthly level in Pakistan and receipts are required to be filed within 30 days, with a grace period of 15 days in exceptional cases.

²¹Monthly rates are relevant here because VAT reporting and payments in Pakistan are conducted on a monthly basis, with filings due within 30 days and a grace period of 15 days allowed under exceptional circumstances.

B. Robustness Checks

Appendix Table A.1 shows that our first stage, reduced form, and 2SLS results are robust across a variety of alternative specifications. Column (1) reproduces the paper’s main estimates from column (1) of Table 2. Column (2) additionally includes firms that moved in or out of an electricity zero-rating multiple times during our sample timeframe. Column (3) includes firm-by-quarter fixed effects to account for potential seasonality in electricity demand. Columns (4)-(6) report results from donut DiD designs, where 1, 2 and 3 month windows around the zero-rating month are excluded from estimation.

In Appendix Table A.2, we use the same DiD regression in equation (4) to show that the zero-rating of sales and non-electricity purchases did not change following a change in zero-rating status for electricity. The outcome variable used is the log of the fraction of sales and non-electricity purchases that are zero-rated. The table indicates that the fraction of zero-rated sales increased by about 2% following a zero-rating (column 1), while the fraction of zero-rated non-electricity purchases remained unchanged (column 2). These results indicate that the increase in VAT input credits and refund requests we observe are likely driven by the electricity zero-rating and not changes in the zero-rating status of other inputs or outputs.

Appendix Tables A.3 and A.4 present the results for VAT and price elasticity using alternative difference-in-differences (DD) estimators, known for their robustness to heterogeneity in treatment effects across cohorts and time. Specifically, we utilize the estimators developed by Callaway and Sant’anna (2021) and Chaisemartin and D’Haultfoeuille (2020). A key concern with traditional two-way fixed effect estimators is the potential aggregation of heterogeneous treatment effects with negative weights, resulting in coefficients that are difficult to interpret. However, the findings in these tables affirm the robustness of our main results, even when accounting for such heterogeneity. They confirm that changes in zero-rating status significantly influence both the VAT rate and electricity expenditure in the anticipated directions. Additionally, they indicate that firms react to changes in electricity prices stemming from the reduced tariff reform. Figures A.1, A.2, and A.3 depict the treatment effects across different event times as derived from these tables. These figures not only corroborate the tabulated results but also provide

further evidence of parallel trends, reinforcing the validity of our findings.

C. Computing Changes in Electricity Prices From Preferential Tariff

This section outlines the process that we use to compute that the preferential tariff regime lowered the unit cost of electricity by 22%. The policy changed the price of electricity to 7.5 cents (USD) per KWh. We first use exchange rate data from the State Bank of Pakistan to compute the rupee equivalent of 7.5 cents per KWh. The electricity price, which is set by the government, has a time of use schedule. We observe the breakdown of electricity consumption by time of use for a limited sample of firms. We first calculate the cost of electricity for these firms under both the preferential tariff and the time of use price schedule. Then, we compute the unit cost under each regime by dividing the total cost by the total units consumed. We then take an average of the percent difference in unit costs under each regime across the firms, which is 22%.

D. Appendix: Tables and Figures

Table A.1: Robustness Checks of the Effect of Zero-Rating on VAT Rates and Electricity Demand

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: VAT rate paid on electricity (first stage)						
Levels (%)	-0.15 (0.01)	-0.15 (0.01)	-0.16 (0.00)	-0.16 (0.00)	-0.16 (0.00)	-0.16 (0.00)
Panel B: Electricity demand (reduced form)						
Levels (Millions of PKR)	0.11 (0.03)	0.16 (0.03)	0.12 (0.03)	0.12 (0.03)	0.12 (0.03)	0.12 (0.03)
Logs	0.13 (0.05)	0.21 (0.04)	0.15 (0.05)	0.16 (0.05)	0.16 (0.05)	0.16 (0.06)
Panel C: 2SLS estimates of VAT rate change on electricity demand						
Electricity demand (log)	-0.81 (0.32)	-1.39 (0.28)	-0.93 (0.32)	-0.93 (0.32)	-1.00 (0.34)	-1.01 (0.36)
Treated Sample						
Allowing for multiple movers	-	✓	-	-	-	-
Firm fixed effects	✓	✓	-	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓
Firm-by-quarter fixed effects	-	-	✓	-	-	-
Donut DiD: 1 month window	-	-	-	✓	-	-
Donut DiD: 2 month window	-	-	-	-	✓	-
Donut DiD: 3 month window	-	-	-	-	-	✓

Notes: This table presents results from a set of robustness checks to the paper's main results in Table 2. Column (1) reproduces the paper's main estimates from column (1) of Table 2. Column (2) includes firms that moved in and out of electricity zero-rating multiple times throughout the sample period in the estimation sample. Column (3) includes firm-by-calendar quarter fixed effects to account for potential seasonality in electricity demand. Columns (4)-(6) contain results from donut DiDs where 1, 2, and 3 month windows around the zero-rating change are excluded from estimation. The number of observations in Columns (1)-(6), respectively, are 294,603; 304,578; 296,395; 290,268; 289,463; 294,908. Standard errors are clustered by firm.

Table A.2: Effect of Electricity Zero-Rating on Fraction of Zero-Rated Sales and Purchases

	Fraction sales (log)	Fraction non-elec purch(log)
	(1)	(2)
Electricity zero-rating	0.02 (0.01)	0.00 (0.00)
Observations	294,586	294,700
R-squared	0.812	0.64

Notes: This table contains estimates of the effect of zero-rated electricity on firms' mix of zero-rating sales and purchase. All regressions include firm and year-by-month fixed effects. Standard errors are clustered by firm.

Table A.3: Robustness Checks - Allowing for Cohort and Time heterogeneity


	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: VAT rate paid on electricity						
Levels (%)	-0.12 (0.01)	-0.12 (0.01)	0.11 (0.01)	0.11 (0.01)	-0.12 (0.00)	-0.12 (0.00)
Panel B: Electricity demand						
Levels (Millions of PKR)	0.14 (0.07)	0.21 (0.08)	-0.10 (0.07)	-0.81 (0.24)	0.74 (0.23)	0.74 (0.23)
Mean (Pre-Treatment)	0.98	0.98	1.76	1.76	1.42	1.42
Pooled	-	-	-	-	✓	✓
Get ZRI	✓	✓	-	-	✓	✓
Lose ZRI	-	-	✓	✓	✓	✓
Only Movers	-	✓	-	✓	-	✓

Notes: This table presents results from a set of robustness checks to the paper's main results in Table 2. Column (1)-(4) contain the DiD estimate of the effect of the zero-rating on VAT rate paid on electricity purchases, and on the PKR amount of electricity purchased, where the estimator used is described in Callaway and Sant'Anna (2021), and the estimates were computed via xthdidregress command in Stata 18. Columns (2) and (4) restrict the sample to those eventually treated, not containing a pure control group. Column (5)-(6) pool the receivers and losers of zero-rating and utilize the DiD estimator in Chaisemartin and D'Haultfoeuille (2020a). The estimates were computed via did_multiplengt command in Stata 18. The pre-treatment means are mean values of the outcome variable in the treated group prior to the treatment, and in case of pooled sample is a weighted average. The number of observations in Columns (1)-(6), respectively, are, 235,839; 7,966; 70,735; 7,792; 281,394; 13,597. Standard errors are clustered by firm.

Table A.4: Robustness: Estimates of the Price Elasticity of Demand for Electricity

	(1)	(2)	(3)	(4)
Elec Demand (1000 KWh)	8.02 (1.91)	9.04 (2.30)	26.0 (5.93)	251 (70.6)
Mean (Pre-Treatment)	80.3	80.3	80.3	310
Firm FE	✓	✓	-	-
MonthxYear FE	-	✓	-	-
Month FE, Year FE	✓	-	-	-
DID C&S (2021)	-	-	✓	✓
Gas	-	-	-	✓

Notes: This table contains estimates of the effect of Pakistan's electricity subsidy on electricity demand. Columns (1)-(2) give DiD estimates from two-way fixed effect regressions and are the levels analog of the estimates shown in table 5. Columns (3)-(4) contain the DiD estimate of the effect of the subsidy on electricity demand, where the estimator used is described in Callaway and Sant'Anna (2021) and was computed via `xthdidregress` command in Stata 18. Column (4) shows the estimates where the sample contains only firms with a known gas connection. Columns (1)-(3) have 51,736 observations whereas column (4) has 4896 observations.



(a) Receive zero-rating

(b) Lose zero-rating

Figure A.1: Effect of Zero-Rating on VAT rate

Notes: This figure presents results from an event study DiD specification estimating the effect of an electricity zero-rating on the VAT rate charged on electricity. Panel (A) represents the effect of receiving zero-rating on the VAT rate. Panel (B) represents the effect of removing zero-rating on the VAT rate. The shaded region around point estimates represent 95 percent confidence intervals for each point. Standard errors are clustered by firm. These figures were generated using `xthdidi` command in Stata 18. The corresponding aggregate (across cohort and time) estimates using Callaway and Santana (2021) estimator is found in Appendix Table A.3, panel A columns (2) and (4).




Figure A.2: Effect of Zero-Rating on Electricity Demand

Notes: This figure presents results from an event study DiD specification estimating the effect of an electricity zero-rating on electricity demand, as measured through purchase value. Panel (A) represents the effect of receiving zero-rating on electricity purchases. Panel (B) represents the effect of removing zero-rating on electricity purchases. The shaded region around point estimates represent 95 percent confidence intervals for each point. Standard errors are clustered by firm. These figures were generated using xthdregress command in Stata 18. The corresponding aggregate (across cohort and time) estimates using Callaway and Santana (2021) estimator is found in Appendix Table A.3, panel B columns (2) and (4).




Figure A.3: Effect of Electricity Subsidies on Electricity Demand

Notes: This figure presents results from a DiD specification of the effect of Pakistan's electricity subsidy on electricity demand. The estimator used is described in Callaway and Sant'Anna (2021) and was computed via `xthdidregress` command in Stata 18 . The shaded region around point estimates represent 95 percent confidence intervals for each point. Standard errors are clustered by firm. The corresponding estimates from a DiD analysis can be found in Table A.4 column (3).