Limits to VAT Self-Enforcement: Role of Network Frictions in Compliance*

Bhanu Gupta† Tejaswi Velayudhan‡ Prince Dhawan§ Ankur Garg ¶

October 30, 2025

Abstract

The value-added tax (VAT) is popular among tax administrations for its "self-enforcing" property: it creates a strategic compliance complementarity in firm-to-firm transactions. Existing work emphasizes the resulting positive compliance spillovers along production chains. This paper shows that in settings with widespread evasion and frictions in production networks, the same strategic complementarity can generate adverse effects. Using administrative tax records and a randomized experiment in New Delhi that nudged firms to pay VAT on time, we find that non-compliance among upstream suppliers can cause targeted firms to remain non-compliant despite stronger enforcement. Even firms that do comply do not induce their non-compliant suppliers to change behavior; instead, they shift evasion to other margins. We provide suggestive evidence that high costs of influencing supplier behavior drive these responses. The results show that the effectiveness and distributional incidence of VAT enforcement depend critically on the structure of production networks.

Keywords: Value-added tax; Tax Filing; Firm Networks; Tax Enforcement

JEL Codes: H25, H26, L14, O12

*The authors would like to thank the Delhi Department of Trade and Taxes for their cooperation in making this project possible. We are grateful to Anne Brockmeyer, Giulia Mascagni, Karthik Muralidharan, David Agrawal, Caroline Weber, Mazhar Waseem and seminar participants at the Zurich Conference in Public Finance, University of Michigan, and UC Irvine for their insightful comments. We thank Devashish Deshpande, Aditya Dhabu, Smriti Saini, and the Center for Effective Governance of Indian States (CEGIS) team for their assistance in data procurement and management. We thank Vedant Vohra and Taha Barwahwala for excellent research assistance. We gratefully acknowledge financial assistance from J-PAL, CEGIS, and the International Center for Tax and Development (ICTD). The project has approval from the IRB at Ashoka University (IRB Reference: 18_21_Gupta). The RCT has been registered on the AEA RCT Registry (AEARCTR-0008357).

†Department of Economics, Ashoka University. Email: bhanu.gupta@ashoka.edu.in

‡Department of Economics, University of California, Irvine. Email: tvelayud@uci.edu

§Department of Trade and Taxes, Government of NCTD

¶Department of Trade and Taxes, Government of NCTD

1 Introduction

The Value-Added Tax (VAT) is one of the largest contributors to government revenue in the world and a cornerstone of tax policy in countries with limited state capacity. A central appeal of the VAT is its self-enforcement property; that the same transaction provides tax credits to the buyer and creates a tax liability for the seller, generating independent reports of the same transaction, asymmetric incentives for the transacting parties, and tax withholding. These are all features that are proven to improve tax compliance (Waseem, 2022).

The first mechanism—generating third-party information—still requires the tax authority to intervene to identify and resolve discrepancies in reports. The infeasibility of comprehensive manual verification was recognized early on as a practical limitation of the VAT that could be overcome with technological advances (Tait, 1991). These advances have largely been realized. Innovations like e-filing, e-invoicing and real-time cross-verification are now widely applied across the world and tie the receipt of credits to the provision of accurately reported information by the seller. If a supplier fails to file, underreports sales on the e-invoice or makes other errors, the buyer's credits are denied, increasing the buyer's liability. This forces buyers to monitor or pressure their suppliers to comply, or to replace them, tasks that would otherwise fall to the tax authority.

In this paper, we show that in a system with comprehensive, real-time cross-checks, even buyers who wish to comply may choose not to if their suppliers are non-compliant. When a buyer's tax credits depend on the supplier's filing, a strategic complementarity is created: a compliant buyer increases the cost of supplier non-compliance, while a non-compliant supplier increases the cost of buyer compliance. Faced with a non-compliant supplier, the buyer must either forgo credits and accept a higher net tax liability, or induce the supplier to comply. If the supplier cannot be influenced, for example due to strong market power, the buyer may choose to lower their liability by underreporting their own sales instead.

We show evidence of such an adverse network effect on compliance in the context of the Goods and Services Tax (GST) in Delhi, India, where buyers can claim input tax credits only for invoices reported by their suppliers. We combine administrative GST registration and filing data from 2020–2022 with invoice-level records that link each buyer to its suppliers. In this setting, late filing is directly observable and equivalent to late payment, allowing us to measure an evasion margin that is usually hidden. We use a randomized voice-message experiment to perturb perceived enforcement among taxpayers with a history of late payment. Buyers with delinquent suppliers do not change their payment timing in response to enforcement but do become more punctual on returns that are not linked to supplier behavior. Crucially, this result cannot be explained by other buyer characteristics and holds when we use experimentally-induced variation in supplier compliance. Even among buyers with fewer delayed credits, who do respond to enforcement, supplier behavior is unchanged. Instead, buyers underreport output and lower their net tax liability. Reducing delinquency decreases compliance on another margin.

We present a framework of the taxpayer's decision-making at the payment deadline when they must choose when to pay, what sales to report and whether to influence their suppliers after considering supplier behavior and associated costs of misreporting, nudging suppliers, and perceived penalties. Supplier behavior determines available input tax credits and hence, the net liability of the taxpayer. Because it is costly to influence non-compliant suppliers' behavior, taxpayers with higher costs of nudging suppliers are more sensitive to supplier compliance. As a result, higher perceived penalties will nudge taxpayers into filing before the deadline only if they overcome the costs associated with misreporting and nudging suppliers. When taxpayers switch to filing on time, they may become non-compliant on other margins such as reported sales. We test these predictions by experimentally varying perceived penalties.

In our experiment, taxpayers were randomly assigned to receive automated voice messages either reminding them of filing deadlines (reminder), warning of penalties (deterrence), or no message (control). These messages increased on-time payment by 8 percent on average. However, this average masks stark heterogeneity: Buyers with high supplier compliance at baseline (above-median share of timely credits) responded strongly to nudges, decreasing late-filing rates by 7.8 percent. Buyers with low supplier compliance decrease late-filing by only around 1 percent which is not a statistically significant change. This difference in treatment effects across the two categories of taxpayers is statistically significant. We address SUTVA violation concerns and estimate spillovers by controlling for actual and expected exposure of suppliers to treatment following Borusyak and Hull (2023). Exploiting experimentally induced variation in supplier compliance, we show that this heterogeneity is likely causal. A 10 percentage point increase in the share of suppliers treated increases taxpayer response by 1.25 percentage points.

We further rule out that network compliance is a proxy for other buyer characteristics such as sorting of taxpayers into compliant production networks by exploiting a feature of the institutional structure. The same taxpayer must file two returns, but only one is linked to payments and therefore to seller behavior. The other return is used to transfer credits and is not dependent on supplier compliance. Our framework predicts that taxpayers should only be affected by upstream compliance when it affects their own tax liability, i.e., for the payment-linked return but not the credit-linked return. In line with our prediction, taxpayers start filing their credit-linked return on time in response to nudges but this response does not vary by supplier compliance. To ensure that this difference is not driven by taxpayer characteristics, we estimate the differential impact of treatment on the two returns within taxpayers by including taxpayer fixed effects. Taxpayers with high network compliance are more likely to file their payment-linked returns on time relative to their credit-linked returns while those with low network compliance are either equally likely or slightly less likely to file their payment-linked returns on time in response to treatment.

Responding taxpayers often reduced reported sales to avoid higher net payments. When choosing to file on time with missing credits, taxpayers are 4.5 percent more likely to report zero output tax liability and 7 percent less likely to report any increase, resulting in a *decrease* in net liability despite claiming fewer credits. A natural explanation for the decrease in output liability is evasion since these responses occur after production decisions have already been made. A nudge to increase compliance on one margin can lead to evasion on a different margin if buyers cannot change their suppliers' behavior.

Across the board, buyers rarely switched away from non-compliant suppliers or changed the behavior of existing suppliers, showing the limits of a key channel through which the

VAT is traditionally thought to propagate compliance. Our study sample consists of buyers who were choosing to trade with non-compliant sellers even in a system where this non-compliance is costly to the buyer. The switching frictions and non-response to enforcement were largest among buyers who face concentrated input markets where the cost to change supplier behavior is plausibly higher.

These results show that when invoice matching links buyers' credits to their suppliers' compliance, weak upstream compliance can blunt or even nullify the effects of downstream enforcement. The findings underscore the role of production networks in shaping the incidence and effectiveness of VAT enforcement, and they highlight the potential for switching frictions and market structure to create persistent pockets of evasion.

A key challenge to empirically examining the effect of supplier non-compliance is the classic difficulty in tax evasion research: that evasion is often hidden. Among registered firms, the impact of an enforcement intervention is generally measured by an increase in reported liability rather than a decrease in evasion because the underlying amount of evasion is unknown.[1] Researchers often cannot distinguish between cases where the firms fail to report higher liability because they are already truthful and cases where they continue to evade. By focusing on an observable margin of non-compliance, delayed payment, we overcome this challenge.

Our work most directly contributes to the literature on the interaction between the VAT and production chains. Seminal work by Pomeranz (2015) showed that the tax authority can "jump-start" compliance all along the supply chain by increasing enforcement at poorly-enforced, down-stream stages. The findings in our paper suggest that this type of positive spillover may not work in all production chains. A contemporaneous paper, Almunia et al. (2023), takes the interaction between firms within a supply chain more seriously and shows that when the two parties make independent reports, nudging just one may be less effective. Unlike in the Ugandan setting they examine, we can clearly identify the non-compliant party, which enables us to demonstrate that the non-response by the targeted taxpayer is due to non-compliance by the supplier. By showing this generalizable dynamic within the

[1] See for example Pomeranz (2015) on VAT evasion, Antinyan and Asatryan (2024) for a survey.

VAT, our paper reveals a limitation to relying on VAT production networks to spread compliance. Related work documents the role of networks VAT registration: firms sort into chains of VAT registration and non-registration (De Paula and Scheinkman, 2010; Gadenne et al., 2023). We demonstrate strategic complementarities in compliance *among* registered firms. More broadly, networks have been shown to be important in tax compliance as taxpayers learn through their peers (Boning et al., 2020; Paetzold and Winner, 2016; Drago et al., 2020). Under the VAT, network effects arise through liability linkages (input-credit claims).

A growing literature emphasizes studying the VAT "as implemented," including frictions in forming and reconfiguring trading relationships (Brockmeyer et al., 2024; Slemrod and Velayudhan, 2022). Our evidence that buyers often do not switch away from non-compliant suppliers aligns with firm-to-firm search and matching frictions (Eaton et al., 2022; Fontaine et al., 2023; Dhyne et al., 2022; Miyauchi, Forthcoming; Bernard et al., 2019). These frictions imply that network position and market structure shape the effectiveness and incidence of VAT enforcement.

Finally, our findings connect to work on tax administration technology. While self-enforcement is powerful, tax gaps persist via over-claimed refunds, ghost firms, and misclassification (Waseem, 2023; Carrillo et al., 2022; Fisman and Wei, 2004); even after detection, collection is hard (Best et al., 2021; Barwahwala et al., 2024). Proposed fixes—e-filing, e-invoicing, and electronic billing machines—can help (Okunogbe and Pouliquen, 2022; Bellon et al., 2022; Fan et al., 2018; Eissa et al., 2014), and invoice-matching is increasingly used to police refunds (Shah, 2020). We show a downside: when upstream firms fail to file, invoice-matching shifts risk and compliance costs onto downstream buyers, dampening the impact of authority-led enforcement.

From a policy perspective, the results suggest that enforcement design should account for taxpayers' position in the supply chain. Downstream-only enforcement may fail when upstream compliance is weak, shifting enforcement costs onto buyers and creating horizontal inequities—especially for firms in concentrated supplier markets or low value-added sectors. Strengthening upstream enforcement or lowering switching costs could preserve the informational advantages of invoice matching while mitigating these unintended burdens.

The rest of this paper is organized as follows. We describe the relevant aspects of the Indian GST in Section 2. We describe our data in Section 3. In Section 4, we present a simple conceptual framework underlying firms' decisions and its dependence on their network's behavior . Our experimental intervention and estimation specifications are described in Section 5. Section 6 presents our results. Finally, Section 7 concludes.

2 Tax Filing and Payment in the GST

India introduced the Goods and Services Tax (GST) in 2017 in a major reform that replaced several commodity taxes, including state-specific VATs, with a nationwide VAT. From the beginning, its administration was digitally sophisticated and mandated e-filing of all returns. The main returns – known as the GSTR-3B and GSTR-1 forms – are filed on either a monthly or quarterly basis based on a taxpayer's annual revenue.[2] The GSTR-3B form is a self-reported summary of their sales, purchases, and tax liability, and is filed simultaneously with tax payment. We refer to this as the *payment-linked* form as late-filing also means late-payment. The GSTR-1 form contains invoice-level details of a firm's transactions with other GST-registered firms. We refer to this form as the *credit-linked* form, as a taxpayer's buyers only receive tax credits for transactions reported on this form.

When a supplier files their credit-linked form (the GSTR-1), the buyers' "GSTR-2" forms are auto-populated with this information. By downloading this form, a buyer can immediately know what credits are available to them and from whom. The buyer then has to nudge their sellers to file to receive their credits in time to make their payment. There is no simple recourse for the buyer to independently report credits owed or to alert the tax authorities to their suppliers'non-compliance. In fact, the role of the buyer in ensuring supplier compliance is so entrenched—despite taxpayers' grievances and litigation—that commercial accounting software specifically advertises the automation of supplier nudges as a selling point to taxpayers.[3]

[2]Registration is mandatory for manufacturers with annual revenue above Rs. 4 million and for service firms with annual revenue above Rs. 2 million. The GST is broad-based with a few commodity and use-specific exemptions. Since January 2021, taxpayers with annual turnover less than Rs. 50 million can file their GSTR-1 and GSTR-3B quarterly and all others must file monthly.

[3]https://cleartax.in/s/max-itc

While such a policy may seem unique to India, all VAT systems place some enforcement burden on trading partners as they must choose how to reconcile different information reports about the same transaction. At one extreme, all taxpayers self-report and any discrepancy must be resolved on a case-by-case basis using the resources of the tax authority. In settings where these resources are limited, digital solutions are already filling the gap in a way that pushes the burden onto taxpayers. For example, Pakistan employs real-time algorithm-based invoice verification to approve or disapprove input tax credit claims (Shah, 2023; Fan et al., 2018). These solutions seem appealing in light of vast, easily detectable discrepancies in other VAT systems with digitized return information (Almunia et al., 2024). In other high-income settings, such as Germany and Canada, tax authorities require buyers to exercise 'due diligence' in verifying the compliance of their suppliers.[4,5] Failure to do so can lead to extensive audits and ITC denial even when no fraud was committed by the buyer. The self-enforcing mechanism shifts the enforcement burden to some degree to the buyer. India's invoice-matching system simply exacerbates this dynamic.[6]

Besides creating incentives within the supply chain to increase compliance, the Indian tax authority employs additional measures to enforce filing compliance. There are several monetary and non-monetary enforcement measures in place to encourage filing, some of which were imperfectly implemented in practice during our study period (see Appendix C for details of penalties). For example, late-filing penalties and interest were only paid by a small fraction of delinquents (see Figure A1). On-time filing rate of both forms are about 70 percent in New Delhi. For comparison, on-time VAT filing rate is on average 77 percent across 117 countries, with even some lower middle-income countries reporting 100 percent on-time filing.[7] Contexts with weaker enforcement environments and high evasion among upstream firms are likely to see stronger adverse network effects than what we observe in India.

One additional institutional feature makes filing compliance important in the Indian con-

[4]https://www.lhp-group.com/news/article/denial-of-input-tax-deduction-even-in-the-case-of-unintentional-participation-in-vat-evasion-in-the-supply-chain/

[5]https://marcil-lavallee.ca/en/bulletin/gst-hst-risks-of-dealing-with-a-shady-supplier-2/

[6]Even in settings where the two parties in a transaction can make separate reports, unilateral action by one party can result in a discrepancy, which could be undesirable for both even if the consequences are not as direct. For example, a seller's failure to report a transaction could be construed as an overreport by the buyer and require documentation or some other onerous participation in an audit.

[7]Source: IMF RA-FIT statistics https://data.imf.org/?sk=4B1DFA79-F9A6-4FFD-8ED0-C41DDA7252F6

text. GST revenue is shared between the central and state governments. As in many federal countries, revenue sharing is a contentious and sensitive issue.[8] Although revenue is collected potentially across many states along the production chain, it is ultimately allocated to the state where final consumption occurs. GST filings provide the supply chain information to accurately apportion revenue across states. Delinquent filing reduces the accuracy of this information and leaves large amounts of revenue unallocated.

Our study takes place within New Delhi, one of the largest metro cities in India, which accounts for about 5-6 percent of total GST revenue. As of July 2021, 32 percent of the Delhi-registered taxpayers made late payments, which is slightly lower than the national average of 35 percent.[9] Taxpayers in New Delhi are typical of large, urban centers – hubs of manufacturing, services, and retail rather than agriculture. Characteristics of taxpayers in our sample are further described in Section 3. In terms of GST administration, Delhi taxpayers face the same rules, use the same filing portal and file the same forms as all other GST taxpayers since policy is set centrally.

3 Data

We use administrative data from registrations and tax filings of all GST taxpayers registered in New Delhi between December 2020 - May 2022. It includes business details provided at registration such as the sector of business, phone number, etc., sales and purchases reported on GSTR-1 and GSTR-3B forms. Crucially, we can construct supply networks from this data including the value of transactions between taxpayers.

Our experimental intervention targeted the subset of these registered taxpayers with a history of filing delinquency, chosen because these taxpayers are likely to be delinquent again. About 78% of taxpayers who file their return late in a given month also file late in the subsequent month. In contrast, only 25% of taxpayers who file on time in a given month, fail to do so in the subsequent month. Of 341,854 taxpayers in our full sample who were registered as

[8]https://www.business-standard.com/article/current-affairs/cag-pulls-up-govt-for-erroneous-process-of-igst-devolution-to-states-121112901302_1.html

[9]These data can be accessed on https://www.gst.gov.in/download/gststatistics

on July 2021, about 60 percent had filed their payment-linked form late or did not file it all, at least once between January and June 2021.[10]

Two additional restrictions were necessitated by the practical logistics of the intervention: 1) Because we worked with the Delhi state commercial tax authority, we limited the intervention to taxpayers directly administered by the state[11], 2) Because the intervention was delivered through automated voice calls to the taxpayer's registered phone numbers we drop any taxpayers that shared a number with another taxpayer in our sample.[12] These restrictions leave us with 284,240 unique taxpayers who are similar in most respects to the full population of taxpayers.

Limiting our sample to state-administered taxpayers implies that the average firm in our sample is slightly smaller than the average GST-taxpayer in Delhi. Appendix Table A2 compares characteristics of taxpayers registered with the state versus the center. State-registered taxpayers are about 7 percentage points more likely to have inter-state sales, and have, on average, been registered for about 6 months longer. Overall differences between the two populations are economically small though statistically significant, particularly in terms of filing compliance. Both groups show similar patterns of late filing – about 50% for GSTR-3B and 78% for GSTR-1. Around 20% taxpayers in both groups have a pending tax return. Propensity to share a phone number with another taxpayer is also similar, at around 38%.

Limiting our sample to taxpayers with unique phone numbers is unlikely to affect the external validity of our findings. Appendix Table A3 compares taxpayers with and and without a shared phone number, among state-administered taxpayers. Again differences here are statistically significant though economically small – even smaller than differences between central government and state taxpayers. The two groups are similar in age, inter-state sales, size of network, and other characteristics. Compliance behavior among the two groups is also very similar with differences of no more than 1 percentage point in ever-filing and on-time

[10] We exclude a few very large taxpayers who sell to a large number of taxpayers (over 200). This was to limit treatment spillovers.

[11] In the GST, taxpayers are partially randomly assigned to be administered either by the state revenue authority or the central revenue authority. Both tax authorities can access data on all taxpayers registered within the state and could even take enforcement action against them but they are the primary adminstrative authority over the assigned subset.

[12] Taxpayers may have the same registered phone number if they listed the number of their accountant, for example, or if they are businesses with a common owner.

filing rates. Our measures of supplier network characteristics are constructed including all suppliers of the targeted taxpayers who are registered in Delhi.[13]

The final data is a taxpayer-by-month dataset of all registered firms who were required to file a GSTR-1 or GSTR-3 form during this time, their filing outcomes, and tax payments. In addition, invoice-level information reported on GSTR-1 forms by taxpayers enables us to link buyers and suppliers, and quantify the strength of their relationship based on the value of transactions between them. We use baseline characteristics of the taxpayer (filing compliance behavior, type of business, turnover, etc.) as measured in July 2021 to select our target population and assign them to treatment groups as described in Section 5. Our outcomes are measured from August 2021 onwards. A key explanatory variable is the compliance level of the supplier network, defined as the share of input tax credits that are available on time. This variable is constructed taking a weighted average of dummies for whether each supplier filed their credit-linked return on time in July 2021, weighted by the amount of credits received from that supplier. Details of variable construction for this and other key variables used in the analysis are given in Appendix A.

In our main analysis, we limit attention to taxpayers whose network pre-treatment is composed mainly of sellers registered in Delhi.[14] We do so because we can only characterize the compliance behavior of taxpayers registered in Delhi on whom we have data. A robustness check re-weighting our analysis sample to match characteristics of the excluded sample reassures us that this sample restriction does not affect external validity of our results. Since our focus is on estimating the impact of network compliance on enforcement efficiency, we further restrict attention to taxpayers with some non-compliance in the seller network. Table 1 describes the characteristics of our main analysis sample, which is similar to the full sample (provided in Appendix Table A1) along several characteristics like average age of registered

[13] Many taxpayers both buy from and sell to other GST-registered entities and as such can be both "buyers" and "sellers" in our analysis. We limit and organize our sample based on the characteristics of taxpayers as buyers, but any outcomes and characteristics of the seller network of these buyers will include all sellers. For example, a taxpayer A who filed all their returns on time but sold to a delinquent taxpayer B is excluded from our analysis sample. But our measure of taxpayer B's (who *is* included in the sample) supplier network compliance will include taxpayer A's compliance.

[14] We restrict the sample to taxpayers who receive at least 70 percent of input tax credits claimed from Delhi-registered sellers. They are registered in Delhi but may be assigned to administered by either the state or central tax authority.

firms, sector of operation, number of sellers in the network, and sales.

Table 1: Summary Statistics of Main Sample

	Mean	Median	Std. Dev.	Min.	Max.	Obs.
Previous late-filer	1.00	1.00	0.00	1.00	1.00	26,877
Share Sellers State-Registered	0.62	0.66	0.30	0.00	1.00	26,877
Network Compliance, Jul 2021	0.63	0.75	0.35	0.00	1.00	26,877
GSTR-1 Filed On Time	0.77	1.00	0.42	0.00	1.00	26,877
GSTR-3B Filed On Time	0.80	1.00	0.40	0.00	1.00	26,877
Reminder	0.15	0.00	0.36	0.00	1.00	26,877
Deterrence	0.60	1.00	0.49	0.00	1.00	26,877
Retailer	0.40	0.00	0.49	0.00	1.00	26,877
Upstream	0.29	0.00	0.45	0.00	1.00	26,877
Services	0.10	0.00	0.30	0.00	1.00	26,877
Months Registered	43.81	50.00	12.96	6.00	50.00	26,877
Has Inter-state sales	0.84	1.00	0.37	0.00	1.00	26,877
Number of Sellers from Delhi	16.76	11.00	21.67	1.00	692.00	26,877
Log(Output Liability)	11.40	11.55	1.86	0.00	20.72	26,027
Zero Net Liability	0.44	0.00	0.50	0.00	1.00	26,877
Higher Net Liability	0.32	0.00	0.47	0.00	1.00	26,877
Zero Output Liability	0.03	0.00	0.18	0.00	1.00	26,877
More Output Liability	0.46	0.00	0.50	0.00	1.00	26,877
More ITC claimed	0.50	1.00	0.50	0.00	1.00	26,877
Log(ITC Claimed)	11.36	11.51	1.87	-1.90	20.61	25,807
Prob. of Repeating Any Late Seller	0.35	0.00	0.48	0.00	1.00	25,382
Prop. Late Seller Repeated	0.28	0.00	0.41	0.00	1.00	25,382
Share Seller On-time	0.86	0.92	0.19	0.00	1.00	26,049
HHI of seller Network in July	0.37	0.29	0.25	0.02	1.00	26,877

Notes: This table shows the summary statistics for our main analysis sample defined as having some history of late-filing, having some non-compliance in the seller network in July 2021 (i.e. at least one seller filed their GSTR-1 late in that month), and receiving at least 70 percent of their input tax credits from Delhi sellers in July 2021. The summary statistics of the full sample are provided in Appendix Table A1.

4 Conceptual Framework

After its production activities have been completed for the month, a firm decides when to make their tax payment (i.e. file their payment return), how much revenue to report on their return, and how much effort to exert to influence their supplier's behavior to minimize their own tax costs.

The firm is owed input tax credits of τM, given the GST rate τ and intermediate input purchases M. Only a share μ_0 of these credits are available on time if they do not exert effort to change their supplier's behavior. The late firm can exert effort e to increase the share of credits available by the payment deadline to $\mu(e)$. Let $\alpha(e)$ represent the cost of exerting effort, which is increasing and convex in e. These costs are general and could include the existence of relationship-specific investments, or of using sophisticated accounting software.

Firms earn true revenue R and choose how much to report, $\hat{R}$. We assume that the marginal cost of underreporting revenue by an amount $x = R - \hat{R}$ is higher if a firm files their return late, reflecting the possibility that late-filing attracts additional attention and a higher probability of detection. That is, taxpayers face evasion cost of $\delta_l(x)$ if they file late and $\delta_o(x)$ if they file on time where both are convex functions with $\delta_l' > \delta_o' \quad \forall x$. It follows trivially therefore that firms will always underreport by more when they file on time.[15]

Total cost if they file on time and exert effort e to get more credits on time is:

$$T_o = \tau(\hat{R} - \mu(e)M) + \delta_o(x) + \alpha(e) \tag{1}$$

Alternatively, total tax cost if they file late and wait for suppliers to provide their credits is:

$$T_l = \tau(\hat{R} - M) + \delta_l(x) + \theta P \tau(\hat{R} - M) \tag{2}$$

where θ is the probability of late penalties P being enforced on the delayed payment $\tau(\hat{R} - M)$.

Firms choose to file on time if $T_o^* < T_l^*$, where the asterix denotes optimized values. Subscripts l and o denotes the optimal choices under late-filing and on-time filing for all choice variables.[16]

[15]Taking the first order conditions with respect to x yields the usual condition that x is set to equate the margnal benefit of evasion, τ to the marginal cost, δ'. Since $\delta_l' > \delta_o'$, $x_l^* < x_o^*$.

[16]Buyers may eventually receive the missing credit and could apply it to future tax revenue. However, non-exporters must carry forward excess credits and cannot receive a refund so future credits may be less valuable if they do not have output tax liability to offset. Working capital constraints also make credit delays costly. Moreover, there is a possibility they may never receive the credits if suppliers never file or underreport sales. If taxpayers believe they will eventually receive and use the credits, the present value of tax costs of filing on time is lower by the discounted value of credits but the tradeoffs remain the same.

Proposition 1: Increasing the likelihood of penalties will only affect firms' compliance if it raises perceived penalties above a tipping point $\bar{\theta}$ that is (1) decreasing in the share of their supplier network that is compliant, $\mu(e)$, and (2) increasing in the cost of exerting effort to increase compliance in the supplier network, $\alpha(e)$.[17]

Proof: $\bar{\theta}$ satisfies the following:

$$\begin{aligned} & T_l^*(\bar{\theta}) = T_o^* \\ \iff & \tau(\hat{R}_l^* - M) + \delta_l(x_l^*) + \bar{\theta}P\tau(\hat{R}_l^* - M) = T_o^* \\ \iff & \bar{\theta} = \frac{T_o^* - \delta_l(x_l^*)}{P\tau(\hat{R}_l^* - M)} - \frac{1}{P} \end{aligned} \tag{3}$$

The tipping point, $\bar{\theta}$ depends positively on T_o^*, which is in turn decreasing in $\mu(e)$ and increasing in $\alpha(e)$ (applying the envelope theorem), which gives the two results. As is intuitive, higher penalties lower the tipping point.

[17] The tipping point $\bar{\theta}$ will be strictly positive as long as the cost of on time filing is greater than the cost of late-filing net of expected penalty i.e. $T_0{}^* > T_l{}^* - \theta P\tau(\hat{R}^* - M)$

Figure 1: Taxpayers with $\bar{\theta} \in (\theta, \theta')$ are induced to become compliant

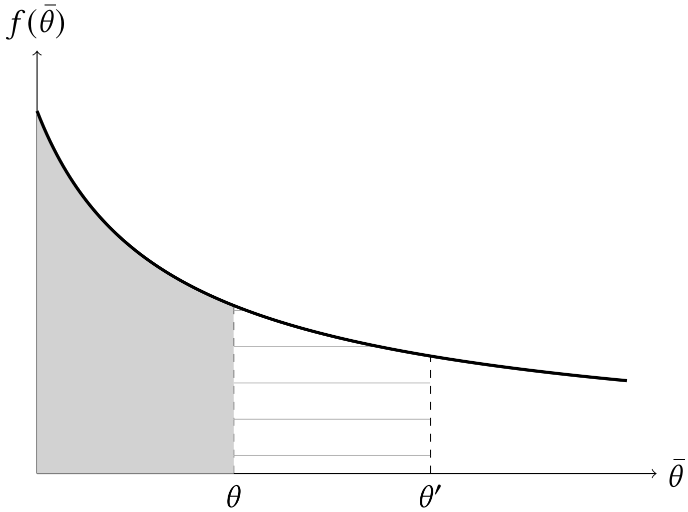

Notes: Figure depicts a hypothetical distribution of threshold $\bar{\theta}$ in the economy, which will vary depending on underlying determinants like non-compliance among suppliers, cost of changing supplier behavior, cost of evasion and credits owed. For expositional purposes, we assume a distribution of $\bar{\theta}$, which has the reasonable properties that some taxpayers comply regardless of the level of enforcement, that increasingly fewer taxpayers have higher values of $\bar{\theta}$, and that some taxpayers may never comply. Those with $\bar{\theta} < \theta$ (gray shaded region) are already filing on time. Raising enforcement from θ to θ' only changes the behavior among taxpayers with sufficiently high $\bar{\theta}$ such that they were not already compliant (i.e. those with more non-compliance among suppliers and sufficiently low evasion costs), but not for those with even higher $\bar{\theta}$ (i.e. those with low evasion costs, high supplier nudging costs, and non-compliance among suppliers).

Figure 1 illustrates a distribution of $\bar{\theta}$ across all taxpayers. For a given level of enforcement, θ, all taxpayers with $\bar{\theta} < \theta$ depicted in the gray shaded region will make on time payments. When the level of enforcement is raised to θ', only taxpayers with $\bar{\theta} \in (\theta, \theta')$ switch to being compliant. These taxpayers that do respond may increase supplier compliance and underreport their revenue by more. In our model, the variation in $\bar{\theta}$ comes from variation in supplier compliance and a taxpayers' own cost of evasion. We investigate these predictions in our data.

5 Experimental Intervention

Nudges from the tax authority in the form of text messages, calls, in-person visits, and letters have been show to successfully encourage tax compliance in several contexts (Mascagni, 2018; Brockmeyer et al., 2019; Castro and Scartascini, 2015). The most successful messages tend to specify penalties for non-compliance – so called "deterrence" messages – which seem to raise the salience of enforcement (Bergolo et al., 2023). We use such messages to perturb perceived enforcement in our context using automated voice-recorded messages aimed at raising filing compliance.

Taxpayers were stratified along certain characteristics to improve power and then randomly assigned within strata to one of three treatment groups or to a control group that received no message. Twenty five percent of taxpayers were assigned to the control group. The remainder were assigned to one of three treatment arms: 1) a simple *reminder* message arm which were reminded of the filing deadlines and provided with a link to the GST portal, 2) a *deterrence* arm that additionally received information about the penalties for late-filing, and 3) a *customized deterrence arm* which additionally was told that the tax authority was aware of their specific past history of non-compliance. The exact content of the messages are provided in Appendix B.3. In our analysis we pool the different treatment arms since we are interested mainly in the differential effect of treatment by network compliance.[18]

Taxpayers in each group received the same message four times each month. The first and second messages were sent on the 9th and the 11th of the month, just before the deadline for filing the credit-linked GSTR-1 form. The third and fourth messages were sent on the 18th and 20th of the month just before the deadline for the payment-linked GSTR-3B form. The calls were made in the months of September - November, corresponding to the filing deadlines for August-October returns. The first round of calls were made on September 9th and the last round of calls on the 20th November. Figure 2 shows the timelines for sample selection and experiment roll out.

We stratified our study population according to three criteria: type of filing non-compliance (late filing or non-filing of either the payment-linked or credit-linked form, or both); revenue (nil-filer, less than Rs. 15 million, above Rs. 15 million)[19]; and nature of business (retail, wholesale, service, or other). See Appendix B for further details of sampling and stratification. Appendix Table A4 shows balance in characteristics among control group and various

[18]The distinct treatment arms were created with the aim of helping the tax authority create an automated messaging enforcement program.

[19]Taxpayers with annual revenue below Rs. 15 million are eligible for the composition scheme which allows them to file returns annually rather than on a quarterly or monthly basis. Composition taxpayers cannot provide input tax credits to their buyers or claim input tax credits from their suppliers.

treatment arms and confirms that randomization was successful.

Figure 2: Intervention Timeline

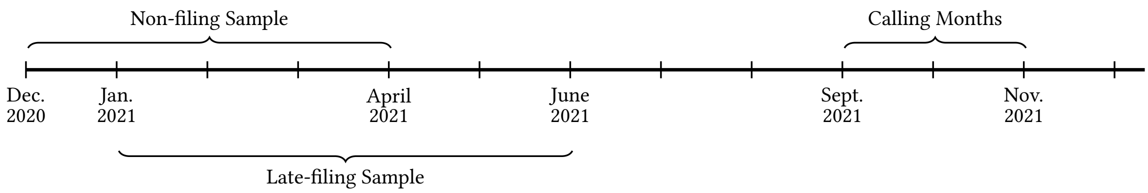

This experimental intervention including the sample stratification and treatment arms were pre-registered.[20] Results from the pre-registered specifications are split between this paper and the companion policy paper, Gupta et al. (2024), which contains results intended for a policy audience such as the effect on filing rates, revenue and deregistration rates. Our pre-analysis plan (PAP) specified that we will analyze network compliance effects and heterogeneity by network compliance. The specific measure of network compliance we use was not specified in the PAP because the detailed invoice level data on transactions between GST-registered taxpayers was provided only after the experiment took place. Variables related to change in composition of buyer's network were not included in the PAP but are included in this paper.

The automated voice-calls allow us to track whether the calls were picked up and how long the receiver listened to the message. Across all the treatment arms, around 80 percent of all firms listened to the voice messages in the first round which declines to around 60 percent in the subsequent rounds (see Appendix Figure A3a). We define a call as successful if the taxpayer listens to the primary content of the message which either reminds or deters against non-compliance. Appendix Figure A3b shows that the success rate is above 50 percent in the first couple of rounds and then declines to around 30 percent in the subsequent rounds.

We are interested in the causal effect of nudges on on-time filing by non-compliant firms. Formally, we define this as:

$$\phi(X) \equiv E[Y_{1i} - Y_{0i} | Y_{0i} = 0, X] = \frac{P(Y_{1i} = 1 \cap Y_{0i} = 0 | X)}{P(Y_{0i} = 0 | X)} \tag{4}$$

where Y_{1i} and Y_{0i} are firm i's potential outcome under treatment (1) or control (0). We denote

[20] Please see study number AEARCTR-0008357 in the AEA RCT Registry.

the compliance level of the supplier network by X, which is the dimension of treatment heterogeneity of interest. Empirically, we define supplier compliance as the share of tax credits available on time based on supplier network and supplier behavior in July 2021 as described in Section 3.

In our main results, we split our sample of taxpayers into two groups with "high" and "low" network compliance, defined as taxpayers having network compliance above and below the sample median (approx. 80 percent), respectively. Within these groups, we estimate the treatment effect using a straightforward specification:

$$Y_{it} = \beta_0 + \beta_D D_i + \text{Controls}_{it} + \eta_{s(i)t} + \varepsilon_{it} \tag{5}$$

where D_i denotes treatment assignment. Our main outcomes are on-time payment and on-time filing of the credit return. We also look at outcomes related to reported revenue, credits, and changes in supplier behavior when we examine mechanisms behind our results. We include strata by month fixed effects $\eta_{s(i)t}$ in all specifications, where $s(i)$ denotes strata of taxpayer i and t denotes month. We also include a set of controls to improve power such as age of the firm, industry fixed effects and on-time filing at baseline. Because we randomized at the taxpayer level, suppliers of a treated or untreated taxpayer may have also been treated. Although treatment status of any individual taxpayer is independent of treatment of the network, we may be concerned about potential spillovers and SUTVA violations. We address these concerns using the method proposed by Borusyak and Hull (2023). Our measure of network exposure to treatment is the share of tax credits coming from treated suppliers, *Share Treated*. Borusyak and Hull (2023) suggest including a measure of expected exposure using simulated treatment assignments as a control to capture variation in treatment exposure due to endogenous characteristics like network centrality. In our setting, this is equivalent to share of credits coming from suppliers in our study sample, multiplied by the probability of treatment, which is 0.75.[21] Therefore, we also include the variable *Share Assigned*, which is the

[21] Let ω_j denote the share of credits received from supplier j such that $\sum_{j\in J_i} \omega_j = 1$ for each taxpayer i, J_i is the set of all suppliers to that taxpayer in July 2021. Share Treated$_i$ = $\sum_{j\in J_i} \omega_j D_j$. $E[\text{Share Treated}_i] = \sum_{j\in J_i} \omega_j E[D_j] = P(D_j) \times \text{Share Assigned}_i$.

share of credits coming from taxpayers in the study sample, as a control in all specifications, along with the actual *Share Treated*.

We report scaled treatment effects ($\frac{\beta_D}{1-\beta_0}$) with standard errors calculated using the delta method. We also test for the difference in treatment effects between the above and below-median network compliance groups using a t-test of differences since the two groups are mutually exclusive.

We estimate the following empirical specification:[22]

$$Y_{it} = \beta_0 + \beta_D D_i + \beta_X X_i + \beta_{DX} D_i X_i + \text{Controls}_{it} + \eta_{s(i)t} + \varepsilon_{it} \tag{6}$$

where the standard errors ε_{it}, are clustered at the level of the taxpayers.

The treatment effect for a given level of X is then given by:

$$\phi(X) = \frac{\beta_D + \beta_{DX} X}{1 - \beta_0 - \beta_X X} \tag{7}$$

6 Results

6.1 Impact of Enforcement on Taxpayer Compliance

We find that the voice-recorded messages make taxpayers more likely to file both the payment-linked and credit-linked forms on time. Table 2 reports the direct effects of each treatment arm on on-time filing in October 2021, the first treatment month when both monthly and quarterly filers are required to file. Both reminder and deterrence messages have statistically significant and positive impacts on payment-linked form filing rates of about 2 percentage points, and on credit-linked form filing rates of about 2.5 percentage points. These translate to a 8 percent and 7.6 percent approximate decrease in late-filing rates, respectively, among the target population of previous late-filers, which is a similar effect size seen for nudges

[22]A version of this specification without the treatment heterogeneity was prespecified in the pre-analysis plan.

in other contexts. These findings indicate that similar to many other contexts, deterrence messages have a statistically significant and positive impact on compliance, even without a corresponding enforcement action.

Table 2: On-time Payment and Filing Response to Nudges

	Previous Late-Filers		With Seller Non-Compliance	
	(1) GSTR-3B	(2) GSTR-1	(3) GSTR-3B	(4) GSTR-1
Reminder	0.016***	0.020***	0.005	0.007
	(0.004)	(0.004)	(0.008)	(0.008)
Deterrence	0.021***	0.025***	0.020***	0.020***
	(0.003)	(0.003)	(0.006)	(0.006)
Constant	0.739***	0.672***	0.783***	0.753***
	(0.002)	(0.002)	(0.005)	(0.005)
Observations	123,280	123,280	26,876	26,876

Notes: Table presents results of specification (5), where treatment is separated into two separate treatment arms: Reminder and Deterrence (which includes "customized deterrence"). The sample in the first two columns is restricted to taxpayers who had filed their return late in any month between April - June 2021. We further restrict the sample to those with at least one non-compliant seller in columns (3) and (4). The outcome variable in columns 1 and 3 is whether the taxpayer filed their payment-linked return (GSTR-3B) on time. The outcome in columns 2 and 4 is whether they filed their credit-linked return (GSTR-1) on time. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This increase was transient and largely limited to the months when taxpayers received a phone call – we see similar impacts in all months of treatment (returns filed for August-October) and then smaller, statistically insignificant treatment effects for the November and December returns once the nudges have stopped (see Figure 3). This pattern occurs despite the fact that taxpayers became less likely over time to listen to the recorded message.

The rest of our analysis will focus on our main sub-sample of taxpayers with mainly Delhi-registered sellers in their network with some history of late-filing. Columns 3 and 4 show the same results for this sub-sample of taxpayers where we find similar effects of about 9 and 8 percent decline respectively on late-filing rates of payment and credit-linked forms. We also restrict attention to the returns filed between September and November 2021 (for August - October returns) when taxpayers received phone calls.

Our conceptual framework predicts that nudges will increase on-time payments on average but not for taxpayers with a high share of delinquent credits. We find that indeed, the overall, positive effect on on-time payment is coming from taxpayers with sufficiently

Figure 3: On-time Filing of GSTR-3B (payment-linked)

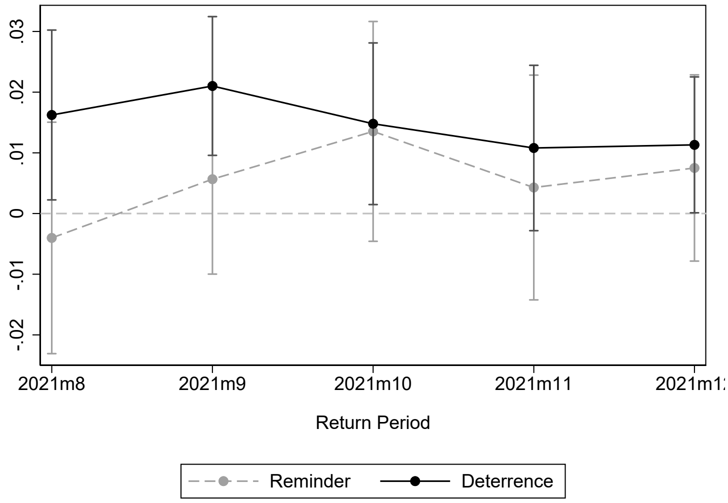

compliant seller networks. We use supplier network compliance one month prior to the intervention to distinguish taxpayers with above and below-median levels of compliance among their supplier network and separately estimate treatment effects. In column 1 of Table 3, we see that treatment raises the likelihood of on-time payment by 2.4 percentage points. In contrast, taxpayers with low levels of compliance show only a small and statistically insignificant increase in on-time payment (Column 2). The last two rows of the table report the implied percent decrease in late filing (scaling the coefficient on treatment by the rate of late payment in the control group), and the difference between high and low network compliance groups. Treatment induces a statistically significant, 7.8 percent decrease in late payment among taxpayers with high compliance among their seller network. This effect is 6.7 percentage points higher than the treatment effect among the low network compliance group in column 2 and this difference is statistically significant at the 10 percent level.

These varying effects of treatment by seller network compliance are illustrated over different values in Figure 4. On the far left, taxpayers who have none of their credits available on time (i.e. a seller network compliance level of zero), do not respond to treatment. At the 25th percentile of network compliance, we see a small decrease in non-compliance of about 5 percent. Treatment effect increases again at the median level of network compliance of about 75 percent, i.e. when 75 percent of input tax credits were available on time in July. Appendix Figure A4 shows that both the reminder and deterrence messages exhibit the same pattern, with the deterrence messages having a stronger effect overall.

Table 3: On-time Payment and Filing Response by Network Compliance

	Payment-Linked Form (Network-Dependent)		Credit-Linked Form (Not Network-Dependent)	
	(1) Above-median	(2) Below-median	(3) Above-median	(4) Below-median
Treated	0.024***	0.007	0.015**	0.013*
	(0.006)	(0.007)	(0.007)	(0.007)
Constant	0.698***	0.443***	0.781***	0.705***
	(0.102)	(0.040)	(0.180)	(0.180)
Observations	31,897	30,139	31,883	30,130
Controls	Yes	Yes	Yes	Yes
Treatment Effect	.078** (.034)	.012 (.012)	.066 (.062)	.044 (.036)
Difference	.067* (.036)		.022 (.072)	

Notes: Table presents results of specification (5). The dependent variable in columns 1 and 2 is a dummy for whether a taxpayer filed the GSTR-3B return (payment-linked) by the deadline in the months of August-October. The dependent variable in columns 3 and 4 is whether they filed the credit-linked return on time in August-October. The sample is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. Columns 1 and 3 restrict attention to taxpayers with above-median compliance in their seller network, which columns 2 and 4 are taxpayers with below-median network compliance. All specifications include strata by month fixed effects. Controls include firm age, industry fixed effects and prior history of filing compliance. We also control for the share of tax credits coming from sellers assigned to study sample and from sellers treated. Standard errors are clustered by taxpayer, and reported in parentheses. The second-to-last row presents the percent decrease in late-filing implied by the coefficients. Standard errors calculated using the delta method in parentheses. The last row shows the difference between above and below-median groups and the standard error of the difference in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

An immediate concern one might have is that the interdependence of seller and buyer compliance reflects matching of trading partners on attitudes towards compliance or that network compliance is proxying for some other correlated dimension of treatment effect heterogeneity. For example, taxpayers in certain industries may be more likely to file late, which would result in higher network non-compliance and lower response to treatment. We address these concerns in two ways.

First, we compare taxpayer response in filing the other mandatory, credit-linked return as a natural placebo check. There is no tax obligation associated with filing this return and so there is no liability linkage between supplier and client compliance associated with the return. However, if taxpayers were sorting into production networks based on general compliance attitudes or the networks reflect correlated compliance costs, we would expect to see the same pattern of correlation between supplier compliance and buyer response in this return. Columns 3 and 4 of Table 3 repeat the same specifications as 1 and 2 but with on-time filing of the credit-linked form as the outcome. Unlike with the payment-return, there difference in treatment effect for the credit-return between high and low supplier compliance networks is 2 percentage points and not statistically significant. Figure 5 shows the treatment effect across different levels of supplier compliance, revealing only a small gradient. These results suggest that the link between supplier compliance and treatment effect on the payment-return must operate through its impact on the tax liability of the buyer and that it is unlikely to be driven by similar tendencies towards compliance among transacting pairs.

As a stronger test, we compare the differences in responses for the two returns by the same taxpayer. We modify our data so that each observation is a return filed by a taxpayer in a month for each of the two returns. The outcome of interest is whether the return was filed on time. Specification 6 is modified to include taxpayer fixed effects (and therefore exclude other controls). The results presented in Table 4 shows the difference in response to nudges for the payment-linked form relative to the credit-linked form among taxpayers with high and low network compliance. In column 1, among high-compliance networks, treatment causes a 6 percentage point higher response for the payment-linked form, as reported in the last row of the table. This estimate is the coefficient on the interaction term between payment-linked form and treatment scaled by rate of late payment in the control group. In contrast, among taxpayers with low-compliance networks, treatment is about 3.6 percentage points less effective for payment returns, though this difference is not statistically significant. However, the difference in response to the two returns between the high and low compliance groups is 9.8 percentage points and statistically significant at the 5 percent level of significance. The results of these tests suggest that even holding constant taxpayer age, industry, and other baseline taxpayer characteristics, network compliance matters for the payment return but not for the credit-linked return. Any characteristic driving this difference must therefore reasonably operate through the tax liability implications on the payment return.

Our second approach to showing the causal link between supplier compliance and buyer's response to enforcement is to use experimentally-induced variation in supplier compliance. In the next section, we explore the spillover effects of treating a large share of the supplier network.

Table 4: Within-Taxpayer Difference in Response of Payment and Credit Provision

	(1) Above-median		(2) Below-median
Treat X Payment-Linked	0.011*		−0.007
	(0.006)		(0.006)
Payment-Linked Form	0.033***		0.081***
	(0.005)		(0.006)
Constant	0.784***		0.734***
	(0.001)		(0.001)
Observations	84,007		83,908
Taxpayer FE	Yes		Yes
Effect on Payment	.062* (.032)		-.036 (.035)
Difference in TE		.098** (.048)	

Notes: The dependent variable is a dummy for whether a taxpayer filed either the payment-linked or credit-linked return on time in August-October. The sample is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. We expand this sample to a taxpayer X month X return level dataset with one observation for each of the two returns that a taxpayer files. Columns 1 restricts attention to taxpayers with above-median compliance in their seller network, while column 2 is taxpayers with below-median network compliance. All specifications include strata by month fixed effects and taxpayer fixed effects. The "Effect on Payment" statistic reported in the second-to-last row in the table is the difference in impact of treatment between the payment-linked and credit-linked returns scaled by the rate of late filing of the payment-linked form among the control group. The "difference in TE" reported in the last row is the difference between the effect on payment between columns 1 and 2. Standard errors calculated using the delta method and clustered by taxpayer are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 4: Heterogeneity in Treatment Effect by Network Compliance

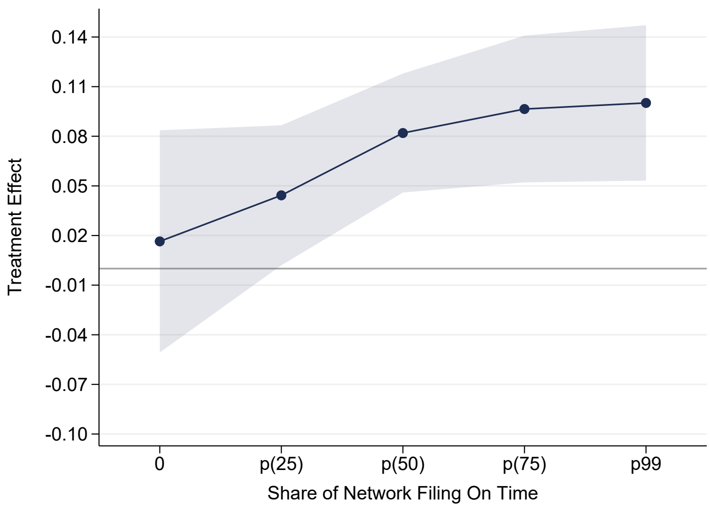

Notes: Figure plots treatment effect as defined in equation (7), evaluated at different quantiles of compliance of the taxpayer's seller network. Standard errors are calculated using the delta method. The treatment effect denotes the percent change in on-time filing among late-filers induced by the treatment.

Before turning to these spillover effects, we address one more potential concern regarding the external validity of our main results. Our main sample excludes taxpayers who source more than 30% of their input tax credits from outside-Delhi sellers, and so we evaluate how

Figure 5: Heterogeneity in Treatment Effect on GSTR-1 Filing by Network Compliance

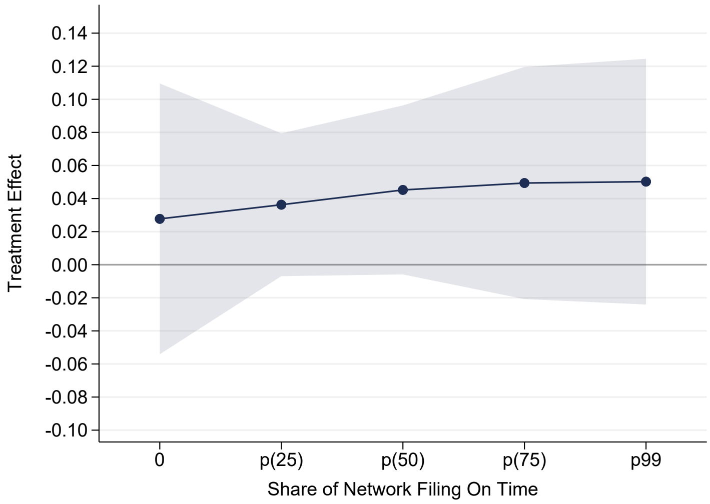

Notes: Figure plots treatment effect as defined in equation (7), evaluated at different quantiles of compliance of the taxpayer's seller network. Standard errors are calculated using the delta method. The treatment effect denotes the percent change in on-time filing among late-filers induced by the treatment.

this exclusion changes our results. In Appendix D, we document that for a few observable characteristics, there are very small but statistically significant differences between our main sample and the excluded group. Following Stuart et al. (2011), we re-estimate our main specification using inverse-propensity-score weights that re-balance the sample to resemble excluded buyers. The results, reported in Table D.2, are qualitatively and quantitatively similar to our main estimates.

6.2 Impact of Experimentally-Induced Compliance in the Network

Our experiment generates an exogenous change in supplier compliance for each firm in the sample. Because we randomized assignment to treatment at the taxpayer level, the share of a taxpayer's seller network in the sample that gets treated is independent of their own treatment assignment. Since treatment induces taxpayers to file their credit-linked forms on time, variation in the share of the network assigned to treatment generates variation in share of on-time credits, μ_0 in our conceptual framework.

We modify our main specification in Equation 5 to include an interaction between treatment and the share of sellers treated, while controlling for the share of sellers in the sample and its interaction with treatment.

$$
\begin{aligned}
Y_{it} &= \beta_0 + \beta_D D_i \\
&\quad + \beta_1 \text{Share Sellers Treated}_i + \beta_2 \text{Share Sellers Treated}_i \times D_i \\
&\quad + \beta_3 \text{Share Sellers Assigned}_i + \beta_4 \text{Share Sellers Assigned}_i \times D_i + \eta_{s(i)t} + \varepsilon_{it}
\end{aligned} \tag{8}
$$

As the seller network becomes large, the share of sellers treated will converge to the share of all taxpayers assigned to treatment (75%) due to the law of large numbers. Therefore, only buyers with small seller networks show substantial variance in the share of sellers treated. This reduction in variance also limits the statistical power of the estimates in this section; while they go in the same direction as our main results, not all are statistically significant.

Column 1 of Table 5 shows our preferred specification where we restrict the sample to firms with fewer than 4 sellers in the network. The coefficient on the interaction between treatment and the share of sellers treated suggests a 10 percentage point increase in the share of sellers treated increases the direct treatment effect by 1.25 percentage points, statistically significant at the 10 percent level. This is a 35 percent decrease in non-compliance. When taxpayers with slightly larger networks of less than 6 suppliers are included in column 2, we still see a positive coefficient of 0.09, which is not statistically significant at the 10 percent level. With no restriction in network size as in column 3, we see an even smaller but still positive coefficient of 0.03.

The effect is smaller, but significant at 10% level, when we include taxpayers with above-median network compliance, where there is of course less room for network compliance to increase (Column 4). These results corroborate our main findings among taxpayers with smaller seller networks. Among those with larger networks, we may not have sufficient power to use the sampling variation so we exercise caution in extending our conclusions to this sample.

We use these estimates to calculate the implied increase in the effect of treating a taxpayer if we increase the share of their suppliers treated from zero to one, conditional on having the mean share of suppliers in the study sample (30%). We report these estimates in the last row of Table 5 for each specification. Increasing the share of suppliers treated from zero to one increases the effect of a nudge by 29 percent among taxpayers with low network compliance and less than 4 suppliers in the network (column 1). Effects are smaller and not statistically significant when we include taxpayers with larger numbers of suppliers (columns 2 and 3). The increase is 15 percent when we pool taxpayers across all levels of network compliance with under 4 suppliers in the network.

Table 5: Impact of Upstream and Direct Nudges

	Below-median network compliance			All
	(1)	(2)	(3)	(4)
Treated	−0.010	0.003	0.007	0.015
	(0.018)	(0.015)	(0.012)	(0.013)
Share July Sellers Treated	−0.086	−0.052	−0.009	−0.057
	(0.054)	(0.049)	(0.046)	(0.041)
Share July Sellers Assigned	0.060	0.060	0.034	0.058
	(0.049)	(0.045)	(0.043)	(0.036)
Treat X Share Sellers Treated	0.124*	0.086	0.028	0.088*
	(0.064)	(0.058)	(0.053)	(0.049)
Treat X Share Sellers Assigned	−0.066	−0.077	−0.038	−0.079*
	(0.059)	(0.054)	(0.049)	(0.044)
Constant	0.625***	0.563***	0.551***	0.474***
	(0.105)	(0.064)	(0.051)	(0.078)
Observations	7,110	10,748	19,213	12,296
Controls	Yes	Yes	Yes	Yes
Number of Sellers	<4	<6	Any	<4
TE if All Sellers Treated	.295* (.157)	.189 (.121)	.063 (.116)	.157* (.086)

Notes: Table presents results of specification (8). The dependent variable is a dummy for whether a taxpayer filed the GSTR-3B return (payment-linked) by the deadline in the months of August-October. The sample is restricted to our main analysis sample. Columns 1-3 restrict the sample to taxpayers with below-median network compliance. We restrict the sample by the number of sellers in the taxpayer's network as given by the indicator at the bottom row. All specifications include strata by month fixed effects and controls as indicated. The "TE if all sellers treated" statistic in the last row presents the difference in treatment effect if *Share Sellers Treated* equals 1, scaled by the rate of late-filing if all sellers were treated but the taxpayer themselves were not and the treatment effect if *Share Sellers Treated and Assigned* equals zero scaled by the rate of late-filing in the control group. Standard errors calculated using the delta method and clustered by taxpayer are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Table 5 and in other specifications in Tables 3 and A5, the spillovers from only treating sellers as captured by the coefficient β_1 are small and statistically insignificant. The effect of treating sellers only seems too small on average to detect in our design.

In the main analysis, we address spillover concerns by controlling for expected exposure to spillovers. To test the robustness of our results, we conduct an additional test by limiting our sample to buyers who had no seller from their July network treated in our experiment.

Columns 1 and 2 of Table A6 show that only taxpayers with above median supplier compliance respond to our nudges even in the absence of any spillover effects from the supply chain.

6.3 Mechanisms: Net Tax Liability and Seller Compliance

Taxpayers with delinquent sellers in their network are likely to have missing tax credits at the tax payment deadline. To make their payment on time, they can do one of the following: (1) accept a larger net liability and file without the missing credits, (2) file on time without the missing credits but report lower output liability to leave net payment unchanged (3) nudge their sellers to file returns on time. Upon deciding to file on time, taxpayers choose how much output to report and whether to nudge their suppliers based on their respective costs. How they respond will depend on these unobserved costs of misreporting output and influencing suppliers.

Table 6 shows the impact on reported output liability, net tax liability, and input tax credits claimed among those with above-median compliance in the seller network, that is, the group that responds to treatment. We measure these outcomes as binary variables showing an increase or decrease relative to their own past value in July 2021 because of the noisiness of these variables. Taxpayers are 2 percentage points more likely to report a net zero tax liability relative to July 2021 in response to a deterrence message, an approximately 4.5 percent increase relative to the mean (Column 1). They are 7 percent less likely to report a larger net tax liability (Column 2). This decrease in net liability at the time of filing is coming from a decrease in reported output liability – taxpayers are 5 percent less likely to report an increase in output tax liability (Column 4), counteracting the fact that they are 3 percent less likely to claim more input tax credits (Column 5).

Table 7 shows the same outcomes for taxpayers with below-median compliance in the seller network, i.e. those who do not increase filing compliance in response to treatment. Notably, we do not see any changes in their reported output tax liability, sales, and credits claimed. This is consistent with their lack of response to the treatment. These results pool outcomes for August - October returns for which the taxpayers received messages. In Appendix Tables A7 - A11, we show month-wise responses for August - December 2021, which confirms

Table 6: Mechanisms: Tax Payment among Above-Median Network Compliance

	August-October 2021				
	(1) Zero Net Liab.	(2) Higher Net Liab.	(3) Zero Output Liab.	(4) More Output Liab.	(5) More ITC claimed
Treated	0.020^{**}	-0.023^{***}	0.003	-0.023^{***}	-0.018^{**}
	(0.008)	(0.008)	(0.002)	(0.008)	(0.009)
Constant	0.445^{***}	0.332^{***}	0.027^{***}	0.462^{***}	0.539^{***}
	(0.008)	(0.007)	(0.002)	(0.008)	(0.009)
Sample Mean	0.46	0.32	0.03	0.45	0.52
Observations	32,419	32,419	32,419	32,419	32,419

Notes: Table presents results of specification (5) Each observation is a taxpayer and return period. *Zero Net Liab.* is a dummy for whether the taxpayer reported net zero liability (i.e. zero payment in cash) in a return period. *Higher Net Liab.* is a dummy for whether the taxpayer reports net liability greater than their net liability in July 2021. *Zero Output Liab.* is a dummy for whether tax liability on sales alone is zero (i.e. zero taxable sales reported). *More Output Liab.* is a dummy for whether output liability is higher than in July 2021. *More ITC Claimed* is a dummy for whether the taxpayer claimed more credits than in July 2021. The sample is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in this table is additionally restricted to those with above-median (i.e. greater than 87%) credits filed by their suppliers in July. All specifications include strata by month fixed effects. Standard errors are clustered by taxpayer, and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Mechanisms: Tax Payment among Below-Median Network Compliance

	August-October 2021				
	(1) Zero Net Liab.	(2) Higher Net Liab.	(3) Zero Output Liab.	(4) More Output Liab.	(5) More ITC claimed
Treated	0.006	−0.007	0.001	0.003	−0.009
	(0.008)	(0.008)	(0.003)	(0.008)	(0.009)
Constant	0.451^{***}	0.323^{***}	0.051^{***}	0.436^{***}	0.516^{***}
	(0.008)	(0.007)	(0.003)	(0.008)	(0.009)
Sample Mean	0.45	0.32	0.05	0.44	0.50
Observations	30,565	30,565	30,565	30,565	30,565

Notes: Table presents results of specification (5). Each observation is a taxpayer and return period. For definitions of various outcome variables, please refer to the footnotes of the previous table. In contrast to the previous table, we now restrict the sample to taxpayers with below-median (i.e. less than 87%) credits filed by their suppliers in July. All specifications include strata by month fixed effects. Standard errors are clustered by taxpayer, and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

that these patterns are mainly observed during months when taxpayers were called. Notably, the effect on lower output liability reported persists past October unlike the decrease in tax credits (Appendix Tables A10 and A11).

This margin of behavioral response shows that taxpayers become more compliant along the timing dimension but become *less* compliant along the liability dimension or by decreasing the size of their operations, both of which are undesirable outcomes from the perspective of raising tax revenue and production efficiency. One way in which this result could arise is as we laid out in the framework where taxpayers perceive a lower cost of evasion if they file on time. We could also get the same result if taxpayers perceive a greater benefit from evasion if they file one time. For example, if it is harder to get excess credits refunded, taxpayers may wish to avoid overpaying today if they anticipate receiving credits in the future. Another possibility is that taxpayers are cash constrained and cannot pay a larger amount than their true tax liability so they make up for the missing credits by underreporting output.

A key dynamic in VAT enforcement is that increasing compliance at weak points (often at the downstream, retail stage) generates positive spillovers through the production network. We examine whether the increase in compliance by buyers translated to greater compliance among their seller network, which we argue depends on the cost to do so. We find that on average, buyers at all levels of supplier non-compliance are unlikely to switch suppliers. Table 8 shows the change in various indicators of supplier non-compliance as a result of buyer's treatment status among those with above-median network compliance. As the sub-group where we see the highest treatment effects, we might expect that the mechanism through which that treatment effect was achieved is most strongly displayed here. On the other hand, supplier compliance is already relatively high among this group, leaving little room or need for change. We find little effect of the treatment on several measures of supplier compliance. Although treated taxpayers are less likely to buy again from a previously delinquent seller (column 1), they continue to have similar shares of credits coming from such delinquent sellers, suggesting they did not simply reduce purchases (column 2). As shown in column 3, there is a marginally significant change in the proportion of late filers that are repeated from the previous month. As a result, the share of credits coming in on time does not change measurably (column 4).

Table 9 looks at the same outcomes among buyers with below-median seller non-compliance where there might be more room for change in supplier behavior. Again, among this group we see no change across the four different measures of supplier non-compliance post-treatment. Appendix Tables A12-A15 show the month-wise impact of nudges on the all the outcome variables for both above and below median seller non-compliance. Overall, we see similar lack of change in the seller network in all return periods.

Table 8: Mechanisms: Seller Compliance among Above-Median Compliance

	September-November 2021			
	(1)	(2)	(3)	(4)
	Prob. of Repeating Any Late Seller	Prop. of ITC from Repeat Late Sellers	Prop. Late Seller Repeated	Share Seller On Time
Treated	−0.016**	0.001	−0.011*	0.000
	(0.008)	(0.002)	(0.006)	(0.003)
Constant	0.357***	0.036***	0.269***	0.849***
	(0.008)	(0.002)	(0.006)	(0.002)
Sample Mean	0.33	0.04	0.25	0.85
Observations	31,152	31,474	31,152	31,474

Notes: Table presents results of specification (5). Each observation is a taxpayer and return period. The outcomes measure change of delinquent sellers from whom the taxpayers have previously purchased inputs. *Prob. of Repeating Any Late Seller* is a dummy for whether a delinquent seller reappears in the taxpayer's network. *Prop. of ITC from Repeat Late Sellers* is the share of a taxpayer's ITC coming from previous delinquent sellers. *Prop. Late Seller Repeated* is the share of delinquent sellers from whom a taxpayer repurchases. *Share Seller On Time* is the share of current period's credits that arrived on time. The sample is restricted to firms with sellers which have some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in this table is additionally restricted to those with above-median (i.e. greater than 87%) credits filed by their suppliers in July. All specifications include strata by month fixed effects.
Standard errors are clustered by taxpayer, and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Mechanisms: Seller Compliance among Below-Median Compliance

	September-November 2021			
	(1)	(2)	(3)	(4)
	Prob. of Repeating Any Late Seller	Prop. of ITC from Repeat Late Sellers	Prop. Late Seller Repeated	Share Seller On Time
Treated	−0.007	−0.001	−0.006	0.001
	(0.008)	(0.003)	(0.007)	(0.004)
Constant	0.385***	0.088***	0.299***	0.774***
	(0.008)	(0.003)	(0.007)	(0.004)
Sample Mean	0.34	0.08	0.27	0.77
Observations	27,849	28,763	27,849	28,764

Notes: Table presents results of specification (5). Each observation is a taxpayer and return period. For definitions of various outcome variables, please refer to the footnotes of the previous table. In contrast to the previous table, we now restrict the sample to taxpayers with below-median (i.e. less than 87%) credits filed by their suppliers in July. All specifications include strata by month fixed effects. Standard errors are clustered by taxpayer, and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These results together suggest that for some firms the cost to switch suppliers leads them to remain with delinquent suppliers even when there are strong incentives to choose compliant suppliers. We provide suggestive evidence of one type of high switching costs that might be relevant here: competitiveness of supplier networks.

An important measure of the strength of a supplier-buyer relationship is the concentration of a buyer's inputs from that supplier. If the input market is concentrated, then it might be harder for buyers to switch suppliers. Dhyne et al. (2022) show using similar VAT trade network data that the concentration of input purchases from a particular supplier is highly predictive of markups even after controlling for industry-level concentration. That is, even if the industry as a whole is competitive, individual suppliers may have market power over their buyers. We measure the concentration of each taxpayer's inputs as the Herfindahl-Hirschman index of their inputs.[23] We find that the treatment effect heterogeneity by supplier compliance comes from taxpayers with concentrated supplier networks. Such taxpayers might find it difficult to switch suppliers in response to nudges. On the other hand, taxpayers in less concentrated supplier network can switch suppliers more easily.

Using this HHI measure, we find that there is strong correlation between the concentration of the supplier network and the persistence of this network measured as the proportion of sellers from the previous month that are repeated. Column 1 of Table A16 shows that as the HHI increases from 0 to 1, the share of repeated sellers increases by 8.4 percentage points or 22.5 percent. Once we include buyer fixed effects in Column 2, the share of repeated sellers more than doubles as we move from least concentrated to more concentrated supply chains.

We then measure the treatment effect heterogeneity by supplier compliance among buyers with more and less-concentrated supplier networks. Figure 6 shows the treatment effect by network compliance among taxpayers with concentrated supplier networks with HHI greater than the median value of 0.3. Among this group, the gradient in treatment effect is steeper. Taxpayers at the 25th percentile of network compliance do not respond at all to treatment while taxpayers with much more compliant networks show a stronger response. Among those with more diffuse supplier networks (Figure 7), we see a slightly negative gradient, likely because the July network is a poorer predictor of subsequent network compliance, introducing measurement error. These results are consistent with our conceptual framework which suggests that the link between supplier compliance and buyer's response is related to the cost of switching suppliers.

[23] That is, for each buyer we calculate the sum of squared shares of each supplier's inputs in a given month.

Figure 6: Effect on Buyers' Payment Filing: Concentrated Supplier Networks

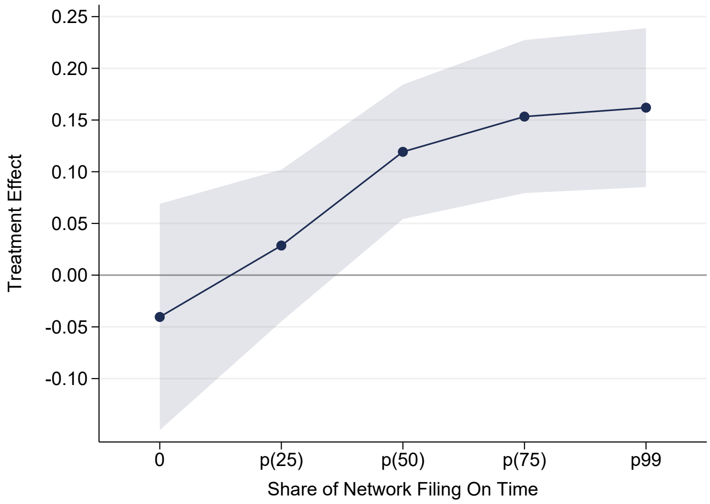

Notes: Figure plots treatment effect as defined in equation (7), evaluated at different quantiles of compliance of the taxpayer's seller network. Standard errors are calculated using the delta method. The treatment effect denotes the percent change in on-time filing among non-filers induced by the treatment. Sample is restricted to buyers with concentrated networks, i.e. HHI of their inputs of at least 0.3.

Figure 7: Effect on Buyers' Payment Filing: Diffuse Supplier Networks

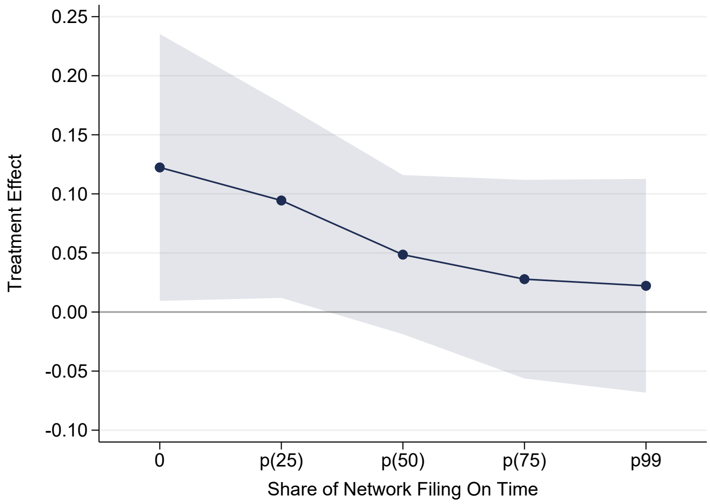

Notes: In this figure, the sample is restricted to buyers with diffuse supplier networks, i.e. HHI of their inputs is less than 0.3. Please also refer to the footnotes of the previous figure.

Indeed, we find suggestive evidence that buyers with less-concentrated supply networks do change their suppliers in response to treatment. Table 10 looks at post-treatment supplier compliance outcomes among buyers with diffuse supplier networks and below-median supplier compliance in July. Treated taxpayers reduce the likelihood of repeating a delinquent seller by around 1 percentage point (column 1), although the coefficient is insignificant. There is a larger albeit statistically insignificant decrease in the share of credits coming from

Table 10: Mechanisms: Change in Suppliers among Buyers with Diffuse Networks and Below-Median Supplier Compliance

	September-November 2021			
	(1) Prob. of Repeating Any Late Seller	(2) Prop. of ITC from Repeat Late Sellers	(3) Prop. Late Seller Repeated	(4) Share Seller On Time
Treated	−0.009	−0.004	−0.007	0.001
	(0.012)	(0.003)	(0.009)	(0.004)
Constant	0.488***	0.076***	0.354***	0.787***
	(0.012)	(0.003)	(0.009)	(0.004)
Sample Mean	0.42	0.06	0.31	0.79
Observations	14,940	15,024	14,940	15,024

Notes: Table presents results of specification (5). Each observation is a taxpayer and return period. Please refer to footnotes of Table 8 for definitions of various outcomes. The sample in this table is additionally restricted to those with below-median (i.e. less than 87%) credits filed by their suppliers in July and the HHI of their inputs less than 0.3 in July . All specifications include strata by month fixed effects.
Standard errors are clustered by taxpayer, and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

previously-delinquent sellers of 0.4 percentage points from a mean of 6 percent (column 2). Surprisingly, there is not much change in the share of credits available on time (column 4). This is possible if the buyers have imperfect information about the compliance behavior of new suppliers. We see no change in suppliers among buyers with below-median network compliance and concentrated supplier networks (Table 11) as their cost of switching suppliers might be significantly higher. The effect sizes among treated taxpayers reduces by almost half compared to buyers with less concentrated networks.

Table 11: Mechanisms: Change in Suppliers among Buyers with Concentrated Networks and Below-Median Supplier Compliance

	September-November 2021			
	(1) Prob. of Repeating Any Late Seller	(2) Prop. of ITC from Repeat Late Sellers	(3) Prop. Late Seller Repeated	(4) Share Seller On Time
Treated	−0.004	0.002	−0.006	0.000
	(0.011)	(0.006)	(0.010)	(0.007)
Constant	0.276***	0.103***	0.242***	0.758***
	(0.011)	(0.006)	(0.010)	(0.007)
Sample Mean	0.25	0.10	0.22	0.75
Observations	12,902	13,734	12,902	13,735

Notes: Table presents results of specification (5). Each observation is a taxpayer and return period. Please refer to footnotes of Table 8 for definitions of various outcomes. The sample in this table is additionally restricted to those with below-median (i.e. less than 87%) credits filed by their suppliers in July and the HHI of their inputs less than 0.3 in July. All specifications include strata by month fixed effects.
Standard errors are clustered by taxpayer, and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The fact that the dependence on network compliance is stronger among more concentrated networks where individual sellers have more market power suggests that such structural features of the economy play a role in the success of self-enforcement.

7 Discussion and Conclusion

The self-enforcement mechanism embedded in the VAT is central to its popularity and success in many developing countries. By design, VAT systems rely on buyers and suppliers to monitor and enforce each other's compliance, reducing the need for tax authorities to deploy their own enforcement resources. In India's GST, this interdependence is amplified by restricting a buyer's credits to invoices the supplier has reported. While such rules can strengthen self-enforcement and raise compliance overall, they can also impose higher costs on taxpayers unable or unwilling to influence non-compliant suppliers.

In practice, many GST taxpayers file returns and remit payments late despite penalties. Our experiment shows that automated voice messages—either simple reminders or penalty warnings—can improve on-time filing, but only where supplier compliance is already high. Buyers with many late-filing suppliers do not respond. Because we directly observe filing compliance, we can show that non-response does not imply full compliance: many buyers continue to pay late when supplier behavior raises their net liability. Rather than prompting buyers to switch toward compliant suppliers and creating positive spillovers up the chain, the liability linkage dampens their own compliance response. Among those who do respond, many do so by lowering reported output. These patterns suggest that persistent ties to non-compliant suppliers are sustained by switching frictions, particularly when suppliers have market power.

The rationale for VAT and real-time enforcement is revenue protection. Fraudulent credit claims are a major source of revenue loss, and invoice matching and related real-time systems can reduce such fraud (Fan et al., 2018; Bellon et al., 2022; Okunogbe and Pouliquen, 2022). As tax authorities shift from post-audit verification to clearance and real-time reporting models, a buyer's ability to deduct VAT increasingly depends on supplier invoices being transmitted or validated through the tax platform, effectively welding buyer compliance to supplier behavior. Similar linkages exist even without explicit invoice matching: under the EU's "knew or should have known" doctrine, authorities can deny input VAT if a purchaser is connected to fraud, and in some member states hold purchasers jointly liable for a supplier's unpaid output VAT. Recent judgments confirm these remedies can be applied cumulatively, subject to proportionality, raising buyer exposure to supplier compliance. Our findings show that in low-compliance environments, such designs have limits and can generate unintended costs.

Two policy lessons follow. First, enforcement should account for the position in the supply chain. While Pomeranz (2015) shows that downstream enforcement can jump-start upstream compliance, our results suggest downstream-only enforcement may leave upstream evasion intact and shift costs onto downstream firms inclined toward compliance. In settings with non-compliant suppliers and sticky input relationships, some downstream firms will rationally remain non-compliant—or substitute towards output under-reporting—raising the marginal cost of compliance for their buyers. Second, the negative spillovers from supplier non-compliance have distributional implications: buyers in concentrated or thin supplier markets face higher switching and bargaining costs. Invoice-dependent crediting can thus create horizontal inequities, adding costs such as loss of working capital, monitoring effort, and enforcement risk. These concerns lead to the second policy implication. Stronger upstream enforcement can be privately desirable: it shifts the monitoring and invoice-chasing burden from buyers to the tax authority and mitigates horizontal inequities when private enforcement costs vary across buyers.

Our findings reveal an important mechanism of non-compliance: when suppliers are delinquent, compliant buyers face a higher net liability if they file on time. Our ability to measure filing/payment delays directly reveals non-response despite underlying non-compliance; similar dampening effects could occur on harder-to-observe margins like under-reporting or non-payment. Therefore, VAT systems need to strike a balance between strengthening self-enforcement incentives through automation while limiting the unintended transfer of enforcement burdens from the tax authority to downstream firms.

References

Almunia, Miguel, David J Henning, Justine Knebelmann, Dorothy Nakyambadde, and Lin Tian, *Leveraging Trading Networks to Improve Tax Compliance: Experimental Evidence from Uganda*, Centre for Economic Policy Research, 2023.

_ , Jonas Hjort, Justine Knebelmann, and Lin Tian, "Strategic or confused firms? evidence from "missing" transactions in Uganda," *Review of Economics and Statistics*, 2024, *106* (1), 256–265.

Antinyan, Armenak and Zareh Asatryan, "Nudging for tax compliance: A meta-analysis," *The Economic Journal*, 2024, p. ueae088.

Barwahwala, Taha, Aprajit Mahajan, Shekhar Mittal, and Ofir Reich, "Is Model Accuracy Enough? A Field Evaluation Of A Machine Learning Model To Catch Bogus Firms," *ACM Journal on Computing and Sustainable Societies*, 2024, *2* (3), 1–43.

Bellon, Matthieu, Era Dabla-Norris, Salma Khalid, and Frederico Lima, "Digitalization to improve tax compliance: evidence from VAT e-Invoicing in Peru," *Journal of Public Economics*, 2022, *210*, 104661.

Bergolo, Marcelo, Rodrigo Ceni, Guillermo Cruces, Matias Giaccobasso, and Ricardo Perez-Truglia, "Tax audits as scarecrows: Evidence from a large-scale field experiment," *American Economic Journal: Economic Policy*, 2023, *15* (1), 110–153.

Bernard, Andrew B, Andreas Moxnes, and Yukiko U Saito, "Production networks, geography, and firm performance," *Journal of Political Economy*, 2019, *127* (2), 639–688.

Best, Michael, Jawad Shah, and Mazhar Waseem, "Detection without deterrence: Long-run effects of tax audit on firm behavior," in "in" 2021.

Boning, William C, John Guyton, Ronald Hodge, and Joel Slemrod, "Heard it through the grapevine: The direct and network effects of a tax enforcement field experiment on firms," *Journal of Public Economics*, 2020, *190*, 104261.

Borusyak, Kirill and Peter Hull, "Nonrandom exposure to exogenous shocks," *Econometrica*, 2023, *91* (6), 2155–2185.

Brockmeyer, Anne, Giulia Mascagni, Vedanth Nair, Mazhar Waseem, and Miguel Almunia, "Does the Value-Added Tax Add Value? Lessons Using Administrative Data from a Diverse Set of Countries," *Journal of Economic Perspectives*, 2024, *38* (1), 107–132.

__ , Spencer Smith, Marco Hernandez, and Stewart Kettle, "Casting a wider tax net: Experimental evidence from costa rica," *American Economic Journal: Economic Policy*, 2019, *11* (3), 55–87.

Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal, "Ghosting the Tax Authority: Fake Firms and Tax Fraud in Ecuador," 2022.

Castro, Lucio and Carlos Scartascini, "Tax compliance and enforcement in the pampas evidence from a field experiment," *Journal of Economic Behavior & Organization*, 2015, *116*, 65–82.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, and Glenn Magerman, "Imperfect competition in firm-to-firm trade," *Journal of the European Economic Association*, 2022, *20* (5), 1933–1970.

Drago, Francesco, Friederike Mengel, and Christian Traxler, "Compliance behavior in networks: Evidence from a field experiment," *American Economic Journal: Applied Economics*, 2020, *12* (2), 96–133.

Eaton, Jonathan, Samuel S Kortum, and Francis Kramarz, "Firm-to-firm trade: Imports, exports, and the labor market," Technical Report, National Bureau of Economic Research 2022.

Eissa, Nada, Andrew Zeitlin, Saahil Karpe, and Sally Murray, "Incidence and impact of electronic billing machines for VAT in Rwanda," *International Growth Centre Report*, 2014.

Fan, Haichao, Yu Liu, Nancy Qian, and Jaya Wen, "Computerizing VAT invoices in China," Technical Report, National Bureau of Economic Research 2018.

Fisman, Raymond and Shang-Jin Wei, "Tax rates and tax evasion: evidence from "missing imports" in China," *Journal of political Economy*, 2004, *112* (2), 471–496.

Fontaine, François, Julien Martin, and Isabelle Mejean, *Frictions and adjustments in firm-to-firm trade*, Centre for Economic Policy Research, 2023.

Gadenne, Lucie, Tushar K Nandi, and Roland Rathelot, "Taxation and Supplier Networks: Evidence from India," *Queen Mary University of London Working Paper No. 947*, 2023.

Gupta, Bhanu, Tejaswi Velayudhan, Prince Dhawan, and Ankur Garg, "Why Don't Firms File Tax Returns on Time? Evidence from the Indian GST," Technical Report, ICTD Working Paper 207, Brighton: Institute of Development Studies 2024.

Mascagni, Giulia, "From the lab to the field: A review of tax experiments," *Journal of Economic Surveys*, 2018, *32* (2), 273–301.

Miyauchi, Yuhei, "Matching and agglomeration: Theory and evidence from japanese firm-to-firm trade," Technical Report Forthcoming.

Okunogbe, Oyebola and Victor Pouliquen, "Technology, taxation, and corruption: evidence from the introduction of electronic tax filing," *American Economic Journal: Economic Policy*, 2022, *14* (1), 341–72.

Paetzold, Jörg and Hannes Winner, "Taking the high road? Compliance with commuter tax allowances and the role of evasion spillovers," *Journal of Public Economics*, 2016, *143*, 1–14.

Paula, Aureo De and Jose A Scheinkman, "Value-added taxes, chain effects, and informality," *American Economic Journal: Macroeconomics*, 2010, *2* (4), 195–221.

Pomeranz, Dina, "No taxation without information: Deterrence and self-enforcement in the value added tax," *American Economic Review*, 2015, *105* (8), 2539–2569.

Shah, Jawad, "Using Computerized Information to Enforce VAT: Evidence from Pakistan," *Available at SSRN 4569607*, 2020.

— , "Using Computerized Information to Enforce VAT: Evidence from Pakistan," 2023.

Slemrod, Joel and Tejaswi Velayudhan, "The VAT at 100: A Retrospective Survey and Agenda for Future Research," *Public Finance Review*, 2022, *50* (1), 4–32.

Stuart, Elizabeth A, Stephen R Cole, Catherine P Bradshaw, and Philip J Leaf, "The use of propensity scores to assess the generalizability of results from randomized trials," *Journal of the Royal Statistical Society Series A: Statistics in Society*, 2011, *174* (2), 369–386.

Tait, Alan A, "Value-added tax: administrative and policy issues," in "Value-Added Tax," International Monetary Fund, 1991.

Waseem, Mazhar, "The role of withholding in the self-enforcement of a value-added tax: Evidence from Pakistan," *Review of Economics and Statistics*, 2022, *104* (2), 336–354.

— , "Overclaimed refunds, undeclared sales, and invoice mills: Nature and extent of noncompliance in a value-added tax," *Journal of Public Economics*, 2023, *218*, 104783.

A Variable Construction

- Share Sellers State-Registered: This variable measures the share of input tax credit (ITC) claimed by a buyer in a given quarter that originates from sellers administered by the state government. We calculate the total ITC received by the buyer from their GSTR-3 form, the amount of ITC received from each seller from the invoice level data in the GSTR-1 form and whether the seller is assigned to the state government or not from the registration form.

- Share ITC from Delhi Sellers : Unlike the previous measure which restricted sellers to those administered by the state government, this variable measures the share of ITC claimed by a buyer that originates from all the sellers i.e. administered by both state and central government.

- Output and Net tax liability: We measure the total output tax liability of a taxpayer using the following equation:

  $$OutputLiability = TotalCash(CGST+SGST+IGST)+TotalCredits(CGST+SGST+IGST)$$

  where, $SGST$ and $CGST$ are the GST paid to the state and central government, respectively. $IGST$ is the tax paid on the inter-state and international transactions. The Net tax liability is the tax paid in only cash after the taxpayer has claimed all the input tax credits.

- **Herfindahl-Hirschman Index of the buyer's input market**: We first calculate the HHI of each supplier's output market and then calculate the HHI of a buyer's input markets as a function of the relative tax credit by each supplier.

  Each firm must list the product-codes (using standard HSN codes) that they intend to sell when they register under GST. The first two digits of the HSN describe the industry category of the product. Multi-product firms list products in order of priority. We calculate the concentration of sellers within each 2-digit industry as follows:

$$HHI_i = \sum_{s=1}^{N_2} \left(\frac{Sales_s}{\sum_{s}^{N_2} Sales_s} \right)^2$$

where N_2 is the total number of sellers s in industry i.

Because we do not know exactly which product is transacted between a buyer and seller, we make the assumption that a buyer is purchasing the supplier's top output and we assign the seller the HHI of the industry of their top output (HHI_s). The average concentration of a buyer's input markets (HHI_b) is then calculated as:

$$HHI_b = \sum_{s=1}^{N_1} \frac{Sales_{sb}}{\sum_{s=1}^{N_1} Sales_{sb}} \times HHI_s$$

where the sum of HHI_s of each seller that a buyer procures from in each quarter is weighted by the sales between buyer b and seller s, $Sales_{sb}$, as a share of total sales from all N_1 suppliers to that buyer in that quarter. The HHI varies from 0 to 1 with 1 implying a perfect monopoly.

B Sample Selection

Our sample is drawn from the universe of GST taxpayers registered with the Delhi Department of Revenue. Within this firm register, we focus on a target population of taxpayers as defined by their past non-compliance behavior.

B.1 Target Population

We target two types of behavior: non-filing and late-filing. The distinction between the two behaviors is important in terms of underlying drivers of each. Filing just a few days or even a month after the deadline is much more likely to be caused by forgetfulness, lack of salience, or disorganization than failing to file a return for 3 months or more. Another reason to distinguish between the two is that tax authorities may want to address both past and future behavior of taxpayers. Nudges could be used to get taxpayers to file pending returns (i.e. non-filers) or it could be used to ensure future on-time filing of taxpayers who filed past returns, but did so after the deadline. These two behaviors may respond differently to nudges. Moreover, identifying the population of future late-filers is not as clear as identifying the population of current non-filers.

For these reasons, we define 2 distinct target populations:

1. Non-filers: These are taxpayers who have a pending return to be filed from any time between December 2020 and March 2021, as of 31st July. We do not consider pending returns between March and July 2021 as there may be many returns filed with a delay, which would be considered late-filing instead.

2. Late-filers: These are taxpayers who filed their returns but only after the deadline in any return period from January 2021 - June 2021. These taxpayers are more likely than others to file their future returns late. Past late-filing behavior is indicative of future late-filing behavior – About 78% of taxpayers who late-file in a given month also file late in the subsequent month. In contrast, only 25% of taxpayers who file on time in a given month, fail to do so in the subsequent month.

Any taxpayer who satisfies the criteria for both "non-filing" and "late-filing" is classified as a non-filer for the purposes of randomization into a treatment group. Because the intervention is to send automated messages to taxpayers by phone, we randomized unique phone numbers into treatment groups, and not unique taxpayers. We drop all taxpayers who share phone number with other taxpayers from the sample.

After these restrictions, we are left with a sampling frame of 283,794 taxpayers who we then stratify according to criteria described below.

B.2 Strata

B.2.1 Type of Filing Non-Compliance

We stratify unique phone numbers according to the following types of filing non-compliance:

1. Non-filer of GSTR 1 and GSTR-3
2. Non-filer of GSTR 1 but not GSTR-3
3. Non-filer of GSTR 3 but not GSTR-1
4. Late-filer of GSTR 1 and GSTR 3
5. Late-filer of GSTR 1 but not GSTR-3
6. Late-filer of GSTR 3 but not GSTR 1

We stratify by filing type because each of these filing patterns can be associated with very different reasons for non-compliance. For example, those who consistently file GSTR-1 but not GSTR-3 are likely to be invoice mills/ fake firms. GSTR-3 filing is associated with payment but GSTR-1 is not, which could mean that GSTR-3 filing is affected by liquidity issues while GSTR-1 is not, and so on.

B.2.2 Imputed Turnover

1. Zero (taxpayers have recorded zero tax payments in FY 2020 including both in cash or GST credit)

2. < Rs. 15 million (threshold for composition scheme) [24]

3. > Rs. 15 million

Turnover data is missing for about 20 percent of our target population. Therefore, we use annual total GST payments as a proxy for firms' taxable turnover, by dividing these payments by the standard GST rate of 18% .

B.2.3 Nature of Business

We also consider the primary nature of businesses and categorize taxpayers into 4 strata:

1. Wholesalers

2. Retailers

3. Services

4. Others

Since different business models might be non-compliant for different reasons, we treat these categories as different strata and randomize within them. Additionally, this approach enhances the external validity of our results across a diverse set of business categories.

B.3 Content of the Nudge Messages

Each month, we nudged taxpayers to submit the return of the preceding calender month. Late filers of GSTR-1 were encouraged to file by the 11th or the applicable due date. We modified the deadline to 20th for GSTR-3 late filers. Lastly, for the customized deterrence messages, we modified the message according to the type of form which the taxpayer had late filed from December 2021 onwards.

[24] Composition scheme reduces the burden of regular filing for small taxpayers. These taxpayers cannot claim input tax credit and have to pay taxes on revenue rather than profits.

Content of the Nudge Messages

Treatment Arm	Message in Hindi	Message in English
1. Reminder	माननीय करदाता, आपको दिल्ली जीएसटी डिपार्टमेंट की तरफ से संपर्क किया जा रहा है। आप अपनी *पिछले महीने* की जीएसटी रिटर्न को *वर्तमान महीने की तिथि* या निर्धारित तारिख तक भरें| फाइलिंग प्रक्रिया और भुगतान के तरीकों के बारे में अधिक जानकारी के लिए जीएसटी की वेबसाइट (www.services.gst.gov.in) पर जाएं।	Respected taxpayer, you are being contacted by the Delhi GST Department. Please file your GST return for *last month* by the *insert date of this month* or by the due date. For more information on the filing process and payment methods, please visit the GST website (www.services.gst.gov.in)
2. Deterrence	माननीय करदाता, आपको दिल्ली जीएसटी डिपार्टमेंट की तरफ से संपर्क किया जा रहा है। आप अपनी *पिछले महीने* की जीएसटी रिटर्न को *वर्तमान महीने की तिथि* या निर्धारित तारिख तक भरें| ऐसा ना करने से आप पर 10,000 रुपया तक का जुरमाना और प्रति वर्ष 18 प्रतिशत ब्याज लगाया जायेगा। फाइलिंग प्रक्रिया और भुगतान के तरीकों के बारे में अधिक जानकारी के लिए, जीएसटी की वेबसाइट (www.services.gst.gov.in) पर जाएं।	Respected taxpayer, you are being contacted by the Delhi GST Department. Please file your GST return for *last month* by the *insert date of this month* or by the due date. Failure to do so may result in a penalty of up to ₹10,000 and interest at 18% per annum. For more information on the filing process and payment methods, please visit the GST website (www.services.gst.gov.in).
3. Custom Deterrence	माननीय करदाता, आपको दिल्ली जीएसटी डिपार्टमेंट की तरफ से संपर्क किया जा रहा है। आप अपनी *पिछले महीने* की जीएसटी रिटर्न को *वर्तमान महीने की तिथि* या निर्धारित तारिख तक भरें| जीएसटी रिकार्ड्स के मुताबिक़ दिसंबर 2020 से आज तक आपने *जीएसटी फॉर्म नंबर* की कम से कम एक रिटर्न निर्धारित समय के बाद जमा की है| समय पर जमा नहीं करने पर आप पर 10,000 रुपया तक का जुरमाना और प्रति वर्ष 18 प्रतिशत ब्याज लगाया जायेगा। फाइलिंग प्रक्रिया और भुगतान के तरीकों के बारे में अधिक जानकारी के लिए, जीएसटी की वेबसाइट (www.services.gst.gov.in) पर जाएं।	Respected taxpayer, you are being contacted by the Delhi GST Department. Please file your GST return for *last month* by the *insert date of this month* or by the due date. Failure to file on time may result in a penalty of up to ₹10,000 and interest at 18% per annum. As per GST records, since December 2020, you have filed at least one return of *GST Form Number* after the due date. For more information about the filing process and payment methods, please visit the GST website (www.services.gst.gov.in).

C Penalties for Filing Non-Compliance

Firms are charged Rs. 200 per day for each day's delay in filing the GSTR-1, capped at a maximum between Rs. 500 and Rs. 10,000 depending on the firm's annual revenue in the previous year. Failure to file the GSTR-3B also carries penalties and additionally interest on the unpaid tax liability at the rate of 18 percent per annum. In addition to penalties, several automated enforcement measures are built in the online portal. First, the return for a subsequent month cannot be filed on the online portal unless all pending returns are filed. Second, no e-way bills [25] can be generated if a GSTR-3B (payment-linked) return is pending. Third, starting in August 2021, no GSTR-1 (credit-linked) return can be filed if a GSTR-3B return is pending, which implies that a taxpayer has to make a payment before providing credits to her buyer, though the payment can be lower than the credits she provides. Finally, firms that haven't filed their returns for over 6 months can be deregistered.[26]

All these enforcement measures do not affect the internal validity of our estimates, as we rely on experimental identification. In fact, in other contexts which lack such sophisticated enforcement, we might find even stronger effects.

We observe that even though the provision of late fees and interest payments were legally available, they were not fully enforced by the tax officials during our study period. Appendix Figure A1 shows that fewer than 60% of taxpayers who filed a late return pay penalties in the subsequent return period. The probability of making any interest payments after making a late payment was even lower. Appendix Figure A2 shows that almost no taxpayers were making interest payments until the start of 2022.[27] This lax enforcement might lower taxpayer's perception of credible enforcement in case of filing non-compliance. Through a communication intervention, we experimentally increase this perceived probability of enforcement without increasing actual enforcement.

[25] GST E-way bill is a document used to track goods in transit. A registered taxpayer transporting goods of value greater than Rs. 50,000 must possess an E-way bill generated on the GST Portal

[26] Deregistration of a firm mandatorily requires a notice issued by a tax officer which can sometimes take longer than 6 months.

[27] Discussions with tax officers revealed that starting in 2022, the filing portal automatically calculated and filled in the interest payment due. This change could account for the sharp increase in interest payments starting in January 2022. Still, no more than 40 percent of late-filers make the payment.

D Additional Robustness Checks

To test whether the restriction of our main sample to taxpayers with at least 70 percent of their credits from Delhi-registered sellers affects our estimated treatment effects, we follow Stuart et al. (2011) and re-estimate our main specifications after re-weighting the sample to match the characteristics of buyers excluded based on the seller registration criteria.

As Table D.1 shows, the main analysis sample differs slightly from the excluded sample. Our main sample is more skewed towards retail, smaller, fewer services, and slightly older. These differences are very small although statistically significant. We reweight our main sample based on these observable differences. We generate weights by estimating a propensity score for inclusion in the sample based on the following characteristics: log output liability, any out-of-state sales, number of Delhi-registered sellers, months registered, and sector of operation. We then re-estimate our main specification weighting by the inverse of these propensity scores.

Table D.2 shows that our results are very similar after re-weighting. Columns 1 and 2 replicate Columns 3 and 4 of Table 2 where the point estimates are only slightly different. Similarly, columns 3 and 4 of Table D.2 replicate columns 1 and 2 of Table 3 where again, the point estimate are very similar. This exercise reassures us that our results are not biased by the sample limitation based on sellers' location.

Table D.1: Comparison of main sample to excluded sample based on seller network location

	Excluded Sample	Main Sample	P-value of Difference
Log(Output Liability)	11.79	11.40	0.000
Has Inter-state sales	0.84	0.84	0.427
No. of Delhi Sellers	15.78	16.75	0.000
Months Registered	41.10	43.81	0.000
Retailer	0.35	0.40	0.000
Upstream	0.26	0.29	0.000
Services	0.14	0.10	0.000
Observations	47483	26877	.

Notes: Comparison of characteristics, in September 2021, of main sample and excluded sample. In our main analysis, we exclude buyers who source over 30 percent of their input tax credits from sellers located outside Delhi.

Table D.2: On-time Filing Behavior: Re-weighting to match taxpayers with out-of-state seller networks

	All		Above-median	Below-median
	(1)	(2)	(3)	(4)
	GSTR-3B	GSTR-1	GSTR-3B	GSTR-3B
Reminder	0.004	0.006	0.018^{*}	−0.009
	(0.010)	(0.010)	(0.010)	(0.011)
Deterrence	0.016^{**}	0.016^{**}	0.025^{***}	0.007
	(0.007)	(0.007)	(0.007)	(0.008)
Share July Sellers Treated			−0.012	0.005
			(0.020)	(0.019)
Share July Sellers Assigned			0.022	0.023
			(0.018)	(0.018)
Constant	0.796^{***}	0.776^{***}	0.798^{***}	0.774^{***}
	(0.006)	(0.006)	(0.007)	(0.008)
Observations	17,596	17,596	27,467	24,216

Notes: Table presents results of specification (6) in columns 1 and 2, and (5) in columns 3 and 4, with inverse propensity score weights to match taxpayer population unrestricted by seller network location. The dependent variable in columns 1,3 and 4 is a dummy for whether a taxpayer filed the GSTR-3B return (payment-linked) by the deadline in the months of August-October. The sample is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. Column 3 restricts attention to taxpayers with above-median compliance in their seller network. Column 4 restricts taxpayers with below-median network compliance. All specifications include strata by month fixed effects. Standard errors are clustered by taxpayer, and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Appendix Figures and Tables

Figure A1: Late Fee or Interest Payment by Late Filers

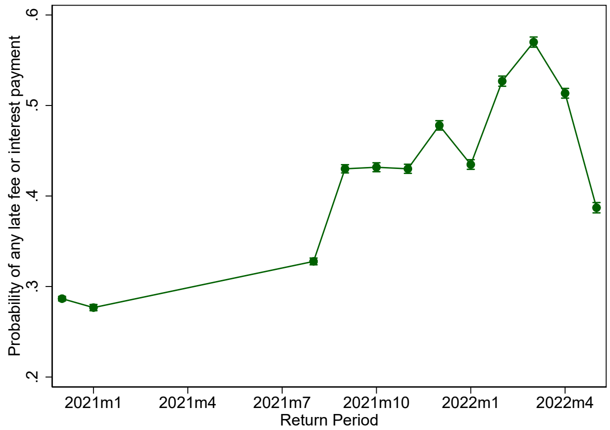

Notes: Figure shows the probability that a taxpayer pays any late fee or interest after filing their previous GSTR-1 or GSTR-3B return late. Sample is restricted to firms that are supposed to file on a monthly basis. We omit Feb - June 2021 return periods because of incomplete data.

Figure A2: Interest Payment by GSTR-3B late filers

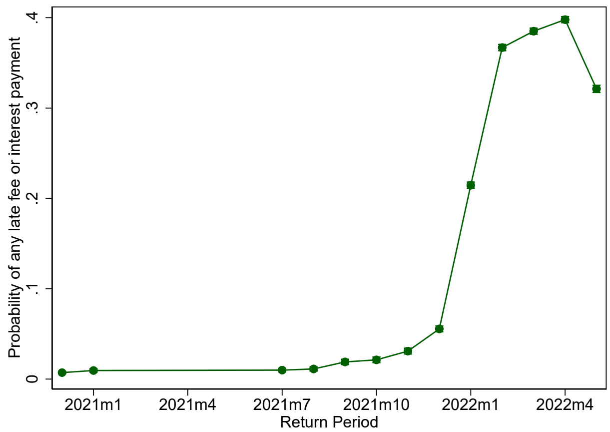

Notes: Figure shows the probability that a taxpayer pays any interest after filing their previous GSTR-3B return late. Sample is restricted to firms that are supposed to file on a monthly basis. We omit Feb - June 2021 return periods because of incomplete data.

Figure A3: Round-wise summary of calls

(a) Response rate of the firms for each round of calling

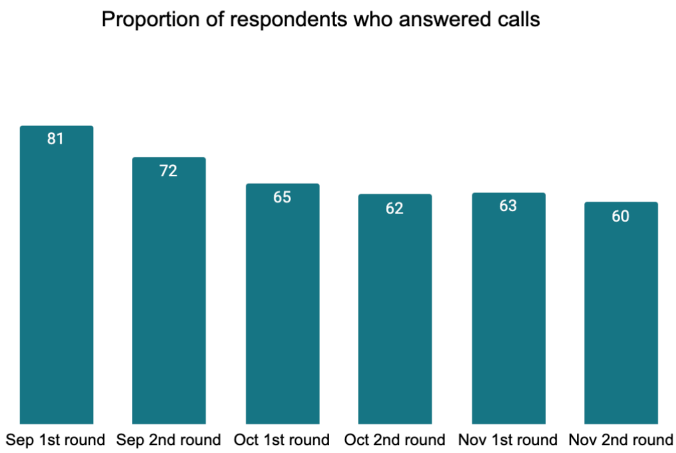

(b) Success rate of calls for each round of calling

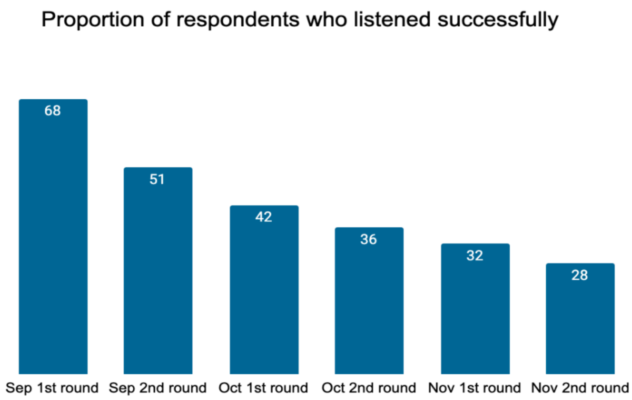

Figure A4: Heterogeneity in Treatment Effect by Network Compliance, Multiple Treatment Arms

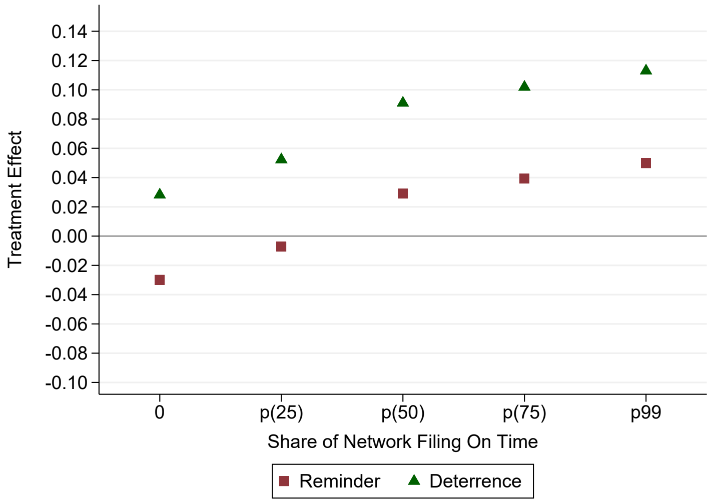

Notes: Figure plots treatment effect as defined in equation (??) for the Reminder and Deterrence messages, evaluated at different quantiles of compliance of the taxpayer's seller network. Standard errors are calculated using the delta method. The treatment effect denotes the percent change in on-time filing among non-filers induced by the treatment.

Table A1: Summary Statistics of Full Sample

	Mean	Median	Std. Dev.	Min.	Max.	Obs.
Previous late-filer	0.60	1.00	0.49	0.00	1.00	341,854
Share Sellers State-Registered	0.61	0.70	0.37	0.00	1.00	233,222
Network Compliance, Jul 2021	0.75	0.98	0.36	0.00	1.00	193,970
Share ITC from Delhi Sellers	11.20	0.77	1,035.30	0.00	317,118.72	189,640
Retailer	0.34	0.00	0.47	0.00	1.00	341,854
Upstream	0.22	0.00	0.41	0.00	1.00	341,854
Services	0.18	0.00	0.38	0.00	1.00	341,854
Months Registered	38.60	49.00	15.11	6.00	50.00	341,853
Has Inter-state sales	0.60	1.00	0.49	0.00	1.00	333,802
Number of Sellers from Delhi	8.65	4.00	14.98	1.00	786.00	233,222
Log(Output Liability)	10.64	11.00	2.53	0.00	20.72	223,109

Notes: This table shows the summary statistics for all the taxpayers registered in Delhi for the month of September, 2021.

Table A2: State Vs. Center-Registered Taxpayers

	Centre	State	P-value of Difference
SGST in FY 2020	586894.37	471256.39	0.000
Has Inter-state sales	0.51	0.58	0.000
No. of Delhi buyers	8.86	9.83	0.000
Months Registered	20.04	25.98	0.000
Retailer	0.34	0.35	0.000
Upstream	0.20	0.23	0.000
Services	0.20	0.18	0.000
Monthly Filer	0.59	0.54	0.000
Pending GSTR 3	0.24	0.21	0.000
Pending GSTR 1	0.23	0.17	0.000
Late Filed GSTR 3	0.50	0.49	0.000
Late Filed GSTR 1	0.75	0.81	0.000
Shared Phone Number	0.39	0.37	0.000
Observations	274861	452551	.

Notes: Comparison of State vs Center-registered taxpayers with some history of filing non-compliance. Sample includes all taxpayers who were GST-registered in Delhi and active as on June 2021.

Table A3: Taxpayers With and Without Shared Registered Phone Numer

	Unique Number	Shared Number	P-value of Difference
SGST in FY 2020	422830.89	553044.07	0.000
Has Inter-state sales	0.60	0.54	0.000
No. of Delhi buyers	10.25	9.09	0.000
Months Registered	26.00	25.94	0.236
Retailer	0.36	0.33	0.000
Upstream	0.23	0.22	0.000
Services	0.17	0.20	0.000
Monthly Filer	0.54	0.54	0.001
Pending GSTR 3	0.21	0.22	0.000
Pending GSTR 1	0.17	0.18	0.000
Late Filed GSTR 3	0.49	0.49	0.621
Late Filed GSTR 1	0.81	0.80	0.000
Observations	284250	168301	.

Notes: Sample consists of population of State-registered taxpayers with some history of non-compliance, active and registered as on June 2021.

Table A4: Balance Table - First Stage

	(1) Control	(2) Reminder	(3) Deterrence	(4) Custom Deterrence
SGST in FY 2020	433,153.82	402,496.40	437,094.12	402,299.11
		(0.34)	(0.88)	(0.20)
Has Inter-state sales	0.60	0.60	0.60	0.60
		(0.72)	(0.77)	(0.74)
No. of Delhi buyers	10.09	10.23	10.25	10.26
		(0.35)	(0.19)	(0.17)
Months Registered	25.98	25.89	26.02	26.00
		(0.35)	(0.65)	(0.81)
Retailer	0.36	0.36	0.36	0.36
		(0.67)	(0.59)	(0.84)
Upstream	0.23	0.23	0.23	0.23
		(0.63)	(0.26)	(0.34)
Services	0.17	0.17	0.17	0.17
		(0.73)	(0.78)	(0.84)
Monthly Filer	0.53	0.54	0.53	0.54
		(0.37)	(0.57)	(0.07)
Pending GSTR 3	0.21	0.21	0.21	0.21
		(0.95)	(0.96)	(0.94)
Pending GSTR 1	0.17	0.17	0.17	0.17
		(0.92)	(0.93)	(0.95)
Late Filed GSTR 3	0.49	0.49	0.49	0.49
		(0.89)	(0.76)	(0.84)
Late Filed GSTR 1	0.81	0.81	0.81	0.81
		(0.89)	(1.00)	(0.99)

Notes: P-value of t-test of difference of mean between treatment group and control in parentheses.

Table A5: Probability of On-time Filing of GSTR-3B by Network Compliance

	Return Month:				
	(1) August	(2) September	(3) October	(4) November	(5) December
Panel A: Buyers Above Median Compliance					
Reminder	0.009	0.011	0.029**	0.007	0.011
	(0.013)	(0.011)	(0.012)	(0.013)	(0.011)
Deterrence	0.024***	0.024***	0.025***	0.006	0.007
	(0.009)	(0.008)	(0.009)	(0.009)	(0.008)
Share July Sellers Treated	-0.008	0.007	-0.002	0.006	0.042*
	(0.026)	(0.022)	(0.025)	(0.026)	(0.022)
Share July Sellers Assigned	0.021	0.013	0.018	0.015	0.007
	(0.023)	(0.020)	(0.023)	(0.023)	(0.020)
Constant	0.781***	0.786***	0.812***	0.819***	0.812***
	(0.009)	(0.007)	(0.008)	(0.008)	(0.007)
Sample Mean	0.80	0.81	0.84	0.83	0.83
Observations	9,709	13,523	9,187	9,185	12,850
Panel B: Buyers Below Median Compliance					
Reminder	-0.019	-0.003	-0.005	-0.001	0.006
	(0.014)	(0.011)	(0.014)	(0.014)	(0.011)
Deterrence	0.007	0.016*	0.002	0.015	0.017**
	(0.010)	(0.008)	(0.010)	(0.010)	(0.008)
Share July Sellers Treated	0.001	-0.027	0.010	0.009	-0.007
	(0.025)	(0.019)	(0.024)	(0.024)	(0.019)
Share July Sellers Assigned	0.025	0.046***	0.022	0.002	0.021
	(0.023)	(0.018)	(0.022)	(0.022)	(0.017)
Constant	0.757***	0.768***	0.800***	0.791***	0.792***
	(0.010)	(0.008)	(0.009)	(0.009)	(0.008)
Sample Mean	0.77	0.78	0.81	0.80	0.81
Observations	8,883	13,349	8,336	8,332	12,691

Notes: Table presents results of specification (5). The dependent variable is a dummy for whether a taxpayer filed the GSTR-3B return (payment-linked) by the deadline in the month listed on top of the column. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: On-time Filing of GSTR-3B by Network Compliance- Direct effects

	Compliance in the Network:	
	(1)	(2)
	Above-median	Below-median
Reminder	0.038**	−0.011
	(0.018)	(0.021)
Deterrence	0.024*	0.003
	(0.014)	(0.015)
Share July Sellers Assigned	−0.034	0.017
	(0.029)	(0.027)
Constant	0.773***	0.455***
	(0.092)	(0.079)
Observations	7,339	7,126
Controls	Yes	Yes

Notes: Table presents results of specification (5). The dependent variable is a dummy for whether a taxpayer filed the GSTR-3B return (payment-linked) by the deadline in the months of August-October. The sample is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. We further only consider buyers with no seller treated in their July network. Columns 1 and 3 restrict attention to taxpayers with above-median compliance in their seller network, which columns 2 and 4 are taxpayers with below-median network compliance. All specifications include strata by month fixed effects. Standard errors clustered by taxpayer, reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Probability of Net Zero VAT Liability by Network Compliance

	Return Month:				
	(1)	(2)	(3)	(4)	(5)
	August	September	October	November	December
Panel A: Buyers Above Median Compliance					
Reminder	0.023	0.010	0.023	0.034**	0.010
	(0.016)	(0.014)	(0.017)	(0.017)	(0.014)
Deterrence	0.026**	0.022**	0.015	0.020	0.009
	(0.012)	(0.010)	(0.012)	(0.012)	(0.010)
Share July Sellers Treated	-0.001	0.031	0.055	0.014	0.052*
	(0.033)	(0.028)	(0.034)	(0.034)	(0.029)
Share July Sellers Assigned	-0.014	-0.036	-0.069**	-0.016	-0.070***
	(0.029)	(0.025)	(0.030)	(0.031)	(0.026)
Constant	0.453***	0.434***	0.453***	0.459***	0.447***
	(0.011)	(0.009)	(0.011)	(0.011)	(0.009)
Sample Mean	0.47	0.44	0.46	0.47	0.45
Observations	9,709	13,523	9,187	9,185	12,850
Panel B: Buyers Below Median Compliance					
Reminder	-0.021	-0.005	-0.011	-0.004	-0.002
	(0.017)	(0.014)	(0.017)	(0.017)	(0.014)
Deterrence	0.002	0.014	0.013	0.011	0.016
	(0.012)	(0.010)	(0.013)	(0.013)	(0.010)
Share July Sellers Treated	0.018	0.034	0.029	0.019	0.001
	(0.029)	(0.023)	(0.030)	(0.030)	(0.024)
Share July Sellers Assigned	-0.030	-0.039*	-0.039	-0.008	-0.023
	(0.027)	(0.021)	(0.028)	(0.028)	(0.022)
Constant	0.469***	0.433***	0.462***	0.483***	0.451***
	(0.011)	(0.009)	(0.012)	(0.012)	(0.010)
Sample Mean	0.46	0.44	0.46	0.49	0.45
Observations	8,883	13,349	8,336	8,332	12,691

Notes: Table presents results of specification (5). The dependent variable is a dummy for whether a taxpayer reported net zero liability (i.e. zero payment in cash) by the deadline in the month listed on top of the column. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median network compliance. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Probability of Higher Net VAT Liability by Network Compliance

	Return Month:				
	(1)	(2)	(3)	(4)	(5)
	August	September	October	November	December
Panel A: Buyers Above Median Compliance					
Reminder	-0.004	-0.021	-0.026	-0.023	-0.005
	(0.015)	(0.013)	(0.016)	(0.016)	(0.014)
Deterrence	-0.023**	-0.021**	-0.030***	-0.013	-0.017*
	(0.011)	(0.010)	(0.012)	(0.011)	(0.010)
Share July Sellers Treated	-0.009	0.002	-0.076**	-0.021	-0.058**
	(0.030)	(0.027)	(0.033)	(0.032)	(0.028)
Share July Sellers Assigned	0.022	0.002	0.071**	0.023	0.043*
	(0.027)	(0.024)	(0.029)	(0.028)	(0.025)
Constant	0.309***	0.337***	0.351***	0.313***	0.347***
	(0.010)	(0.009)	(0.011)	(0.010)	(0.009)
Sample Mean	0.30	0.32	0.33	0.30	0.34
Observations	9,709	13,523	9,187	9,185	12,850
Panel B: Buyers Below Median Compliance					
Reminder	0.007	-0.011	0.020	-0.002	0.015
	(0.016)	(0.013)	(0.017)	(0.016)	(0.014)
Deterrence	-0.004	-0.014	-0.008	-0.002	-0.006
	(0.012)	(0.010)	(0.012)	(0.012)	(0.010)
Share July Sellers Treated	-0.010	-0.040*	-0.018	-0.017	-0.007
	(0.028)	(0.022)	(0.029)	(0.028)	(0.023)
Share July Sellers Assigned	0.031	0.032	0.026	0.010	0.016
	(0.025)	(0.020)	(0.026)	(0.026)	(0.021)
Constant	0.305***	0.336***	0.323***	0.300***	0.329***
	(0.011)	(0.009)	(0.011)	(0.011)	(0.009)
Sample Mean	0.31	0.33	0.33	0.30	0.33
Observations	8,883	13,349	8,336	8,332	12,691

Notes: Table presents results of specification (5). The dependent variable is an indicator for whether the taxpayer reports a net tax liability greater than their net liability in July 2021 by the deadline in the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Probability of Zero Output Liability by Network Compliance

	Return Month:				
	(1) August	(2) September	(3) October	(4) November	(5) December
Panel A: Buyers Above Median Compliance					
Reminder	-0.001	0.006	-0.001	-0.003	-0.002
	(0.006)	(0.004)	(0.005)	(0.006)	(0.004)
Deterrence	0.006	0.005*	-0.001	-0.002	-0.000
	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
Share July Sellers Treated	0.010	0.002	0.000	-0.013	-0.007
	(0.012)	(0.008)	(0.011)	(0.012)	(0.009)
Share July Sellers Assigned	-0.014	-0.010	-0.017*	-0.009	-0.002
	(0.011)	(0.008)	(0.010)	(0.011)	(0.008)
Constant	0.033***	0.021***	0.030***	0.040***	0.024***
	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
Sample Mean	0.03	0.02	0.02	0.03	0.02
Observations	9,709	13,523	9,187	9,185	12,850
Panel B: Buyers Below Median Compliance					
Reminder	0.006	-0.007	-0.010	0.004	-0.001
	(0.008)	(0.006)	(0.008)	(0.009)	(0.005)
Deterrence	0.006	0.001	0.002	0.005	0.005
	(0.006)	(0.004)	(0.006)	(0.007)	(0.004)
Share July Sellers Treated	-0.009	0.003	-0.018	-0.015	-0.010
	(0.014)	(0.009)	(0.013)	(0.015)	(0.009)
Share July Sellers Assigned	-0.003	-0.013	0.005	0.012	-0.001
	(0.013)	(0.009)	(0.012)	(0.014)	(0.008)
Constant	0.060***	0.044***	0.052***	0.065***	0.037***
	(0.006)	(0.004)	(0.005)	(0.006)	(0.004)
Sample Mean	0.06	0.04	0.05	0.07	0.04
Observations	8,883	13,349	8,336	8,332	12,691

Notes: Table presents results of specification (5). The dependent variable is a dummy for whether tax liability on sales alone is zero (i.e. zero taxable sales reported) for the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: Probability of Higher Output Liability by Network Compliance

	Return Month:				
	(1)	(2)	(3)	(4)	(5)
	August	September	October	November	December
Panel A: Buyers Above Median Compliance					
Reminder	-0.012	-0.043***	-0.018	-0.016	-0.033**
	(0.016)	(0.014)	(0.017)	(0.017)	(0.014)
Deterrence	-0.007	-0.026**	-0.034***	-0.025**	-0.028***
	(0.012)	(0.010)	(0.012)	(0.012)	(0.010)
Share July Sellers Treated	0.003	0.014	-0.001	0.051	0.030
	(0.032)	(0.028)	(0.035)	(0.034)	(0.030)
Share July Sellers Assigned	0.008	0.020	0.037	-0.015	0.004
	(0.029)	(0.025)	(0.031)	(0.030)	(0.026)
Constant	0.397***	0.480***	0.504***	0.414***	0.514***
	(0.011)	(0.009)	(0.011)	(0.011)	(0.010)
Sample Mean	0.39	0.47	0.49	0.40	0.50
Observations	9,709	13,523	9,187	9,185	12,850
Panel B: Buyers Below Median Compliance					
Reminder	-0.000	0.017	0.008	-0.012	-0.006
	(0.017)	(0.014)	(0.018)	(0.017)	(0.014)
Deterrence	0.004	-0.002	0.002	-0.014	-0.019*
	(0.012)	(0.010)	(0.013)	(0.013)	(0.010)
Share July Sellers Treated	0.014	-0.017	0.026	-0.024	-0.020
	(0.029)	(0.024)	(0.031)	(0.030)	(0.024)
Share July Sellers Assigned	-0.016	0.008	-0.005	-0.008	-0.002
	(0.026)	(0.021)	(0.028)	(0.027)	(0.022)
Constant	0.392***	0.448***	0.463***	0.419***	0.516***
	(0.011)	(0.009)	(0.012)	(0.012)	(0.010)
Sample Mean	0.39	0.45	0.47	0.40	0.50
Observations	8,883	13,349	8,336	8,332	12,691

Notes: Table presents results of specification (5). The dependent variable is a dummy for whether output liability is higher than in July 2021 for the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: Probability of Higher ITC by Network Compliance

	Return Month:				
	(1) August	(2) September	(3) October	(4) November	(5) December
Panel A: Buyers Above Median Compliance					
Reminder	-0.041**	-0.026*	-0.017	-0.016	-0.012
	(0.016)	(0.014)	(0.017)	(0.017)	(0.014)
Deterrence	-0.017	-0.017*	-0.012	-0.012	-0.006
	(0.012)	(0.010)	(0.012)	(0.012)	(0.010)
Share July Sellers Treated	-0.025	-0.002	-0.030	-0.030	-0.010
	(0.033)	(0.028)	(0.034)	(0.034)	(0.029)
Share July Sellers Assigned	-0.008	0.004	0.043	0.043	0.030
	(0.030)	(0.025)	(0.031)	(0.031)	(0.026)
Constant	0.542***	0.525***	0.556***	0.556***	0.547***
	(0.011)	(0.009)	(0.011)	(0.011)	(0.009)
Sample Mean	0.52	0.51	0.55	0.55	0.55
Observations	9,709	13,523	9,187	9,185	12,850
Panel B: Buyers Below Median Compliance					
Reminder	0.011	0.008	-0.000	-0.001	0.002
	(0.017)	(0.014)	(0.017)	(0.017)	(0.014)
Deterrence	-0.005	-0.011	-0.023*	-0.023*	-0.015
	(0.012)	(0.010)	(0.013)	(0.013)	(0.010)
Share July Sellers Treated	0.041	0.015	0.008	0.008	-0.005
	(0.029)	(0.024)	(0.030)	(0.030)	(0.024)
Share July Sellers Assigned	-0.048*	-0.032	-0.024	-0.025	-0.013
	(0.027)	(0.021)	(0.028)	(0.028)	(0.022)
Constant	0.505***	0.502***	0.549***	0.549***	0.539***
	(0.011)	(0.009)	(0.012)	(0.012)	(0.010)
Sample Mean	0.50	0.49	0.53	0.53	0.53
Observations	8,883	13,349	8,336	8,332	12,691

Notes: Table presents results of specification (5). The dependent variable is a dummy for whether the taxpayer claimed more credits than in July 2021 for the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A12: Probability of Repeating Any Late Seller by Network Compliance

	Return Month:				
	(1) August	(2) September	(3) October	(4) November	(5) December
Panel A: Buyers Above Median Compliance					
Reminder	0.001	-0.026**	-0.009	-0.028*	-0.024*
	(0.016)	(0.013)	(0.016)	(0.015)	(0.012)
Deterrence	0.001	-0.010	-0.011	-0.025**	-0.009
	(0.012)	(0.010)	(0.012)	(0.011)	(0.009)
Share July Sellers Treated	0.031	0.019	0.036	0.032	-0.021
	(0.032)	(0.027)	(0.033)	(0.032)	(0.026)
Share July Sellers Assigned	-0.078***	-0.076***	-0.085***	-0.053*	-0.036
	(0.029)	(0.025)	(0.030)	(0.028)	(0.023)
Constant	0.409***	0.370***	0.379***	0.315***	0.264***
	(0.011)	(0.009)	(0.011)	(0.010)	(0.008)
Sample Mean	0.40	0.34	0.36	0.29	0.24
Observations	9,547	13,140	9,019	8,993	12,527
Panel B: Buyers Below Median Compliance					
Reminder	0.007	-0.000	-0.008	0.009	0.006
	(0.017)	(0.014)	(0.017)	(0.016)	(0.013)
Deterrence	-0.024*	-0.009	-0.011	-0.005	-0.004
	(0.013)	(0.010)	(0.013)	(0.012)	(0.009)
Share July Sellers Treated	0.030	-0.004	0.002	-0.040	-0.014
	(0.031)	(0.024)	(0.031)	(0.030)	(0.022)
Share July Sellers Assigned	-0.191***	-0.111***	-0.131***	-0.077***	-0.098***
	(0.028)	(0.022)	(0.028)	(0.027)	(0.020)
Constant	0.565***	0.393***	0.417***	0.340***	0.294***
	(0.012)	(0.009)	(0.012)	(0.011)	(0.009)
Sample Mean	0.50	0.35	0.37	0.30	0.26
Observations	8,387	12,237	7,849	7,763	11,841

Notes: Table presents results of specification (5). The dependent variable is a dummy for whether a seller from the taxpayer's network who had late filed GSTR-1 in the previous month appears again in the taxpayer's network in the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Proportion of ITC from Repeat Late Seller, by Network Compliance

	Return Month:				
	(1)	(2)	(3)	(4)	(5)
	August	September	October	November	December
Panel A: Buyers Above Median Compliance					
Reminder	0.001	0.003	0.007	0.002	-0.002
	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)
Deterrence	-0.000	0.002	0.002	-0.004	0.000
	(0.002)	(0.003)	(0.003)	(0.003)	(0.002)
Share July Sellers Treated	0.010	0.016**	-0.003	0.005	-0.003
	(0.007)	(0.007)	(0.008)	(0.008)	(0.005)
Share July Sellers Assigned	-0.020***	-0.016***	-0.005	-0.004	-0.004
	(0.006)	(0.006)	(0.008)	(0.007)	(0.005)
Constant	0.035***	0.037***	0.037***	0.032***	0.023***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.002)
Sample Mean	0.03	0.04	0.04	0.03	0.02
Observations	9,547	13,315	9,108	9,051	12,712
Panel B: Buyers Below Median Compliance					
Reminder	0.011	0.004	-0.006	0.012*	0.000
	(0.010)	(0.006)	(0.007)	(0.007)	(0.005)
Deterrence	-0.001	-0.003	-0.007	0.005	-0.002
	(0.008)	(0.004)	(0.005)	(0.005)	(0.003)
Share July Sellers Treated	0.030	0.005	0.003	-0.012	-0.001
	(0.019)	(0.010)	(0.013)	(0.013)	(0.008)
Share July Sellers Assigned	-0.090***	-0.029***	-0.037***	-0.015	-0.016**
	(0.017)	(0.009)	(0.012)	(0.011)	(0.007)
Constant	0.198***	0.089***	0.100***	0.074***	0.057***
	(0.007)	(0.004)	(0.005)	(0.005)	(0.003)
Sample Mean	0.18	0.08	0.08	0.07	0.05
Observations	8,387	12,729	8,059	7,977	12,300

Notes: Table presents results of specification (5). The dependent variable is the input tax credits available to the taxpayer from delinquent sellers in the previous month, as a share of their total ITC available in the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A14: Proportion of Late Sellers Repeated, by Network Compliance

	Return Month:				
	(1) August	(2) September	(3) October	(4) November	(5) December
Panel A: Buyers Above Median Compliance					
Reminder	0.005	-0.018	-0.010	-0.013	-0.020*
	(0.014)	(0.012)	(0.013)	(0.013)	(0.011)
Deterrence	0.009	-0.004	-0.008	-0.020**	-0.007
	(0.010)	(0.008)	(0.009)	(0.010)	(0.008)
Share July Sellers Treated	0.029	0.017	0.021	0.022	-0.020
	(0.028)	(0.024)	(0.026)	(0.027)	(0.023)
Share July Sellers Assigned	-0.054**	-0.053**	-0.048**	-0.036	-0.023
	(0.025)	(0.021)	(0.023)	(0.024)	(0.020)
Constant	0.313***	0.292***	0.261***	0.245***	0.219***
	(0.009)	(0.008)	(0.009)	(0.009)	(0.007)
Sample Mean	0.31	0.28	0.25	0.23	0.20
Observations	9,547	13,140	9,019	8,993	12,527
Panel B: Buyers Below Median Compliance					
Reminder	-0.007	-0.005	-0.004	0.007	0.002
	(0.015)	(0.012)	(0.014)	(0.014)	(0.011)
Deterrence	-0.029***	-0.007	-0.012	-0.003	-0.006
	(0.011)	(0.009)	(0.010)	(0.010)	(0.008)
Share July Sellers Treated	0.039	-0.007	0.007	-0.027	-0.018
	(0.027)	(0.021)	(0.025)	(0.025)	(0.020)
Share July Sellers Assigned	-0.169***	-0.080***	-0.092***	-0.065***	-0.070***
	(0.025)	(0.019)	(0.022)	(0.023)	(0.018)
Constant	0.459***	0.320***	0.296***	0.269***	0.248***
	(0.010)	(0.008)	(0.009)	(0.010)	(0.008)
Sample Mean	0.40	0.29	0.26	0.24	0.22
Observations	8,387	12,237	7,849	7,763	11,841

Notes: Table presents results of specification (5). The dependent variable is the share of all delinquent sellers in their previous month's network that appear again in their seller network in the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A15: Proportion of Sellers Filing On-time, by Network Compliance

	Return Month:				
	(1) August	(2) September	(3) October	(4) November	(5) December
Panel A: Buyers Above Median Compliance					
Reminder	0.007	-0.012***	0.003	0.005	-0.001
	(0.006)	(0.004)	(0.006)	(0.006)	(0.004)
Deterrence	-0.002	-0.002	0.003	0.004	0.004
	(0.005)	(0.003)	(0.004)	(0.005)	(0.003)
Share July Sellers Treated	-0.005	-0.014	-0.016	-0.006	-0.006
	(0.013)	(0.009)	(0.012)	(0.013)	(0.008)
Share July Sellers Assigned	-0.023**	0.025***	-0.030***	-0.018	0.019***
	(0.011)	(0.008)	(0.011)	(0.011)	(0.007)
Constant	0.798***	0.883***	0.822***	0.825***	0.894***
	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
Sample Mean	0.79	0.88	0.81	0.82	0.90
Observations	9,547	13,315	9,108	9,051	12,712
Panel B: Buyers Below Median Compliance					
Reminder	-0.010	-0.002	-0.007	0.003	0.003
	(0.010)	(0.006)	(0.009)	(0.009)	(0.006)
Deterrence	-0.004	0.005	-0.000	-0.001	0.006
	(0.007)	(0.005)	(0.007)	(0.007)	(0.004)
Share July Sellers Treated	-0.007	0.005	0.021	0.017	0.010
	(0.017)	(0.011)	(0.016)	(0.017)	(0.010)
Share July Sellers Assigned	-0.060***	0.042***	-0.069***	-0.074***	0.024***
	(0.016)	(0.010)	(0.015)	(0.015)	(0.009)
Constant	0.684***	0.821***	0.728***	0.744***	0.850***
	(0.007)	(0.004)	(0.006)	(0.006)	(0.004)
Sample Mean	0.66	0.84	0.71	0.72	0.86
Observations	8,387	12,729	8,060	7,977	12,300

Notes: Table presents results of specification (5). The dependent variable is the share of their sellers in the current return period that filed their GSTR-1 on-time in the month specified in the column header. The sample in both panels is restricted to our main analysis sample of firms with some history of non-compliance, and at least 70 percent of tax credits from Delhi-registered sellers prior to treatment. The sample in Panel A is additionally restricted to those buyers with above-median (i.e. greater than 87%) credits filed by their suppliers in July, while Panel B consists of below-median buyers. All specifications include strata fixed effects. Standard errors reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A16: Correlation between Persistence and Competitiveness of the Supplier's Network

	Proportion of Repeated Sellers	
	(1)	(2)
HHI of the Sellers	0.084^{***}	0.309^{***}
	(0.001)	(0.002)
Constant	0.373^{***}	0.259^{***}
	(0.001)	(0.001)
Observations	1,109,599	1,011,461
Month FEs	Yes	Yes
Taxpayer FEs	No	Yes

Notes: Each observation is a taxpayer and return period. The dependent variable is the proportion of sellers that are repeated from the previous month. We measure the competitiveness of a buyer's input market by the Herfindahl-Hirschman Index (HHI) of their inputs. The sample consists of all the registered Delhi buyers from May-July 2021. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.