

Detection Without Deterrence: Long-Run Effects of Tax Audit on Firm Behavior

Michael Best, Jawad Shah, and Mazhar Waseem*

July 2021

Abstract

We exploit a program of randomized audits covering the entire population of VAT filers from Pakistan to study how much evasion audit detects and how much evasion it deters by changing behavior. We document substantial evasion at the baseline: almost one-third of firms engage in tax evasion, and conditional on some evasion, the average evasion rate exceeds 40 percent. We find remarkable heterogeneity in evasion by firm size with the evaded amount exceeding the reported liability in the bottom three quartiles but is merely 7 percent in the top. Despite detecting substantial liabilities, audit does not deter tax evasion. Examining more than ten outcomes, we detect no effect of audit on proximate or distant behavior. We offer an explanation of the detection-without-deterrence result and discuss its optimal policy implications.

Keywords: VAT, Tax Evasion, Firm Behavior

JEL Classification: H25, H26, H32

*We thank Miguel Almunia, Zareh Asatryan, Pierre Bachas, Anne Brockmeyer, Giacomo Brusco, Michael Devereux, Bill Hoyt, Anders Jensen, Matthias Kasper, Carlos Lamarche, Joana Naritomi, Maximilian Todtenhaupt, Christian Traxler, and Ron Zimmer for helpful insights. We also thank the seminar participants at the Institute of Fiscal Studies, the University of Mannheim, CESifo Public Economics, IEB, University of Barcelona, NHH Norwegian School of Economics, the University of Kentucky, and the National Tax Association for helpful suggestions and comments. Financial support from the TAXDEV, Institute of Fiscal Studies is gratefully acknowledged. Email Addresses: Michael Best mcb2270@columbia.edu; Jawad Shah jawad.shah@uky.edu; Mazhar Waseem mazhar.waseem@manchester.ac.uk.

I Introduction

Modern tax systems are based on the principle of self-assessment. Taxpayers assess their tax liability and report it to the government through their tax returns. The reported liability is considered final unless the return is selected for audit. Audit thus is typically the sole point of contact between the government and the taxpayer and therefore the sole instrument through which the government can detect and punish noncompliance and create deterrence against it. How much revenue a country collects for the given tax policy hence hinges critically on how effective its revenue authority's audit processes are. [Sarin & Summers \(2019\)](#) estimate that by investing in IRS's audit capacity the US can shrink its tax gap by around 15 percent, generating additional revenue of around \$1 trillion in a decade. Despite the centrality of audit to the tax collection of a country, it has received little attention from public finance researchers. Importantly, we still do not understand fully how effective tax audits are in uncovering tax evasion and creating deterrence against it.

The central challenge in identifying the causal effect of tax audit is that it is endogenous to a taxpayer's evasion choice. Modern tax administrations use sophisticated, risk-based algorithms to target audits toward more egregious occurrences of tax evasion. While such targeting helps them deploy their scarce resources optimally, it makes the job of empirical researchers harder. To the extent that firms select into audit based on their reporting histories, the audit effects are intricately intertwined with the selection effects and there is no natural way to disentangle the two. The causal effect of audit thus cannot be identified unless one is willing to make strong and often arbitrary assumptions. Risk-based assignment further means that audit results do not represent unbiased estimates of the average level of noncompliance in the population. In this paper, we overcome these challenges by exploiting a national program of randomized audits from Pakistan. In distinction to small-scale randomized audit programs many tax administration run to estimate the tax gap, our program covers the entire population of tax filers in the country, and it ran for three consecutive years (2013–2015). Our empirical setting is therefore akin to a randomized controlled trial done at scale (at the national level) and repeated in three consecutive years—an incredibly compelling setup. Using this variation, we are able to identify cleanly the average treatment effect of audit as well as the extent and distribution of tax evasion at the baseline.

The randomized audit program we exploit began in 2013. Before that Pakistan's

revenue authority (FBR) used to assign audits following the standard risk-based approach. This practice, however, was challenged before the superior courts of the country *inter alia* on the ground that the risk criteria used to assign audits were confidential and likely discriminatory against some taxpayers. While these challenges were pending, the FBR could not assign audits using the standard approach and instead had to assign them using random computer ballots. These random ballots were held in public, in the presence of taxpayer representatives. The program covered both income tax and VAT, but in this paper we focus only on VAT audits. The VAT return is filed every month, and the availability of high-frequency data allow us to distinguish between immediate and long-run effects of audit in transparent event-study research designs.

Why should audit have a long-run effect on behavior? In the standard tax compliance model, a taxpayer chooses the amount of tax liability to report trading off the benefit and cost of tax evasion (Allingham & Sandmo, 1972). The cost of evasion here is that with some probability the government would detect evasion and would recover the evaded amount along with a penalty. The true detection probability, however, is unknown to firms, although they would have formed beliefs on it based on their history of interactions with the government and other items in their pre-audit information set. Audit exposes the government's detection technology to firms, forcing them to revise their beliefs on the detection probability. In the simple case, where both the prior belief and the signal contained in audit are Gaussian, the posterior belief is just the weighted sum of the two with the weights provided by the precision of each distribution. This posterior would be different from the prior if (1) the precision of the priors distribution is finite and (2) the precision of the signal is not zero. Under these two very mild and intuitive conditions, audit would lead to a revision of firm priors on the detection probability and hence its future behavior, in particular its tax evasion choice.

This paper has two distinct parts. In the first part, we use audit results to map out the contours of tax evasion in Pakistan—a representative emerging economy. Given random assignment, our estimates are not contaminated by selection and represent unbiased estimates of the level of evasion at the baseline. We have access to administrative tax records that comprise (1) a long panel of tax returns spanning 120 months (July 2008 – June 2018); (2) the tax register that contains important firm characteristics; and (3) audit data that contain both audit characteristics and findings. We find that roughly one-third of firms engage in tax evasion. Conditional on engaging in some

evasion, the evaded tax amounts to nearly 40 percent of the true tax liability. We uncover remarkable heterogeneity in evasion by firm size. The evasion rate is only 7 percent among large firms (in the top quartile of the baseline size distribution) but more than 50 percent among the rest.¹ Strikingly, the distribution of tax payment is even more unequal than the distribution of tax evasion. Our most granular analysis shows that top 1 percent firms contribute more than 94 percent of the tax remitted; the evasion rate among these firms is virtually zero. The data thus point to a peculiar shape of the evasion-size gradient: tax evasion is extremely high at the bottom, decreases monotonically with firm size in the middle, and collapses to zero at the very top. No other determinant of tax evasion discussed in literature exhibits as strong an association with tax evasion as firm size does.

In the second part of the paper, we identify the causal effect of audit on firm behavior. We compare more than ten outcomes reported on the tax return, including sales, input costs, and tax liability, across audited and unaudited firms. None of these outcomes, however, shows any impact of audit either in the short or in the long run. Trajectories of these outcomes continue to evolve on the preexisting trends with virtually no difference between the two groups. Nor is there any heterogeneity in this result. We use both the standard approach and a more flexible, machine-learning-based approach developed in [Athey *et al.* \(2019\)](#) but find no heterogeneity in the treatment effect along any of the more than ten firm characteristics we explore. Importantly, we find a precisely-estimated null effect holds even for the subsample of firms audit found a positive liability against.

Our estimates are robust to multiple specification checks and have broad applicability. Pakistan's revenue authority could not audit all firms drawn in the random ballots. To account for this violation of the experimental protocol, we estimate the LATE parameters corresponding to each specification we mention above. Unsurprisingly, these LATE estimates are indistinguishable from their ITT counterparts. To show that these estimates are not relevant only to compliers but rather broadly to the whole population, we use the marginal treatment effects (MTEs) framework popularized by [Heckman & Vytlacil \(2005, 2007\)](#). Since we have access to a binary instrument only, we identify a linear version of the model following [Brinch *et al.* \(2017\)](#) and [Kowalski \(2016\)](#). The MTE functions we estimate are flat, which rules out significant treatment effect heterogeneity and selection on unobserved gains, showing that our LATE es-

¹Note that an evasion rate in excess of 50 percent means that the evaded amount exceeds the reported tax liability.

estimates have global external validity. The scale of our intervention implies that our estimates are also robust to other external validity concerns small-scale interventions commonly face (see for example [Muralidharan & Niehaus, 2017](#); [Al-Ubaydli *et al.*, 2017](#); [Deaton & Cartwright, 2018](#)). Our sample frame is the universe of VAT filers and our randomized sample includes all audits done in a year. Our results therefore apply to a typical firm with the audit done under typical conditions (managerial oversight, intensity of audit, political economy, etc.).

We find that despite detecting substantial amounts of unpaid liabilities audit does not cause any change in behavior. We propose a simple explanation of this detection-without-deterrence puzzle. Following recent empirical literature ([Pomeranz, 2015](#); [Naritomi, 2019](#); [Waseem, 2020b](#)), we argue that in a VAT the detection probability associated with a transaction depends on how much information it generates for the government. Transactions between arm's length firms create information trails and are thus difficult to hide. On the other hand, transactions between colluding firms or between firms and consumers do not create any information trail and are thus easy to hide. Arranging transactions by the hiding cost they entail, one obtains an S-shaped detection probability function, with the probability of detection on most transactions being close either to zero or to one ([Kleven *et al.*, 2011](#)). In this world, firms in general are too far away from their indifference point of reporting or not reporting a transaction and a marginal audit even when it forces some revision in the firm's belief on the detection probability does not create an observable change in behavior. This process is reinforced further by the fact that amounts detected by audit are not automatically recoverable but rather are subject to adjudication and appeal processes. In weak-state-capacity settings, these process are manipulable and thus firms may not revise their priors on the detection probability significantly even when a large amount is detected against them by audit. We present two pieces of evidence consistent with this explanation. First, we show that both the detected amount and the detection probability fall with the share of final sales (firm to consumer transactions) in a firm's turnover. Second, we show that only a small fraction of the detected amount (2 percent) is paid voluntarily at the time of audit, the rest becoming subject to the adjudication and appeal processes.

Modern tax instruments—personal income tax and VAT, as we note above, rely heavily on audits to deter noncompliance. Ineffective audits can lower the revenue efficiency of these instruments with serious consequences for the optimal tax policy. Importantly, the welfare maximizing instrument mix in settings characterized by in-

effective audits may involve some use of distortionary instruments such as tariffs that though distort production have superior revenue efficiency, a point made in great detail by [Best *et al.* \(2015\)](#). In terms of optimal enforcement policy, our results suggest that the revenue authority should conduct fewer but more intense audits. The focus of the audit program should be to forge future behavior by shifting firm beliefs on the detection probability outwards, thus creating abiding gains in revenue. Audits that fails to do so create large deadweight loss with little upside.

Tax evasion has received renewed research interest in recent years. This revival in large part is driven by the strong association between the fiscal capacity of a state and its economic development reemphasized recently by [Besley & Persson \(2013\)](#) and others. Exploiting random assignment of audits, the first part of this paper contributes a comprehensive analysis of tax evasion in a representative emerging economy. We document the average level of tax evasion in the population as well as its distribution by firm size and other important determinants of tax compliance. In this effort, the paper is similar to [Kleven *et al.* \(2011\)](#) and [Waseem \(2021, 2020a\)](#) who do so in other contexts. The second part of the paper estimates the causal effect of audit on future behavior, contributing to a small strand of literature that includes [Gemmell & Ratto \(2012\)](#), [DeBacker *et al.* \(2013, 2018\)](#), and [Advani *et al.* \(2019\)](#).² These papers exploit the randomized audit programs the IRS and HMRC run to identify the dynamic effect of audit in the US and the UK. Audit's effect on compliance has also been examined by laboratory studies and deterrence message experiments (see [Antinyan & Asatryan, 2020](#) for a recent survey). Importantly, however, in none of the three strands of literature does any consensus exist on the question, and both the sign and the magnitude of the audit effect are open questions. For example, of the four random audits based studies, the latter two find significant positive effects of audit, while the former two report a null effect. Similarly, in some laboratory studies evasion decreases after audit (see [Kirchler, 2007](#) for a survey) but in others it increases (for example [Maciejovsky *et al.*, 2007](#)).³

²The scale and scope of the programs exploited by these studies, however, differ substantially from ours. Random audits are in general not an optimal use of resources for a revenue authority and hence randomized audits exploited by these studies are a small subset of the overall audit program. Furthermore, because the chief purpose of these programs is to estimate the tax gap, randomized audits carried out by the IRS and HMRC are extensive audits. In distinction, our program is a national program that covers all audits done in a year and our audits are routine audits carried out by the revenue authority as a part of their operations.

³The negative effect of audit arises either because it causes a downward revision of the perceived detection probability or because taxpayers irrationally believe that current audit makes them less likely to face future audit, a phenomenon known as the gambler's fallacy ([Gilovich, 1983](#)) or the bomb crater

II Conceptual Framework

This section outlines a simple model that links audit to firm behavior, highlighting the channel through which it may deter noncompliance in future. The framework is based on a version of the canonical tax compliance model (Allingham & Sandmo, 1972) presented in Kleven *et al.* (2011).

II.A Firm Behavior to Taxation

Consider a firm that uses taxable inputs valuing $c(s)$ and nontaxable inputs valuing $\psi(s)$ to produce an amount s of output. The firm is subject to the standard VAT whereby it charges tax at the rate τ of its sales and adjusts tax paid on inputs, facing a tax liability of $T(\tau) = \tau (s - c)$. We assume that the enforcement is imperfect so that the firm can underreport sales $\hat{s} < s$ and overreport input costs $\hat{c} > c$, evading an amount e of its tax liability $e = \hat{T} - T$, where $\hat{T} = \tau (\hat{s} - \hat{c})$.

The government runs an audit program to detect tax evasion, imposing a proportional penalty at the rate θ of the evaded tax liability. The probability the government detects evasion with is $p(e)$ with $p'(e) > 0$ and $p''(e) > 0$. The firm does not know this *true* detection probability and its belief on the probability is denoted by $\tilde{p}(e)$. Based on this belief and other parameters of the tax system, the risk-neutral firm decides how much tax to evade solving the following program

$$(1) \quad \max_e \tilde{p}(e) \cdot \pi^A + (1 - \tilde{p}(e)) \cdot \pi^{NA}.$$

Here $\pi^A = s - c(s) - \psi(s) - \theta\tau e$ and $\pi^{NA} = s - c(s) - \psi(s) + \tau e$ denote the after-tax profits of the firm in the detected and undetected states. The FOC of the problem

$$(2) \quad [\tilde{p}(e) + e \cdot \tilde{p}'(e)] (1 + \theta) = 1$$

implicitly defines the mapping between the perceived detection probability and the evasion choice $e(\tilde{p}, \theta)$. The comparative statics of the problem with respect to $\tilde{p}(e)$ are unambiguous: the evaded amount decreases as the perceived detection probability increases $e(\tilde{p}', \theta) < e(\tilde{p}, \theta)$ for $\forall \tilde{p}' > \tilde{p}$.⁴

effect (Mittone, 2006).

⁴See, for example, (Kleven *et al.*, 2011).

II.B Audit and Belief Updating

Audit is a rare event. We show later that in a typical year the government audits only around five percent of the population, a rate at which a typical firm will experience audit once every twenty years.⁵ Audit thus represents a rare opportunity for the firm to learn the efficacy of government's detection technology, update its beliefs on it, and tailor its future behavior in accordance with the revised beliefs. To see how this process works, assume that the firm's prior belief on the detection probability is a draw from the normal distribution with mean \tilde{p}_t and variance $\sigma_{\tilde{p}_t}^2$. The firm undergoes audit at time t , receiving a noisy signal x_t of the real detection probability

$$(3) \quad x_t = p_t + \epsilon_t.$$

For simplicity, we assume that ϵ_t is also a normal process with $\epsilon_t \sim \mathcal{N}(0, \sigma_{\epsilon_t}^2)$. When both the prior and the signal are Gaussian, the posterior belief is also Gaussian with mean

$$(4) \quad \tilde{p}_{t+1} = \alpha x_t + (1 - \alpha) \tilde{p}_t,$$

and standard deviation $\sigma_{\tilde{p}_{t+1}}^2 = \frac{\sigma_p^2 \sigma_{\epsilon}^2}{\sigma_{\epsilon}^2 + \sigma_p^2}$. The mean posterior belief is a weighted average of the signal and the mean of the prior, with weights provided by the precision of each distribution

$$(5) \quad \alpha \equiv \frac{\frac{1}{\sigma_{\epsilon}^2}}{\frac{1}{\sigma_{\epsilon}^2} + \frac{1}{\sigma_p^2}}.$$

Intuitively, the weight $\alpha \in [0, 1]$ depends on the noise to signal ratio of audit with a more precise signal receiving a higher weight. In the extreme case, when the precision of the signal approaches infinity ($\sigma_{\epsilon}^2 \rightarrow 0$), its weight tends to one and prior beliefs play no role in the formation of posterior on the detection probability. This simple learning model provides intuitive formulation to two conditions under which audit leads to a significant revision of firm priors on the detection probability.

Condition 1. The distribution of prior beliefs is not degenerate $\sigma_{\tilde{p}_t}^2 \neq 0$.

⁵Note that the likelihood of a firm facing an audit is endogenous to firm behavior if the authority runs a parametric, risk-based system of audit selection. The raw audit probability is for illustrative purpose only, showing that on average the authority can only audit one-twentieth of the population each year.

Condition 2. The signal contains some useful information $\sigma_{\epsilon_t}^2 < \infty$.

The first of these condition requires that the firm does not know beforehand the detection probability with certainty. As long as there is some randomness to the audit process, this condition must be satisfied trivially. The second condition requires that the firm gleans some new, credible information from audit. Given that audit is such a rare and intrusive process (see details in the following section), this condition must also hold. To the extent that these conditions are satisfied, they lead to the following result.

Result. If conditions 1 and 2 hold, audit causes a revision in firm beliefs on the detection probability $\tilde{p}_{t+1} \neq \tilde{p}_t$.

The revision of beliefs will in turn reflect in the firm's future behavior via the mapping $e(\tilde{p}, \theta)$. For example, in case of upward revision $\tilde{p}_{t+1} > \tilde{p}_t$, tax evasion will go down $e(\tilde{p}_{t+1}, \theta) < e(\tilde{p}_t, \theta)$ and the firm will remit more tax. To quantify the direction and magnitude of these movements, we define the deterrence value of audit (DV) as the proportional change in tax evasion caused by a marginal audit

$$(6) \quad DV = \frac{e(\tilde{p}_{t+1}, \theta) - e(\tilde{p}_t, \theta)}{e(\tilde{p}_t, \theta)}.$$

We call it the deterrence value because any revision of firm beliefs will impact its behavior not only in the next period but all future periods. In our empirical application, we use the variation created by the randomized audit program to estimate this deterrence value directly from the data.

Note that the functional form of the learning model we assume above plays little role in our key result, although the Gaussian case simplifies the exposition considerably. Importantly, the result will hold in a general setting with $\tilde{p}_{t+1} = f(\tilde{p}_t, x_t)$ as long as the intuitive and trivial Conditions 1 and 2 are satisfied. Neither is it necessary for the firm to be a rational Bayesian learner for the above result to hold. Biased learning due either to mechanical failures of inference (bounded rationality, limited attention, etc.) or to motivated thinking of owners and managers would only mean that the updating may exceed or fall short of the rational benchmark (Bénabou & Tirole, 2016). In either case, it would reflect in the firm's future behavior, albeit to a different degree.

Heterogeneity. Our analysis so far is from the standpoint of a single firm. Our data, however, contain many firms which may be heterogeneous in terms of their prior

beliefs as well as in how they acquire and process information or how this information maps on to their future behavior. Given that audits are randomly assigned in our sample, our empirical results capture an unbiased estimate of the *average* deterrence effect of audit. We, however, run multiple subgroup analyses to uncover any heterogeneity in the treatment effect along these dimensions.

II.C Audit Rate and Detection Probability

In the tax compliance literature, for simplicity the detection probability $\tilde{p}(e)$ is commonly modeled in a reduced form way. But it is important to emphasize that this probability is a composite term comprising the audit rate (the probability that a given firm will be picked for audit) and the detection probability conditional on audit (the probability that the firm’s evasion will be uncovered by audit). Denoting these two terms by $\tilde{p}_a(e)$ and $\tilde{p}_d(e)$, the detection probability $\tilde{p}(e)$ featuring in the behavioral rule (2) can be written as

$$(7) \quad \tilde{p}(e) \equiv \tilde{p}_a(e) \cdot \tilde{p}_d(e).$$

This distinction is particularly important in our setup. Pakistan’s revenue authority explicitly announced before each wave of audits the fraction of the population it intended to audit. With this announcement, the perceived audit rate in the population would converge toward its true value $\mathbb{E}[\tilde{p}_a(e)] \rightarrow p_a(e)$. The second component of the detection probability, however, remains unknown and only firms that undergo audit learn it from their interaction with auditors. Our empirical results therefore isolate the effect of learning the probability of detection conditional on audit, $p_d(e)$, which captures the government’s detection technology on firms’ future behavior.

Of the two components of the detection probability, the existing empirical literature primarily focuses on the first. Several studies manipulate through randomized interventions the firm’s real or perceived likelihood of facing an audit and examine its effects on future tax payments (see for example [Bérgolo et al., 2017](#) or [Slemrod, 2019](#) for a survey). In our setup, however, all firms know the audit likelihood $p_a(e)$, but only a random subset of them learns how likely the audit is to detect their tax evasion $p_d(e)$. This learning as we describe above would lead to updating of their priors, thus shaping the trajectory of their future tax payments.

III Institutional Background

In this section, we document institutional features of the Pakistani environment that are important for our empirical analysis.

III.A Randomized Audit Program

Like all tax authorities, the FBR conducts the audit of a fraction of taxpayers each year. Before 2010, the selection for audit used to take place at the local level with each regional tax office picking taxpayers from their jurisdiction for audit. In 2010, the FBR centralized this process, giving it the power to pick audits for all regional offices using a computer ballot, which could be either random or risk-based (parametric). Exercising these new powers, the FBR picked the first batch of audits using parametric criteria in 2012. The selection, however, was challenged before the superior courts mainly on the grounds that the selection criteria, which were confidential, could be discriminatory against some taxpayers. While these challenges were pending, the FBR could not pick audits using parametric criteria. The legal challenge was not resolved till the end of 2015, and during the intervening period the FBR was constrained to pick audits using random computer ballots. Importantly, random audits in our setting are not a small subsample of total audits, but for three consecutive years (2013–2015) the entire audit program of the country was randomized.

Before each random ballot, the FBR issued an audit policy setting out the fraction of the population to be audited and the criteria for exclusion from the draw. The first information, as we note above, anchors firms' expectations on the true audit rate they face $\mathbb{E}[\tilde{p}_a(e)] \rightarrow p_a(e)$. The criteria for exclusion from the draw were fairly minor in the first two draws, which only excluded government departments and taxpayers already under audit. The third draw, however, also excluded firms under fixed and withholding type regimes of VAT. The required number of cases were picked randomly from the eligible sample (population minus exclusions) after stratifying it by business organization (corporate vs. noncorporate).⁶ The ballots were held in public in the presence of taxpayer representatives, and the list of drawn cases was posted on the FBR portal. The whole process was anonymous and in case was any personal information such as the name or address revealed.⁷

⁶Please see [FBR \(2015\)](#) for details of the randomization procedure, including the set of exclusions.

⁷Both audit policies and lists of drawn cases are public information and have been available on the FBR portal for view and download.

The drawn cases were promptly communicated to local tax offices for initiating audits. Although these audits were conducted by the local offices, the FBR maintained central oversight through the newly developed Taxpayers' Audit Monitoring System (TAMS).⁸ In addition to the centrally assigned audits, local tax offices could initiate audits on their own. But they could do so only in exceptional circumstances, such as when they received specific information on tax evasion, and only after informing the taxpayer in writing the grounds for doing so.

Table I reports descriptive statistics of the five audit waves in our sample. For our empirical analysis we use the first three only, where audit was assigned through the random ballot. The fraction of population picked (p_a) varied across audit waves, ranging between 5 percent and 12 percent. Audits are by nature backward looking. Firms picked in year t are audited for the returns filed in year $t - 1$, although auditors can go back up to five years if necessary.⁹ The FBR did not have the capacity to take up audits of all picked cases, and the actual audit rate in all years remained below 100 percent (70 percent for the first wave and significantly lower in the later). As we note above, local tax offices initiated a small number of audits on their own. These audits are listed in the last column of the table. Our empirical framework takes into account these two violations of the experimental protocol namely that the audit rate remained below 100 percent and that some audits not assigned through random ballots were conducted.

Table II shows audits were initiated soon after assignment. For example, almost 65 percent of those assigned through the first ballot were initiated within one month of the draw. This ratio was even higher for the later waves. Significant underpayment was detected by audits. The distribution of the detected amount, however, is strongly skewed rightward, and the median detection in all three waves is zero. We present a more detailed analysis of the audit findings in section V of the paper.

III.B Pakistani VAT System

Pakistani VAT largely follows the standard design. Firms charge VAT on their sales (output tax) and adjust the VAT paid on inputs (input tax). They remit the tax due

⁸TAMS was the new audit portal of the FBR. All processes related to audit, including all communications to taxpayers, were to be handled through it. This meant the FBR could monitor the progress of audits, compare it across regional offices, and take action in case of delinquency.

⁹For example, firms picked through the 13th September 2013 draw would be audited normally for the twelve returns filed in the tax year 2012 for months July 2011 to June 2012.

(output tax minus input tax) through the tax return, which is filed every month.¹⁰ The filing is based on the principle of self-assessment. Firms assess their own tax liability, which is considered final unless the return is picked for audit. Audit, thus, is the sole instrument through which the revenue authority can detect noncompliance and create deterrence against it.

Pakistan's revenue authority, FBR, is composed of a head office, located in Islamabad, and multiple regional office located throughout the country. These regional offices include four Large Taxpayers Units, two Corporate Regional Tax Offices and twenty Regional Tax Offices. Random audits in our sample were assigned by the head office and were completed at the regional offices. An audit team typically consists of two auditors who report to the local hierarchy. The central audit office, located at the FBR headquarter, exercises overall oversight through the online monitoring system (TAMS). Importantly, all written communications with taxpayers have to be routed through it and are considered invalid unless they contain a bar code issued by the TAMS (FBR, 2015).

Revenue authorities conduct multiple types of audits, which vary in terms of their intrusiveness, such as desk audits or comprehensive audits. All random audits in our sample are comprehensive audits. In each case, the taxpayer was notified, the records were called and examined, and the results were entered into the TAMS.

Like other developing economies, tax evasion is a major issue in Pakistan. In a recent paper, Waseem (2021) estimates an evasion rate of 35-40 percent among the VAT filers of the country. The tax evasion occurs through both undeclared sales and overclaimed tax credits. Given a nontrivial amount is evaded, tax audits have the potential to shift firms' beliefs on the probability of detection outward, creating deterrence against future noncompliance.

In terms of tax evasion and quality of its institutions, Pakistan is not different from other emerging economies. Gómez Sabaini & Jiménez (2012), for example, estimate the VAT evasion rate among a host of Latin American economies. These rates are quite similar to the Pakistan's.¹¹ Similarly, Pakistan's score on the Ease of Doing Business (59.51) is indistinguishable from the average (59.06) of all countries excluding the High Income ones (World Bank, 2020).¹² Nor is Pakistan an atypical country in

¹⁰Some small firms in some of the periods included in our sample were allowed to file on a quarterly rather than monthly frequency.

¹¹For example, the VAT evasion rates of Guatemala, Nicaragua, Panama, and Peru are 37.5 percent, 38.1 percent, 33.8 percent, and 37.7 percent. These are within the range for the Pakistan's estimate.

¹²The Ease of Doing Business score is widely used as a measure for the quality of institutions of a

terms of its tax morale: its score on the tax morale question in the World Value Survey is in fact better than the world average (Haerpfer *et al.*, 2020).¹³

III.C Data

We use administrative data from Pakistan that include the universe of VAT returns filed between July 2008 and June 2018. The VAT return consists of three main sections. In the first section, firms report the value of their sales, decomposing it into its foreign (exports) and domestic components. In the second section, the value of purchased inputs are reported, divided likewise in the two parts. In the final section, firms compute their tax liability, indicating the tax charged on sales, the tax credited on inputs, and the difference between the two—the tax payable. Since 2011, firms also report the transaction-level details of their sales and purchases. Each firm is assigned a unique ID and is required to file every month. The data, therefore, have a panel structure.

In addition to the return data, we use information on firm characteristics from the tax register. This information includes the business organization of the firm (corporate vs. noncorporate etc.), its date of registration, and other variables we use in our heterogeneity analysis. Appendix A.1 provides a complete list of these variables.

Finally, we use audit data available on the FBR portal and the TAMS. As we note above, the list of cases drawn in each computer ballot is publicly available. We download it from the FBR portal and merge it with our VAT return data using the unique firm ID. We are able to merge 43,465 out of 43,625 audits in our sample. For the remaining 218 cases, the firm ID mentioned in the list is incorrect. We add the audit information from the TAMS to this dataset. This information includes the date the audit was initiated, the type of audit (randomly assigned vs. locally assigned), and the amount detected.

country (see for example Besley & Persson, 2014).

¹³We refer to the Question 180 on the World Value Survey 2017-2021. The question asks respondents if “Cheating on taxes if you have a chance” is justified, with responses varying from 1 (never justifiable to 10 always justifiable). Pakistan’s average score on the question is 1.967, which is better than the world’s average of 2.197.

IV Empirical Strategy

One of our empirical goals in this paper is to estimate the deterrence value of audit defined in equation (6). Since the VAT can be evaded by underreporting sales ($\hat{s} < s$) or overreporting input costs ($\hat{c} > c$), the DV in our setup takes the following form

$$(8) \quad DV = \frac{\hat{s}(\tilde{p}_t, \theta) - \hat{s}(\tilde{p}_t, \theta)}{\hat{s}(\tilde{p}_t, \theta)} - \frac{\hat{c}(\tilde{p}_t, \theta) - \hat{c}(\tilde{p}_t, \theta)}{\hat{c}(\tilde{p}_t, \theta)}.$$

We can compute the two terms on the RHS by estimating how reported sales and input costs respond to a tax audit, running regressions of the following type

$$(9) \quad y_i = \alpha + \beta \text{assign}_i + \text{corporate}_i + \epsilon_i,$$

where y_i is the log of reported sales or input costs, assign_i denotes that firm i 's audit was assigned through a random ballot, and corporate_i is a dummy indicating that the firm is a corporation. For space consideration, we sometimes denote the assign_i dummy simply as Z_i . Since audits in our sample are assigned randomly on stratified corporate and noncorporate samples, $\hat{\beta}$ from these regressions identifies the causal effect of interest. But most of our results are from the parallel difference-in-differences model

$$(10) \quad y_{it} = \mu_i + \gamma \text{assign}_i \times \text{after}_t + \lambda_t + \epsilon_{it}.$$

Note that the *corporate* dummy—being time invariant—is absorbed by the firm fixed effect here.¹⁴ This DD model offers us greater transparency (visual event-study results) and precision. We cluster standard errors at the firm level, but in some specifications we cluster at the tax office level as robustness check.

The coefficient $\hat{\gamma}$ from above model identifies the intention-to-treat effect (ITT). We also estimate the corresponding LATE parameter by instrumenting audit with initial random assignment. With treatment effect heterogeneity and selection on the unobserved gain, the LATE is informative only about the average effect on compliers (Imbens & Angrist, 1994). Compliers are an interesting population in our setup. They are the firms the tax authority would audit whenever they have spare audit capacity available. Notwithstanding the policy-relevance of LATE, we are also inter-

¹⁴The tax code requires a firm that changes its business organization from non-corporate to corporate and vice versa to re-register. Upon re-registration, a new identifier is issued to the firm.

ested to know the average effect among the population. For this reason, we estimate the marginal treatment effect (MTE) of audit following the framework developed in Heckman & Vytlacil (2005, 2007). Because we have access to a binary instrument only, we cannot identify the MTE nonparametrically and do so assuming a linear functional form (Brinch *et al.*, 2017; Kowalski, 2016).

Table III runs balance tests on our baseline data. We compare ten VAT outcomes and ten firm characteristics across firms drawn in a given random ballot ($Z_i = 1$) with others using model (9). The compared groups are very similar for the first two waves: the difference in means is almost always insignificant or trivial. This, however, is not true for the third wave. Firms drawn in this wave, for example, are on average larger and more likely to be manufacturers. These differences are unlikely to have arisen by chance. We have noted in section III.A that exclusions from the draw were significantly expanded for the third wave. Importantly, firms under fixed and withholding regimes were excluded from audit. We do not identify these firms in our data and are thus unable to replicate the sample used for the random ballot of the third wave. For this reason, we focus solely on the first two waves for our empirical results. Nonetheless, for the sake of completeness we always present our main results for the third wave as well.

V Tax Evasion at the Baseline

Audits we consider are randomly assigned. The amount detected by them therefore represents an unbiased estimate of the extent of tax evasion at the baseline. In this section, we document the tax evasion rate implied by audit findings, examining in particular its relationship with firm observables.

Table IV presents the results. All amounts in this table are in PKR billions. The top row shows that 3,482 firms were audited in the first wave. These firms reported a total turnover of around 500 billion in the baseline year. The audits detected 2.15 billion of unpaid liability against them, which constitutes around 0.45 percent of the turnover. These firms remitted 28.16 billion of VAT at the baseline with an average effective tax rate of 5.65 percent (columns 5–6). The unpaid revenue therefore amounts to nearly 7 percent of the *true* tax liability (column 7).¹⁵

¹⁵We report the evasion rate as a fraction of the true tax liability. Specifically, for a given bin b the evasion rate we report is calculated as $Evasion\ Rate_b = \sum_{i \in b} \frac{T_i^d}{T_i^d + T_i}$, where T_i^d is the amount detected

The next five rows of the table explore any variation in the evasion rate across different subpopulations. The second row shows that positive liability is detected against 28 percent of firms. The detected amount equals around 67 percent of the VAT remitted by these firms in the baseline year, which translates into an evasion rate of 40 percent. The next four rows divide firms into four quartiles based on their annual turnover in the baseline year. Strikingly, the detected amount exceeds the reported tax liability for all the bottom three quartiles, implying an evasion rate in excess of 50 percent. In contrast, the evasion rate is only 6 percent in the top quartile. The top-quartile firms also contribute disproportionately to the tax revenue. Of the 28.16 billion VAT remitted by all audited firms included in the sample here, more than 99 percent (27.91 billion) was remitted by them. We find qualitatively similar results for the second audit wave, although the evasion at the top is even lower for this wave.

Figure I examines the relationship between tax evasion and firm size in more detail. We divide audited firms into equal-sized bins based on their annual turnover in the baseline year and see how the evasion rate and tax payments vary along the size distribution. Panels A through D successively decrease the bin size, showing progressively granular analysis. Figure A.I repeats this analysis but shows along the right y-axis the average effective tax rate paid by firms at the baseline instead of their relative contribution to the revenue. We find a curious shape of the evasion-size gradient: tax evasion is particularly high at the bottom, with the evasion rate exceeding 80 percent at the 20th percentile of the distribution; it then declines almost monotonically before collapsing discontinuously to a trivial level at the very top. One other striking feature of the analysis is the remarkable concentration of government revenue at the top (see the right y-axis). Almost entire government revenue is contributed by few large firms (in the top 2.5 percentile of the size distribution), relative to whom the contribution of other firms is negligible.

Table A.I explores heterogeneity in evasion rate across other firm characteristics commonly discussed in literature in relation to tax evasion. The key takeaway from this exercise is that firm size remains by far the most important determinant of tax evasion. In comparison to it, other firm characteristics, including the firm's position in the production chain, its location, age, or industry, are either not associated with tax evasion or their association with evasion is much weaker than that of firm size. These results are consistent with the recent models of tax compliance in weak enforcement

against firm i and \hat{T}_i is the reported tax liability of firm i .

setting, which predict a strong, decreasing relationship of tax evasion with firm size;¹⁶ although to our knowledge we are the first to document this curious relationship empirically. Large firms tend to have transparent accounting mechanisms within the firm. These mechanisms let them operate at their economically optimal scale but render commonly used strategies to evade taxes—such as cash payments or keeping double books of account—*infeasible*.¹⁷ Tax evasion as a result decreases with firm size, with large firms ending up remitting a disproportionately oversized share of revenue.

In the audit data, the detected amount is reported in six heads. Table A.II decomposes the detected amount into its major heads. Less than 2 percent of the detected amount is recovered at the time of audit either by direct payment (column 2) or by curtailing the taxpayer’s refund claim (column 7). The rest of the amount is not paid voluntarily by taxpayers and can be recovered by the revenue authority only after quasi-judicial determination of the payable amount. Under these proceedings, taxpayers are given the option to contest audit findings and the decisions are subject to appeal before tax tribunals. We do not have data on the outcome of these processes but anecdotal evidence suggests they are cumbersome and inefficient so that the detected amount remains stuck in litigation for a long time.¹⁸

Although audits in our sample were randomly assigned, the audit rate for both waves remained below 100 percent. If audits were targeted toward specific firm types, selection resulting from it could bias the evasion rates we report above. Figure A.II explores such selection, examining if firms audited earlier were systematically different from those audited later. We find no systematic correlation between the amount detected and the order in which audits were taken up. Nor is the order correlated with other firm observables (see Table A.III). A much detailed analysis of selection appears later in the paper. We find no evidence of such selection: within the randomly assigned sample, audits do not appear to target any specific subgroup. To this extent, our estimates represent unbiased estimates of noncompliance at the baseline.

Tax audits are unlikely to uncover all tax evasion. For this reason, revenue author-

¹⁶See for example Kleven *et al.* (2016); Gordon & Li (2009); Kopczuk & Slemrod (2006).

¹⁷Without strong internal controls, firms cannot grow beyond a given scale as they may worry about pilferage and stealing by local managers.

¹⁸According to a recent press report a total of 76,700 cases involving a recoverable amount of PKR 1.77 trillion are stuck in litigation. Nearly two-thirds of the litigated amount (PKR 1.1 trillion) is pending internally (at the two appeal fora available within the FBR) and the rest with the superior courts of the country. For details of these numbers see [here](#).

ities that use random audits to estimate the tax gap multiply the detected amount by a scale factor to convert it into their official estimate of the tax gap. IRS, for example, used to apply a scale factor of 3.28 for this purpose.¹⁹ This factor was derived from a direct survey of taxpayers on tax compliance (see IRS, 1996; Kleven *et al.*, 2011 for details). We do not have access to such a multiplying factor for the case of VAT in Pakistan. Nor are audits in our sample *extensive* audits, done for the express purpose of estimating noncompliance. They rather are routine audits revenue authorities conduct during the course of their normal operation. Our estimates therefore likely represent a conservative lower bound on the true evasion rate in Pakistan.

VI Audit and Firm Behavior

We now examine the effects of audit on firm behavior, assessing in particular if they deter tax evasion in future periods.

VI.A ITT Estimates

We begin by presenting nonparametric evidence. Figure II plots the coefficients δ_j s from the following regression

$$(11) \quad y_{it} = \mu_i + \sum_{j=2}^N \delta_j \cdot 1.(month=j)_t + u_{it},$$

where y denotes the log of variable indicated in the title of each panel. The regression is run separately for firms drawn in the random ballot ($assign_i = 1$) and other firms in the sample ($assign_i = 0$).²⁰ We drop the dummy for the first month (July 2008) and plot coefficients on the other month dummies (up to June 2018). Figure III illustrates the DD version of these plots, where we add interactions of the month and *assign* dummies into (11) and plot the coefficients on these interactions along with the 95 percent confidence intervals around them. Given the drawn firms are a random sam-

¹⁹The IRS has now replaced the scale factor with a more sophisticated econometric-based algorithm called Detection Controlled Estimation (DCE). The algorithm adjusts the amount detected by random audits to account for income that taxpayers do not report on their tax returns and the income IRS auditors do not detect (IRS, 2012).

²⁰The sample here includes all firms other than government departments and firms already under audit. Both categories of excluded firms together constitute a small (<5 percent) fraction of the $assign = 0$ sample.

ple of the population, it is unsurprising that the trajectory of treated and untreated outcomes is indistinguishable from each other in the 62 pre-draw months. Table A.IV shows this formally by estimating baseline trends using model (10).

Strikingly, however, the outcomes continue to evolve on the common, preexisting trend even in the post-draw period. The relative difference between the two groups remains indistinguishable from zero in the 70 post-draw months we consider. Figures IV and V replicate this analysis for the second draw, showing similar results. Initial evidence thus suggests that audit does not cause significant revision in firm priors on the detection probability and thus does not induce a significant change in behavior. Below, we examine this result in more details by running formal, regression-based tests.

The top panel of Table V reports our ITT estimates from model (10). We examine both short- (one-year) and medium-run (three-year) impacts produced by the audits assigned in the first wave. Consistent with the visual evidence none of the ten coefficients differs significantly from zero at the conventional level. Nor is there any systematic difference between the proximate and distant responses. Table VI repeats the exercise for the second wave. Tables VII–VIII examine six other VAT outcomes, and Table A.V clusters at the tax office level. All these 46 specifications—covering ten intensive margin outcomes, one extensive margin outcome, and two audit waves—tell a consistent story: audit does not have a meaningful impact on firm behavior, either in the short or in the long run.

VI.B LATE Estimates

Since the FBR did not conduct audit of all cases drawn in the random ballots, the above estimates capture the average effect of getting *picked* for audit rather than the average effect of audit. To compute the latter parameter, we estimate the 2SLS models corresponding to (10), instrumenting the endogenous variable *audit* by the initial random assignment.²¹ Table A.VI reports the first stage of these regressions, illustrating that a strong first stage exists in this setting. The bottom panels of Tables V–VIII and A.V report the LATE estimates for the 46 specifications we run. The results are similar. The majority of the LATE estimates are of negative sign, statistically insignificant, and economically trivial.

Figures A.III–A.IV and Table A.VII examine the third wave of audits, reporting

²¹For brevity, we sometimes denote *audit*_{*i*} variable simply as *D*_{*i*} in the subsequent sections.

parallel results comprising the ITT and LATE estimates. Recall that for this wave the balance tests reveal significant differences between $Z_i = 1$ and $Z_i = 0$ groups (see Table III). We therefore do not draw any conclusion from these results and produce them only for the sake of completeness.

VI.C ATE Estimates

When treatment effects are heterogeneous and there is selection into treatment on the unobserved gain, the LATE is informative on the average effect of the treatment on compliers only (Imbens & Angrist, 1994; Abadie, 2003). Compliers, in our setting, are firms that are pushed into audit by the instrument (being drawn in the random computer ballot). The LATE we identify therefore may not reflect the average effect in the population unless the impact of audit does not vary across firms or auditors do not target specific firms, using information we do not observe.

We first explore the latter point, examining if auditors target selective types of firms. Table IX compares audited and unaudited firms.²² Audited firms here include both that were picked by a random draw ($Z_i = 1$) and that were picked by local tax offices based on their information ($Z_i = 0$). Tables X-XI separate the analysis for the two subgroups. A typical audited firm indeed differs from the unaudited in terms of observables we examine (Table IX). But these differences are almost entirely driven by the small group of firms local tax offices picked for audit on their own ($Z_i = 0$). Within the random-assignment group ($Z_i = 1$), audits do not seem to target any selected subgroup. Figures VI-VII compare audited and unaudited firms in our event study framework (11). Since the specification includes firm fixed effects, the results capture any residual selection into audit which is not explained by the firm's fixed characteristics, such as size or industry. There does not appear to be any such residual selection as the reporting histories of both groups are similar. Table A.VIII establishes this rigorously by running formal tests on the baseline data. Parallel trends for a long preaudit period mean our DD estimator remains internally valid and applies to all audited firms rather than compliers only.

The above result is supported by our two previous results. First, the compliance rate falls from 70 percent in the first audit wave to 30 percent in the second, yet we see no meaningful difference between the corresponding LATE estimates (compare

²²Since audits were done at the local tax office, we need to compare audited and unaudited firms within a tax office to rule out selection. We therefore include tax office fixed effects into these regressions.

Tables V and VI). This suggests that the marginal firm pushed into audit may not be significantly different from others within the randomly assigned ($Z_i = 1$) sample. Second, the amount detected and other firm observables bear no correlation with the order in which audits were taken up (Figure A.II and Table A.III). This suggests that audits are not systematically targeted toward specific group of firms. Auditors do not seem to possess any privileged information to do so.

Continuing our effort to go beyond LATE, we next exploit the marginal treatment effect (MTE) framework popularized by Heckman & Vytlačil (1999). Since our instrument is binary, we cannot identify the MTE function nonparametrically and instead identify a linear version of it following Brinch *et al.* (2017) and Kowalski (2016). Figures VIII–IX show the MTEs we estimate using the two randomization waves as instruments. The technical details of the estimations are in Appendix A.2. Importantly, the MTEs from all specifications are flat. The change in the unaudited outcomes as the potential fraction audited increases reflects selection. On the other hand, the gradient in the audited outcomes reflects selection and audit effect heterogeneity. That both these curves are flat rules out these factors in our setup. Note that the functional form assumption we make is not too restrictive. We have access to two randomized experiments and therefore can exploit more information than is typically available in an RCT. Specifically, because the compliance rate varies between the two waves, both audited and unaudited outcomes in our setup are identified at four rather than two points. The flat MTEs we obtain from all specifications therefore suggest that our LATEs have global external validity.

VI.D Heterogeneity

To strengthen the above conclusion, we also examine treatment heterogeneity directly. We do so using two nonparametric approaches. First, we estimate triple-difference versions of model (10), interacting the DD term with firm traits. We explore eight traits introduced into the model as dummies indicating (i) firm size; (ii) firm age; (iii) firm location; (iv) local tax office having jurisdiction over the firm; (v) the type of local tax office (LTU vs. RTO etc.); (vi) firm’s position in the supply chain (manufacturer vs. wholesaler etc.); (vii) firm’s business organization; and (viii) industry the firm operates in. All these traits are measured at the baseline before the announcement of ballot results, and we estimate the model separately for the two audit waves. Figures A.V–A.XII display the results. We do not find any systematic treatment effect

heterogeneity across the subgroups we compare. The 95 percent confidence interval almost always includes zero, showing that the response of each subgroup is statistically indistinguishable from that of the omitted category.

In addition to the predetermined firm traits, we also explore heterogeneity by the timing and outcome of audit. Figure A.XIII divides audited firms into ten groups, depending upon the time lag between the assignment and initiation of audit. If auditors have hidden information they use to target specific subgroups, it would be reflected in the order they took up the assigned audits in. We, however, do not see any differences along this dimension. Audited firms in all deciles appear to be very similar. Table A.IX stratifies the audited sample by the detected amount, looking for any differential effect upon firms auditors did find an underpaid amount against. Here also we do not find any differential effect.

Finally, we explore treatment heterogeneity using a more flexible machine-learning approach. We ask if the audit effect varies with the firm's predetermined traits using the Generalized Random Forest algorithm developed in *Athey et al. (2019)*.²³ To reduce the computational demands of the algorithm, we use the simple difference-in-means model (9) as the baseline rather than the DD model (10) we have been using so far. The results are in Figures A.XIV-A.XXIII. The first four of these figures show the audit effect does not vary with firm size or age. The rest of the figures explore binary traits. Again, we do not find any systematic heterogeneity in the audit effect along any of the eight traits we look at.

VII Detection Without Deterrence

We present extensive evidence above showing that despite detecting a substantial amount of tax evasion audits have no significant impact on firm behavior. Not only does this finding hold on average but also among more than 20 subgroups we define based on firm observables. In terms of our model, it suggests either that audit does not cause any revision in firm priors or that the revision it causes does not affect firm behavior. The former possibility is unlikely given that the two conditions (Conditions 1 and 2) under which audit must cause some revision in firm priors are

²³In the approach, individual trees are grown by greedy recursive partitioning of the sample space, with each split chosen to improve the model fit. The trees are then randomized using bootstrap aggregation, whereby each tree is grown on a different random subset of the training data, and random split selection that restricts the variable available at each step of the algorithm.

quite trivial and unlikely to fail in this setting. The failure of the first condition, for example, implies that prior beliefs of all firms are concentrated on the true detection probability with no variance around the mean.²⁴ The failure of the second condition is also equally unlikely. Not only is audit a rare event,²⁵ it is quite intrusive as well. Auditors spend considerable time with taxpayers going through their records, visiting their premises, and discussing their audit findings. It is therefore highly unlikely that taxpayers do not glean any useful information on the government’s detection technology during this lengthy and intrusive interaction. No updating in either directions and consequently no reoptimization of future behavior is puzzling. In this section, we make sense of this detection-without-deterrence puzzle.

We begin by tweaking the model we presented in section II slightly. Following [Basri et al. \(2019\)](#), the revised model treats evasion as a discrete rather than the continuous choice. Discretizing the choice variable brings the model closer to our VAT setting, leading to simpler and more intuitive exposition. The firm engages in L transactions, indexed by $l = 1 \dots L$, and decides separately for each transaction whether to report or hide it. It would report a transaction and remit the VAT due if the cost of hiding the transaction exceeds the benefit from doing so

$$(12) \quad \left[\tilde{p}_l(e_l) + e_l \cdot \tilde{p}'_l(e_l) \right] (1 + \theta) > 1.$$

This inequality is a discrete version of the behavioral rule (2), showing that the firm’s choice critically hinges on the detection probability hiding a transaction entails. Ordering transactions in terms of the detection probability and hence the hiding cost, we can define L^* as the first transaction for which the above inequality holds. The firm will accordingly report transactions $L^* \dots L$ and will remit the tax due, amounting to $\int_{L^*}^L \tau(s_l - c_l) d(l)$. Note that L^* could be the first transaction, in which case the firm does not evade at all, or it could be the last, in which case the firm evades the entire tax due. In general, L^* would be idiosyncratic to firms, depending on their scale, production technology, trading network, and other characteristics.

²⁴This is extremely unlikely as existing evidence shows that taxpayers misperceive even the most simple and accessible details of the tax system, such as the marginal tax rates ([Taubinsky & Rees-Jones, 2017](#)). How can then they be expected to know something that has not been revealed and that too with certitude.

²⁵During the ten-year period we consider, the FBR could not audit more than 5 percent of firms a year, a rate at which a typical firm would experience audit once every twenty years. Note that the likelihood of a firm facing an audit is endogenous to firm behavior if the authority runs a parametric, risk-based system of audit selection. The raw audit probability is for illustrative purpose only, showing that on average the authority can only audit one-twentieth of the population each year.

Given the input-output linkages the VAT creates across firms, the costs of hiding a transaction would vary substantially depending upon who the two parties to the transaction are. Specifically, hiding a transaction would be easier for the reporting firm if the other party is (1) a consumer, (2) an unregistered firm, or (3) a firm willing to collude. In these cases, the firm can cover its tracks, making it harder for the government to detect evasion. On the other hand, hiding a transaction would be difficult if the other party is unwilling to collude, such as a firm that cannot handle unaccounted cash and therefore cannot keep a transaction out of books.²⁶ The $\tilde{p}_l(e_l)$ faced by the firm on different transactions is therefore likely to have the shape shown in Figure X. It would be typically low for the former type of transactions but would rise sharply once transactions of the latter type begin. Such an S-shape detection probability function was first suggested by Kleven *et al.* (2011) and has since then been confirmed in other empirical settings (see Waseem, 2020a for one such example). The shape reflects that the probability of detection to a first order depends on the external information an economic transaction generates for the government.

The discrete choice model predicts a simple behavioral rule. The firm will report transactions entailing high detection probability $[L^*, L]$, hiding the rest. In this world, it is easy to see why audit may not cause any *observable* change in future behavior. For this purpose, let us characterize a marginal audit as the following.

Definition. An audit is *pivotal* if it leads to the flipping of inequality (12).

A pivotal audit causes sufficiently large revision in the firm's perceived detection probability so that the LHS of inequality (12) exceeds the RHS after the audit if it did not do so earlier and vice versa. For example, indexing the post- and pre-audit variables by $t + 1$ and t , an audit will be pivotal if

$$(13) \quad \begin{aligned} & \left[\tilde{p}_{l,t}(e_{l,t}) + e_{l,t} \cdot \tilde{p}'_{l,t}(e_{l,t}) \right] (1 + \theta) < 1 \\ & \left[\tilde{p}_{l,t+1}(e_{l,t+1}) + e_{l,t+1} \cdot \tilde{p}'_{l,t+1}(e_{l,t+1}) \right] (1 + \theta) > 1. \end{aligned}$$

In this case, the transaction l will not be reported prior to audit but will be reported after it. Thus, a necessary condition for audit to cause an observable change in firm behavior is that it is pivotal.

²⁶These consideration can lead to segmentation of firms into good and bad VAT chains with compliant firms dealing with compliant firms only and vice versa. See de Paula & Scheinkman (2010); Gadenne *et al.* (2019); Gerard *et al.* (2019) for empirical evidence on market segmentation caused by a VAT.

If the structure of the detection probability is of the form shown in Figure X with the probability on most transactions being close either to zero or one and therefore too far away from the indifference point of inequality (12), it is highly unlikely that a given audit will be pivotal. Indeed, in this world even when Conditions 1 and 2 are satisfied so that audit does cause a revision in firm priors, it still may not cause an observable change in firm behavior if such revision is not large enough to flip inequality (12). A less obvious but related point is that detection alone is insufficient to cause a revision in firm priors. The above model implicitly assumes that uncovering a transaction (detection) automatically leads to the recovery of the tax due and penalty from the taxpayer. But we have seen in section V that it is not the case as only 2 percent of the detected amount is recovered upon audit. The rest of the amount is contested by firms and can be recovered only once it passes the adjudication and appeal processes stipulated in the tax code. Thus even large detection may not force firms to revise their priors on the detection probability if they feel that courts are unlikely to enforce the detected amount against them.

Our preferred explanation of the detection-without-deterrence puzzle is therefore the following. The S-shaped detection probability function means that small changes in firm priors are less likely to be consequential, triggering an observable change in reporting behavior. This process is further reinforced by weak state capacity, whereby the amount detected by audit is subject to leaky, unpredictable judicial processes and hence less likely to create deterrence against future noncompliance. Note that these two factors are flip sides of the same coin. A corollary of the S-shaped detection probability is that the government possesses little information on the unreported tax liability and therefore has limited ability to enforce it through courts.

What evidence can we bring to bear to support this explanation? We begin by showing that consistent with our reasoning (see Figure X) both the detected amount and the probability of detection fall with the share of final sales in a firm's turnover (see Tables A.X and A.XI). Final sales are transactions where the other party is either a consumer or an informal firm and the negative correlations we document in the two tables are accordingly a direct confirmation of the S-shaped detection probability function. Importantly, these correlations are robust to controlling for other major determinants of tax compliance including firm size. Our preferred explanation is further confirmed by the breakdown of the detected amount we presented in Table A.II. As we note above, 98 percent of the detected amount is not paid by firms at the time of audit, becoming subject to adjudication and appeal processes. According to our

model above, the government possesses little external information on this component of the tax base and it is therefore not surprising if firms believe they can get the detected liability reversed in these proceedings.²⁷

What consequences do our results have for optimal policy design? A straightforward implication of our results is that the revenue authority should conduct fewer but more deterrence-focused audits. Detection that does not lead to deterrence serves little purpose. A vast chunk of the detected amount remains unpaid and the resulting litigation wastes precious resources of the revenue authority and taxpayers—a pure deadweight loss from the society’s point of view. Any upward shift in firm priors as a result of fewer but more-intense audits would generate revenues in both current and future periods not only among the audited firms (direct effect) but also among the unaudited firms (spillovers), creating thereby an abiding deterrence against noncompliance. Modern tax instruments—personal income tax and VAT, as we note above, rely heavily on audits to deter noncompliance. Ineffective audits can lower the revenue efficiency of these instruments relative to the first-best substantially. This would have important consequences for the optimal tax mix in settings characterized by ineffective audits. Importantly, the welfare maximizing instrument mix may involve relying on some distortionary instruments such as tariffs that though distort production have better revenue efficiency, a point made in great detail in *Best et al. (2015)*.

VIII Conclusion

We exploit a national program of randomized audits from Pakistan to identify the causal effects of audit as well as the extent and distribution of tax evasion at the baseline. Nearly one-third of firms engage in some tax evasion. Conditional on evading, the evaded tax amounts to nearly 40 percent of the true tax liability. There exists remarkable heterogeneity in evasion by firm size. The evaded amount exceeds the voluntarily-reported tax liability for firms in the bottom three size quartiles but is trivial (6 percent) in the top. No other determinant of tax evasion matters as much as firm size does. Despite detecting substantial amounts of tax evasion, audit has no effect on future behavior. We examine more than ten outcomes but in no case does

²⁷Note that paying the detected amount forthwith strictly dominates delaying it if the firm believes the government would be able to prove its case before courts. In addition to incurring legal costs, contesting the detected amount exposes the firm to the payment of interest, which is charged at a higher rate than the usual rate at which firms can get credit from the financial sector.

the evolution of the outcomes undergo any change at the time of audit. Our empirical setup is compelling (random assignment at the national level) and our data rich. The null result is therefore cleanly identified and is robust to usual identification and inference concerns.

Our results suggest that firms face an S-shaped detection probability function. The probability of detection on most transactions is close either to zero or to one. For such transactions, firms are in a corner solution, being too far away from the point where the marginal cost of reporting a transaction equals the marginal benefit. In this world, a small revision in firm beliefs on the detection probability is unlikely to produce an observable change in outcomes reported on the tax return. We show that both the detection probability and the amount detected fall with the share of final sales in a firm's turnover, thus providing direct evidence in support of the S-shaped detection probability. Our results suggest that audits which detect large recoverable liabilities but still do not create any deterrence are unlikely to be optimal because there is no dynamic gain in revenue to compensate the effort that goes into recovering liabilities pointed by audit. Optimal audit plan potentially involves fewer but more intense audits that shift firm priors to a degree that reporting transactions that otherwise would have gone unreported becomes optimal. We do not observe the deadweight loss created by audit, in particular arising from the legal costs incurred by both parties in prosecuting the contested liabilities, and thus cannot determine key features of the optimal audit plan in such setting, an exercise left for future work.

References

- ABADIE, ALBERTO. 2003. Semiparametric instrumental variable estimation of treatment response models. *Journal of econometrics*, **113**(2), 231–263.
- ADVANI, ARUN, ELMING, WILLIAM, & SHAW, JONATHAN. 2019. The dynamic effects of tax audits. IFS working paper.
- AL-UBAYDLI, OMAR, LIST, JOHN A., & SUSKIND, DANA L. 2017. What Can We Learn from Experiments? Understanding the Threats to the Scalability of Experimental Results. *American Economic Review*, **107**(5), 282–286.
- ALLINGHAM, MICHAEL G., & SANDMO, AGNAR. 1972. Income Tax Evasion: A Theoretical Analysis. *Journal of Public Economics*, **1**, 323–338.


- ANTINYAN, ARMENAK, & ASATRYAN, ZAREH. 2020. Nudging for tax compliance: A meta-analysis.
- ATHEY, SUSAN, TIBSHIRANI, JULIE, WAGER, STEFAN, *et al.* 2019. Generalized random forests. *The Annals of Statistics*, **47**(2), 1148–1178.
- BASRI, M CHATIB, FELIX, MAYARA, HANNA, REMA, & OLKEN, BENJAMIN A. 2019. *Tax Administration vs. Tax Rates: Evidence from Corporate Taxation in Indonesia*. Tech. rept. National Bureau of Economic Research.
- BÉNABOU, ROLAND, & TIROLE, JEAN. 2016. Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, **30**(3), 141–164.
- BÉRGOLO, MARCELO L, CENI, RODRIGO, CRUCES, GUILLERMO, GIACCOBASSO, MATIAS, & PEREZ-TRUGLIA, RICARDO. 2017. *Tax audits as scarecrows: Evidence from a large-scale field experiment*. Tech. rept. National Bureau of Economic Research.
- BESLEY, TIMOTHY, & PERSSON, TORSTEN. 2013. Taxation and Development. *In*: ALAN J. AUERBACH, RAJ CHETTY, MARTIN FELDSTEIN, & SAEZ, EMMANUEL (eds), *handbook of public economics*, vol. 5.
- BESLEY, TIMOTHY, & PERSSON, TORSTEN. 2014. Why Do Developing Countries Tax So Little? *Journal of Economic Perspectives*, **28**(4), 99–120.
- BEST, MICHAEL CARLOS, BROCKMEYER, ANNE, KLEVEN, HENRIK JACOBSEN, SPINNEWIJN, JOHANNES, & WASEEM, MAZHAR. 2015. Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan. *Journal of Political Economy*, **123**(6), 1311–1355.
- BRINCH, CHRISTIAN N, MOGSTAD, MAGNE, & WISWALL, MATTHEW. 2017. Beyond LATE with a discrete instrument. *Journal of Political Economy*, **125**(4), 985–1039.
- DE PAULA, AUREO, & SCHEINKMAN, JOSE A. 2010. Value-Added Taxes, Chain Effects, and Informality. *American Economic Journal: Macroeconomics*, **2**(4), 195–221.
- DEATON, ANGUS, & CARTWRIGHT, NANCY. 2018. Understanding and misunderstanding randomized controlled trials. *Social Science & Medicine*, **210**, 2–21.

- DEBACKER, JASON, HEIM, BRADLEY T, TRAN, ANH, & YUSKAVAGE, ALEXANDER. 2013. The Impact of Legal Enforcement: An Analysis of Corporate Tax Aggressiveness after an Audit. *Middle Tennessee State, Indiana University, and US Department of Treasury working paper*.
- DEBACKER, JASON, HEIM, BRADLEY T, TRAN, ANH, & YUSKAVAGE, ALEXANDER. 2018. Once bitten, twice shy? The lasting impact of IRS audits on individual tax reporting. *Journal of Law and Economics*, **61**, 1–35.
- FBR. 2015. Audit Policy 2015. Federal Board of Revenue, Pakistan, Taxpayer’s Audit Wing.
- GADENNE, LUCIE, NANDI, TUSHAR K., & RATHELOT, ROLAND. 2019. *Taxation and Supplier Networks: Evidence from India*. Mimeo.
- GEMMELL, NORMAN, & RATTO, MARISA. 2012. Behavioral responses to taxpayer audits: evidence from random taxpayer inquiries. *National Tax Journal*, **65**(1), 33.
- GERARD, FRANÇOIS, NARITOMI, JOANA, SEIBOLD, ARTHUR, & ZULIAN, BRUNO. 2019. Two-Tier Tax Systems and Firms: Evidence from Brazil.
- GILOVICH, THOMAS. 1983. Biased evaluation and persistence in gambling. *Journal of personality and social psychology*, **44**(6), 1110.
- GÓMEZ SABAINI, JUAN CARLOS, & JIMÉNEZ, JUAN PABLO. 2012. Tax structure and tax evasion in Latin America. *Macroeconomics of Development Series 118*.
- GORDON, ROGER, & LI, WEI. 2009. Tax structures in developing countries: Many puzzles and a possible explanation. *Journal of Public Economics*, **93**(7-8), 855–866.
- HAERPFER, C, INGLEHART, R, MORENO, A, WELZEL, C, KIZILOVA, K, DIEZ-MEDRANO, J, LAGOS, M, NORRIS, P, PONARIN, E, PURANEN, B, *et al.* 2020. World Values Survey: Round Seven–Country-Pooled Datafile. *Madrid, Spain & Vienna, Austria: JD Systems Institute & WWSA Secretariat*.
- HECKMAN, JAMES J, & VYTLACIL, EDWARD. 2005. Structural equations, treatment effects, and econometric policy evaluation 1. *Econometrica*, **73**(3), 669–738.
- HECKMAN, JAMES J, & VYTLACIL, EDWARD J. 1999. Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the national Academy of Sciences*, **96**(8), 4730–4734.

- HECKMAN, JAMES J., & VYTLACIL, EDWARD J. 2007. Econometric evaluation of social programs, part I: Causal models, structural models and econometric policy evaluation. *Handbook of econometrics*, **6**, 4779–4874.
- IMBENS, GUIDO W., & ANGRIST, JOSHUA D. 1994. Identification and Estimation of Local Average Treatment Effects. *Econometrica*, **62**(2), 467–475.
- IRS. 1996. Federal Tax Compliance Research, Individual Income Tax Gap Estimates for 1985, 1988, and 1992. *Department of the Treasury Internal Revenue Service*. IRS Publications 1415 (Rev. 4-96), Washington DC.
- IRS. 2012. Federal Tax Compliance Research: Tax Year 2006 Tax Gap Estimation. *Department of the Treasury Internal Revenue Service*. IRS Research, Analysis and Statistics Working Paper.
- KIRCHLER, ERICH. 2007. *The economic psychology of tax behaviour*. Cambridge University Press.
- KLEVEN, HENRIK J., KNUDSEN, MARTIN B., KREINER, CLAUS THUSTRUP, PEDERSEN, SØREN, & SAEZ, EMMANUEL. 2011. Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark. *Econometrica*, **79**(3), 651–692.
- KLEVEN, HENRIK JACOBSEN, KREINER, CLAUS THUSTRUP, & SAEZ, EMMANUEL. 2016. Why can modern governments tax so much? An agency model of firms as fiscal intermediaries. *Economica*, **83**(330), 219–246.
- KOPCZUK, WOJCIECH, & SLEMROD, JOEL. 2006. Putting Firms into Optimal Tax Theory. *American Economic Review Papers and Proceedings*, **96**(2), 130–134.
- KOWALSKI, AMANDA E. 2016. *Doing more when you're running LATE: Applying marginal treatment effect methods to examine treatment effect heterogeneity in experiments*. Tech. rept. National Bureau of Economic Research.
- MACIEJOVSKY, BORIS, KIRCHLER, ERICH, & SCHWARZENBERGER, HERBERT. 2007. Misperception of chance and loss repair: On the dynamics of tax compliance. *Journal of Economic Psychology*, **28**(6), 678–691.
- MITTONE, LUIGI. 2006. Dynamic behaviour in tax evasion: An experimental approach. *The Journal of Socio-Economics*, **35**(5), 813 – 835.

- MURALIDHARAN, KARTHIK, & NIEHAUS, PAUL. 2017. Experimentation at Scale. *Journal of Economic Perspectives*, **31**(4), 103–24.
- NARITOMI, JOANA. 2019. Consumers as tax auditors. *American Economic Review*, **109**(9), 3031–72.
- POMERANZ, DINA. 2015. No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax. *American Economic Review*, **105**(8), 2539–2569.
- SARIN, NATASHA, & SUMMERS, LAWRENCE H. 2019. *Shrinking the Tax Gap: Approaches and Revenue Potential*. Tech. rept. National Bureau of Economic Research.
- SLEMROD, JOEL. 2019. Tax compliance and enforcement. *Journal of Economic Literature*, **57**(4), 904–54.
- TAUBINSKY, DMITRY, & REES-JONES, ALEX. 2017. Attention variation and welfare: theory and evidence from a tax salience experiment. NBER Working Paper.
- WASEEM, MAZHAR. 2020a. Does Cutting the Tax Rate to Zero Induce Behavior Different from Other Tax Cuts? Evidence from Pakistan. *The Review of Economics and Statistics*, **102**(3), 426–441.
- WASEEM, MAZHAR. 2020b. The Role of Withholding in the Self-Enforcement of a Value-Added Tax: Evidence from Pakistan. *The Review of Economics and Statistics*, 1–44. https://doi.org/10.1162/rest_a_00959.
- WASEEM, MAZHAR. 2021. Overclaimed Refunds, Undeclared Sales, and Invoice Mills: Nature and Extent of Noncompliance in a Value-Added Tax. CEPR. Discussion Paper No. 14601.
- WORLD BANK. 2020. *Doing Business 2020 : Comparing Business Regulation in 190 Economies*. Tech. rept. World Bank, Washington DC.

FIGURE I: EVASION RATE BY FIRM SIZE



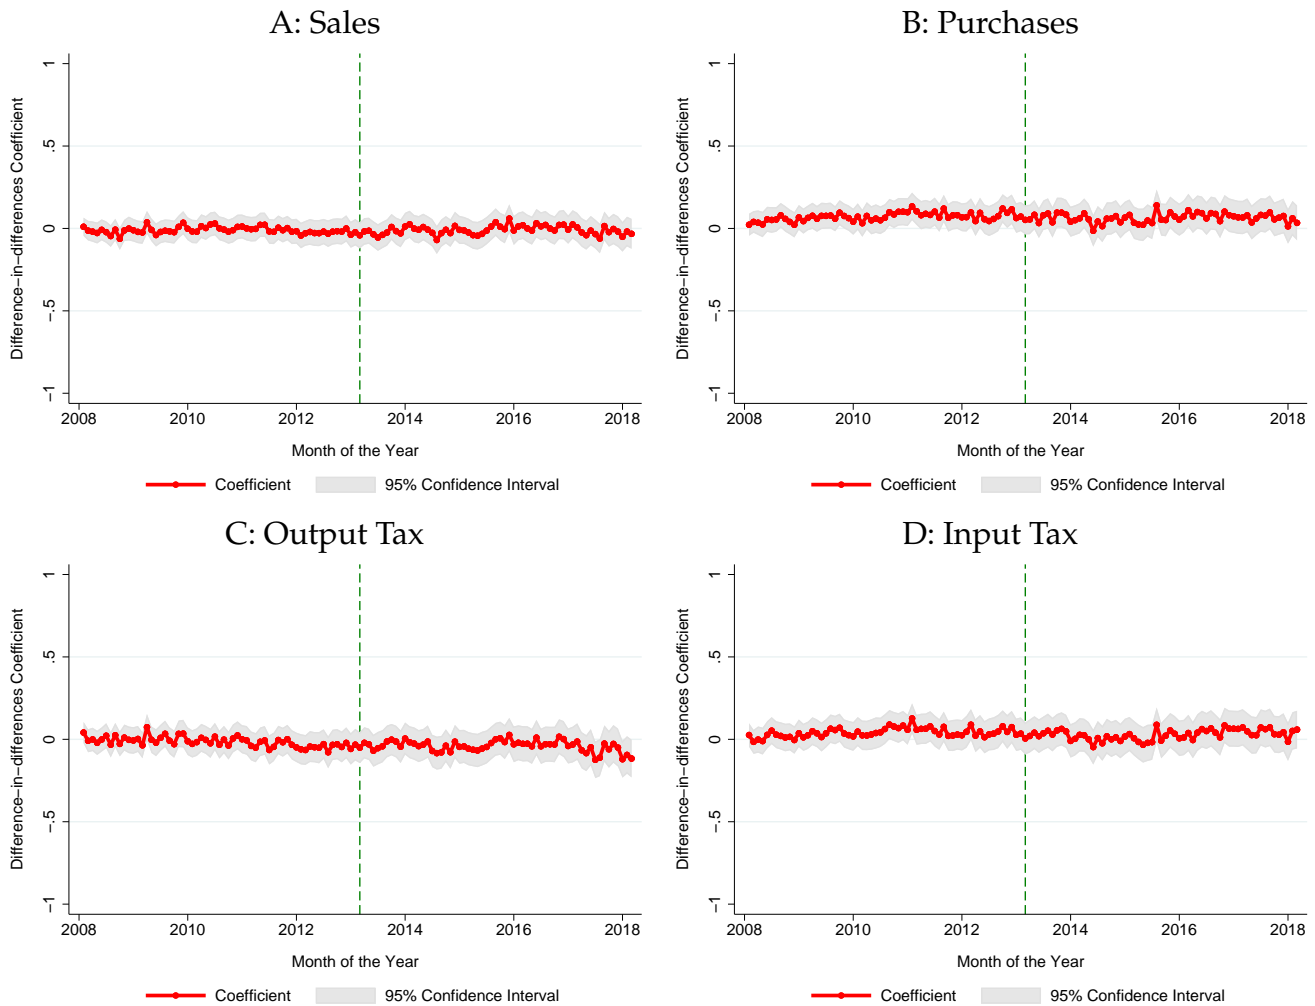
Notes: The figure plots the tax evasion rate by baseline firm size. We divide firms into equal-sized bins based on their annual turnover in the baseline year. We then calculate the evasion rate in each bin as the total amount detected by audit against all firms in the bin as a fraction of the total *real* VAT liability of these firms at the baseline. The real VAT liability is calculated as the sum of total VAT remitted by these firms at the baseline plus the total amount detected by audit against them. We winsorize the amount detected by audit at the 99th percentile of the distribution to account for outliers. The estimated evasion rate is shown by the red curve with the y-axis on the left. To increase statistical power, we pool together firms audited in the first two audit waves. We superimpose a series indicating the total VAT remitted by firms in each bin as a fraction of total VAT remitted by all audited firms in this sample. This series is shown by the blue curve with the y-axis on the right. The top two panels divide firms into 10 and 20 bins and the bottom-two into 50 and 100 bins. All plots begin from the 20th percentile because firms below this threshold remit no VAT at the baseline so that their evasion rate is not defined.

FIGURE II: INTENTION TO TREAT EFFECTS OF AUDIT – FIRST WAVE




Notes: The figure explores the impacts of audit on future firm behavior. We compare the evolution of four VAT outcomes across the treatment and control groups. The treatment groups consists of firms whose audit was assigned through the first random ballot held on September 13, 2013. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm and month fixed effects, dropping the dummy for July 2008. We then plot the coefficients on the time dummies of these regressions. The sample includes all tax periods from July 2008 to June 2018. The regressions are run separately for the two groups of firms. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines demarcate the date the random computer ballot was held on.

FIGURE III: INTENTION TO TREAT EFFECTS OF AUDIT – FIRST WAVE



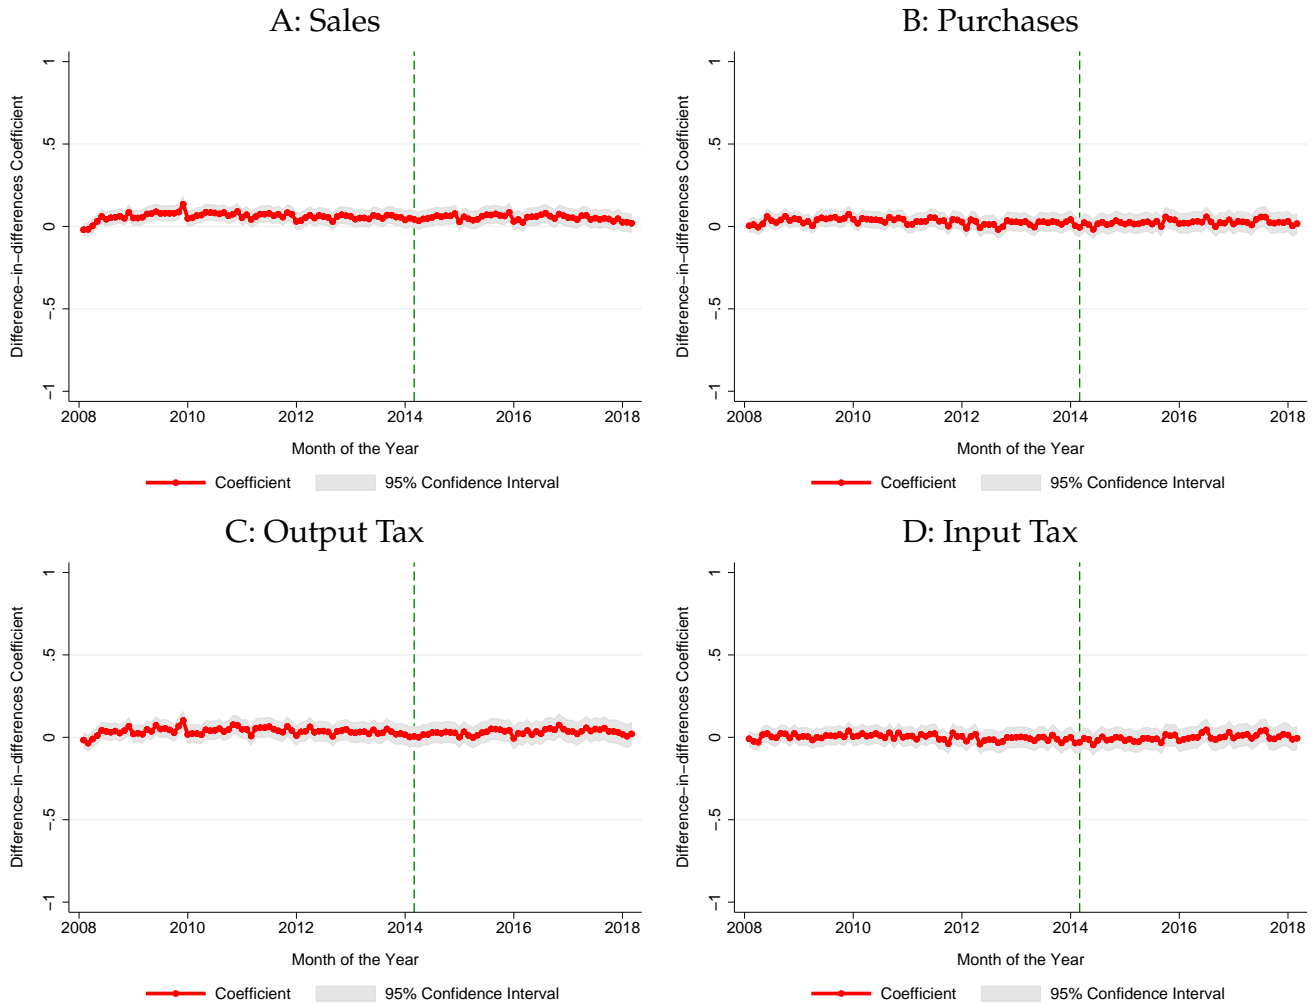
Notes: The figure shows the difference-in-differences version of the plots in Figure I. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm, month, and month \times treat dummies, dropping the dummies for July 2008. We then plot the coefficients on the month \times treat dummies from these regressions. The gray surface plot shows the 95 percent confidence interval around the coefficient. The treatment groups consists of firms whose audit was assigned through the first random ballot held on September 13, 2013. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. We cluster standard errors at the firm level. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines demarcate the date the random computer ballot was held on.

FIGURE IV: INTENTION TO TREAT EFFECTS OF AUDIT – SECOND WAVE




Notes: The figure explores the impacts of audit on future firm behavior. We compare the evolution of four VAT outcomes across the treatment and control groups. The treatment groups consists of firms whose audit was assigned through the first random ballot held on September 25, 2014. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm and month fixed effects, dropping the dummy for July 2008. We then plot the coefficients on the time dummies of these regressions. The sample includes all tax periods from July 2008 to June 2018. The regressions are run separately for the two groups of firms. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines demarcate the date the random computer ballot was held on.

FIGURE V: INTENTION TO TREAT EFFECTS OF AUDIT – SECOND WAVE



Notes: The figure shows the difference-in-differences version of the plots in Figure IV. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm, month, and month \times treat dummies, dropping the dummies for July 2008. We then plot the coefficients on the month \times treat dummies from these regressions. The gray surface plot shows the 95 percent confidence interval around the coefficient. The treatment groups consists of firms whose audit was assigned through the first random ballot held on September 25, 2014. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. We cluster standard errors at the firm level. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines demarcate the date the random computer ballot was held on.

FIGURE VI: AUDITED VS. UNAUDITED FIRMS – FIRST AUDIT WAVE




Notes: The figure compares the evolution of outcomes across audited and unaudited firms. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm, month, and month×audit dummies, dropping the dummies for July 2008. We then plot the coefficients on the month×audit dummies from these regressions. The gray surface plot shows the 95 percent confidence interval around the coefficient. The audit dummy indicates firms whose audit was conducted during the first wave. These includes firms whose audit was assigned through the random computer ballot ($Z_i = 1$) and firms whose audit was initiated by the local tax office on their own accord ($Z_i = 0$). The unaudited firms are all other firms in the population of VAT filers. We cluster standard errors at the firm level. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines denotes September 13, 2013—the date first random computer ballot was held on.

FIGURE VII: AUDITED VS. UNAUDITED FIRMS – SECOND AUDIT WAVE




Notes: The figure compares the evolution of outcomes across audited and unaudited firms. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm, month, and month \times audit dummies, dropping the dummies for July 2008. We then plot the coefficients on the month \times audit dummies from these regressions. The gray surface plot shows the 95 percent confidence interval around the coefficient. The audit dummy indicates firms whose audit was conducted during the second wave. These includes firms whose audit was assigned through the random computer ballot ($Z_i = 1$) and firms whose audit was initiated by the local tax office on their own accord ($Z_i = 0$). The unaudited firms are all other firms in the population of VAT filers. We cluster standard errors at the firm level. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines denotes September 25, 2014—the date first random computer ballot was held on.

FIGURE VIII: MARGINAL TREATMENT EFFECTS – FIRST AUDIT WAVE




Notes: The figure plots the $MTE(p)$ curve for four outcomes using random assignment in the first audit wave as instrument. Please see Appendix A.2 for technical details. The fraction treated $p \equiv P(D = 1|Z)$ is shown along the horizontal axis. It increases from 0 (no treatment) to 1 (full treatment). We also indicate the baseline treatment probability $p_B \equiv P(D = 1|Z = 0)$ and the intervention treatment probability $p_I \equiv P(D = 1|Z = 1)$ along this axis. The green solid curve shows the marginal treated outcomes curve $MTO(p)$. It is identified at two points indicated in the plot by circular markers. The blue, dashed curve depicts the marginal untreated outcomes curve $MUO(p)$. It is also identified at two points indicated in the plot with square markers. For both curves, we extrapolate between the two points using linearity assumption. The difference between the two curves represents the $MTE(p)$. Since in our setting all three curves sit above each other, we lift both $MTO(p)$ and $MUO(p)$ up by adding the constant from the corresponding regression to distinguish them from the primary object of our interest $MTE(p)$.

FIGURE IX: MARGINAL TREATMENT EFFECTS – SECOND AUDIT WAVE



Notes: The figure plots the $MTE(p)$ curve for four outcomes using random assignment in the second audit wave as instrument. Please see Appendix A.2 for technical details. The fraction treated $p \equiv P(D = 1|Z)$ is shown along the horizontal axis. It increases from 0 (no treatment) to 1 (full treatment). We also indicate the baseline treatment probability $p_B \equiv P(D = 1|Z = 0)$ and the intervention treatment probability $p_I \equiv P(D = 1|Z = 1)$ along this axis. The green solid curve shows the marginal treated outcomes curve $MTO(p)$. It is identified at two points indicated in the plot by circular markers. The blue, dashed curve depicts the marginal untreated outcomes curve $MUO(p)$. It is also identified at two points indicated in the plot with square markers. For both curves, we extrapolate between the two points using linearity assumption. The difference between the two curves represents the $MTE(p)$. Since in our setting all three curves sit above each other, we lift both $MTO(p)$ and $MUO(p)$ up by adding the constant from the corresponding regression to distinguish them from the primary object of our interest $MTE(p)$.

FIGURE X: PROBABILITY OF DETECTION



Notes: The figure plots the probability of detection faced by a typical firm. We arrange L transactions carried out by the firm in term of the detection probability they entail $p_l(e_l)$ in ascending order. The probability of transaction is low if the other party to the transaction is (1) a consumer, (2) an unregistered firm, or (3) a firm willing to collude. In all these case, the transaction does not create any third-party information for the government. The probability of detection is high otherwise. The curve accordingly turns sharply once transactions between arm-length parties unwilling to collude begin. The transaction L^* represents the first transaction for which the detection probability is so high that inequality (12) fails. The firm would accordingly report transactions $[L^*, L]$, hiding the rest. Note that the threshold L^* would vary across firms depending among other things on their size, industry, and trading network.

TABLE I: DESCRIPTIVE STATISTICS OF AUDIT I

Audit Wave	Tax Year	Ballot Date	Audits Assigned			Audits Conducted	
			Mode	Number	Percent	Assigned	Unassigned
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	2013	Sep 13, 2013	Random	4,926	5%	3,482	521
2	2014	Sep 25, 2014	Random	12,447	12%	3,612	293
3	2015	Sep 14, 2015	Random	8,372	7.5%	1,122	164
4	2016	Jan 05, 2017	Parametric	8,935	7.5%	884	332
5	2017	Apr 12, 2018	Parametric	8,785	7.5%	852	352

Notes: The table reports some descriptive statistics of the five audit waves in our sample. Column (2) reports the tax year during which the computer ballot to draw audit cases was held. Column (3) reports the exact ballot date. The ballot was random for the first three waves and parametric for the next two. The volume of cases picked during the ballot is mentioned in Column (5) in numbers and in Column (6) as the proportion of population. Column (7) reports the number of audits completed out of those assigned through the computer ballot. Column (8), on the other hand, reports the number of audits initiated by the local tax office on their own accord. During the five audit waves, a total of 43,625 cases were picked for audit through computer ballots. Out of these, the tax identifiers of 218 were inaccurate. We were therefore unable to merge these 218 cases with VAT and audit records. We accordingly drop these 218 cases from the sample and focus instead on the 43,465 audits assigned through the computer ballot as reported in Column (5).

TABLE II: DESCRIPTIVE STATISTICS OF AUDIT II

Audit Wave	Audits Initiated			Amount Detected		
	Within 1 Month	Within 3 Months	Within 6 Months	Mean	Median	90th Percentile
(1)	(2)	(3)	(4)	(5)	(6)	(7)
1	0.646	0.942	0.950	617	0	165
2	0.925	0.993	0.998	619	0	100
3	0.852	0.945	0.964	4,098	0	158

Notes: The table presents a few descriptive statistics of randomly assigned audits during the first three audit waves. Columns (2)-(4) report the time lag between the assignment and initiation of audit. Column (2), for example, shows that around 65 percent of audits assigned in the first random ballot were initiated with the first month of assignment. This ratio was 93 percent and 85 percent for the next two audit waves. Columns (5)-(7) report the amount detected during each wave of audit. Column (5) reports the mean amount detected in PKR thousands. The US\$-PKR exchange rate during this time (2013) was around 100. The next columns of the table report the median and the 90th percentile of the amount detected, illustrating that it is highly skewed toward right with the mean significantly larger than the median for all three audit waves.

TABLE III: RANDOMIZATION TEST

	First Wave				Second Wave				Third Wave			
	Mean ($Z_i = 0$)	Mean ($Z_i = 1$)	Diff. in Means	SE	Mean ($Z_i = 0$)	Mean ($Z_i = 1$)	Diff. in Means	SE	Mean ($Z_i = 0$)	Mean ($Z_i = 1$)	Diff. in Means	SE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
<u>A: VAT Outcomes</u>												
1. Sales	14.251	14.282	0.031	0.043	14.278	14.298	0.020	0.026	14.335	14.831	0.496	0.029
2. Purchases	14.081	14.095	0.014	0.047	14.234	14.186	-0.048	0.029	14.264	14.248	-0.015	0.035
3. Output Tax	11.671	11.707	0.036	0.049	11.791	11.768	-0.024	0.030	11.969	11.953	-0.017	0.035
4. Input Tax	11.768	11.802	0.033	0.052	11.990	11.911	-0.079	0.031	12.149	11.886	-0.263	0.037
5. Tax Payable	10.200	10.300	0.100	0.063	10.392	10.360	-0.032	0.041	10.570	10.830	0.260	0.045
6. Tax Paid	9.532	9.607	0.076	0.058	9.805	9.785	-0.020	0.034	9.850	10.338	0.488	0.039
7. Exports	15.288	15.169	-0.119	0.114	14.904	15.145	0.241	0.068	14.619	15.655	1.036	0.064
8. Imports	14.905	14.887	-0.018	0.078	14.858	14.843	-0.015	0.048	14.878	15.902	1.024	0.076
9. Refund	12.037	11.884	-0.153	0.152	12.214	12.188	-0.026	0.089	12.086	12.424	0.338	0.093
10. Carry Forward	11.642	11.667	0.026	0.078	12.010	12.160	0.150	0.046	12.162	12.248	0.086	0.050
<u>B: Firm Characteristics</u>												
11. Manufacturer	0.339	0.350	0.010	0.010	0.314	0.339	0.025	0.006	0.215	0.786	0.572	0.006
12. Importer	0.111	0.109	-0.003	0.006	0.124	0.118	-0.006	0.004	0.159	0.019	-0.140	0.002
13. Exporter	0.025	0.019	-0.005	0.003	0.040	0.025	-0.016	0.002	0.050	0.021	-0.029	0.002
14. Distributor	0.028	0.030	0.001	0.003	0.031	0.034	0.003	0.002	0.036	0.011	-0.025	0.002
15. Wholesaler	0.240	0.241	0.001	0.008	0.229	0.240	0.011	0.005	0.251	0.046	-0.205	0.003
16. Service Provider	0.193	0.192	-0.002	0.008	0.193	0.185	-0.009	0.005	0.208	0.099	-0.110	0.005
17. Major City	0.640	0.636	-0.004	0.010	0.631	0.639	0.008	0.006	0.625	0.650	0.024	0.007
18. LTU	0.013	0.013	0.000	0.004	0.012	0.008	-0.004	0.002	0.005	0.042	0.037	0.003
19. Years Registered	12.987	13.680	0.694	0.109	11.745	12.967	1.222	0.070	10.496	13.607	3.111	0.091
20. Textile	0.162	0.163	0.001	0.008	0.143	0.152	0.009	0.005	0.108	0.266	0.157	0.006

Notes: The table runs balance tests on the three randomization waves in our sample. For each outcome, we estimate model (9) restricting the sample to the baseline period only. The baseline period is June 2012 for the first, June 2013 for the second, and June 2014 for the third randomization wave. The last two columns for each randomization wave report the coefficient $\hat{\beta}$ and its standard error from the model. The details of the variables used here are provided in Appendix A.1.

TABLE IV: TAX EVASION RATE AT THE BASELINE

	# Audits	Sales	Amount Detected		VAT Paid at the Baseline		Evasion Rate
			PKR	% of Sales	PKR	% of Sales	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<u>A: First Audit Wave</u>							
All Audited Firms	3,482	498.4	2.15	0.43	28.16	5.65	7.1
Amount Detected > 0	986	137.0	2.15	1.57	3.20	2.33	40.2
Size Quartile 1	1,057	0.0	0.06	684.76	0.00	8.78	98.7
Size Quartile 2	824	1.7	0.07	3.94	0.04	2.52	61.0
Size Quartile 3	809	12.3	0.22	1.75	0.21	1.67	51.1
Size Quartile 4	792	484.3	1.80	0.37	27.91	5.76	6.1
<u>B: Second Audit Wave</u>							
All Audited Firms	3,612	2200.0	2.24	0.10	88.37	4.02	2.5
Amount Detected > 0	1,220	264.6	2.24	0.84	7.52	2.84	22.9
Size Quartile 1	1,007	0.4	0.04	10.21	0.02	3.81	72.8
Size Quartile 2	892	4.9	0.17	3.37	0.11	2.15	61.0
Size Quartile 3	862	24.4	0.22	0.89	0.30	1.24	41.8
Size Quartile 4	851	2170.2	1.81	0.08	87.95	4.05	2.0

Notes: The table estimates the evasion rate at the baseline using audit outcomes data. The first column reports the number of audits conducted for the group of firms indicated in the corresponding row. Aggregate turnover of this group in the baseline year in PKR billions is reported in the next column. The next two columns report the amount detected by audit, in PKR billions in column 3 and as a percent of aggregate sales in column 4. Columns 5-6 report the VAT paid at the baseline by the group of firms indicated in the corresponding row, in PKR billions in column 5 and as a percent of aggregate sales in column 6. The last column presents the evasion rate implied by the detected amount. It is calculated as the ratio of columns 4 and the sum of columns 4 and 6 and represents the evaded amount as a fraction of the *real* VAT liability of firms.

TABLE V: IMPACT OF AUDIT ON FIRM BEHAVIOR – FIRST WAVE

	Impacts After One Year					Impacts After Three Years				
	Sales	Purchases	Output Tax	Input Tax	Tax Payable	Sales	Purchases	Output Tax	Input Tax	Tax Payable
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<u>A: ITT Estimates</u>										
assign × after	-0.010 (0.016)	-0.010 (0.018)	-0.016 (0.021)	-0.017 (0.023)	-0.037 (0.028)	-0.007 (0.017)	-0.021 (0.019)	-0.025 (0.022)	-0.036 (0.023)	-0.016 (0.028)
Observations	2,831,140	2,468,502	2,086,889	2,099,210	1,415,795	3,839,502	3,328,628	2,884,225	2,906,045	1,913,096
<u>B: LATE Estimates</u>										
audit × after	-0.014 (0.023)	-0.014 (0.027)	-0.023 (0.030)	-0.024 (0.033)	-0.051 (0.039)	-0.010 (0.024)	-0.030 (0.028)	-0.035 (0.031)	-0.051 (0.033)	-0.022 (0.039)
Observations	2,831,140	2,468,502	2,086,889	2,099,210	1,415,795	3,839,502	3,328,628	2,884,225	2,906,045	1,913,096
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table estimates the impact of audit on firms' future behavior. In the top panel, the coefficient $\text{assign} \times \text{after}$ shows $\hat{\gamma}$ from model (10), where the dummy variable assign_i denotes that firm i 's audit was assigned through the first random ballot held on September 13, 2013. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable after_t indicates that month t falls after the date of the ballot. The sample includes periods up to October 2014 for the first five columns and periods up to October 2016 for the rest. Panel B shows the corresponding results from 2sls regressions, where the endogenous variable audit_i is instrumented by the initial random assignment. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE VI: IMPACT OF AUDIT ON FIRM BEHAVIOR – SECOND WAVE

	Impacts After One Year					Impacts After Three Years				
	Sales	Purchases	Output Tax	Input Tax	Tax Payable	Sales	Purchases	Output Tax	Input Tax	Tax Payable
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<u>A: ITT Estimates</u>										
assign × after	-0.021 (0.010)	-0.021 (0.012)	-0.030 (0.012)	-0.026 (0.013)	-0.022 (0.016)	-0.010 (0.010)	-0.009 (0.012)	-0.013 (0.013)	-0.007 (0.013)	0.006 (0.016)
Observations	3,133,061	2,725,243	2,343,583	2,357,343	1,568,363	4,159,404	3,587,740	3,088,403	3,137,794	2,034,932
<u>B: LATE Estimates</u>										
audit × after	-0.071 (0.033)	-0.073 (0.043)	-0.109 (0.043)	-0.091 (0.046)	-0.081 (0.058)	-0.032 (0.035)	-0.033 (0.043)	-0.044 (0.044)	-0.025 (0.045)	0.022 (0.057)
Observations	3,133,061	2,725,243	2,343,583	2,357,343	1,568,363	4,159,404	3,587,740	3,088,403	3,137,794	2,034,932
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table estimates the impact of audit on firms' future behavior. In the top panel, the coefficient $\text{assign} \times \text{after}$ shows $\hat{\gamma}$ from model (10), where the dummy variable assign_i denotes that firm i 's audit was assigned through the first random ballot held on September 25, 2014. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable after_t indicates that month t falls after the date of the ballot. The sample includes periods up to October 2015 for the first five columns and periods up to October 2017 for the rest. Panel B shows the corresponding results from 2sls regressions, where the endogenous variable audit_i is instrumented by the initial random assignment. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE VII: IMPACTS OF RANDOM AUDITS ASSIGNED IN THE FIRST WAVE – OTHER OUTCOMES

	Impacts After One Year					Impacts After Three Years				
	Exports	Imports	Tax Paid	Refund	Carry Forward	Exports	Imports	Tax Paid	Refund	Carry Forward
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<u>A: ITT Estimates</u>										
assign × after	0.013 (0.037)	0.047 (0.028)	-0.052 (0.031)	-0.116 (0.092)	-0.049 (0.040)	0.027 (0.038)	0.035 (0.027)	-0.025 (0.033)	-0.070 (0.091)	-0.085 (0.040)
Observations	317,130	570,949	1,161,513	234,207	1,594,740	450,661	838,590	1,723,448	287,241	2,490,894
<u>B: LATE Estimates</u>										
audit × after	0.018 (0.051)	0.073 (0.043)	-0.072 (0.043)	-0.175 (0.138)	-0.071 (0.058)	0.037 (0.053)	0.054 (0.042)	-0.035 (0.046)	-0.102 (0.134)	-0.124 (0.059)
Observations	317,130	570,949	1,161,513	234,207	1,594,740	450,661	838,590	1,723,448	287,241	2,490,894
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table estimates the impact of audit on firms' future behavior. In the top panel, the coefficient $\text{assign} \times \text{after}$ shows $\hat{\gamma}$ from model (10), where the dummy variable assign_i denotes that firm i 's audit was assigned through the first random ballot held on September 13, 2013. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable after_t indicates that month t falls after the date of the ballot. The sample includes periods up to October 2014 for the first five columns and periods up to October 2016 for the rest. Panel B shows the corresponding results from 2sls regressions, where the endogenous variable audit_i is instrumented by the initial random assignment. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE VIII: EXTENSIVE MARGIN IMPACT OF RANDOM AUDITS

Outcome:	$1(\text{Return Filed}_{it})$					
	September 13, 2013		September 25, 2014		September 14, 2015	
Random Draw Held On:	One Year	Three Years	One Year	Three Years	One Year	Three Years
Impacts After:	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: ITT Estimates</u>						
assign \times after	0.002 (0.002)	0.004 (0.002)	0.008 (0.001)	0.009 (0.001)	0.010 (0.001)	0.008 (0.001)
Observations	7,097,120	9,852,941	8,129,498	11,062,795	8,502,891	11,171,180
<u>B: LATE Estimates</u>						
audit \times after	0.002 (0.002)	0.006 (0.003)	0.027 (0.004)	0.029 (0.004)	0.075 (0.010)	0.058 (0.009)
Observations	7,097,120	9,852,941	8,129,498	11,062,795	8,502,891	11,171,180
Mean of the Dependent Variable	0.955	0.955	0.956	0.956	0.956	0.956
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table estimates the impact of audit on firms' extensive margin behavior. We estimate model (10) using an indicator that the firm filed its VAT return for the period (month) t as the outcome variable. In the top panel, the coefficient assign \times after shows $\hat{\gamma}$ from the model. The dummy variable assign $_i$ denotes that firm i 's audit was assigned through the random ballot held on the date indicated in the heading of each column. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable after $_t$ indicates that month t falls after the date of the ballot. The sample for odd columns includes periods up to one year after the ballot and for even columns up to three years after the ballot. Panel B shows the corresponding results from 2sls regressions, where the endogenous variable audit $_i$ is instrumented by the initial random assignment. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE IX: SELECTION IN COMPLIANCE? AUDITED VS. NON-AUDITED FIRMS

	First Wave				Second Wave			
	Mean ($D_i = 0$)	Mean ($D_i = 1$)	Difference in Means	Standard Error	Mean ($D_i = 0$)	Mean ($D_i = 1$)	Difference in Means	Standard Error
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>A: VAT Outcomes</u>								
1. Sales	14.547	14.816	0.269	0.044	14.553	14.776	0.222	0.043
2. Purchases	14.311	14.567	0.255	0.048	14.438	14.519	0.080	0.049
3. Output Tax	11.936	12.214	0.279	0.050	12.031	12.199	0.168	0.051
4. Input Tax	12.006	12.328	0.322	0.051	12.196	12.230	0.033	0.051
5. Tax Payable	10.537	10.866	0.328	0.068	10.698	10.870	0.172	0.075
6. Tax Paid	10.039	10.435	0.397	0.062	10.221	10.359	0.137	0.063
7. Exports	15.752	15.705	-0.047	0.114	15.353	15.793	0.440	0.096
8. Imports	15.183	15.261	0.078	0.075	15.096	15.235	0.139	0.074
9. Refund	12.410	12.673	0.263	0.139	12.578	12.667	0.089	0.130
10. Carry Forward	11.926	12.192	0.266	0.081	12.276	12.446	0.170	0.083
<u>B: Firm Characteristics</u>								
11. Manufacturer	0.383	0.448	0.064	0.010	0.361	0.418	0.056	0.009
12. Importer	0.105	0.087	-0.018	0.006	0.116	0.111	-0.005	0.006
13. Exporter	0.023	0.016	-0.007	0.003	0.036	0.013	-0.023	0.003
14. Distributor	0.027	0.026	-0.001	0.004	0.029	0.028	-0.001	0.004
15. Wholesaler	0.214	0.196	-0.018	0.008	0.206	0.219	0.012	0.008
16. Service Provider	0.190	0.174	-0.016	0.008	0.189	0.166	-0.023	0.008
17. Major City	0.661	0.661	0.000	0.000	0.654	0.654	0.000	0.000
18. LTU	0.045	0.045	-0.000	0.000	0.039	0.039	-0.000	0.000
19. Years Registered	13.499	14.729	1.230	0.117	12.388	14.221	1.833	0.119
20. Textile	0.165	0.171	0.005	0.007	0.148	0.160	0.012	0.006

Notes: The table explores selection in audit, comparing audited and unaudited firms. We estimate a version of model (9), regressing the outcome in each row on two dummy variables (D_i and $corporate_i$) and tax office fixed effects. We restrict the sample to the baseline period only. The dummy variable D_i takes the value 1 for all audited firms including those whose audit was assigned through the random ballot and those whose audit was taken up by the local tax office of its own accord. The unaudited firms ($D_i = 0$) include all other firms in the eligible sample. The baseline period is June 2012 for the first and June 2013 for the second audit wave. The last two columns for each wave report the coefficient $\hat{\beta}$ and its standard error from the model. The details of the variables used here are provided in Appendix A.1.

TABLE X: SELECTION IN COMPLIANCE? AUDITED VS. NON-AUDITED FIRMS (WITHIN $Z_i = 1$ GROUP)

	2013 Draw				2014 Draw			
	Mean ($D_i = 0$)	Mean ($D_i = 1$)	Difference in Means	Standard Error	Mean ($D_i = 0$)	Mean ($D_i = 1$)	Difference in Means	Standard Error
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>A: VAT Outcomes</u>								
1. Sales	14.567	14.569	0.001	0.095	14.560	14.633	0.073	0.061
2. Purchases	14.360	14.312	-0.048	0.108	14.393	14.410	0.017	0.066
3. Output Tax	11.885	12.005	0.120	0.104	11.982	12.094	0.112	0.069
4. Input Tax	11.944	12.075	0.131	0.117	12.131	12.115	-0.017	0.070
5. Tax Payable	10.524	10.666	0.142	0.132	10.648	10.715	0.067	0.101
6. Tax Paid	9.935	10.175	0.240	0.129	10.206	10.173	-0.033	0.083
7. Exports	15.602	15.678	0.076	0.285	15.476	15.897	0.422	0.226
8. Imports	15.131	15.150	0.018	0.178	15.057	15.154	0.097	0.101
9. Refund	11.650	12.482	0.832	0.385	12.502	12.681	0.179	0.257
10. Carry Forward	11.833	12.023	0.190	0.173	12.424	12.331	-0.093	0.108
<u>B: Firm Characteristics</u>								
11. Manufacturer	0.364	0.406	0.042	0.022	0.378	0.397	0.019	0.013
12. Importer	0.115	0.096	-0.019	0.016	0.107	0.120	0.013	0.010
13. Exporter	0.018	0.017	-0.001	0.006	0.022	0.022	0.000	0.004
14. Distributor	0.030	0.027	-0.003	0.008	0.033	0.029	-0.003	0.005
15. Wholesaler	0.228	0.210	-0.018	0.020	0.218	0.215	-0.003	0.012
16. Service Provider	0.186	0.188	0.001	0.017	0.185	0.170	-0.015	0.011
17. Major City	0.655	0.655	-0.000	0.000	0.659	0.659	-0.000	0.000
18. LTU	0.043	0.043	0.000	0.000	0.035	0.035	0.000	0.000
19. Years Registered	13.865	14.313	0.448	0.258	13.175	14.222	1.047	0.167
20. Textile	0.163	0.167	0.005	0.017	0.158	0.154	-0.004	0.009

Notes: The table explores selection in audit, comparing audited and unaudited firms within the sample drawn for audit in the corresponding random ballot. We estimate a version of model (9), regressing the outcome in each row on two dummy variables (D_i and $corporate_i$) and tax office fixed effects. We restrict the sample to the baseline period only. The dummy variable D_i takes the value 1 for firms whose audit was conducted. The unaudited firms ($D_i = 0$) include all other firms in the randomly drawn sample. The baseline period is June 2012 for the first and June 2013 for the second audit wave. The last two columns for each wave report the coefficient $\hat{\beta}$ and its standard error from the model. The details of the variables used here are provided in Appendix A.1.

TABLE XI: SELECTION IN COMPLIANCE? AUDITED VS. NON-AUDITED FIRMS (WITHIN $Z_i = 0$ GROUP)

	2013 Draw				2014 Draw			
	Mean ($D = 0$)	Mean ($D = 1$)	Difference in Means	Standard Error	Mean ($D = 0$)	Mean ($D = 1$)	Difference in Means	Standard Error
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>A: VAT Outcomes</u>								
1. Sales	14.548	15.693	1.145	0.086	14.556	15.776	1.220	0.149
2. Purchases	14.312	15.330	1.018	0.095	14.446	15.549	1.103	0.169
3. Output Tax	11.936	12.958	1.022	0.098	12.040	13.046	1.006	0.170
4. Input Tax	12.008	13.045	1.037	0.097	12.208	13.053	0.844	0.171
5. Tax Payable	10.537	11.704	1.167	0.148	10.708	11.979	1.270	0.235
6. Tax Paid	10.040	11.220	1.180	0.132	10.227	11.390	1.163	0.190
7. Exports	15.750	16.009	0.258	0.216	15.330	16.019	0.689	0.372
8. Imports	15.183	15.473	0.290	0.129	15.101	15.689	0.588	0.198
9. Refund	12.425	13.168	0.743	0.291	12.585	13.168	0.583	0.381
10. Carry Forward	11.927	12.857	0.930	0.164	12.255	13.138	0.883	0.279
<u>B: Firm Characteristics</u>								
11. Manufacturer	0.384	0.622	0.239	0.021	0.359	0.531	0.172	0.030
12. Importer	0.105	0.049	-0.055	0.010	0.117	0.110	-0.007	0.019
13. Exporter	0.023	0.017	-0.006	0.004	0.037	0.026	-0.011	0.002
14. Distributor	0.026	0.020	-0.006	0.007	0.028	0.017	-0.012	0.014
15. Wholesaler	0.214	0.133	-0.081	0.012	0.205	0.182	-0.023	0.022
16. Service Provider	0.190	0.111	-0.079	0.016	0.190	0.103	-0.087	0.025
17. Major City	0.661	0.661	0.000	0.000	0.652	0.652	-0.000	0.000
18. LTU	0.045	0.045	-0.000	0.000	0.040	0.040	0.000	0.000
19. Years Registered	13.493	16.379	2.886	0.288	12.266	14.597	2.331	0.437
20. Textile	0.165	0.187	0.021	0.017	0.147	0.208	0.061	0.025

Notes: The table explores selection in audit, comparing audited and unaudited firms excluding from the sample firms drawn for audit in the corresponding random ballot. We estimate a version of model (9), regressing the outcome in each row on two dummy variables (D_i and $corporate_i$) and tax office fixed effects. We restrict the sample to the baseline period only. The dummy variable D_i takes the value 1 for firms whose audit was conducted. The baseline period is June 2012 for the first and June 2013 for the second audit wave. The last two columns for each wave report the coefficient $\hat{\beta}$ and its standard error from the model. The details of the variables used here are provided in Appendix A.1.

A Online Appendix

A.1 Definition of Variables

- (i) **Sales.** The value of all goods and services supplied by the firm in the given tax period (month) including exports.
- (ii) **Purchases.** The value of all taxable intermediates acquired by the firm in the given tax period (month).
- (iii) **Output Tax.** The value of VAT charged on sales made by the firm in the given tax period (month). It equals $\tau \cdot (\hat{s}_{it} - \hat{s}_{E,it})$, where τ is the applicable VAT rate and $(\hat{s}_{it} - \hat{s}_{E,it})$ is the value of non-export sales reported by firm i in period t . Because exports are zero-rated, they do not appear in the output tax.
- (iv) **Input Tax.** The value of VAT credit claimed on intermediates acquired by the firm in the given tax period (month). It equals $\tau \cdot \hat{c}_{it}$, where τ is the applicable VAT rate and \hat{c}_{it} is the value of purchases of intermediates claimed by firm i in period t .
- (v) **Tax Payable.** The VAT payable by the firm in the given tax period (month). By definition, it equals the output tax minus the input tax.
- (vi) **Tax Paid** The VAT actually paid by the firm in the given tax period (month). It may differ from Tax Payable if the firm has any carry-forward from previous months.
- (vii) **Exports.** The value of all goods and services exported by the firm in the given tax period (month).
- (viii) **Imports.** The value of all goods and services imported by the firm in the given tax period (month).
- (ix) **Refund.** The amount of refund claimed by the firm in the given tax period (month). The refund arises when the firm's input tax exceeds its output tax. In this case, the firm has the option to carry forward the balance amount or seek its refund. Because exports are zero-rated, firms the majority of whose output is exported are likely to claim refund every tax period.

- (x) **Carry Forward.** The amount of carry forward claimed by a firm. The carry forward arises when the firm's input tax exceeds its output tax and it does not opt to seek the refund of the balance amount.
- (xi) **Manufacturer.** A firm whose principal business activity is the manufacture of goods. Manufacturing is the process whereby a firm converts inputs into a distinct article capable of being put to use differently than inputs and includes any process incidental or ancillary to it.
- (xii) **Importer.** A firm whose principal business activity is the import of goods for sale in the local market without carrying out any manufacturing process on them.
- (xiii) **Exporter.** A firm whose principal business activity is the export of goods. These firms may supply in the local market, but a majority of their output is exported out of country.
- (xiv) **Distributor.** Distributor means a person appointed by a manufacturer, importer or any other person for a specified area to purchase goods from him for further supply and includes a person who in addition to being a distributor is also engaged in supply of goods as a wholesaler or a retailer.
- (xv) **Wholesaler.** Wholesaler' includes a dealer and means any person who carries on, whether regularly or otherwise, the business of buying and selling goods by wholesale or of supplying or distributing goods, directly or indirectly, by wholesale for cash or deferred payment or for commission or other valuable consideration or stores such goods belonging to others as an agent for the purpose of sale; and includes a person supplying taxable goods to a person who deducts income tax at source under the Income Tax Ordinance, 2001.
- (xvi) **Retailer.** A person, supplying goods to general public for the purpose of consumption.
- (xvii) **Industry.** The Pakistani tax administration uses 4-digit Harmonized Commodity Description and Coding System (HS code) to classify firms into industry. The code, used by customs administrations throughout the world, divides all goods and services into 99 chapters (the first two digits in the code) and 21

sections. The sections broadly correspond to major industries in the country. I take the section a firm falls in as its industry.

- (xviii) **Major City** The dummy variable takes the value 1 if the firm’s head office is in one of the three major cities of Pakistan—Karachi, Lahore, and Islamabad.
- (xix) **LTU** The dummy variable takes the value 1 if the firm is administered by on of the four Large Taxpayer Centers in the country located in Karachi, Lahore, and Islamabad.

A.2 Marginal Treatment Effects

In this section, we describe how we estimate the $MTE(p)$ curves shown in Figures VIII and IX. Because we have access to a binary instrument only, full nonparametric identification (see Heckman & Vytlacil, 2005, 2007) is not feasible in our setup, and instead we identify MTEs under a functional structure following the approach developed in Kowalski (2016) and Brinch *et al.* (2017).

As in the paper, Z here denotes the instrument (random assignment) and D the treatment (actual audit). Following the standard terminology in this literature, we refer to $p \equiv P(D = 1|Z)$ as the potential fraction treated. For any outcome Y , The $MTE(p)$ is defined as

$$MTE(p) \equiv \mathbb{E}(Y_T - Y_U|U_D = p)$$

where Y_T represents the potential outcome in the audited state ($D = 1$) and Y_U the potential outcome in the unaudited state ($D = 0$). The unobserved cost and benefit of audit are represented by U_D and p . The MTE therefore captures the treatment effect on a unit marginal to selecting into treatment. Using the above definition, it can be written as the difference between the marginal treated outcome (MTO) and the marginal untreated outcome (MUO)

$$\begin{aligned} MTO(p) &\equiv \mathbb{E}(Y_T|U_D = p) \\ MUO(p) &\equiv \mathbb{E}(Y_U|U_D = p) \end{aligned}$$

These curves are defined for every value of $p(Z)$ but given our binary instrument only two values of p are observed: the baseline treatment probability $p_B \equiv P(D = 1|Z = 0)$ and the intervention treatment probability $p_I \equiv P(D = 1|Z = 1)$. We therefore assume that both these curves are linear. The $MTO(p)$ is identified at two

points

$$\begin{aligned} BTTO &= \mathbb{E}(Y|X = x, D = 1, Z = 0) \\ LATO &= \frac{1}{p_I - p_B} [p_I ITTO - p_B BTTO], \end{aligned}$$

where $ITTO = \mathbb{E}(Y|X = x, D = 1, Z = 1)$. We use the linearity assumption to extrapolate between these two points. Similarly, the $MUO(p)$ is identified at

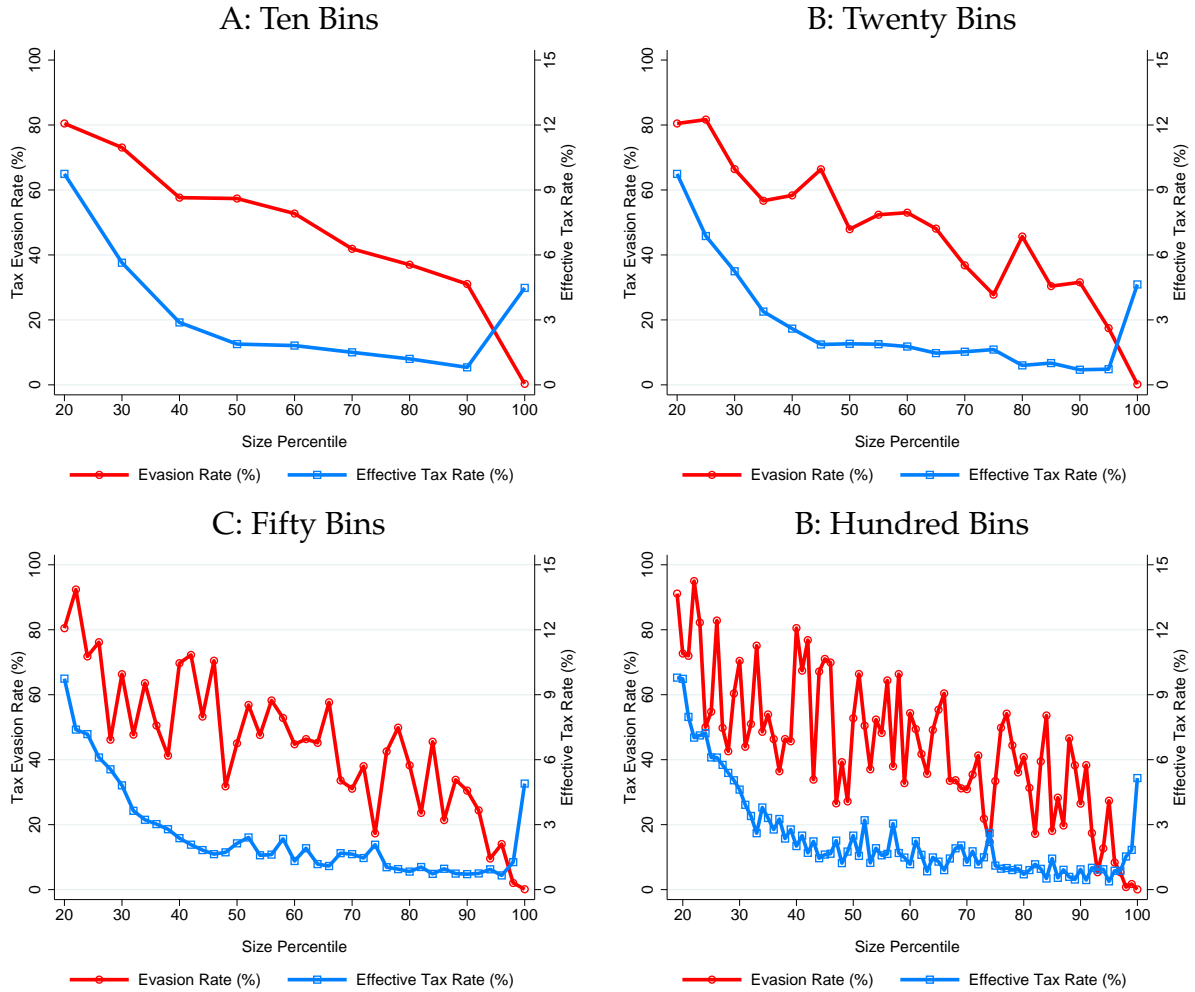
$$\begin{aligned} IUUO &= \mathbb{E}(Y|X = x, D = 0, Z = 1) \\ LAUO &= \frac{1}{p_I - p_B} [(1 - p_B)BUUO - (1 - p_I)IUUO] \end{aligned}$$

where $BUUO = \mathbb{E}(Y|X = x, D = 0, Z = 0)$.²⁸

To plot the $MTO(p)$ curve, we regress the outcome variable on a full set of firm and period fixed fixed effects and an interaction term of the audit (D) and post dummies, restricting the sample to firms randomly selected for audit ($Z = 1$). The regression gives us estimates of ITTO and IUUO. Running a similar regression on a sample of firms not drawn in the random ballot ($Z = 0$) delivers the estimates of BTTO and BUUO. We then find LATO and LAUO using the definitions above. The $MTO(p)$ curve is identified at two points $(BTTO, \frac{p_B}{2})$ and $(LATO, \frac{p_B + p_I}{2})$. We extrapolate between the two using the linearity assumption. Similarly, $MUO(p)$ is identified at $(LAUO, \frac{p_B + p_I}{2})$ and $(IUUO, \frac{p_I + 1}{2})$, and we extrapolate using linearity. The $MTO(p)$ curve is the difference between the two. We draw these curves for four outcomes and two audit waves separately. Since in our setting all these curves sit above each other, we lift both $MTO(p)$ and $MUO(p)$ up by adding the constant from the corresponding regression to distinguish them from the primary object of our interest $MTE(p)$.

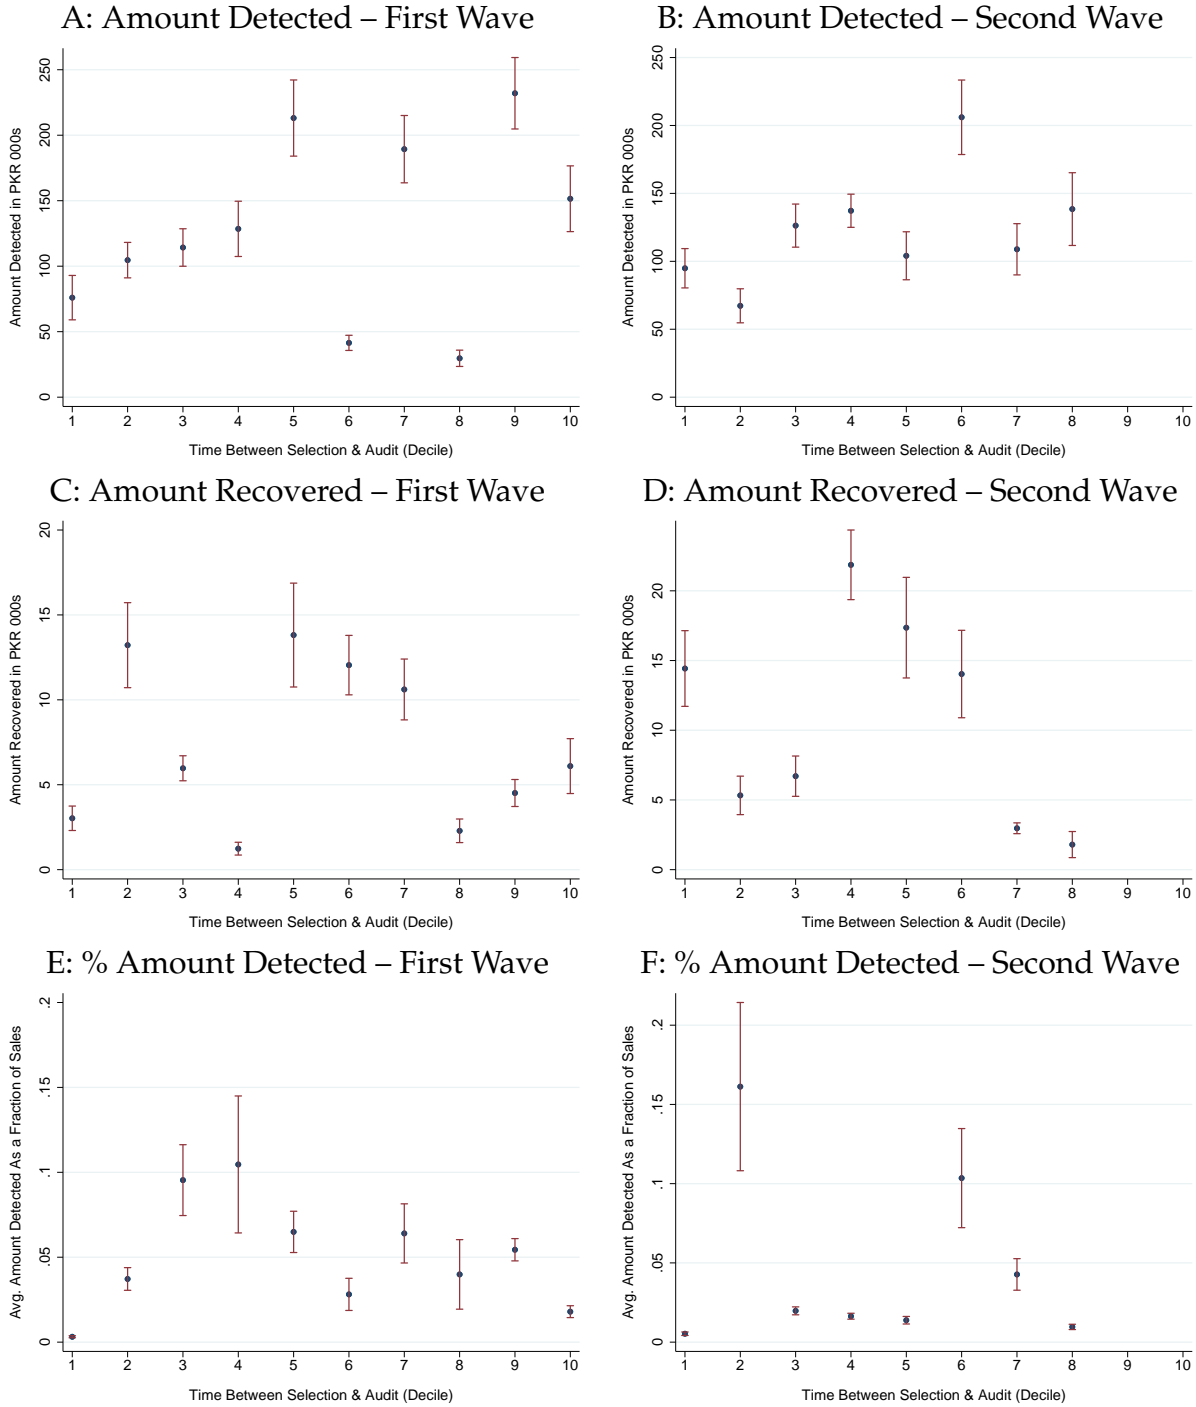
²⁸In all these definitions, O stands for outcomes, T for treated, U for untreated, B for baseline, I for intervention, and LA for local average. Please see [Kowalski \(2016\)](#) for detail of these terms.

FIGURE A.I: EVASION RATE BY FIRM SIZE




Notes: The figure plots the tax evasion rate by baseline firm size. We divide firms into equal-sized bins based on their annual turnover in the baseline year. We then calculate the evasion rate in each bin as the total amount detected by audit against all firms in the bin as a fraction of the total *real* VAT liability of these firms at the baseline. The real VAT liability is calculated as the sum of total VAT remitted by these firms at the baseline plus the total amount detected by audit against them. We winsorize the amount detected by audit at the 99th percentile of the distribution to account for outliers. The estimated evasion rate is shown by the red curve with the y-axis on the left. To increase statistical power, we pool together firms audited in the first two audit waves. We superimpose a series indicating the average effective tax rate in each bin at the baseline. We calculate the average effective tax rate as the ratio between the total VAT remitted by firms in the bin and total sales reported by these firms, both in the baseline year. The average effective tax rate series is shown by the blue curve with the y-axis on the right. The top two panels divide firms into 10 and 20 bins and the bottom-two into 50 and 100 bins. All plots begin from the 20th percentile because firms below this threshold remit no VAT at the baseline so that their evasion rate is not defined.

FIGURE A.II: AMOUNT DETECTED BY TIMING OF AUDIT



Notes: The figure examines if the order in which audits were taken up is correlated with audit outcomes, exploring thereby if audits were systematically targeted toward specific firms. We divide the time between assignment and initiation of audit into ten deciles and then plot the average audit outcome and the 95 percent confidence interval around it for each decile. The top panels look at the average amount detected by audit in PKR thousands, the middle panels at the average amount recovered in PKR thousands, and the bottom panels at the average amount detected as a ratio of annual baseline turnover of the firm. To take care of outliers, we drop observations where the amount detected is more than the 99th percentile of the distribution. This affects the top and bottom panels only. The LHS panels plot outcomes for the first randomized ballot and the RHS for the second.

FIGURE A.III: INTENTION TO TREAT EFFECTS OF THIRD AUDIT WAVE



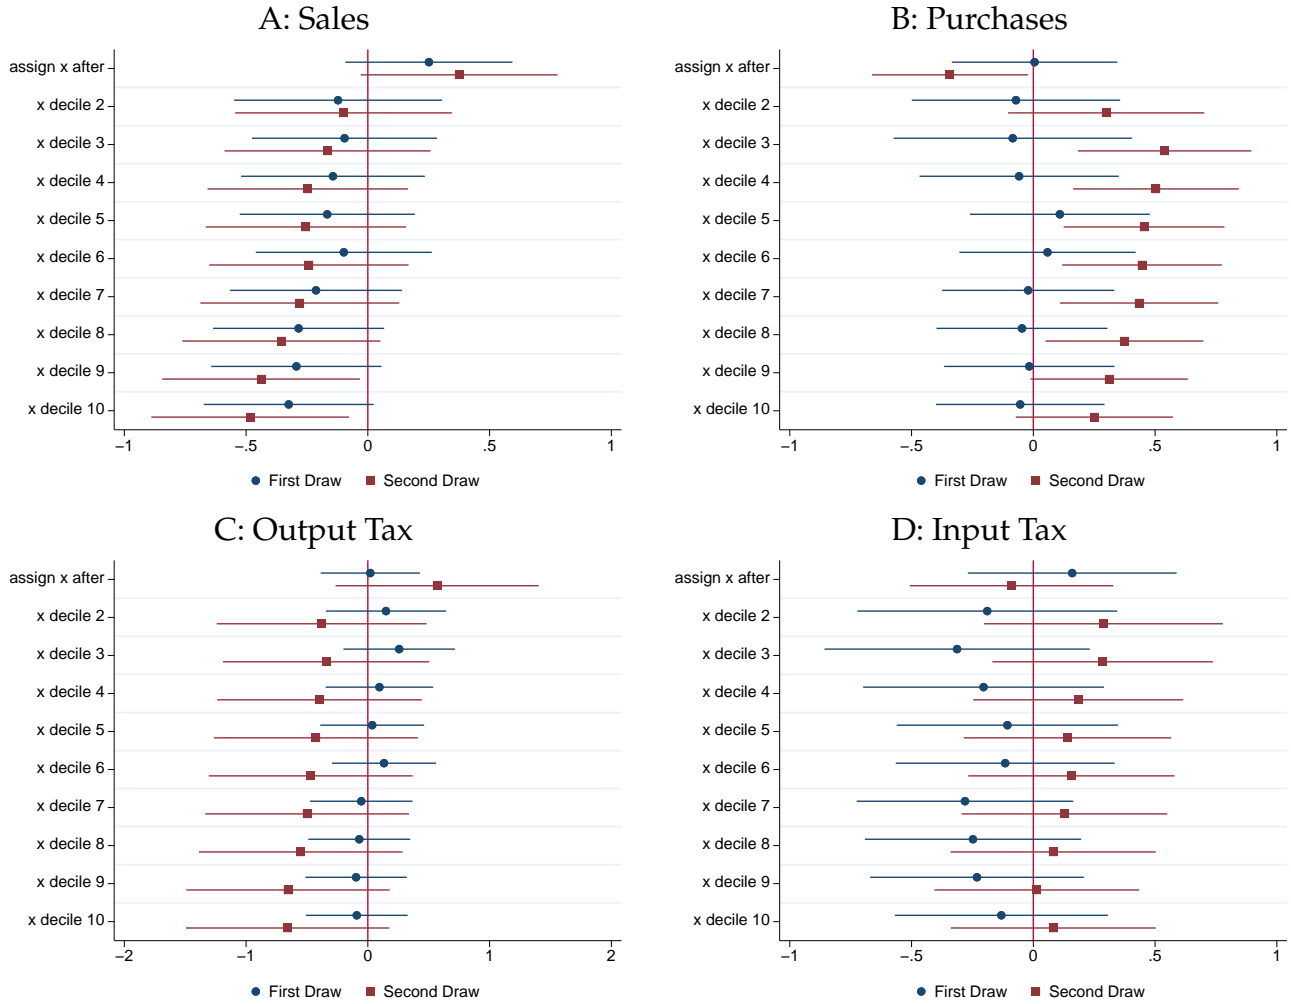
Notes: The figure explores the impacts of audit on future firm behavior. We compare the evolution of four VAT outcomes across the treatment and control groups. The treatment groups consists of firms whose audit was assigned through the first random ballot held on September 14, 2015. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments, firms already under audit, and firms subject to fixed and withholding tax regimes. We do not identify the last type of firms and therefore are unable to exclude them from the eligible sample. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm and month fixed effects, dropping the dummy for July 2008. We then plot the coefficients on the time dummies of these regressions. The sample includes all tax periods from July 2008 to June 2018. The regressions are run separately for the two groups of firms. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines demarcate the date the random computer ballot was held on.

FIGURE A.IV: INTENTION TO TREAT EFFECTS OF THIRD AUDIT WAVE



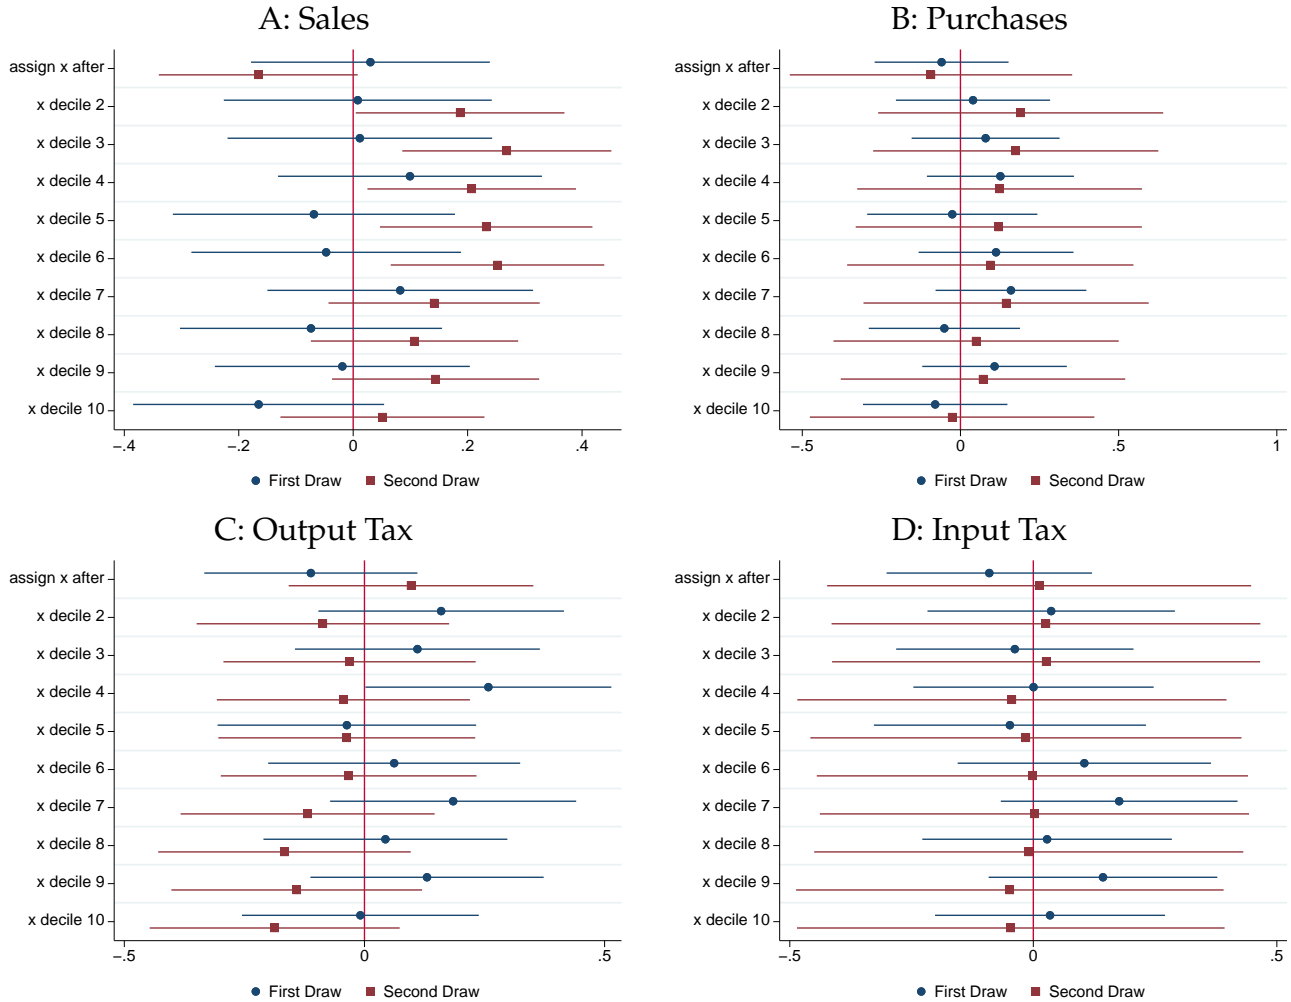
Notes: The figure shows the difference-in-differences version of the plots in Figure A.III. To construct these charts, we regress the log of the outcome variable shown in the title of each panel on the full set of firm, month, and month \times treat dummies, dropping the dummies for July 2008. We then plot the coefficients on the month \times treat dummies from these regressions. The gray surface plot shows the 95 percent confidence interval around the coefficient. The treatment groups consists of firms whose audit was assigned through the first random ballot held on September 14, 2015. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments, firms already under audit, and firms subject to fixed and withholding tax regimes. We do not identify the last type of firms and therefore are unable to exclude them from the eligible sample. We cluster standard errors at the firm level. Year t on the horizontal axis indicates July of the corresponding year. Vertical dashed lines demarcate the date the random computer ballot was held on.

FIGURE A.V: HETEROGENEITY IN RESPONSE BY FIRM SIZE



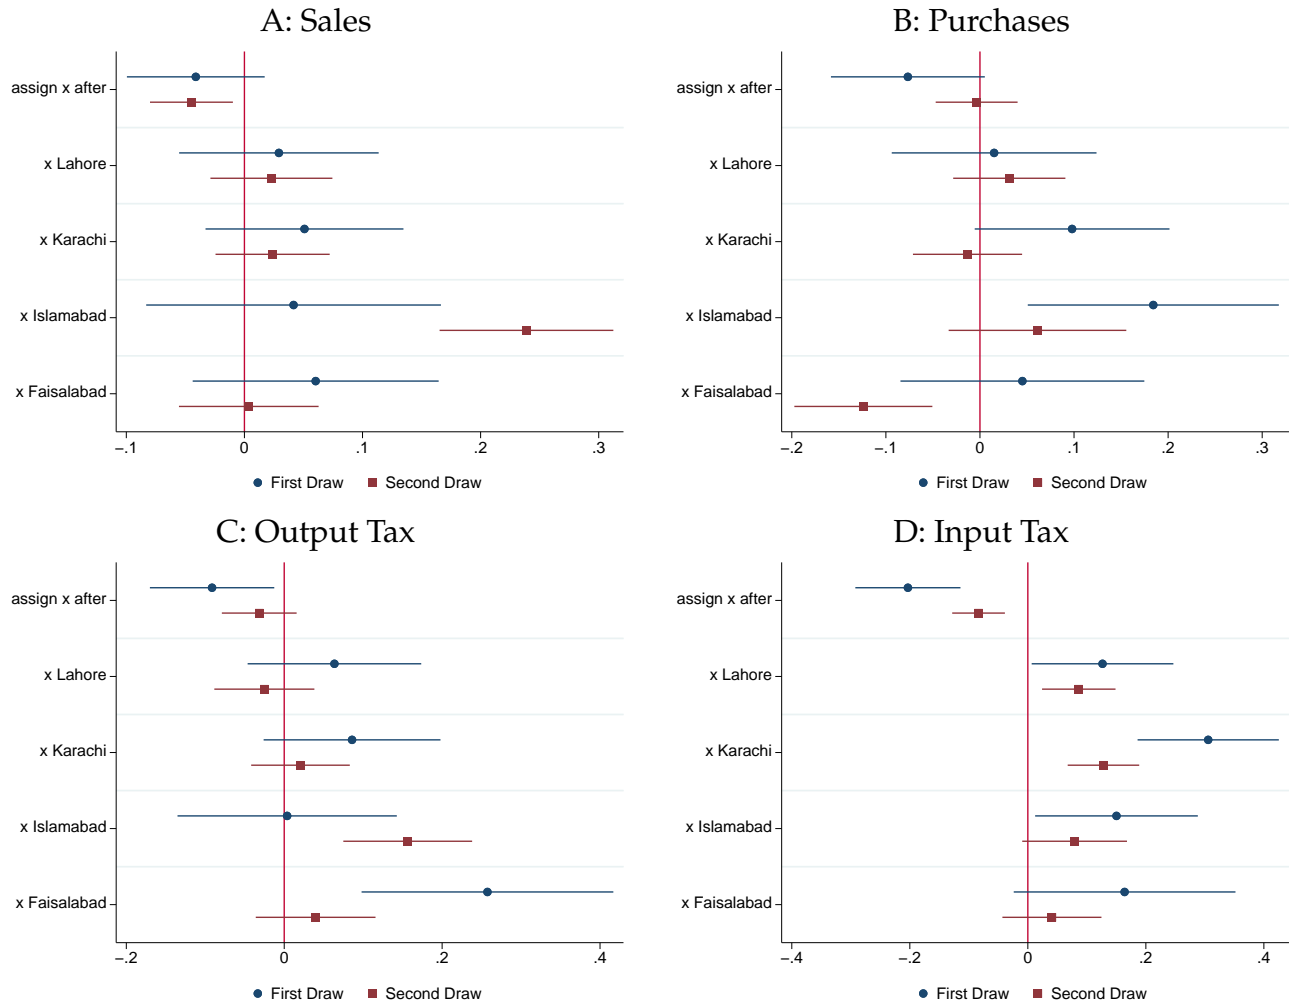
Notes: The figure explores heterogeneity in the audit effect. We divide firms into ten deciles based on their annual turnover in the baseline year. We then estimate a triple-difference version of model (10). The model includes interactions of the firm decile dummy with the $assign \times after_{it}$ dummy. The $assign_i$ dummies takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. We drop the triple-interaction term involving the first decile. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.VI: HETEROGENEITY IN RESPONSE BY FIRM AGE




Notes: The figure explores heterogeneity in the audit effect. We divide firms into ten deciles based on their age, defining age as the number of days between July 1, 2013 and the date of registration of the firm. We then estimate a triple-difference version of model (10). The model includes interactions of the firm decile dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. We drop the triple-interaction term involving the first decile. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.VII: HETEROGENEITY IN RESPONSE BY LOCATION



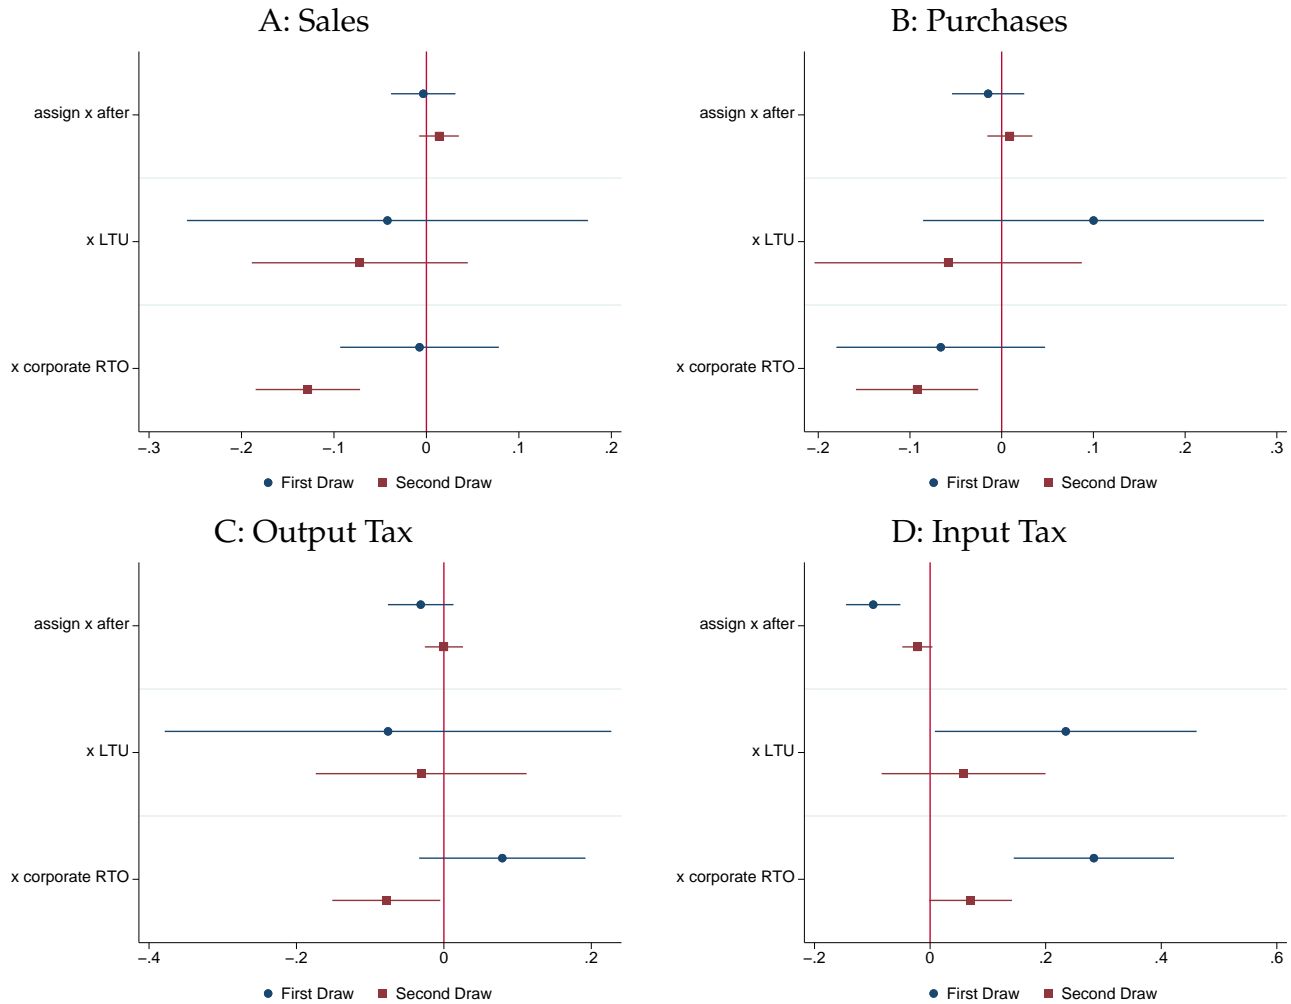
Notes: The figure explores heterogeneity in the audit effect. We divide firms into five groups depending upon the city their head office is located in. Firms not located in the four major cities of the country—Lahore, Karachi, Islamabad, and Faisalabad—are included in the baseline category. We then estimate a triple-difference version of model (10). The model includes interactions of the firm location dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm’s audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.VIII: HETEROGENEITY IN RESPONSE BY TAX OFFICE




Notes: The figure explores heterogeneity in the audit effect. We divide firms into eleven groups based on the local tax office they are subject to. Firms not in the ten major tax offices are included in the baseline category. We then estimate a triple-difference version of model (10). The model includes interactions of the tax office dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.IX: HETEROGENEITY IN RESPONSE BY TAX OFFICE TYPE



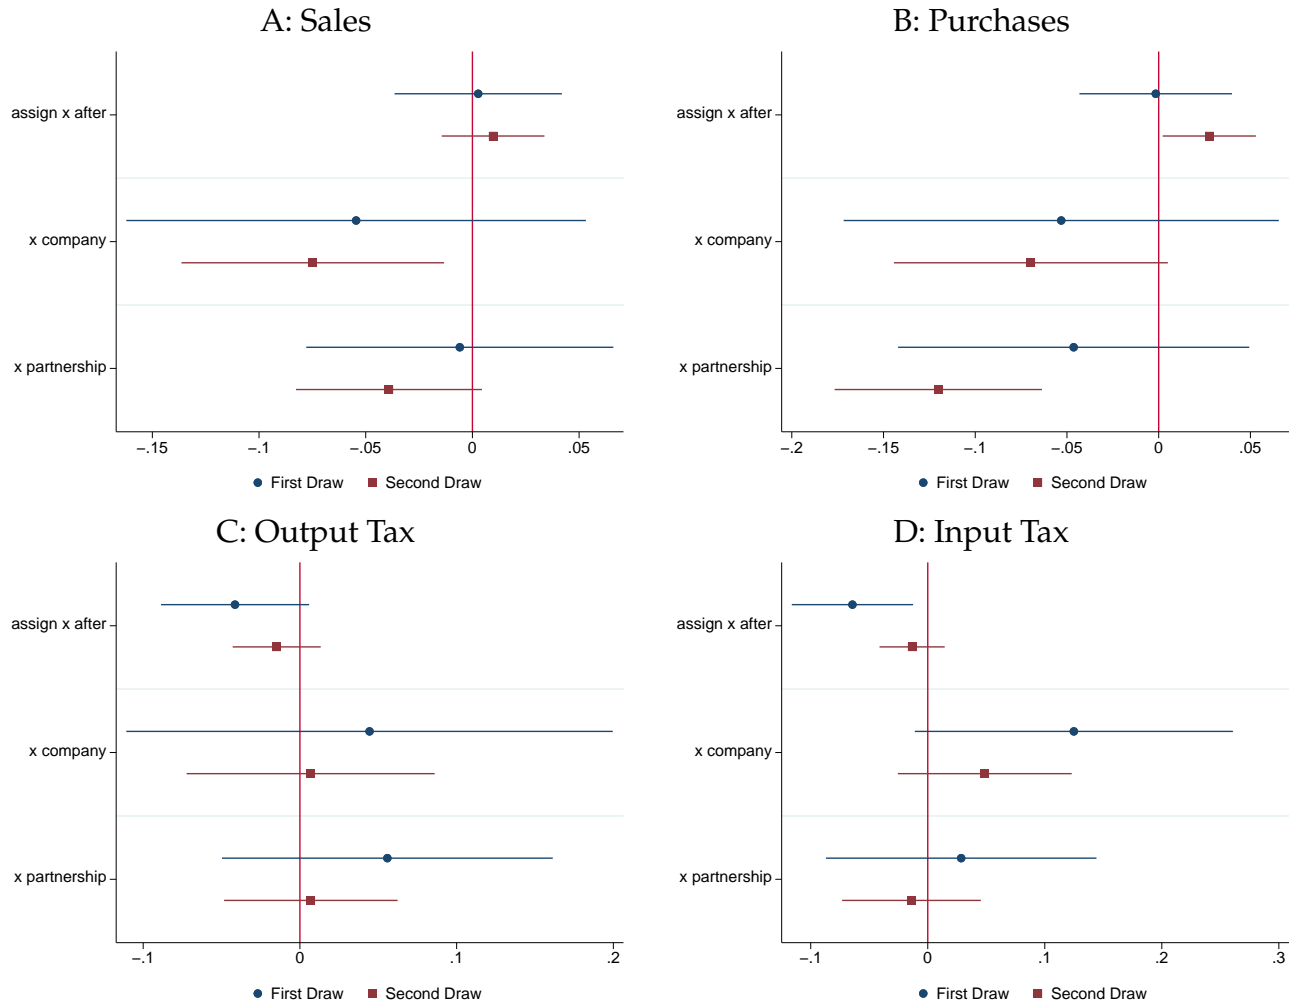
Notes: The figure explores heterogeneity in the audit effect. We divide firms into three groups based on the type of tax office they are subject to. Firms in four Large Taxpayer Units of the country are included in the first group (LTU), firms in the two Corporate Regional Tax Offices are included in the second group, and the rest of the firms are included in the baseline category. These firms are subject to a normal Regional Tax Office. We then estimate a triple-difference version of model (10). The model includes interactions of the tax office type dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.X: HETEROGENEITY IN RESPONSE BY PRODUCTION STAGE



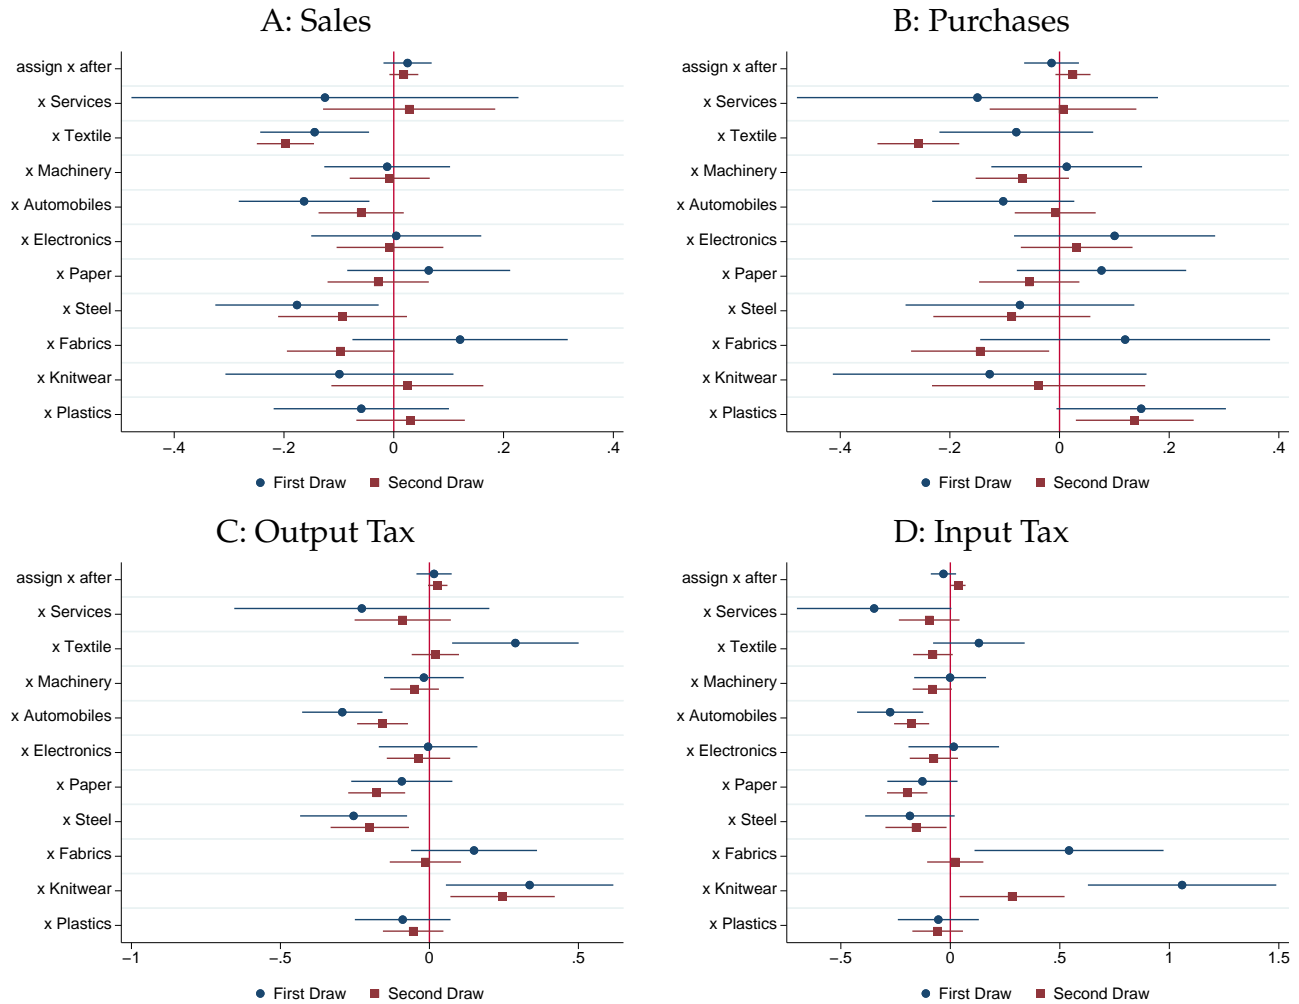
Notes: The figure explores heterogeneity in the audit effect. We divide firms into seven groups based on their principle business activity. The baseline category are retailers. These activities roughly capture the position of the firm in the supply chain. We then estimate a triple-difference version of model (10). The model includes interactions of the production stage dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.XI: HETEROGENEITY IN RESPONSE BY BUSINESS ORGANIZATION




Notes: The figure explores heterogeneity in the audit effect. We divide firms into three groups based on their business organization. The baseline category are sole proprietors. We then estimate a triple-difference version of model (10). The model includes interactions of the business organization dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.XII: HETEROGENEITY IN RESPONSE BY INDUSTRY



Notes: The figure explores heterogeneity in the audit effect. We divide firms into 12 groups based on the industry they operate in. We separate firms in 11 major industries of the country and club the rest into the baseline category. We then estimate a triple-difference version of model (10). The model includes interactions of the industry dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.


FIGURE A.XIII: HETEROGENEITY IN RESPONSE BY TIMING OF AUDIT




Notes: The figure explores heterogeneity in the audit effect. We divide firms into ten deciles based on the time lag between the assignment and initiation of audit in days. We then estimate a triple-difference version of model (10). The model includes interactions of the firm decile dummies with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. We drop the triple-interaction term involving the first decile. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. The first wave results are in blue and the second wave results are in red. Standard errors are clustered at the firm level.

FIGURE A.XIV: HETEROGENEITY IN RESPONSE BY FIRM SIZE (FIRST WAVE)


A: Sales




B: Purchases



C: Output Tax




D: Input Tax




Notes: The figure explores heterogeneity in the audit effect. We use firm-size as a continuous variable. We then use a generalized random forest model to estimate the treatment effects of the audit for all values within the feasible range based on the available data. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The estimated treatment effects and 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for each outcome variable.

FIGURE A.XV: HETEROGENEITY IN RESPONSE BY FIRM SIZE (SECOND WAVE)


A: Sales




B: Purchases



C: Output Tax




D: Input Tax




Notes: The figure explores heterogeneity in the audit effect. We use firm-size as a continuous variable. We then use a generalized random forest model to estimate the treatment effects of the audit for all values within the feasible range based on the available data. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The estimated treatment effects and 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for each outcome variable.

FIGURE A.XVI: HETEROGENEITY IN RESPONSE BY FIRM AGE (FIRST WAVE)


A: Sales




B: Purchases



C: Output Tax




D: Input Tax




Notes: The figure explores heterogeneity in the audit effect. We use firm-age as a continuous variable. We then use a generalized random forest model to estimate the treatment effects of the audit for all values within the feasible range based on the available data. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The estimated treatment effects and 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for each outcome variable.

FIGURE A.XVII: HETEROGENEITY IN RESPONSE BY FIRM AGE (SECOND WAVE)


A: Sales




B: Purchases



C: Output Tax




D: Input Tax




Notes: The figure explores heterogeneity in the audit effect. We use firm-age as a continuous variable. We then use a generalized random forest model to estimate the treatment effects of the audit for all values within the feasible range based on the available data. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The estimated treatment effects and 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for each outcome variable.

FIGURE A.XVIII: HETEROGENEITY IN RESPONSE BY LOCATION



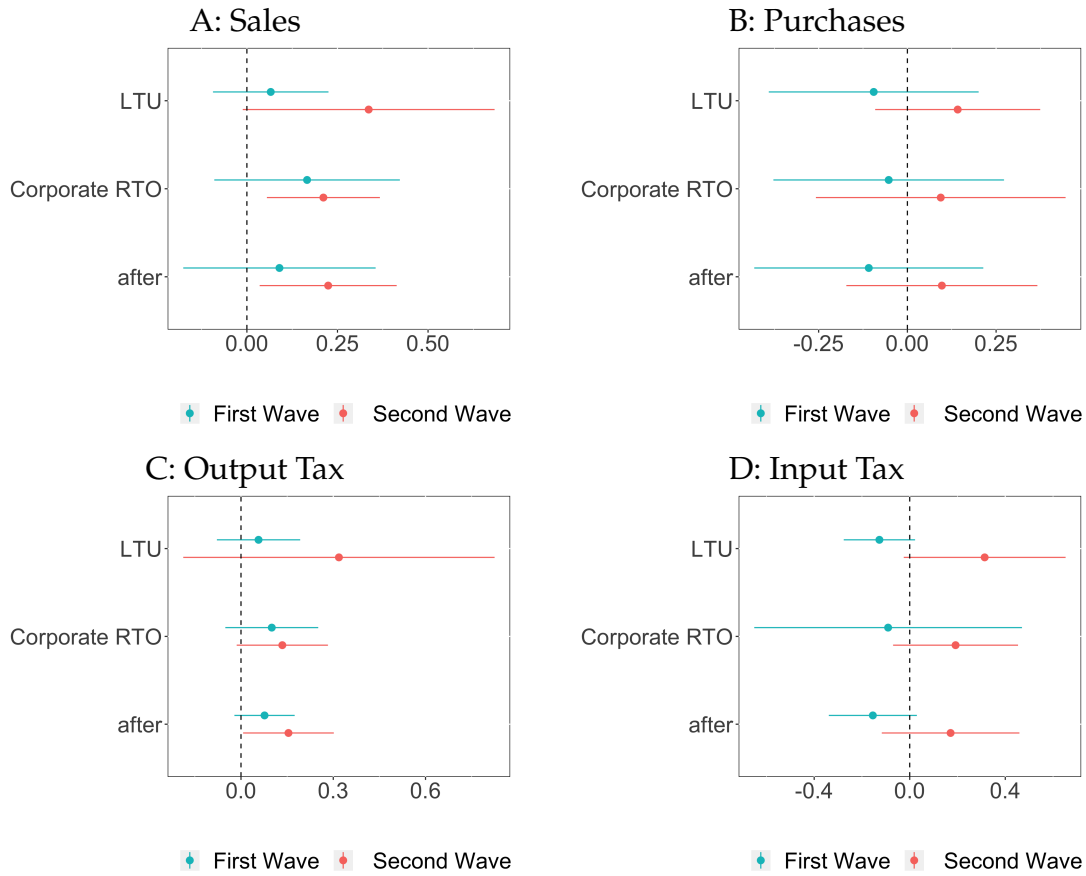
Notes: The figure explores heterogeneity in the audit effect. We divide firms into five groups depending upon the city their head office is located in. Firms not located in the four major cities of the country—Lahore, Karachi, Islamabad, and Faisalabad—are included in the baseline category. We then use a generalized random forest model to estimate the treatment effects of the audit. The model includes dummy variables for each group along with a dummy variable for "after" - indicating the time period after the date of the ballot. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The coefficients and the 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for the first and the second audit waves and for each outcome variable. The first wave results are in blue and the second wave results are in red.

FIGURE A.XIX: HETEROGENEITY IN RESPONSE BY TAX OFFICE



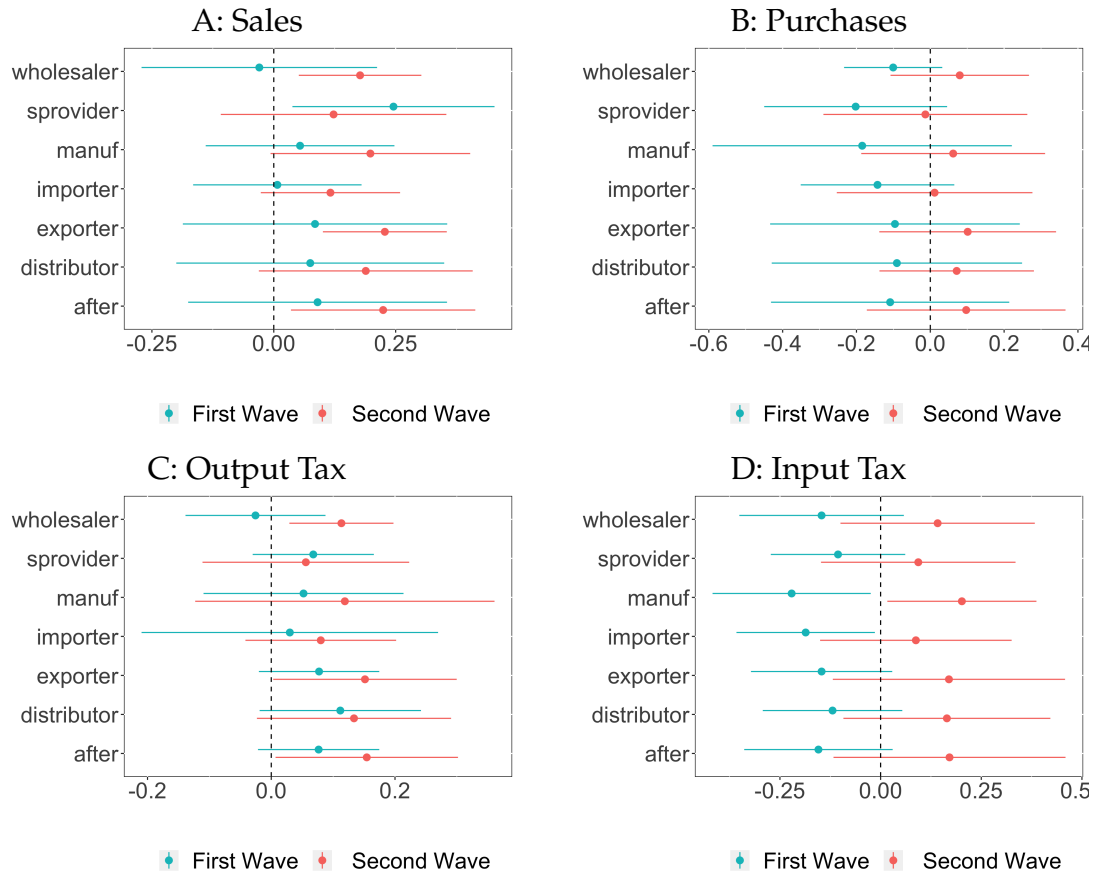
Notes: The figure explores heterogeneity in the audit effect. We divide firms into eleven groups based on the local tax office they are subject to. Firms not in the ten major tax offices are included in the baseline category. We then use a generalized random forest model to estimate the treatment effects of the audit. The model includes dummy variables for each group along with a dummy variable for "after" - indicating the time period after the date of the ballot. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The coefficients and the 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for the first and the second audit waves and for each outcome variable. The first wave results are in blue and the second wave results are in red.

FIGURE A.XX: HETEROGENEITY IN RESPONSE BY TAX OFFICE TYPE




Notes: The figure explores heterogeneity in the audit effect. We divide firms into three groups based on the type of tax office they are subject to. Firms in four Large Taxpayer Units of the country are included in the first group (LTU), firms in the two Corporate Regional Tax Offices are included in the second group, and the rest of the firms are included in the baseline category. These firms are subject to a normal Regional Tax Office. We then use a generalized random forest model to estimate the treatment effects of the audit. The model includes dummy variables for each tax office type along with a dummy variable for "after" - indicating the time period after the date of the ballot. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The coefficients and the 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for the first and the second audit waves and for each outcome variable. The first wave results are in blue and the second wave results are in red.

FIGURE A.XXI: HETEROGENEITY IN RESPONSE BY PRODUCTION STAGE




Notes: The figure explores heterogeneity in the audit effect. We divide firms into seven groups based on their principle business activity. The baseline category are retailers. These activities roughly capture the position of the firm in the supply chain. We then use a generalized random forest model to estimate the treatment effects of the audit. The model includes dummy variables for each group along with a dummy variable for "after" - indicating the time period after the date of the ballot. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The coefficients and the 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for the first and the second audit waves and for each outcome variable. The first wave results are in blue and the second wave results are in red.

FIGURE A.XXII: HETEROGENEITY IN RESPONSE BY BUSINESS ORGANIZATION



Notes: The figure explores heterogeneity in the audit effect. We divide firms into three groups based on their business organization. The baseline category are sole proprietors We then use a generalized random forest model to estimate the treatment effects of the audit. The model includes dummy variables for each group along with a dummy variable for "after" - indicating the time period after the date of the ballot. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The coefficients and the 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for the first and the second audit waves and for each outcome variable. The first wave results are in blue and the second wave results are in red.

FIGURE A.XXIII: HETEROGENEITY IN RESPONSE BY INDUSTRY



Notes: The figure explores heterogeneity in the audit effect. We divide firms into 11 groups based on the industry they operate in. We separate firms in 10 major industries of the country and club the rest into the baseline category. We then use a generalized random forest model to estimate the treatment effects of the audit. The model includes dummy variables for each group along with a dummy variable for "after" - indicating the time period after the date of the ballot. We consider a firm as treated (audited) if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The coefficients and the 95 percent confidence intervals on the estimated treatment effects are plotted. Models are estimated separately for the first and the second audit waves and for each outcome variable. The first wave results are in blue and the second wave results are in red.

TABLE A.I: EVASION RATE AND FIRM CHARACTERISTICS

	Outcome: Tax Evasion Rate					
	(1)	(2)	(3)	(4)	(5)	(6)
Log Firm Size	-4.065*** (0.481)	-3.995*** (0.450)	-4.020*** (0.454)	-3.863*** (0.457)	-4.552*** (0.484)	-4.416*** (0.516)
Share Manufacturers	0.091 (0.282)					0.004 (0.281)
Share Retailers		0.142 (0.176)				0.182 (0.175)
Share Major City			-0.116 (0.669)			-0.694 (0.730)
Share Young Firms				0.783* (0.447)		0.950* (0.489)
Share Textile					0.502*** (0.169)	0.503*** (0.170)
Constant	100.526*** (7.466)	99.989*** (7.483)	102.595*** (12.735)	87.577*** (10.578)	105.866*** (7.575)	99.519*** (13.414)
Observations	818	818	818	818	818	818

Notes: The table explores heterogeneity in evasion rate with respect to important firm characteristics. We divide firms into 1000 size bins based on their reported annual turnover in the baseline year. We then regress the evasion rate in each bin on log firm size and other firm characteristics. As earlier, the evasion rate here is calculated as the total amount detected by audit against all firms in the bin as a fraction of total *real* VAT liability of these firms at the baseline. The details of firm characteristics used here are given in Appendix A.1. To increase statistical power, we pool together firms audited in the first two audit waves. We winsorize the amount detected at the 99th percentile of the distribution to account for outliers. Robust standard errors are in parenthesis. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

TABLE A.II: BREAKDOWN OF THE DETECTED AMOUNT

	Amt. Detected		Amt. Recovered		Amt. Recoverable		Refund Curtailed	
	PKR	%	PKR	%	PKR	%	PKR	%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<u>A: First Audit Wave</u>								
All Audited Firms	2.147	0.431	0.023	0.005	2.118	0.425	0.004	0.001
Amount Detected > 0	2.147	1.567	0.023	0.017	2.118	1.546	0.004	0.003
Size Quartile 1	0.062	684.756	0.001	11.221	0.061	673.534	0.000	0.000
Size Quartile 2	0.067	3.936	0.003	0.186	0.064	3.750	0.000	0.000
Size Quartile 3	0.215	1.746	0.008	0.067	0.203	1.648	0.003	0.021
Size Quartile 4	1.802	0.372	0.011	0.002	1.790	0.370	0.002	0.000
<u>B: Second Audit Wave</u>								
All Audited Firms	2.235	0.102	0.040	0.002	2.191	0.100	0.003	0.000
Amount Detected > 0	2.235	0.845	0.040	0.015	2.191	0.828	0.003	0.001
Size Quartile 1	0.045	10.205	0.002	0.473	0.042	9.649	0.000	0.000
Size Quartile 2	0.166	3.367	0.009	0.179	0.157	3.188	0.000	0.000
Size Quartile 3	0.217	0.889	0.009	0.036	0.205	0.840	0.003	0.012
Size Quartile 4	1.808	0.083	0.020	0.001	1.786	0.082	0.000	0.000

Notes: The table breaks down the total amount detected by audit (columns 1-2) into its three major components (columns 3-8). The odd-number columns report the amounts in PKR billions and the even-number columns the amount as a ratio of the aggregate annual turnover of the corresponding group of firm. Amount Recovered is the amount paid by the taxpayer as a result of audit. Amount Recoverable, on the other hand, is unpaid amount out of the total detected by audit. This amount is subject to quasi-judicial determination and appeal processes. Refund Curtailed indicates the amount by which the firm agreed to reduce its refund claim pending with the department.

TABLE A.III: SELECTION IN SEQUENCING OF AUDITS

	Outcome: Days between assignment and initiation			
	(1)	(2)	(3)	(4)
Sales	-1.785 (7.492)	-4.301 (7.489)	2.542 (2.749)	2.679 (2.657)
Purchases	-0.727 (8.569)	-3.636 (8.568)	-2.583 (5.626)	0.588 (5.433)
Output Tax	9.936 (30.624)	8.030 (30.012)	-2.929 (12.057)	0.718 (11.651)
Input Tax	-4.050 (14.118)	1.030 (13.982)	3.229 (11.034)	-2.713 (10.648)
Tax Paid	-6.673 (23.011)	-3.513 (22.538)	-1.108 (4.718)	-2.919 (4.554)
Exports	-0.550 (1.560)	-0.126 (1.540)	1.836 (1.002)	2.399 (0.974)
Imports	-0.201 (1.884)	-0.223 (1.916)	-0.370 (0.643)	-0.264 (0.624)
Refund	1.382 (1.395)	1.662 (1.377)	-1.847 (0.866)	-2.325 (0.840)
Carry Forward	1.734 (3.374)	1.132 (3.355)	-0.143 (0.569)	-0.300 (0.549)
Manufacturer	-13.271 (5.331)	-11.003 (5.298)	-1.860 (1.615)	-1.986 (1.581)
Importer	-0.785 (6.230)	-0.614 (6.190)	-3.302 (1.833)	0.310 (1.791)
Exporter	1.834 (9.390)	6.001 (9.301)	-1.649 (2.282)	-1.134 (2.295)
Distributor	7.098 (9.143)	9.746 (8.977)	-0.251 (2.469)	-1.645 (2.395)
Wholesaler	-5.548 (5.391)	-2.847 (5.315)	-1.848 (1.669)	0.958 (1.624)
Service Provider	-7.959 (5.332)	-4.111 (5.247)	0.109 (1.661)	1.141 (1.606)
Constant	46.995 (4.843)	44.436 (4.768)	18.961 (1.490)	17.830 (1.443)
Observations	3,482	3,481	3,612	3,611
Corporation FEs	Yes	Yes	Yes	Yes
Tax Office FEs	No	Yes	No	Yes

Notes: The table explores selection in audit. We regress the time lag measured in number of days between the assignment and initiation of audit on baseline firm characteristics. We standardize the first nine variables in this table by subtracting the mean and dividing by the standard deviation of the variable. Since audits were taken up by local tax offices, we include the tax office fixed effects in even-numbered columns. The first two columns report results for the first audit wave and the last two for the second audit wave. Standard errors are in parenthesis.

TABLE A.IV: PREEXISTING TRENDS

	First Wave					Second Wave				
	Sales	Purchases	Output	Input	Tax Payable	Sales	Purchases	Output	Input	Tax Payable
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$assign \times year \in [s - 1, s]$	-0.018 (0.015)	-0.005 (0.017)	-0.039 (0.020)	-0.004 (0.021)	-0.033 (0.025)	-0.016 (0.010)	-0.018 (0.011)	-0.027 (0.013)	-0.030 (0.014)	0.002 (0.017)
$assign \times year \in [s - 3, s]$	0.001 (0.014)	0.021 (0.016)	-0.031 (0.018)	0.021 (0.020)	-0.006 (0.022)	-0.006 (0.010)	-0.014 (0.012)	-0.005 (0.014)	-0.020 (0.014)	0.012 (0.017)
$assign \times year \in [s - 5, s]$	-0.004 (0.021)	0.042 (0.022)	-0.019 (0.022)	0.040 (0.024)	0.051 (0.033)	0.028 (0.011)	0.007 (0.012)	0.020 (0.013)	-0.002 (0.014)	0.056 (0.017)
Observations	2,324,186	2,025,380	1,672,095	1,681,583	1,154,574	2,628,878	2,290,848	1,934,273	1,945,733	1,312,928
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores if the preexisting trends for the five outcomes indicated in the heading of each column were parallel between firms who were picked for audit in a random ballot and other firms in the eligible sample. We estimate a model similar to (10) replacing the $assign \times after_{it}$ dummy with three dummies shown in the top three rows. The dummy variable $assign_i$ denotes that firm i 's was picked for audit in the random ballot indicated in the heading of the column. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The sample for these regressions include the baseline periods only, from July 2008 to August 2013 for the first wave and from July 2008 to August 2014 for the second. The dummy variable $year \in [s - 1, s]$ indicates that the period is one of the last twelve months included in the regression and so on. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.V: IMPACTS OF RANDOM AUDITS ASSIGNED IN THE FIRST WAVE

	Impacts After One Year					Impacts After Three Years				
	Sales	Purchases	Output Tax	Input Tax	Tax Payable	Sales	Purchases	Output Tax	Input Tax	Tax Payable
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<u>A: ITT Estimates</u>										
assign × after	-0.009 (0.016)	-0.009 (0.019)	-0.016 (0.021)	-0.017 (0.026)	-0.037 (0.027)	-0.007 (0.014)	-0.021 (0.019)	-0.025 (0.019)	-0.036 (0.023)	-0.015 (0.030)
Observations	2,802,387	2,456,864	2,061,472	2,089,489	1,393,541	3,809,614	3,315,994	2,857,885	2,895,330	1,890,220
<u>B: LATE Estimates</u>										
treat × after	-0.013 (0.022)	-0.014 (0.027)	-0.022 (0.029)	-0.024 (0.037)	-0.051 (0.036)	-0.010 (0.019)	-0.030 (0.027)	-0.035 (0.026)	-0.051 (0.031)	-0.021 (0.041)
Observations	2,802,387	2,456,864	2,061,472	2,089,489	1,393,541	3,809,614	3,315,994	2,857,885	2,895,330	1,890,220
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table estimates the impact of audit on firms' future behavior. In the top panel, the coefficient $\text{assign} \times \text{after}$ shows $\hat{\gamma}$ from model (10), where the dummy variable assign_i denotes that firm i 's audit was assigned through the first random ballot held on September 13, 2013. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable after_t indicates that month t falls after the date of the ballot. The sample includes periods up to October 2014 for the first five columns and periods up to October 2016 for the rest. Panel B shows the corresponding results from 2sls regressions, where the endogenous variable audit_i is instrumented by the initial random assignment. Standard errors are in parenthesis, which have been clustered at the tax office level.

TABLE A.VI: AUDIT IMPACTS – FIRST STAGE

Outcome:	<i>audit</i> × <i>after_{it}</i>					
	September 13, 2013		September 25, 2014		September 14, 2015	
Random Draw Held On:	One Year	Three Years	One Year	Three Years	One Year	Three Years
Post Sample:	(1)	(2)	(3)	(4)	(5)	(6)
<i>assign</i> × <i>after</i>	0.704 (0.007)	0.703 (0.007)	0.294 (0.004)	0.296 (0.004)	0.133 (0.004)	0.134 (0.004)
Observations	6,893,186	9,681,146	7,894,004	10721371	8,241,185	10829729
F Statistic	10,353	10,071	4,751	4,658	1,120	1,102

Notes: The table reports the first stage of our 2sls models. We estimate model (10) using the dummy $treat \times after_{it}$ as the outcome variable, where $treat_i$ takes the value 1 if firm i was audited in the corresponding audit wave indicated in the heading of each column. The coefficient $assign \times after$ shows $\hat{\gamma}$ from these regressions. The dummy variable $assign_i$ denotes that firm i 's audit was assigned through the random ballot indicated in the heading of each column. The sample includes the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. We report results for two Post Samples: One Year specifications include twelve $after_t$ periods and Three Years specifications include 36 $after_t$ periods. In each case, the samples includes all months from July 2008 to the last $after_t$ period. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.VII: IMPACTS OF RANDOM AUDITS ASSIGNED IN THE THIRD WAVE

	Impacts After One Year					Impacts After Three Years				
	Sales	Purchases	Output Tax	Input Tax	Tax Payable	Sales	Purchases	Output Tax	Input Tax	Tax Payable
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
assign × after	-0.034 (0.011)	-0.024 (0.013)	-0.039 (0.014)	-0.009 (0.014)	0.004 (0.014)	-0.050 (0.011)	-0.040 (0.014)	-0.071 (0.015)	-0.076 (0.015)	-0.093 (0.015)
Observations	3,007,568	2,590,734	2,256,294	2,265,080	2,758,303	3,910,133	3,341,025	2,879,242	2,930,477	3,577,794
B: LATE Estimates										
treat × after	-0.261 (0.083)	-0.185 (0.102)	-0.296 (0.106)	-0.063 (0.105)	0.033 (0.108)	-0.376 (0.087)	-0.297 (0.106)	-0.487 (0.110)	-0.527 (0.108)	-0.652 (0.112)
Observations	3,007,568	2,590,734	2,256,294	2,265,080	2,758,303	3,910,133	3,341,025	2,879,242	2,930,477	3,577,794
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table estimates the impact of audit on firms' future behavior. In the top panel, the coefficient $\text{assign} \times \text{after}$ shows $\hat{\gamma}$ from model (10), where the dummy variable assign_i denotes that firm i 's audit was assigned through the first random ballot held on September 14, 2015. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable after_t indicates that month t falls after the date of the ballot. The sample includes periods up to October 2016 for the first five columns and periods up to October 2018 for the rest. Panel B shows the corresponding results from 2sls regressions, where the endogenous variable audit_i is instrumented by the initial random assignment. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.VIII: PREEXISTING TRENDS – AUDITED VS. NOT AUDITED

	First Wave					Second Wave				
	Sales	Purchases	Output	Input	Tax Payable	Sales	Purchases	Output	Input	Tax Payable
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\text{treat} \times \text{year} \in [s - 1, s]$	0.019 (0.016)	0.038 (0.018)	-0.016 (0.021)	0.022 (0.020)	-0.046 (0.026)	0.001 (0.015)	0.020 (0.019)	0.022 (0.023)	-0.024 (0.022)	-0.007 (0.028)
$\text{treat} \times \text{year} \in [s - 3, s]$	0.070 (0.016)	0.074 (0.017)	0.006 (0.019)	0.071 (0.019)	0.029 (0.024)	0.003 (0.015)	0.011 (0.019)	0.037 (0.023)	0.029 (0.022)	-0.006 (0.027)
$\text{treat} \times \text{year} \in [s - 5, s]$	0.089 (0.024)	0.066 (0.022)	0.011 (0.022)	0.066 (0.024)	0.098 (0.033)	0.034 (0.018)	0.028 (0.019)	0.054 (0.022)	0.064 (0.022)	0.025 (0.028)
Observations	2,324,186	2,025,380	1,672,095	1,681,583	1,154,574	2,628,878	2,290,848	1,934,273	1,945,733	1,312,928
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table explores if the preexisting trends for the five outcomes indicated in the heading of each column were parallel between audited and unaudited firms. We estimate a model similar to (10) replacing the $\text{assign} \times \text{after}_{it}$ dummy with three dummies shown in the top three rows. The dummy variable treat_i denotes that firm i was audited in the wave indicated in the heading of the column. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The sample for these regressions include the baseline periods only, from July 2008 to August 2013 for the first wave and from July 2008 to August 2014 for the second. The dummy variable $\text{year} \in [s - 1, s]$ indicates that the period is one of the last twelve months included in the regression and so on. Standard errors are in parenthesis, which have been clustered at the firm level.

TABLE A.IX: HETEROGENEITY IN RESPONSE WITH RESPECT TO AMOUNT DETECTED

	Sales	Purchases	Output Tax	Input Tax	Tax Payable
	(1)	(2)	(3)	(4)	(5)
<u>A: First Wave</u>					
assign × after	-0.009 (0.019)	-0.016 (0.021)	-0.020 (0.025)	-0.029 (0.026)	0.004 (0.031)
assign × after × trait	0.009 (0.040)	-0.023 (0.048)	-0.022 (0.052)	-0.031 (0.054)	-0.089 (0.070)
Observations	3,839,502	3,328,628	2,884,225	2,906,045	1,913,096
<u>B: Second Wave</u>					
assign × after	-0.014 (0.011)	-0.019 (0.013)	-0.016 (0.013)	-0.009 (0.013)	0.005 (0.017)
assign × after × trait	0.040 (0.031)	0.119 (0.041)	0.053 (0.042)	0.038 (0.039)	0.010 (0.048)
Observations	4,390,478	3,791,277	3,262,221	3,313,664	2,151,912
Firm FEs	Yes	Yes	Yes	Yes	Yes
Period FEs	Yes	Yes	Yes	Yes	Yes

Notes: The table explores heterogeneity in the audit effect. We divide firms into two groups. Firms against whom a positive amount was detected by audit are included in one group (indicated by the dummy variable $trait_i$); the rest of the firms are included in the baseline category. We then estimate a triple-difference version of model (10). The model includes interactions of the $trait_i$ dummy with the $assign \times after_{it}$ dummy. The $assign_i$ dummy takes the value 1 if the firm's audit was assigned in the corresponding random computer ballot. The control group comprises the rest of the firms in the eligible sample. The eligible sample consists of the population of VAT filers excluding government departments and firms already under audit. The dummy variable $after_t$ indicates that month t falls after the date of the ballot. The coefficients and the 95 percent confidence intervals on the double and triple-interaction terms from these regressions are plotted. Regressions are run separately for the first and the second audit waves. Standard errors are clustered at the firm level.

TABLE A.X: DETECTED AMOUNT BY SHARE FINAL SALES

	Outcome: Amount Detected (Std. Deviations)					
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Share Final Sales</u>						
2nd Quartile	-0.100*	-0.099**	-0.097*	-0.098*	-0.105**	-0.096**
	(0.051)	(0.051)	(0.050)	(0.050)	(0.052)	(0.048)
3rd Quartile	-0.094*	-0.091*	-0.085*	-0.090*	-0.098*	-0.086*
	(0.051)	(0.050)	(0.047)	(0.050)	(0.052)	(0.046)
4th Quartile	-0.101**	-0.097*	-0.090**	-0.085*	-0.108*	-0.085*
	(0.051)	(0.049)	(0.046)	(0.044)	(0.056)	(0.045)
Observations	6,561	6,561	6,561	6,560	6,548	6,547
<u>B: Share (Final Sales + Purchases from Unregistered Sector)</u>						
2nd Quartile	-0.085	-0.082	-0.076	-0.081	-0.088	-0.074
	(0.052)	(0.051)	(0.048)	(0.051)	(0.053)	(0.046)
3rd Quartile	-0.108**	-0.087*	-0.094**	-0.083*	-0.113**	-0.074*
	(0.052)	(0.045)	(0.045)	(0.043)	(0.057)	(0.042)
4th Quartile	-0.113**	-0.086**	-0.095**	-0.086**	-0.118**	-0.076*
	(0.052)	(0.044)	(0.043)	(0.043)	(0.059)	(0.043)
Observations	6,561	6,561	6,561	6,560	6,548	6,547
Size FEs	No	Yes	No	No	No	Yes
Production Stage FEs	No	No	Yes	No	No	Yes
Tax Office FEs	No	No	No	Yes	No	Yes
Industry FEs	No	No	No	No	Yes	Yes

Notes: The table shows that the amount detected by audit falls with the baseline share of final sales in a firm's turnover. The outcome variable here is the amount detected by audit, normalized by its standard deviation. To maximize power, we pool together audits conducted in the first two waves. Final sales are defined as sales where the other party to the transaction does not possess a national tax number: they are either consumers or informal firms. We divide firms into four quartiles based on the share of final sales in their turnover at the baseline. We regress the outcome variable on the three quartile dummies, omitting the first quartile as the reference group. We successively introduce the controls indicated in the last four lines. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

TABLE A.XI: DETECTION PROBABILITY BY SHARE FINAL SALES

	Outcome: 1 (Amount Detected > 0)					
	(1)	(2)	(3)	(4)	(5)	(6)
<u>A: Share Final Sales</u>						
2nd Decile	0.004 (0.016)	0.002 (0.016)	0.009 (0.016)	-0.016 (0.014)	0.004 (0.016)	-0.019 (0.014)
3rd Decile	-0.016 (0.016)	-0.019 (0.016)	-0.002 (0.016)	-0.024* (0.014)	-0.002 (0.016)	-0.029* (0.015)
4th Decile	0.022 (0.016)	0.026 (0.016)	0.038** (0.016)	-0.025* (0.014)	0.040** (0.016)	-0.017 (0.015)
Observations	6,561	6,561	6,561	6,560	6,548	6,547
<u>B: Share (Final Sales + Purchases from Unregistered Sector)</u>						
2nd Decile	-0.014 (0.017)	-0.017 (0.017)	-0.001 (0.017)	-0.019 (0.014)	0.000 (0.017)	-0.019 (0.015)
3rd Decile	-0.014 (0.016)	0.005 (0.017)	0.008 (0.017)	-0.053*** (0.014)	-0.001 (0.017)	-0.032** (0.015)
4th Decile	-0.053*** (0.016)	-0.035** (0.017)	-0.023 (0.017)	-0.056*** (0.014)	-0.040** (0.017)	-0.048*** (0.016)
Observations	6,561	6,561	6,561	6,560	6,548	6,547
Size FEs	No	Yes	No	No	No	Yes
Production Stage FEs	No	No	Yes	No	No	Yes
Tax Office FEs	No	No	No	Yes	No	Yes
Industry FEs	No	No	No	No	Yes	Yes

Notes: The table shows that the probability of detection falls with the baseline share of final sales in a firm's turnover. The outcome variable here is a dummy indicating that audit detects a positive amount against the firm. To maximize power, we pool together audits conducted in the first two waves. Final sales are defined as sales where the other party to the transaction does not possess a national tax number: they are either consumers or informal firms. We divide firms into four quartiles based on the share of final sales in their turnover at the baseline. We regress the outcome variable on the three quartile dummies, omitting the first quartile as the reference group. We successively introduce the controls indicated in the last four lines. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.