

Firm Networks and Tax Compliance: Experimental Evidence from Uganda*

Miguel Almunia

CUNEF Universidad

David Henning

UCLA

Justine Knebelmann

Sciences Po

Dorothy Nakymbadde

Uganda Revenue Authority

Lin Tian

INSEAD

February 2025

Abstract

How do tax enforcement interventions diffuse through firm-to-firm networks? We explore this question with a randomized trial in Uganda. Using transaction-level VAT data, we map seller-buyer networks and identify discrepancies in the amounts reported by trading partners. Enforcement letters highlighting these discrepancies are sent to either the seller, the buyer, or both. The correction rate in the treatment group is 23.8%, fourteen times higher than in the control group. This response is asymmetric: corrections are primarily made by sellers, even when only buyers receive letters, providing novel evidence that firms can induce changes in their partners' tax reporting. Spillover effects extend to transactions not listed in the letters, including those involving other trading partners. The intervention also results in sustained improvements in reporting behavior over subsequent months. Our study sheds light on firm-to-firm communication within networks and offers policy-relevant insights for fighting tax evasion.

Keywords: firm networks; tax enforcement; spillovers; value-added tax (VAT); Uganda

JEL codes: O12, L14, H26, H25.

*Almunia: CUNEF Universidad, CEPR, IGC, IFS, miguel.almunia@cunef.edu. Henning: UCLA, djhenning@ucla.edu. Knebelmann: Sciences Po, IFS, J-PAL, justine.knebelmann@sciencespo.fr. Nakymbadde: Uganda Revenue Authority, dnakymbadde@ura.go.ug. Tian: INSEAD, CEPR, lin.tian@insead.edu. We thank the staff at the Uganda Revenue Authority (URA) for their invaluable collaboration throughout this project. Stefano D'Angelo, Pablo García-Guzmán, Diana C. León, and Claude Raisaro provided excellent research assistance. We also benefited from many insightful discussions with Jonas Hjort. We gratefully acknowledge funding from the International Growth Centre (IGC), Fundación Ramón Areces, the British Academy, the Leverhulme Foundation, and the Spanish Ministry of Science (Agencia Estatal de Investigación) through a Ramón y Cajal Fellowship. This RCT was pre-registered as AEARCT-0002958 (Almunia et al., 2018). An earlier version of this paper was circulated under the title "Leveraging trading networks to improve tax compliance: Experimental evidence from Uganda" (CEPR DP 18151).

1 Introduction

Tax evasion by firms remains a persistent challenge for governments worldwide, particularly in developing countries with limited administrative capacity (Besley and Persson, 2013, 2014; Best et al., 2015). Firms operate within highly interconnected networks, where information flows and spillovers may amplify, or diffuse, the effects of government policies. To what extent tax enforcement measures—a core government policy—spread through transaction networks, and whether these networks can be leveraged to improve enforcement effectiveness, remain open questions. Answering these questions is challenging because the most granular unit in a network—and the ideal unit of treatment—is a *pair* of firms, whereas enforcement policies typically target *individual* firms (Carrillo et al., 2017; Holz et al., 2023).

We study these questions in the context of the Value-Added Tax (VAT), which is relevant for two reasons. First, VAT is the main source of tax revenue in developing countries, yet significant compliance gaps persist (Brockmeyer et al., 2024). Second, its structure inherently creates conflicting reporting incentives for sellers and buyers (Keen and Lockwood, 2010; Pomeranz, 2015; Slemrod and Velayudhan, 2022), which shape how compliance behaviors propagate through transaction networks (Gadenne et al., 2024). Understanding how these network effects shape compliance is crucial for designing effective enforcement policies, as it requires accounting for both the direct impact of interventions and the spillovers they produce.

In this paper, we present experimental evidence on how information from firms’ transaction networks can be used to improve tax compliance among firms and their trading partners. We design a novel two-stage randomized controlled trial (RCT) in Uganda—implemented with the Uganda Revenue Authority (URA)—that separately estimates the direct and spillover effects of an enforcement intervention. Leveraging an administrative dataset of firm-to-firm transactions reported in VAT returns, the study is, to our knowledge, the first RCT to use such fine-grained data *ex ante* both to map taxpayer networks and to design a targeted intervention, fully accounting for the high interconnectedness of firms. We see this as a major innovation because failing to account for the transaction network in the experimental design makes it challenging to disentangle the direct and spillover effects (Vazquez-Bare, 2023).¹

¹Our approach builds on prior studies examining firms’ social or geographical networks (Boning et al., 2020; Drago et al., 2020; Cruces et al., 2024), or assessing effects of interventions on firms’ clients and suppliers without incorporating network data into the intervention design (Pomeranz, 2015; Carrillo et al., 2023; Garriga and Tortarolo, 2024). Firms’ transaction networks are uniquely complex due to high interconnectedness. This explains why, despite their central role in shaping economic outcomes (McMillan and Woodruff, 1999), randomized trials involving firms and their transaction networks are underrepresented in the experimental literature, with exceptions such as Atkin et al. (2017) and Hjort et al. (2021).

The experiment exploits the full network of VAT-registered firms in Uganda, constructed from transaction data spanning March to December 2017. This dataset comprises 115,856 unique seller-buyer links, allowing us to identify potential tax evasion cases based on discrepancies between the amounts reported by sellers and buyers in bilateral trade. A key reason for VAT’s persistent compliance gaps in low-income countries is firms’ belief that revenue authorities cannot effectively cross-check reports from sellers and buyers, weakening VAT’s self-enforcing properties (Mascagni et al., 2022; Almunia et al., 2024). The intervention aims to update these beliefs and assess how improved perceptions of enforcement affect compliance and propagate through firm networks. Specifically, we focus on instances where sellers report a lower amount than buyers, referred to as “seller shortfall” henceforth.² The intervention targets a randomly selected subset of these seller-buyer links by sending official letters from the URA. These letters notify firms of the URA’s newly developed method for detecting misreporting, list up to three discrepancies with the trading partner in the selected link, and request prompt amendments of past returns. The detailed information communicated in these letters allows us to identify exactly what information was passed on to trading partners. To interpret the results from our experiment, we lay out a simple conceptual framework.

Our experimental design incorporates three innovations to rigorously account for the network structure of the data. First, treatment is randomized at the seller-buyer link level, rather than the firm level, to prevent violations of the stable-unit treatment value assumption (SUTVA) due to potential spillovers through trading partners (Rosenbaum, 2007). Second, an iterative sampling procedure ensures that no firm is exposed to multiple treatments by selecting 1,235 seller-buyer links separated by at least one degree in the firm network. Specifically, once a seller-buyer link is selected to be in the study sample, all directly connected eligible links are excluded from further selection.³ This approach enables—for the first time—a convincing identification of the letter’s effects on treated firms’ links with their untreated partners. Third, treatment intensity varies by changing the letter recipient(s) at the link level: in one arm, only the seller receives the letter (20%), in another, only the buyer (20%), and in the third arm, both (20%).⁴ This design allows us to evaluate whether targeting pairs

²Seller shortfall is widespread in the Ugandan VAT data, and suggests tax-evading behavior, as it can potentially reduce net VAT liability for at least one or both firms involved (Almunia et al., 2024). Similar patterns have been documented in other low-income countries (Brockmeyer et al., 2024).

³The average seller in the final study sample has 45 clients and 17 suppliers over the ten months prior to the intervention.

⁴The remaining 40% of links form the control group, receiving no communication from the revenue authority. This design parallels two-stage cluster randomization or partial population designs commonly used to study spillover effects at different treatment intensities (Cruces et al., 2024), with clusters defined as seller-buyer links.

of linked firms in a network is more effective than treating them individually.

The intervention significantly increases corrections of past discrepancies, supporting our prediction that the letters update firms’ beliefs about the administration’s detection capacity. In treated links, the correction rate for listed discrepancies rises by 22.3 percentage points (pp), a fourteen-fold increase over the 1.6% in the control group. Corrections of discrepancies not listed in the letters increase by 10.6pp within treated links relative to the control group. Additionally, the intervention leads to sustained improvements in reporting behavior, including a 15.0pp (19%) reduction in seller shortfall instances and an 13.0pp (115%) increase in matching reports over ten months after the treatment. These findings indicate that the intervention strengthens firms’ perceptions of the tax administration’s monitoring and enforcement capacity, driving both immediate and lasting compliance gains.

The intervention also reveals marked differences between sellers and buyers, highlighting asymmetry in reporting behaviors and providing evidence of widespread firm-to-firm communication. Sellers account for most corrections, with 19.8% of discrepancies corrected by sellers in the treatment group compared to only 2.6% by buyers. Sellers also drive the spillover effects to untreated links. Even more striking is the finding that, when the letters are only sent to buyers, they prompt a 7.9pp, or fivefold, increase in corrections made by *sellers*, underscoring the influence firms have on their trading partners’ compliance. Furthermore, the strongest effects occur when both firms in the link receive the letter.⁵ Interestingly, however, the combined effect of targeting both firms is not greater than the sum of treating each individually, showing that firm-to-firm communication amplifies single-firm interventions.

The asymmetric responses shed light on the mechanisms underlying VAT evasion. Sellers’ corrections modestly increase their monthly VAT liability but are accompanied by adjustments on other margins: sellers increase their reported sales to other VAT firms while reducing reported final sales, which lack third-party verification. This reclassification potentially explains why sellers are more responsive to the intervention—they can address discrepancies without substantially increasing their overall tax liability. Our findings reveal a novel result: VAT evasion in this context is primarily driven by the underreporting of sales rather than the overreporting of costs, contrasting with existing models of business tax evasion in developing countries (Best et al., 2015; Carrillo et al., 2017; Waseem, 2023), which emphasize cost inflation.⁶

⁵The finding is consistent with Deserranno et al. (2022), which shows that incentives in Sierra Leone’s healthcare sector are most effective when distributed equally between frontline workers and supervisors.

⁶While Waseem (2023) also studies the VAT, Carrillo et al. (2017) and Best et al. (2015) focus on corporate or turnover taxation. Unlike our randomized intervention, these studies rely on quasi-experimental evidence.

Spillover effects extend to untreated links, with corrections in treated firms’ other trading relationships increasing by 2.7pp, or sixfold. Comparing across treatment arms, sending letters to both firms leads to significantly larger spillover effects on untreated links compared to targeting sellers alone. Indeed, if letters are only sent to the seller, the effect is not significantly different from zero, underscoring that notifying both trading partners is crucial for these spillovers to materialize. We find no evidence that the intervention affects the probability or volume of subsequent trade in treated links, indicating that the treatment did not alter the network structure.⁷

Overall, our back-of-the-envelope calculations find that the intervention uncovered \$368,681 in previously unreported business-to-business (B2B) trade and proved cost-effective for the revenue authority due to its low implementation cost. We estimate that the direct effect—the uncovered B2B sales from discrepancies mentioned in the letter—is \$114,986 when letters are sent only to sellers, compared to \$108,154 when letters are sent to both firms. Although this might imply that targeting sellers alone is sufficient, accounting for spillover effects raises the total uncovered B2B sales increase to \$151,452 for letters sent to both firms, versus \$138,139 for letters sent only to sellers. Given the low marginal cost of sending one extra letter, sending letters to both firms is the most effective treatment for promoting compliance.

Our first contribution lies in the creation of a randomized experiment within a network of firms—using transaction-level data—enabling us to distinguish between the direct and (part of) the spillover effects with precision. Understanding direct and spillover effects of an intervention is essential for designing optimal policies. The reflection problem—whereby a firm’s trading partners’ behavior affects the firm’s behavior—makes it hard to distinguish between direct and spillover effects (Manski, 1993; Vazquez-Bare, 2023). This is especially a concern in a setting where the network is dense, something that is typically the case in business transaction networks. While network experimental design has been explored, it often focuses on social or geographic settings (Baird et al., 2018; Vazquez-Bare, 2023; Cruces et al., 2024).⁸ Existing studies document spillovers through tax preparers (Battaglini et al., 2019; Boning et al., 2020), across tax types (Lopez-Luzuriaga and Scartascini, 2019), and among geographically proximate entities (Drago et al., 2020; Lediga et al., 2022).⁹ We instead focus on spillovers through the transaction network. The study closest to ours, Pomeranz (2015), analyzes VAT enforcement spillovers in Chile through an experiment that randomizes

⁷This also suggests that reduced discrepancies result from improved compliance rather than collusion, which would likely reduce the reported trading volume.

⁸Cruces et al. (2024) provide an application to property tax compliance in Argentina (and a methodological contribution to account for heterogeneity across clusters), with a specific focus on geographic spillovers.

⁹Lediga et al. (2022) also explore spillovers within industries and value chains using input-output tables.

treatment at the firm level. While also leveraging (partial) data from the transaction network, given the reflection problem, it is hard in that setting to disentangle the direct from the spillover effects when randomizing at the firm level.

Second, we contribute to the literature on firm networks by highlighting that they are instrumental in amplifying the effect of direct tax enforcement interventions. Previous literature has underscored the importance of transaction networks to understand firm performance (Alfaro-Ureña et al., 2022; Dhyne et al., 2022; Adão et al., 2022; Bernard et al., 2022; Demir et al., 2022). We move this literature forward by documenting direct evidence that firms communicate extensively with their trading partners about actions from the government, such as tax enforcement measures. These spillovers are quantitatively important—increasing the impact of our intervention from \$257,817 in undeclared B2B transaction to \$368,681 (a 43% increase)—and matter for firm behavior: the spillover effects to outside-link trading partners only materialize when the trading partner communicates with the treated firm.

Finally, we contribute to a rapidly expanding literature in public finance on tax enforcement and firms’ evasion behavior by highlighting the potential of leveraging transaction networks for designing enforcement measures. While some interventions have successfully deterred tax evasion (see, e.g., Shimeles et al. (2017) on Ethiopian firms and Holz et al. (2023) on Dominican businesses), growing evidence suggests that enforcement interventions often fail to increase revenue collection, despite the availability of third-part reporting, if firms can easily adjust their declarations on other margins (Carrillo et al., 2017; Slemrod et al., 2017; Mascagni et al., 2021; Best et al., 2021; Hoy et al., 2022). Our results are mixed in this regard: we find a positive, though small, effect on tax liability through amendments of past returns but no significant effect on subsequent liabilities. However, the strong effects on reporting behavior suggest that leveraging firm network information could be a promising strategy in low-capacity settings, especially when complemented with additional monitoring resources (Almunia and Lopez-Rodriguez, 2018; Basri et al., 2021). Indeed, following our experiment, the URA expanded the use of data crosschecks and enforcement messages central to our intervention, underscoring the practical relevance of our findings. si

The rest of the paper is organized as follows. Section 2 provides the context and describes the data. Section 3 outlines the framework and experimental design. Section 4 presents the results, followed by a discussion of their implications and relation to existing evidence in Section 5. Section 6 concludes.

2 Background

This section starts by describing the institutional context of the VAT in Uganda. We then present the data that allows us to map firm networks and observe discrepancies in VAT reporting.

2.1 The VAT in Uganda

Uganda is a low-income country with a per-capita income of \$2,140 in PPP (World Bank, 2021). Its tax-to-GDP ratio—14.4% in 2020/21 (IMF, 2022)—is slightly below the 16% average in Sub-Saharan Africa (OECD/ATAF, 2022) and substantially lower than the 34.1% average in OECD countries (OECD, 2022). The VAT was introduced in 1996 and currently contributes 30% of total tax revenue (IMF, 2022), a figure similar to the average in Sub-Saharan Africa. However, the IMF estimates that the VAT compliance gap—the difference between the potential revenue and actual collections—is large in Uganda, at around 60% of potential VAT and equivalent to 6% of GDP (IMF, 2014).

The Ugandan VAT has a standard design: a general rate of 18% applies to all domestic sales, with the usual exemptions for necessities and some services.¹⁰ Firms with annual turnover above 150 million UGX (\$45,000) are required to register for VAT, while smaller firms can choose to pay a simplified turnover tax.¹¹

VAT-registered firms have to submit monthly VAT returns to the URA, reporting all their sales (on which VAT is due), and their inputs, on which they can claim tax credits if purchased from another VAT firm. The tax base is computed as the difference between sales and creditable purchases. Payments of positive tax liabilities are due within 15 days of filing a return. As in other low-income countries, there are restrictions on VAT refund applications when firms report negative liabilities.¹² Firms may amend their monthly returns at any time after the initial filing, as we explain in more detail below.

Since 2012, all VAT returns are filed electronically. VAT firms are required to submit detailed transaction-level records along with their monthly VAT return, covering all domestic

¹⁰For instance, unprocessed agricultural products and medical, educational and financial services are exempted from VAT. As in other countries, exports are zero-rated, but the VAT applies to imports.

¹¹In our sample, 51.68% of firms are above the VAT registration threshold. Our analysis focuses on VAT-registered firms, and we do not find that the intervention affects the likelihood of staying in the VAT system.

¹²When the stock of negative VAT liabilities is above 5 million UGX (\$1,500), firms can claim a refund, but they have to agree to an audit by the revenue authority. If negative liabilities are less than that amount, they can only be carried over as an offset against future VAT liabilities. The strict regulation of VAT refunds is common practice in other low-income countries (Lemgruber et al., 2015).

sales and purchases with other VAT firms. This system is designed such that the URA receives two reports for each firm-to-firm transaction, one from the seller and one from the buyer.¹³ This dual-reporting mechanism enables us to map firm-to-firm trading networks, detect potential misreporting, and define types of VAT misreporting.

2.2 Data

Our data covers the universe of monthly VAT returns filed by Ugandan firms to the URA in the period between March 2017 and December 2018. Firms' monthly VAT return includes their tax identification number (TIN), the period covered by the return, the filing date, total sales and purchases to/from other VAT-registered firms, total sales to final consumers, total VAT liability, and VAT credit carried over from previous months. As noted above, firms also report their individual transactions with other VAT-registered companies. For each entry, firms report the TIN of the trading partner, the amount and the date of the transaction. Final sales, which include sales to final consumers and to non VAT-registered firms, are reported as a single aggregate figure in each monthly return.

We verify the consistency of the data along several dimensions. First, about 80% of transactions are reported within one month of the transaction date, as required by law, with another 15% reported within two months. Second, the transaction-level data align closely with the monthly summaries: in 97.4% of declarations, total output VAT matches the sum of VAT collected on individual transactions and sales to final consumers, while input VAT matches in 99% of declarations. This consistency confirms that the transaction-level records provide reliable paper trails for firms' VAT declarations and liabilities.

Amendments of past returns. To submit an amendment, the starting point is the initial monthly VAT return. Firms tick a box of the new version of the return indicating that it is an amendment of a previously filed return. All entries of the return can be modified. We refer to *amendments* as the action of filing a new return for a past period, and to *corrections* as the change in the discrepancies which occur due to these amended filings. The return period reflects the date when the transaction was initially filed and thus does not change for amended returns, while the filing date indicates when the amendment was submitted.¹⁴

¹³This reporting requirement is stricter than what is commonly observed in advanced countries, where transaction-level information is only requested during tax audits.

¹⁴At the time of the study, the data extraction process used by the URA eliminated the initial return when an amendment was filed. However, thanks to the frequent data extractions we set up with the administration, we are able to observe both initial and amended returns for our analysis.

Amendments are relatively rare: of all monthly VAT returns filed in 2016, about 10% were amended in the subsequent 12 months.

Seller-buyer links and networks. We define two firms as forming a link in a given month if (i) both firms are VAT-registered, and (ii) at least one of them lists the other as a client or a supplier in the transaction-level entries of its VAT return for that month. In our analysis, we aggregate all transactions data at the link-month level. There are 21,548 unique firms that submit a VAT return for the March to December 2017 period, and 115,856 distinct seller-buyer links. We define a firm’s network as all its direct clients and suppliers. VAT-registered firms have an average (respectively, median) of 10.51 (2.00) unique trading partners over the year, corresponding to 5.38 (resp., 0.00) unique clients and 5.38 (resp., 1.00) unique suppliers.¹⁵

2.3 Discrepancies in VAT Reporting

VAT reporting requirements mandate two reports per transaction—one from the seller and one from the buyer—creating paper trails considered vital to the system’s effectiveness (Agha and Haughton, 1996; Keen and Lockwood, 2010). However, at the time of the study, the URA analyzed transaction-level data only during audits. The process was not automated and was hindered by technological bottlenecks and lack of qualified staff. We document important limitations to VAT performance, potentially attributed in part to this weak administrative capacity.

There are two types of reporting discrepancies that we can identify with the transaction-level information. We denote as “seller shortfall” cases in which the seller reports a lower amount than the buyer. Conversely, “buyer shortfall” cases are those in which the buyer reports a lower amount. We define these concepts more formally in the next section. Cross-checking the amounts reported by sellers and buyers at the monthly level in our study period, we find widespread discrepancies. Specifically, we observe seller shortfall in 41.37% of the link-month observations.¹⁶ More than 92% of these are extensive-margin discrepancies, meaning that only one of the firms in the link reported trading with the other firm.

¹⁵The number of unique clients does not include final consumers or client firms that are not registered in the VAT, because those transactions are not reported in a disaggregated way.

¹⁶We observe buyer shortfall in 51.62% of the link-month observations, and matching amounts in the remaining 6.53% of cases. In earlier work, we find similar rates of discrepancies over the 2013-2019 period and document that these discrepancies lead to a substantial loss in VAT revenue for Uganda (see Almunia et al., 2024, for details). These patterns are not unique to Uganda and have been documented in other low-income countries where similar data exists (Brockmeyer et al., 2024).

While only seller shortfall leads to a lower tax liability (whereas buyer shortfall could even increase it), the prevalence of these discrepancies strongly suggests that VAT evasion is rampant. Furthermore, the distribution of firms’ reported value added (total sales minus purchases) raises concerns, as about 30% of firms report a negative or zero value added amount over the entire fiscal year.

Against this backdrop, the objective of the letter intervention we designed in collaboration with the Uganda Revenue Authority is to reduce VAT evasion, by targeting seller shortfall discrepancies and leveraging the information available about the network of firms’ transactions.

3 Experimental Design and Implementation

In this section, we lay out a simple framework to characterize firms’ tax evasion behavior and how it may be affected by the letter intervention. We then present the two-stage randomization procedure that allows to identify the direct and spillover effects of the intervention, before describing the experiment implementation.

3.1 A Conceptual Framework

We define a link as a seller-buyer pair where at least one party reports trading with the other.¹⁷ Let \hat{s}_{jk} denote the sales *reported* by firm j to their client k , \hat{c}_{kj} denote the purchases *reported* by firm k from their supplier j , and the *true* amounts are denoted by $c_{kj} = s_{jk}$. Formally, a seller shortfall occurs when seller j declares less than buyer k on the same transactions, i.e., $\hat{s}_{jk} < \hat{c}_{kj}$.¹⁸ In general, a seller shortfall reduces the total tax remitted.¹⁹ However, note that a seller shortfall can occur if the seller underreports its sales ($\hat{s}_{jk} < s_{jk} = c_{kj}$), if the buyer overreports its purchases ($\hat{c}_{kj} > c_{kj} = s_{jk}$), or if both happen simultaneously. Thus, when we observe this type of discrepancy, it is not clear *ex ante* which of the two parties of the transaction is liable for misreporting.

In either case, a firm engaging in seller shortfall implicitly assumes that the risk of detection by the tax authority is low, because it will not cross-check its declaration against that of its trading partner, i.e., comparing \hat{s}_{jk} and \hat{c}_{kj} . Our letter intervention aims to increase

¹⁷We ignore the time dimension for the purposes of describing the conceptual framework.

¹⁸Conversely, in the case of a buyer shortfall, the seller declares more than the buyer on the same transactions, i.e., $\hat{s}_{jk} > \hat{c}_{kj}$. There is no discrepancy if $\hat{s}_{jk} = \hat{c}_{kj}$.

¹⁹The discrepancies may arise from either intentional evasion or mistakes in filing VAT returns, as noted in [Almunia et al. \(2024\)](#). In our context, the key point is that these discrepancies ultimately lead to lower tax liabilities.

the perceived risk of detection by informing firms of specific discrepancies (the *listed* discrepancies) identified by the administration and warning of potential penalties if tax evasion is confirmed. This direct exposure to detection is likely to change firms' prior belief of detection risk and influence firms' reporting behavior, leading to our first set of testable predictions:

Set of Predictions #1: Effects on Treated Links.

- (1a) Firms in treated links amend past VAT returns to correct the seller shortfall discrepancies ($\hat{s}_{jk} < \hat{c}_{kj}$) that are *listed* in the letter.
- (1b) Firms in treated links amend past VAT returns to correct other seller shortfall discrepancies that are *not listed* in the letter.
- (1c) Firms in treated links change their behavior and engage in fewer instances of seller shortfall in the monthly VAT returns following the treatment.

Predictions #1a and #1b focus on the contemporary effects of the letter on past returns. Prediction #1a represents a direct effect, as firms amend the listed discrepancies. Prediction #1b extends this impact, suggesting that firms may correct other discrepancies not mentioned in the letter, reflecting an updated belief about the increased risk of detection. Prediction #1c, which refers to a potential change in reporting behavior in subsequent months, indicates a more persistent shift in beliefs.

The predictions above do not distinguish between sellers and buyers in our analysis. The symmetry between underreporting sales and overreporting purchases breaks down once we take final sales into account. While all of a firm's sales are subject to VAT, including the final sales, only purchases from VAT-registered firms count towards the calculation of input tax credits. Critically, declared final sales cannot be verified through cross-checking, unlike transactions between VAT-registered firms. This structural difference creates an asymmetry in firms' reporting incentives and strategies, as final sales offer sellers plausible deniability when evading taxes on sales.

To illustrate this asymmetry, consider a firm that sells to both VAT-registered firms and final consumers, with total sales given by $s_j \equiv \sum_k s_{jk} + s_j^F$, where s_j^F denotes final sales. If the firm misreports by underreporting on sales ($\hat{s}_j < s_j$), and the tax authority only cross-checks a *subset* of its transactions, the firm can claim that some transactions were misclassified as final sales. Then, the firm can relabel those transactions as sales to VAT-registered firms, without altering the overall reported sales and hence tax liability. Because cross-checking is not feasible for \hat{s}_j^F , the authority can only detect evasion if it identifies enough discrepancies

such that the total purchases declared by firm j 's trading partners (indexed by i) exceed its declared total sales, i.e., $\sum_i \hat{c}_{ij} > \hat{s}_j^F + \sum_k \hat{s}_{jk}$. This level of verification demands significant administrative resources, particularly in data analysis and computational capacity, which is typically lacking in low-income settings like Uganda. On the input side, such plausible deniability does not exist: if the declared purchases exceed the reported sales for a given link, i.e., $\hat{c}_{kj} > \hat{s}_{jk}$, it raises immediate red flags, providing a direct indication of potential evasion.

In summary, this difference in the risk of detection means that, holding all else constant, sellers have more incentives to engage in seller shortfall than buyers. Consequently, the letter intervention may prompt asymmetric responses between buyers and sellers, leading to our next set of testable predictions:

Set of Predictions #2: Asymmetric Responses within Treated Links.

(2a) Sellers are more likely to be responsible for seller shortfall instances than buyers, i.e., $\hat{s}_{jk} < s_{jk} = c_{kj}$ is more prevalent than $\hat{c}_{kj} > c_{kj} = s_{jk}$. Hence, sellers are more likely to react to the letter by correcting past returns.

(2b) If the letter is sent to a firm that is not responsible for the seller shortfall discrepancy, they communicate with their trading partner to induce them to make appropriate corrections.

In addition to the effects on treated links, our setup also provides insights into how the information delivered to the firms in a specific link might influence behavior in other trade links. With updated beliefs about the tax authority's ability to cross-check transactions and detect discrepancies, treated firms may adjust their reporting behaviors for untreated links as well. This brings us to the third testable prediction:

Set of Predictions #3: Effects on Untreated Links.

(3a): Treated firms are more likely to correct seller shortfall discrepancies with other trading partners.

Finally, all of these predicted effects at the link level may have an impact on net reported VAT liability at the firm level, which leads to our last set of predictions:

Set of Predictions #4: Firm-level Effects on Net VAT Liability.²⁰

²⁰Here we focus on describing the reporting behavior of those firms treated as sellers in the study sample, since they are the ones expected to react more to the letter treatment. In the empirical analysis, we report the results for buyers on these dimensions as well.

- (4a)** Since the tax administration cannot cross-check final sales s_j^F , sellers react to the letter by relabeling part of s_j^F as s_{jk} to keep their total sales s_j unchanged.
- (4b)** As a result of the correction of past discrepancies (Predictions #1a and #1b), the net VAT liability declared by treated sellers increases, despite some relabeling (prediction 4a).
- (4c)** The improvement in reporting behavior in post-treatment months (Prediction #1c), leads to a net increase in reported VAT liability in periods after the treatment.

In summary, the existence of firm networks implies that an intervention informing firms of specific reporting discrepancies could generate a range of outcomes. Beyond the direct effects of the letter interventions on treated links, including some within-link spillover effects driven by communication between firms, there are also potential spillover effects to other trading links. We summarize the predictions in Table 1, which classifies the effects based on whether they occur within the treated links, outside treated links, or at the firm level.

Table 1
Summary of Predictions on the Effects of the Letter Intervention

Within-Link	<ul style="list-style-type: none"> ◇ Correct listed discrepancies (1a) ◇ Correct unlisted discrepancies (1b) ◇ Reduce future seller shortfalls (1c) ◇ Sellers more likely to correct (2a) ◇ Communication from a non-responsible partner leads to corrections within the link (2b)
Outside-Link	<ul style="list-style-type: none"> ◇ Correct transactions with other trading partners (3a)
Firm-Level	<ul style="list-style-type: none"> ◇ Treated sellers relabel final sales (4a) ◇ Treated sellers' net VAT liability increases due to corrections (4b) ◇ Treated sellers declare higher VAT liability in post-treatment month (4c)

Notes: This table summarizes the predictions laid out in Section 3.1.

3.2 Sample Selection for Identification in Firm Networks

In order to empirically test the predictions in Table 1, we leverage the structure of the firm network to design a novel experiment. The unit of analysis is the seller-buyer link, rather than the individual firm, for two reasons: first, this is the level at which we observe seller shortfall

discrepancies. Second, randomizing at the firm level would potentially violate the stable-unit treatment value assumption (SUTVA): in the presence of spillover effects, more connected firms would be more exposed to a randomly-assigned treatment (Rosenbaum, 2007). This is true even if the randomization is stratified by the baseline number of trading partners.

However, a “naive” randomization at the link level also introduces other issues in experimental design. The main one is that a given firm (seller or buyer) could potentially be assigned to multiple treatment arms through its links with different trading partners. This overlap would preclude identification of causal effects, as the combined impact of different treatments on the same firm could interfere with one another.

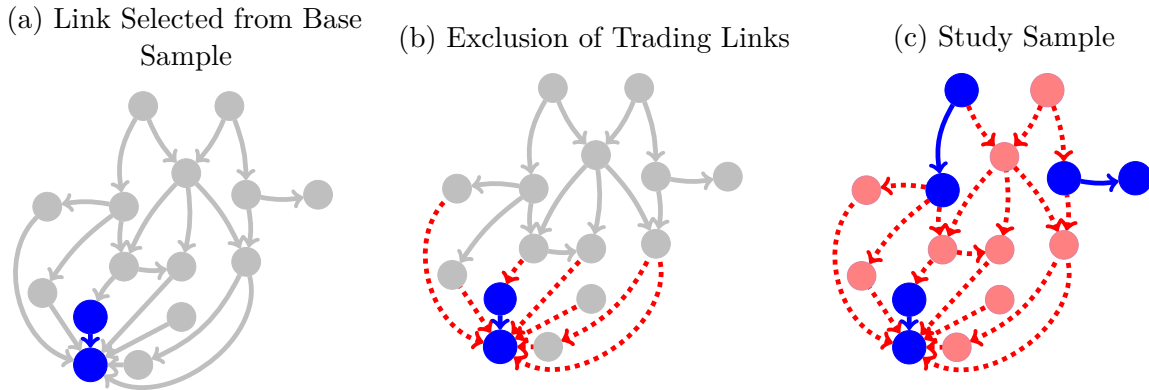
Our sample selection strategy is designed to address this challenge in several steps. We first select a time window of interest, March to December 2017, for which we have a total of 115,856 seller-buyer links formed by all VAT-registered firms. Next, we define a Base Sample of links with reporting discrepancies that are suggestive of potential tax evasion—a behavior that the administration wants to target. More precisely, we select links that meet the following criteria as of January 2018: both firms in the link must be registered for VAT, must have filed a VAT return in the previous month, and the link must have an accumulated seller shortfall with each other greater than 1 million UGX (\$303) over the previous 10 months.²¹ After applying these filters, the Base Sample consists of 11,036 seller-buyer links and 4,514 unique firms.

Next, we draw our Study Sample from the Base Sample, applying an iterative sample selection procedure, as illustrated in Figure 1. Panel (a) represents an example of a firm network of targeted discrepancies: each solid line represents a link between a seller and a buyer in the Base Sample. First, we randomly select a seller-buyer link (in blue). Second, we identify all other relevant links of that seller and buyer with other firms, e.g., trading relations where the same Base Sample selection criteria are met, as shown in red dotted line in Panel (b), and remove these links from the eligible pool of links. Third, we randomly select another link from the remaining sample of eligible links, as shown in Panel (c). Similarly, we remove all links that involve the selected seller and buyer. We repeat the procedure until there are no links left in the Base Sample. Following this iterative approach, we obtain a set of 1,235 seller-buyer links (and hence 2,470 firms) in the Study Sample, as shown in Figure 2.²² Unsurprisingly, firms in the Study Sample are more connected than the typical firm in

²¹Additionally, we remove 2,352 firms that were part of the URA’s annual audit plan for the financial year 2017/18 to avoid interfering with their normal operations.

²²An error in one data extraction, in which all input transactions filed between 19/9/2017 and 15/10/2017 showed up as duplicates in our data, caused a mechanical inflation in the occurrences and amounts of seller

Figure 1
Sample Selection Strategy



Note: This figure illustrates the iterative sample selection procedure for the study. Each dot represents a firm and each arrow a seller-buyer link with potential tax evasion, as defined in Section 3.2. In Panel (a), we start from the Base Sample and randomly select one seller-buyer link. In Panel (b), we remove all other links involving the selected seller and buyer with other firms (red dashed lines). The algorithm is repeated among the remaining eligible links, until no links are left. Panel (c) displays the resulting study sample.

Uganda, with an average of 55.80 distinct trading partners, compared to 10.51 in the full sample.


Through this procedure, we ensure that links in the Study Sample are at least one degree of separation apart within the firm network. In other words, no firm within a given link in the Study Sample has a direct trading relationship that results in a seller shortfall above the threshold with a firm from another link.

Identification assumptions

This sample selection design minimizes direct trading overlaps among firms, allowing us to isolate and identify the direct and spillover effects of the intervention. To ensure the validity of our analysis, we rely on two key assumptions. First, we assume the letter primarily impacts firm links that meet the Base Sample selection criteria, with an important condition being that the seller-buyer links exhibit seller shortfall. This is plausible since the letter directly targets discrepancies related to potential tax evasion within these links. We provide further evidence validating this assumption by showing that the treatment does not significantly

shortfalls between links. As a result, the original Study Sample included some links whose actual seller shortfall did not meet our discrepancy criteria (amounting to 28% of the original Study Sample). We correct this in our analysis by excluding all pairs that did not feature discrepancies after correcting for this error. We discuss the data correction in more detail in Appendix B, and show that it is balanced across treatment arms, indicating that correcting for it is unlikely to bias our results.

Figure 2
Samples and Experimental Design



Note: This figure shows the number of firms and unique seller-buyer links included in the study. The universe of transactions refers to the entire transaction-level data for VAT registered firms between March and December 2017. The base sample refers to the links with potential tax evasion as described in Section 3.2. The study sample refers to the links included in the study after applying the iterative sample selection procedure described in Section 3.2. Finally, the links from the study sample are randomly assigned to the control group or one of the treatment arms. *Source:* Data from monthly VAT returns submitted to the URA.

affect other reporting behaviors, i.e., instances of buyer shortfall or matching reports, in Appendix Table A1.

The second assumption is that spillover effects across links do not differentially affect the control and treated groups. Since all links in the Study Sample are separated by at least one degree in the network, a sufficient condition to validate this assumption is that spillovers do not extend beyond direct trading relationships. In other words, untreated firms connected to treated links should remain unaffected. Appendix Table A2 supports this condition by showing that nearly all corrections are made by treated firms themselves, i.e., those directly involved in treated links, rather than by their untreated partners outside these links. Furthermore, even if second-degree spillovers were present, we find no evidence

that second-degree exposure to treatment—defined as treatment received by a firm’s direct trading partners—differentially affects the reporting behaviors of treated and control firms. This additional evidence, presented in Appendix Table A3, further validates the assumption.

3.3 Treatment Arms

We implement a stratified randomization to allocate seller-buyer links of the Study Sample into three treatment arms and a control group, as shown at the bottom of Figure 2. In the “Buyer only” arm, a letter is sent exclusively to the buyer, while the “Seller only” arm targets the seller, and the “Seller and Buyer” arm targets both. 741 pairs are divided up equally across these three treatment arms, while the control group consists of the 494 remaining pairs (40% of the Study Sample).²³

Our experimental design can be thought of as an extreme version of the two-stage experiments used in the literature to identify spillover effects (Baird et al., 2018; Duflo and Saez, 2003; Hudgens and Halloran, 2008; Vazquez-Bare, 2023). In such experiments, the researchers randomly select clusters and then vary across clusters which individuals and how many of them get treated. In our setting, seller-buyer links are the clusters.

In all arms, we limit the listed discrepancies to three per letter, by selecting the three largest discrepancies of the link within the past ten months. Within each seller-buyer link, we can therefore further examine the effects on unlisted discrepancies and contrast them with the direct effects on listed discrepancies (Predictions #1a and #1b). We apply the same criteria to define counterfactual listed and unlisted discrepancies for the control group.

Having separate “Seller only” and “Buyer only” arms allows us to examine whether the direct effects differ between two sides of the trade and to identify which side is more responsive to the treatment (Predictions #2a and #2b). It also facilitates the detection of spillover effects, where untreated firms within the treated link adjust their reporting when their trading partners receive a letter, thereby revealing how information spreads within firm networks through firm communication (Prediction #3a).

The “Seller and Buyer” arm evaluates whether treating both trading partners simultaneously is more effective and, if so, to what extent. The results are not obvious ex ante. Treating both firms could amplify effects by prompting corrections from both sellers and buyers or by reducing the ability of one firm to shift blame onto the other when both are aware of their treatment. Conversely, joint treatment might encourage collusion, leading to

²³The randomization is stratified by three variables: the ratio of final sales over total sales (above vs. below median), firm size measured by total output VAT (above vs. below median), and the location of the firms’ headquarters (in Kampala, the capital city of Uganda, vs. not).

weaker responses. Furthermore, this treatment arm tests the cost-effectiveness of targeting both firms within a link compared to treating unrelated firms, testing whether the additional impact justifies the extra cost of sending two letters.

3.4 Implementation of the Letter Experiment

After carrying out the sample selection and randomization procedures described in Sections 3.2 and 3.3, we focus on the Study Sample and check that our randomization generates balanced groups. Descriptive statistics are displayed in Table 2, which shows that observable characteristics in each of the treatment arms are not significantly different from those in the control group. Additionally, the attrition rates in treatment and control groups are very similar, alleviating a common concern in RCTs (see Appendix Figure C1).

The letters were sent to a total of 1,002 firms in the treatment arms, notifying them that the URA has developed new analytical methods to detect discrepancies in VAT declarations. A template is displayed in Appendix Figure C2. Each letter lists up to three examples of monthly “seller shortfall” discrepancies between the recipient firm and its partner. The letter instructs the firm to resolve the issue by filing the necessary amendments, and warns them that the URA will carry out similar cross-checks in the future. It also reminds the taxpayer of the fines and prosecution it is exposed to in case of tax evasion. All the letters are officially signed by the Assistant Commissioner for Compliance Management, to ensure that they are reliable and credible in the eyes of taxpayers.

The timeline of the experiment was as follows: between February 28th and April 17th, 2018, physical copies of the letters were delivered to treated firms by a private courier company, using firms’ postal addresses from the URA’s taxpayer register. Because a small percentage of letters (less than 10%) could not be delivered in person, the URA emailed a copy of the letters to all treated firms on April 6th, 2018. During the implementation period and the following months, URA staff kept track of all communications from treated firms with the administration. Table C1 and Figure C3 in the Appendix provide information on how firms reacted to the letters. The implementation was successful: 92% of firms selected for treatment either confirmed reception of the physical letter by a signature, or by contacting the URA.

Table 2
Descriptive Statistics and Randomization Balance

Link characteristics	Mean	Difference with respect to control			
	Control (1)	Any treatment (2)	Buyer only (3)	Seller only (4)	Seller and Buyer (5)
Monthly transactions amount	1.11 (3.35)	-0.03 [0.08]	-0.16 [0.12]	-0.01 [0.11]	0.09 [0.11]
Months with seller shortfall	3.79 (3.06)	-0.12 [0.18]	-0.29 [0.25]	-0.06 [0.25]	-0.02 [0.23]
Share of extensive margin	0.91 (0.29)	0.00 [0.02]	-0.01 [0.02]	-0.01 [0.02]	0.02 [0.03]
Buyer characteristics					
Total input	14.61 (57.19)	-2.62 [5.16]	-1.13 [4.30]	-11.39 [8.72]	3.98 [3.76]
Total output	11.47 (48.39)	-1.63 [3.18]	-1.62 [3.68]	-6.83 [5.18]	3.11 [3.10]
Share of final sales	0.51 (0.44)	-0.02 [0.03]	-0.03 [0.03]	-0.01 [0.04]	-0.02 [0.03]
Audited in 2016	0.06 (0.24)	0.00 [0.01]	0.01 [0.02]	-0.01 [0.02]	0.01 [0.02]
Unique trading partners	46.38 (80.18)	-1.45 [4.68]	-2.84 [5.86]	-0.33 [6.19]	-1.18 [6.70]
Unique clients	18.90 (67.45)	0.11 [3.96]	0.88 [4.72]	-1.11 [5.16]	0.49 [5.87]
Unique suppliers	28.39 (32.79)	-1.78 [1.92]	-4.19 [2.65]	0.59 [2.56]	-1.71 [2.49]
Seller characteristics					
Total Input	27.90 (63.65)	1.64 [3.79]	-0.88 [5.51]	3.06 [4.81]	2.67 [4.77]
Total Output	24.49 (56.52)	2.43 [3.13]	0.48 [4.51]	3.37 [4.24]	3.36 [4.15]
Share of final sales	0.56 (0.41)	0.00 [0.02]	0.03 [0.03]	-0.05 [0.03]	0.01 [0.03]
Audited in 2016	0.04 (0.21)	-0.01 [0.01]	-0.03* [0.02]	0.00 [0.02]	-0.01 [0.02]
Unique trading partners	60.57 (152.93)	-6.37 [9.48]	-5.68 [10.94]	-4.50 [14.02]	-8.72 [12.03]
Unique clients	44.80 (144.30)	-5.24 [9.08]	-2.54 [10.25]	-3.93 [13.51]	-8.92 [11.41]
Unique suppliers	17.37 (23.92)	-1.29 [1.41]	-3.38* [2.02]	-0.61 [1.83]	0.03 [1.77]
Observations	2358	3790	1278	1183	1329
Links	494	741	242	238	261

Notes: This table displays descriptive statistics for firms of the study sample and their balance across treatment arms and the control group. Column 1 reports the mean and standard deviation (in parentheses) for the control group. Columns 2-5 report the differences between treatment arms and the control group, and the standard errors from a t-test of the difference in means (in square brackets). The top panel reports characteristics at the seller-buyer link level. The buyer (respectively seller) characteristics in the second and third panels are at the firm level. We compute monthly transactions amount as the highest amount reported by either trading partner (seller or buyer). Months with seller shortfall is the number of months for which seller shortfall is observed. Share of extensive margin indicates the share of seller shortfall instances where the seller doesn't declare any trade with the buyer. Share of final sales is the ratio of sales to final consumers plus sales to non-VAT firms over total sales. All monetary values are in thousands of US\$. *Source:* Data from monthly VAT returns submitted to the URA between March and December 2017.

4 Results

We analyze the effects of the letter on firms’ behavior in three steps. First, we study how the treatment affects the probability that past discrepancies are corrected. We test our predictions from Section 3.1 by studying both within-link and outside-link effects, as well as direct responses by firms having received the letter and communication-driven spillovers. Second, we analyze the impact of the treatment on subsequent reporting discrepancies in the post-treatment period, which goes until December 2018 (ten months after the letters were sent). This sheds light on longer-term changes in firms’ behavior induced by the treatment. Third, we analyze the firm-level effects of the intervention, focusing on changes in the aggregate amounts reported in tax returns and the resulting impact on net VAT liability.

4.1 Impact on the Correction of Past Reporting Discrepancies

Effect on corrections in treated links


To analyze the effect of the letters on the correction of past discrepancies, we estimate the following regression:

$$Y_{it} = \alpha + \sum_{h=1}^3 \beta_h T_{ih} + \delta_t + \epsilon_{it}, \quad (1)$$

where Y_{it} is a dummy variable indicating whether the discrepancy for month t of seller-buyer link i has been reduced at any point in the ten months after treatment, through corrections observed in firms’ amendments. T_{ih} denotes a set of dummy variables capturing the three mutually exclusive treatment arms, δ_t is a month fixed effect, and ϵ_{it} is the error term. The coefficients of interest are the β_h ’s, which capture the intent-to-treat estimates. Observations are at the link-month level, and standard errors are clustered at the link level.

Panel (a) of Figure 3 reports the results for Prediction #1a on the sample of listed discrepancies. The bars indicate the raw share of discrepancies corrected in each treatment arm by either the buyer or the seller, and the β_h coefficient estimates and standard errors displayed to the right of each bar. As shown in the bottom bar, the share of corrected discrepancies for listed months is 22.3pp higher (significant at the 1% level) for any treated link than in the control group (at 1.6%), which corresponds to a fourteen-fold increase in the correction rate. Analyzing the effect for each treatment arm separately, we find that the effect is strongest (30.3pp) for the “Seller and Buyer” treatment, followed by the “Seller only” treatment (24.3pp), and lastly the “Buyer only” treatment (11.6pp).

Figure 3
 Within-Link Effects of the Letter on the Correction of Past Discrepancies



Note: This figure reports the effect of the letter on the correction of past discrepancies within links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The bars represent the share of seller shortfall discrepancies corrected through amendments. The first bar shows the share for control links, the second for the “Buyer only” treatment group, the third for the “Seller only” treatment group, and the fourth for the “Seller and Buyer” treatment group. The bottom bar reports the share aggregating all three treatment groups. To the right of each bar, we report the $\hat{\beta}$ coefficients from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced. $\hat{\tau}_T$ is the coefficient for the aggregate effect of any treatment. Panel (a) focuses on discrepancies mentioned in the letter (*listed* discrepancies), while Panel (b) focuses on unlisted discrepancies. The sample size is slightly smaller in Panel (b) since we drop links for which all discrepancies observed in the pre-treatment period were listed on the letter. Standard errors are clustered at the link level. * $p < 0.10$; ** $p < .05$; *** $p < .01$. All regression results are shown in Appendix Table D1. *Source:* Data from monthly VAT returns submitted to the URA.

Appendix Table D1 reports the full regression results and tests for the significance of the difference in treatment effects across arms. The effect for the “Buyer only” treatment is significantly smaller than the “Seller and Buyer” treatment and the “Seller only” treatment (at the 1% level), while the coefficients for the difference between latter two are not significantly different.²⁴

In panel (b) of Figure 3, we test Prediction #1b, looking at effects for unlisted discrepancies. The sample is restricted to the 528 seller-buyer links for which there were more than three discrepancies in the pre-treatment period. Despite not being mentioned in the letter, we observe a significant effect on the correction of these discrepancies. The overall effect is a 10.6pp increase in the share of reduced discrepancies (bottom bar), reaching up to 13.8pp in the “Seller and Buyer” group and 14.4pp in the “Seller only” group, compared to 1.1% in the control group.²⁵

These findings confirm that the letter has significant effects on the correction of within-link discrepancies, even those that were not mentioned in the letters, in line with Predictions #1a and #1b.

Asymmetric effects for sellers and buyers

In Figure 4, we break down the results by distinguishing whether it is the seller or the buyer in a treated link who corrects the discrepancies. This enables us to test for asymmetric responses as laid out in Predictions #2a and #2b. In panel (a), the outcome variable is a dummy taking value one if the discrepancy is corrected by the seller while, in panel (b), the outcome is a dummy taking value one if the discrepancy is corrected by the buyer, using the same specification as equation (1). We focus on listed discrepancies.

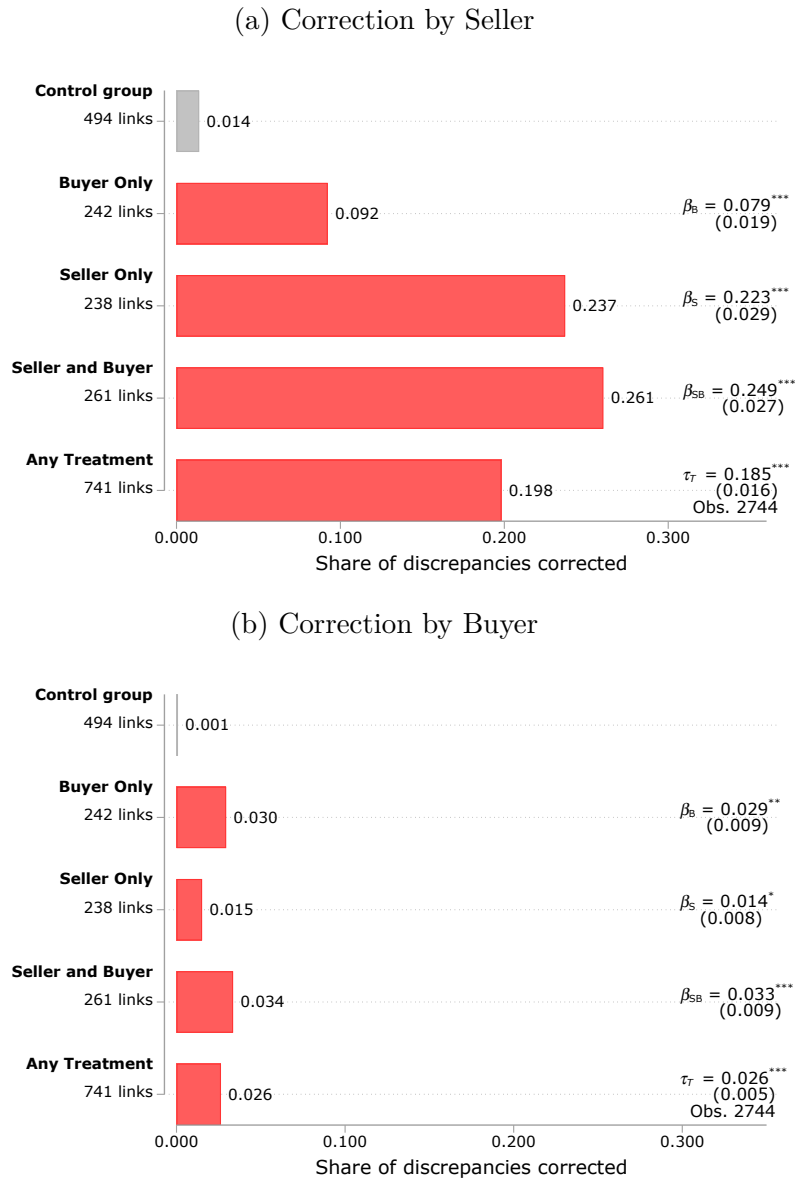
When both seller and buyer are treated, sellers are substantially more likely to react than buyers: the share of reduced discrepancies is 24.9pp larger than in the control group due to seller corrections (panel a), versus 3.3pp due to buyer corrections (panel b).

In the treatment arms in which only one firm receives the letter, we find that the seller is always more likely to correct than the buyer. Sellers are 22.3pp more likely to correct compared to the control group when they receive the letter themselves, corresponding to a

²⁴In an alternative specification, the outcome variable is defined more narrowly and takes value one only if the discrepancy has been fully resolved through the corrections (columns 5-6 of Appendix Table D1). The same pattern of results holds, with slightly muted effects: the correction rate is 14.0pp higher for any treatment, compared to a mean of 0.8% in the control group.

²⁵Columns 7 and 8 of Appendix Table D1 report results for full corrections: the share is 5.2pp higher in the treated group than in control, reaching 7.6pp for the “Seller and Buyer” group, significant at the 5% level.

Figure 4
 Within-Link Direct and Spillover Effects: Who Corrects Discrepancies?



Note: This figure reports the effect of the letter on corrections made by each firm (seller and buyer) within links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The bars represent the share of seller shortfall discrepancies corrected through amendments. The first bar shows the share for control links, the second for the “Buyer only” treatment group, the third for the “Seller only” treatment group, and the fourth for the “Seller and Buyer” treatment group. The bottom bar reports the share aggregating all three treatment groups. To the right of each bar, we report the β coefficients from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced. $\hat{\tau}_T$ is the coefficient for the aggregate effect of any treatment. Panel (a) focuses on discrepancies corrections made by the seller, while Panel (b) focuses on corrections made by the buyer. Standard errors are clustered at the link level. * $p < 0.10$; ** $p < .05$; *** $p < .01$. All regression results are shown in Appendix Table D2. *Source:* Data from monthly VAT returns submitted to the URA.

sixteenfold increase in corrections. Remarkably, sellers are also 7.9pp more likely to correct when the buyer—their trading partner—receives the letter (a fivefold increase in the correction rate). In contrast, the effects on the correction rates of buyers are small regardless of whether the letter is sent to them (2.9pp), or to their seller, (1.4pp).²⁶

In general, while the effects are strongest when we treat both firms in the link, the magnitude of the effect is always lower than the sum of the effects for the “Seller only” and “Buyer only” arms. This novel result highlights that, while treating both firms in a pair simultaneously generates stronger responses, it may not be more cost effective than sending letters to only one firm within the pair.

In line with Prediction #2a, these results reveal that sellers react significantly more strongly to the treatment than buyers, suggesting that sellers are more likely to be responsible for the discrepancies. This is an important finding because it is not obvious *ex ante* that seller shortfall discrepancies are due to sellers underreporting (rather than buyers overreporting). Moreover, this asymmetry has not been documented in the prior literature.

Furthermore, the strong effect of the “Buyer only” treatment on sellers’ behavior indicates that firms engage in communication with their trading partners upon receiving the letters, subsequently prompting these partners to submit amended returns (Prediction #2b). This provides unique evidence of the flow of information between firms, a dynamic that has been challenging to capture in empirical research. Moreover, this illustrates the ability of firms to influence their trading partners’ tax reporting behavior (communication-driven spillovers), even in ways that may not align with the partners’ immediate interests.


Effects on corrections outside treated links

In Figure 5, we test for outside-link effects between treated firms and their other partners within the Base Sample, as laid out in Prediction #3a. We focus on sellers, since our results from Figure 4 show that the effects of the letters are primarily driven by sellers. The sample consists of seller shortfall discrepancies of the sellers from the Study Sample with all their buyers, excluding the buyer from the Study Sample link. We further restrict the sample to months matching the discrepancies listed on the letters. This leads to 11,530 link-month observations between March and December 2017.²⁷ The outcome variable is a dummy variable taking value one when the discrepancy is reduced by the seller.

²⁶Full regression results are reported in Appendix Table D2.

²⁷We show regression results for all months in Appendix Table D3, where the sample includes 12,783 observations. In all cases, we exclude observations with buyers who have a tax ID number (TIN) but are not

Figure 5
Outside-Link Effects of the Letter on the Correction of Past Discrepancies



Note: This figure reports the effect of the letter on corrections outside the seller-buyer links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. To identify outside-link effects, the sample includes all seller shortfall discrepancies of the sellers from the study sample with all their buyers, excluding the buyer from the study sample link. We focus on corrections made by sellers, for months listed on the letter. The bars represent the share of seller shortfall discrepancies corrected through amendments. The first bar shows the share when the seller is in the control group, the second when the seller is in the “Buyer only” treatment group, the third when the seller is in the “Seller only” treatment group, and the fourth when the seller is in the “Seller and Buyer” treatment group. The bottom bar reports the share aggregating all three treatment groups. To the right of each bar, we report the $\hat{\beta}$ coefficients from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced by the seller. $\hat{\tau}_T$ is the coefficient for the aggregate effect of any treatment. Standard errors are clustered at the seller level. * $p < 0.10$; ** $p < .05$; *** $p < .01$. All regression results are shown in Appendix Table D4. *Source:* Data from monthly VAT returns submitted to the URA.

We find a 2.7pp (sixfold) increase in the outside-link correction rate, significant at the 5% level, when considering all treatment arms together. The effects are driven by the “Seller and Buyer” treatment which leads to a 7.2pp increase in corrections (significant at the 10% level), while the coefficients for “Buyer only” and “Seller only” are not significant. The resulting correction rate is still low (7.6%) since these corrections are extremely rare in the control group (0.4%). The full regression results are shown in Appendix Table D4.

While there is no significant difference between sending the letter only to the seller or to both firms in terms of direct effects (p-value = 0.149, see Appendix Table D1), sending letters to both firms leads to significantly larger spillover effects on untreated links compared to targeting sellers alone (p-value = 0.078, see Appendix Table D4). In other words, it is crucial that both trading partners receive the letter for the spillover effects on untreated links

registered for VAT.

to materialize. Consequently, accounting for these spillovers may change our assessment of the most effective treatment.

In columns 3 and 4 of Appendix Table D4, we show further the results for unlisted months: the effects are of similar magnitudes (a 3.2pp increase), and only significant at the 10% level, and when aggregating all treatment arms together. The effects are also less precisely estimated when including all months (Appendix Table D3), the 1.0pp increase being only significant at the 10% level and when considering full corrections.²⁸ Nonetheless, taken together, these results on outside-link effects show that sellers anticipate the revenue authority’s cross-checking capacity to extend beyond trade with the particular partner indicated in the letter. This results in the propagation of belief updating within firms’ networks.

4.2 Impact on Post-treatment Reporting Behavior

Prediction #1c states that if the letter results in updated beliefs about the administration’s enforcement capacity, there should be changes in firms’ reporting behavior in the months subsequent to the treatment. We examine whether the intervention leads to a sustained improvement in reporting behavior beyond the correction of past discrepancies.

Effect on reporting in treated links

To study the tax reporting behavior of firm pairs after the intervention, we run the following event-study specification:

$$Y_{it} = \sum_{\substack{j=-10 \\ j \neq -1}}^{10} \beta_j(m_j \cdot D_i) + \delta_t + \gamma_i + \epsilon_{it} \quad (2)$$

where Y_{it} is the outcome variable of interest for link i in month t , D_i is an indicator for Any treatment for link i , m_j is a dummy for month j (defined relative to the treatment period, $j = 0$), and parameters δ_t and γ_i represent month and link fixed effects, respectively. Our event-study coefficients of interest are β_j ’s.²⁹ The analysis period runs between April 2017 and December 2018, with January 2018 taken as the reference month.³⁰ The sample includes

²⁸In Appendix Table D5, we report results when considering the same sellers acting as buyers in other trading links. We find a very small (0.1pp or 33%) reduction in the correction rates significant at the 10% level, driven exclusively by the “Seller and Buyer” treatment. In Tables D6 and D7 we report results when considering the buyers of the study sample acting as sellers, there is no detectable effect of the intervention.

²⁹We find no evidence of differential attrition, as shown in Appendix Figure C1.

³⁰The letters started being delivered in the last week of February 2018. Hence, the first tax return affected by the intervention is the one referring to February 2018, which firms filed between March 1st and 15th.

all monthly trades reported by a given seller-buyer link. Additionally, we implement an alternative difference-in-differences specification in which we pool together all post-treatment months.³¹

The first set of results is reported in Figure 6, which displays the event-study results for reporting discrepancies. The outcome variable is, in turn, a dummy for whether the link displays seller shortfall, buyer shortfall, or matching reports. Observations are conditional on there being trade between the two firms, so the underlying sample is an unbalanced panel. In addition to the event-study coefficients β_j , we report the difference-in-differences coefficients at the bottom of the Figure.

We find that the intervention leads to a sustained reduction of 15.0pp (19%) in the probability of reporting seller shortfall with the trading partner mentioned in the letter in the ten months following the intervention, in line with Prediction #1c.³² This decline is accompanied by a 13.0pp or 115% increase in the probability of matching reports (no discrepancy). There is a small increase in the buyer shortfall rate, but it is not statistically significant. Overall, the results suggest that the intervention leads to sustained improved reporting behavior within treated links. Furthermore, the seller shortfall amount (conditional on having seller shortfall) does not appear to change significantly, as shown in Figure 7a. This suggests that the change in reporting behavior is mainly in the number of transactions reported, rather than in the amounts reported.³³

One potential concern is that this improved reporting behavior may mask collusion between the firms, if the seller and buyer agree to report lower amounts following the intervention. We provide evidence suggesting that this is not the case: in Figure 7b, we show that the intervention does not change transaction size as reported by the seller-buyer link.³⁴ In Figure 7c, we show further that there is no impact on the likelihood that the treated links continue to report trading with each other after the intervention. This suggests that the intervention

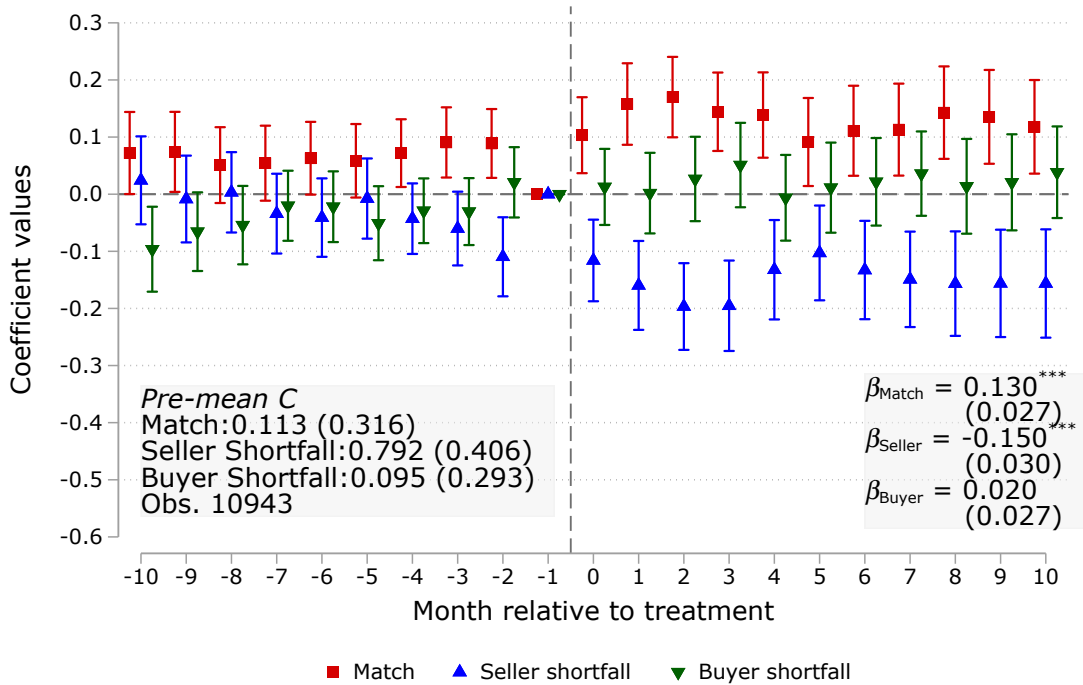
³¹Our difference-in-differences specification is: $Y_{it} = \beta (D_i \cdot \mathbb{1}_{j>0}) + \delta_t + \gamma_i + \sum_{j=-10}^{-1} \beta_j (m_j \cdot D_i) + \epsilon_{it}$, where D_i is an indicator taking value one for treated links. This is interacted with $\mathbb{1}_{j>0}$, a “post” dummy taking value one for all post-treatment months. The coefficient estimate for β reveals the cumulative treatment effect for the entire post-treatment period between March and December 2018. Following Roth and Sant’Anna (2023), we include a separate control for each of the pre-treatment periods.

³²While the last pre-treatment coefficient on seller-shortfall is statistically different from zero, Appendix Figure E1 shows that this is due to a change in the share of seller shortfall in the control group, rather than the treatment group. Hence, we rule out the possibility that this is due to anticipation effects.

³³More than 90% of the seller shortfall discrepancies in the Study Sample are on the extensive margin, with the seller not reporting the transaction at all. Hence, the intervention mostly affects the most frequent type of misreporting.

³⁴Transaction size is defined as the maximum amount reported by the buyer or the seller.

Figure 6
 Post-Treatment Effect of the Letter on Within-Link Discrepancies




Note: This figure reports the effect of the letter on subsequent reporting discrepancies. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). The figure plots $\hat{\beta}_j$ coefficients estimated in the event-study laid out in equation (2), with three different outcomes: probability of there being no discrepancy (“Match”) in red; probability of seller shortfall in blue; and probability of buyer shortfall in green. Outcomes are conditional on trade occurring within the link. In the bottom right corner, we report the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months. Standard errors are clustered at the link level and the bars report 95% confidence intervals. * $p < 0.10$; ** $p < .05$; *** $p < .01$. See Appendix Tables E1 and E2 for the full regression results. Appendix Figure E1 shows the raw proportion of seller shortfall occurrences in treatment and control groups. *Source:* Data from monthly VAT returns submitted to the URA.

did not disrupt trading relationships in a significant way, nor push firms to start trading “under the radar.”³⁵

Similar to what we observe for the effects on the correction of past VAT returns, we find asymmetric effects of the treatment on future reporting too. In Figure 7d, we show that, conditional on trade being reported within the link, there is a 17.6pp (62%) higher probability of it being reported by the seller, with no significant change for the buyer. This shows that the effects on changes in subsequent reporting behavior are mostly driven by sellers as well.

³⁵The regression coefficients underlying panels (a)-(c) of Figure 7 are reported in Appendix Table E1.

Figure 7
Post-Treatment Effect of the Letter on Within-Link Trade



Note: This figure reports the effect of the letter on subsequent trade. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). The figure plots $\hat{\beta}_j$ coefficients estimated in the event-study laid out in equation (2), with four different outcomes: seller shortfall amount conditional on there being seller shortfall, in Panel (a); transaction size conditional on trade in Panel (b), defined as the maximum amount reported by the buyer or seller; probability of any trade being reported in Panel (c); conditional on trade, whether it is reported by the seller and/or the buyer, in Panel (d). In the bottom right corner, we report the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months. Standard errors are clustered at the link level and the bars report 95% confidence intervals. * $p < 0.10$; ** $p < .05$; *** $p < .01$. All amounts are in thousands of USD and winsorized at the 0.5% level. See Appendix Table E1 for full regression results. *Source:* Data from monthly VAT returns submitted to the URA.

Effect on reporting outside treated links

To study whether the intervention affects the subsequent behavior of firms with other trading partners, we perform a similar event-study analysis but focusing on the sample of links formed by the sellers in our Study Sample with all their buyers, except for the one from the Study Sample link. Results are shown in Appendix Figure E2. We do not find significant effects on these outside-link reporting discrepancies subsequent to the treatment.³⁶

³⁶See Appendix Table E4 for detailed regression results and Appendix Table E5 for the diff-in-diff coefficients.

4.3 Firm-Level Effects

In this section, we consider the effects of the intervention on outcomes at the firm-level, instead of link-level. We focus on discussing the results for sellers, since they are shown to be the ones driving the observed responses. The corresponding results for buyers are displayed in Appendix F.

Effect of corrections of past discrepancies on VAT liability

To analyze how corrections of past discrepancies change the amounts reported in the monthly VAT returns, we estimate the following equation:

$$Y_{jt} = \alpha + \sum_{h=1}^3 \beta_h T_{jh} + \delta_t + \epsilon_{jt}, \quad (3)$$

where Y_{jt} denotes firm-level outcomes for seller j in month t , T_{jh} denotes a set of dummy variables capturing the three mutually exclusive treatments j arms, δ_t is a month fixed effect, and ϵ_{jt} is the error term.

Table 3 shows the results for four outcomes after the amended monthly returns are submitted: changes in the sales to other VAT-registered firms (labelled “B2B sales”), in final sales, in taxable inputs, and in the overall VAT liability.³⁷ The sample is restricted to sellers’ outcomes in the months listed in the letters.³⁸ All amounts are in thousands of US dollars (\$), and we winsorize all outcomes at 0.5% level.³⁹

In columns 1-2, we look at the change in reported B2B sales after amendments are filed. We find that all treatments combined lead to an increase of \$512 in reported monthly B2B sales (col. 1). Analyzing each treatment arm separately, we find that the “Seller and Buyer” treatment has the largest effect (\$689), consistent with the previous results on corrections. In columns 3-4, we examine the impact on reported final sales and find a decrease of \$309 (\$446 in the “Seller and Buyer” group). In columns 5-6, we find no significant impact on reported taxable inputs. These coefficients indicate that amendments increase the amount of B2B sales reported, but about 60% of this increase is offset by a reduction in reported final sales.⁴⁰ Finally, in columns 7-8 we estimate the overall effect on the VAT liability reported in

³⁷To calculate the VAT liability for each monthly return, we calculate the difference between output VAT charged and input VAT paid. We do not consider the potential application of tax credits carried forward from past negative liabilities.

³⁸In Table F1, we run the same analysis for the buyers in our Study Sample.

³⁹Our results are similar when we winsorize at different levels, namely 1%.

⁴⁰The 60% is calculated as follows: $309/512 \approx 446/689 \approx 0.6$.

Table 3
Effect of Corrections on Firm-Level VAT Liability (Sellers)

Dependent variable:	Δ B2B Sales		Δ Final Sales		Δ Taxable Inputs		Δ VAT Liability	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	0.512*** (0.059)		-0.309*** (0.037)		0.002 (0.003)		0.013*** (0.004)	
Buyer only		0.151*** (0.057)		-0.081*** (0.029)		0.004 (0.005)		0.003 (0.005)
Seller only		0.684*** (0.116)		-0.390*** (0.073)		0.002 (0.004)		0.019** (0.008)
Buyer and Seller		0.689*** (0.111)		-0.446*** (0.075)		0.000 (0.005)		0.016*** (0.006)
R-squared	0.036	0.056	0.035	0.057	0.007	0.007	0.010	0.013
Observations	2735	2735	2735	2735	2735	2735	2735	2735
No. of Firms	1233	1233	1233	1233	1233	1233	1233	1233
Mean of Dep. in Control	44.384	44.384	134.837	134.837	154.857	154.857	4.386	4.386
Mean of Diff. in Control	0.018	0.018	-0.003	-0.003	0.001	0.001	0.006	0.006
Median of Dep. in Control	1.661	1.661	11.656	11.656	27.837	27.837	0.255	0.255
P-value of $\beta_S = \beta_B$		0.000		0.000		0.718		0.060
P-value of $\beta_{SB} = \beta_B$		0.000		0.000		0.469		0.054
P-value of $\beta_{SB} = \beta_S$		0.978		0.594		0.640		0.753
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effects of the letters on past VAT liability for sellers, estimated using equation (3). The sample includes VAT returns of the sellers for all months mentioned in the letters. The outcome variables are defined as the change in a given entry of the VAT return before and after the treatment, where the changes occur through amendments. Columns 1-2 report results for *B2B Sales* (sales to other VAT firms), columns 3-4 for *Final Sales* (sales to final consumers or non-VAT firms), columns 5-6 for *Taxable Inputs* (purchases from VAT firms) and columns 7-8 for *VAT Liability* (total output tax minus total input tax). *Mean* (resp. *Median*) of *Dep. in Control* reports the average (resp. median) value of each entry for the firms in the control group. *Mean of Diff. in Control* reports the average differences in the values before and after treatment for the firms in the control group. Standard errors are clustered at the seller level. *p < 0.10; **p < .05; ***p < .01. Appendix table F1 reports the equivalent results for buyers. *Source:* Data from monthly VAT returns submitted to the URA.

the tax returns referring to the ten pre-treatment months. VAT liability increases slightly by \$13 for all treatments combined, which is about 5% of the monthly VAT liability of the median firm in our Study Sample. Correspondingly, the results are \$19 for the “Seller only” group and \$16 for the “Seller and Buyer” group.

Our findings indicate that while firms amend past returns to report higher B2B sales, they partially offset this increase by reducing their reported final sales, consistent with Prediction #4a. As a result, the overall rise in reported tax liability for the months mentioned in the letter is smaller than it would have been if firms had only corrected the listed discrepancies. Nonetheless, despite the modest magnitudes, the intervention successfully increased VAT liability for the targeted periods, in line with Prediction #4b.⁴¹

⁴¹These small but positive effects on tax liability are similar to those observed by Carrillo et al. (2017), where firms reacted to a similar tax compliance intervention by adjusting other margins of reporting in order to leave total tax liability almost unchanged.

Effect on subsequent VAT liability

Next, we analyze the impact of the letters on reported VAT liability in post-treatment months. We rely on a specification similar to equation (2), and additionally the difference-in-differences specification, except that here, observations are at the firm level. Results are shown in Figure 8. Panels (a), (b), and (c) show that there are no significant effects in the reporting of B2B sales, final sales, and total inputs after the treatment. Figure 8d shows that most of the monthly event-study coefficients for the net change in VAT liability after the treatment are positive, but none of them is significantly different from zero. Using the diff-in-diff specification, we obtain an increase of \$628 in monthly VAT liability, which is not statistically significant either. Despite the lack of statistical significance, this coefficient is economically relevant, amounting to 14% of the average monthly VAT liability.

In summary, although the intervention led to higher tax liabilities for amended past returns (Table 3) and triggered sustained changes in firm behavior in subsequent months, our study lacks sufficient statistical power to precisely estimate an increase in reported VAT liability.

Heterogeneity analysis

To better understand which firms are driving our results, we estimate the average characteristics of compliers—firm pairs who are induced into reducing at least one discrepancy—following the methodology of Pinotti (2017). The analysis is run at the seller-buyer link level, with pre-treatment baseline characteristics drawn from the seller within each pair. To estimate the average characteristics of the compliers we run a two-stage least squares regression whose first and second stage are given by the following:

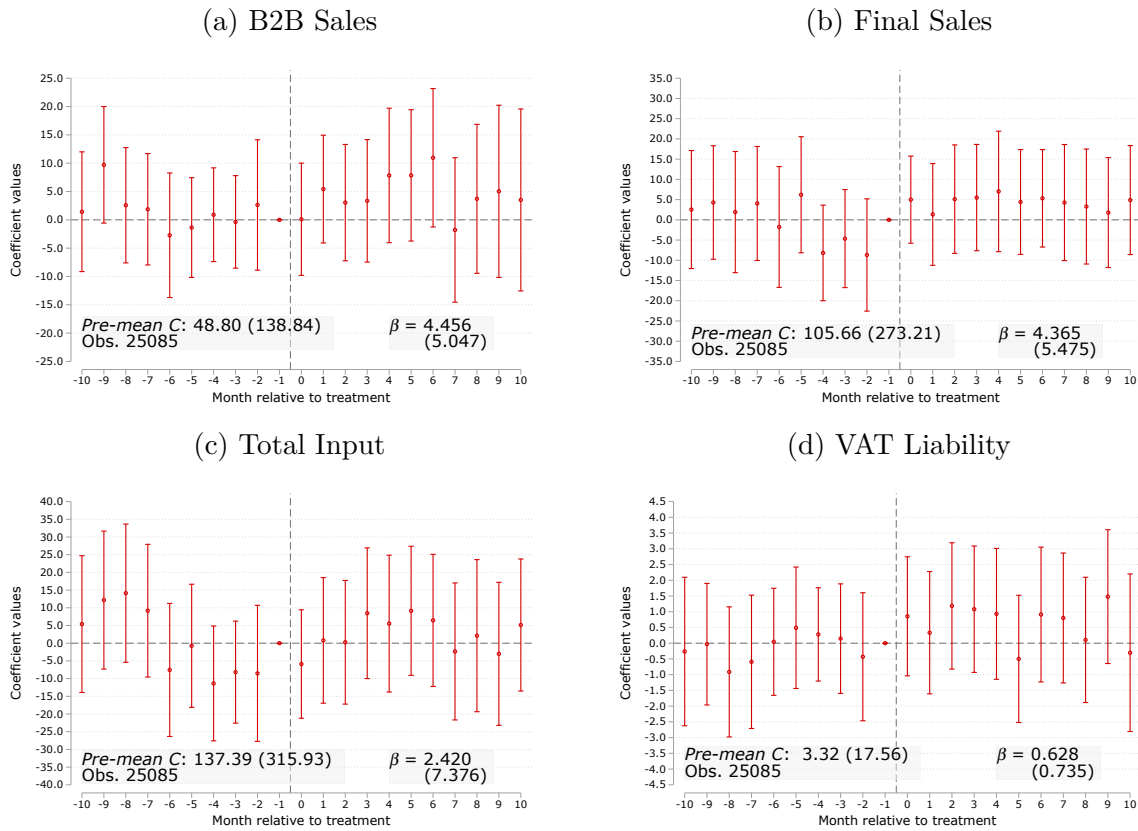
$$h_{i(j)} = \lambda T_{i(j)} + \epsilon_i; \tag{4}$$

$$h_{i(j)} \times k_j = \theta T_{i(j)} + \nu_i, \tag{5}$$

where $h_{i(j)}$ is an indicator variable equal to one if firm pair i that seller j is a part of corrected at least one discrepancy, k_j is the characteristic of interest for seller j , and $T_{i(j)}$ indicates whether pair i is treated (received a letter). The seller characteristics of the complying firm pairs are given by θ . Intuitively, θ captures the average baseline value of the characteristics of interest for the set of sellers who are part of firm pairs induced into reducing at least one discrepancy. Again, we focus on the characteristics of sellers in each pair.

Table 4 reports the average baseline characteristics of the compliers, in addition to the

Figure 8
Post-treatment Effects on Firm-Level VAT Liability (Sellers)



Note: This figure reports the effect of the letter on subsequent VAT liability of sellers. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-month level (April 2017 to December 2018). The figure plots $\hat{\beta}_j$ coefficients estimated in an event-study similar to equation (2), but at the firm level, with four different outcomes: amount of B2B sales (sales to other VAT firms) in Panel (a), final sales (sales to final consumers or non-VAT firms) in Panel (b), total inputs in Panel (c), and VAT liability (total output tax minus total input tax) in Panel (d). In the bottom right corner, we report the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months. All amounts are in thousands of USD and winsorized at the 0.5% level. Standard errors are clustered at the link level and the bars report 95% confidence intervals. * $p < 0.10$; ** $p < .05$; *** $p < .01$. See Appendix Table F2 for full regression results. *Source:* Data from monthly VAT returns submitted to the URA.

sample averages and the p-values for the difference between the complier means and sample averages. Compliers have a significantly higher share of sales to final consumers (at 61.8% vs. 56.5% in the whole sample, with the difference being significant at 10% level), consistent with our conceptual framework: sellers responding to the treatment by reducing discrepancies tend to have a larger share of final sales that they can relabel as B2B sales to limit increases in VAT liability. Compliers are also less connected to other firms (significant at the 10% level). We further examine whether compliers report higher total sales, have a greater share of tax returns with negative VAT liability, are older, or were audited the year before the intervention; none of these appear significant. Finally, we check for differences across a full

set of sector indicators and find no significant differences.

Table 4
Heterogeneity Analysis: Baseline Characteristics of the Compliers

	All Months		
	Letter	P-value	Sample
Ln(Sales)	18.733	0.931	18.720
Final sales ratio	0.618	0.088	0.564
Share of negative returns	0.404	0.299	0.434
Existence length	90.865	0.254	97.315
Degree	44.106	0.070	53.883
Audited in 2016	0.041	0.410	0.053
Sector: Agriculture	0.016	0.672	0.012
Sector: Construction	0.070	0.789	0.065
Sector: Manufacturing	0.111	0.565	0.125
Sector: Retail	0.187	0.624	0.172
Sector: Service	0.319	0.210	0.274
Sector: Wholesale	0.298	0.134	0.349

Notes: This table reports heterogeneity results estimated in equation (5). We run regressions where the outcome is an indicator equal to 1 if there is a correction for any discrepancy of a given link, interacted in turn with characteristics of the seller. The seller characteristics are computed using VAT returns from 10-months prior to the intervention. The characteristics are (i) the natural logarithm of sales, (ii) the ratio of final sales to total sales, (iii) the share of monthly VAT returns with negative Value-Added, (iv) number of months since firm exists, (v) the seller's degree of connection within the network computed as the number of unique partners, (vi) whether the firm was audited in 2016, (vii) a categorical variable indicating the sector (Agriculture, Construction, Retail, Service, and Wholesale). The column "Letter" reports the θ coefficient on the indicator for receiving the letter, from equation (5). The "P-value" column reports results of a significance test of the difference in the mean value of each characteristic between the compliers and the whole sample. The sample average of each characteristic is reported in the last column. *Source:* Data from monthly VAT returns submitted to the URA.

We next investigate whether specific firm characteristics predict the likelihood of correcting discrepancies in response to the letter. In columns 1-3 of Table F3, we estimate an OLS regression where the outcome variable—whether the seller in a treated link reduced a discrepancy—is regressed on a large set of baseline seller, buyer, and link characteristics. Sellers in the service sector and those that received the letter are more likely to respond. Older firms with a lower share of returns with negative VAT liability also respond more strongly, possibly indicating an effort to appear compliant to the URA. However, the variables overall have limited predictive power, with the R-squared for the OLS regression at 0.09. In column 4, we use a LASSO estimator to select the baseline characteristics that are most predictive of the pair reducing at least one discrepancy. Despite including all covariates from columns 1-3, only one variable—whether the seller receives a letter—is picked by the LASSO estimator. These findings suggest a limited heterogeneity in firms' responses to the letter

based on observable characteristics.

5 Discussion

The results presented in the previous section show that the intervention led to a substantial increase in the correction of past discrepancies, demonstrating that the letters effectively updated firms' beliefs about the revenue authority's ability to detect reporting discrepancies. We have documented spillover effects both within treated links, where firms correct discrepancies not listed in the letter, and across untreated links, where firms improve their reporting behavior in other trading relationships. These effects are strongest when both firms in a treated link receive the letter, although sellers are the primary drivers of corrections. Sellers respond not only when directly targeted, but also when their buyers receive the letter, emphasizing the role of firm-to-firm communication in shaping compliance. The intervention also reveals that VAT evasion is predominantly driven by sellers underreporting sales, with adjustments often achieved through reclassification of final sales.

These findings underscore the potential of leveraging firm network data to improve tax compliance, particularly by targeting sellers in enforcement interventions. However, a substantial proportion of treated firms (more than 70%) do not file any amendments in response to the intervention, suggesting that the threat of a letter from the revenue authority is not strong enough in this context for the majority of taxpayers. In this sense, our results resonate with the findings of [Carrillo et al. \(2017\)](#), despite the very different setting of their study, which takes place in a middle-income country (Ecuador) and refers to corporate income tax rather than VAT. Our results are also consistent with those of [Hoy et al. \(2022\)](#).

Mechanisms

In terms of mechanisms, the results indicate that the amending firms increase their reported B2B sales, but this is partially offset by a decrease in their reported final sales. We highlight that this is a different mechanism from the one documented in [Carrillo et al. \(2017\)](#), where firms compensate for additional reported sales by increasing reported input costs to leave tax liability unchanged. In our setting, we do not find any significant increases in reported inputs. Instead, firms manipulate their reported sales by relabeling some sales transactions to attain the same goal. This underscores the limitation of VAT's self-enforcing property: Even though the reporting rules of the VAT provide additional tools to enforce compliance, there are still margins of evasion for firms that report both B2B sales and final sales.

The asymmetric responses of sellers and buyers are consistent with the findings of Pomeranz (2015), providing evidence that enforcement interventions focusing on VAT tend to affect upstream trading partners (suppliers) more strongly than downstream ones (clients). However, when it comes to spillovers, our results differ in two key ways. First, while Pomeranz’s framework predicts spillovers only in cases of collusive evasion—where firms jointly misreport and thus cannot be detected with simple cross-checking—we find evidence of spillovers even in cases of unilateral evasion, such as seller shortfall, where a single firm misreports and discrepancies can be detected through cross-checking. In fact, we find no evidence suggesting that the intervention leads to increased collusive evasion, even when both firms within the treated links receive the letter. Second, unlike Pomeranz (2015), who finds significant spillovers only in suppliers, we observe spillovers in both directions, affecting both upstream suppliers and downstream clients, as shown in Figure 5.

Cost effectiveness

Despite the modest size of the increase in VAT liability due to the amendments, the intervention was cost effective. Summing over all the corrections, the intervention led to a net increase of \$368,681 in B2B trade reported by firms in the Study Sample.⁴² Once we take into account that some firms relabeled originally reported final sales as B2B sales, we estimate that the revenue gain from the amendments of transactions listed in the letters is \$35,555.⁴³ The total cost of sending the letters—including hiring a courier company and the time spent by URA officers on the project—was \$5,716 (see Appendix G for details on the calculation). Thus, the additional revenue obtained was more than six times higher than the cost of the intervention. Therefore, we believe that a scale-up of this type of intervention could potentially raise a significant amount of additional revenue. This is especially the case because there is a limit to the reclassification strategy used by firms: if the amendments that increase reported B2B sales are always offset by a reduction in reported final sales, at some point the latter will be too small. If the revenue authority identifies all reporting discrepancies and notifies firms about them systematically, firms will be forced to file amendments in which they cannot apply that reclassification strategy.

Comparing treatment arms, we estimate that the direct effect—the uncovered B2B sales from discrepancies mentioned in the letter—is \$114,986 when letters are sent only to sellers,

⁴²There were a total of 1249 corrections in the treated group revealing \$441,792 of trade, compared to 137 corrections adding up to \$73,111 of additional reported trade in the control group.

⁴³This results from multiplying the change in tax liability estimated in column 7 of Table 3 times the number of firm-month observations: $\$13 \times 2,735 = \$35,555$.

compared to \$108,154 when letters are sent to both firms. On the basis of direct effects alone, one might conclude that targeting sellers is sufficient. However, when spillover effects are included, the total uncovered B2B sales rise to \$151,452 for letters sent to both firms, compared to \$138,139 for letters sent only to sellers. This evidence indicates that sending letters to both firms may be the most effective treatment to promote compliance.

Furthermore, the results on the propagation of corrections in the firm network imply that the aggregate effects of such interventions depend heavily on the number of buyers connected to treated firms. As an illustration, we find that the seller making the most corrections sells to 353 buyers and makes 295 corrections, uncovering \$5,800 previously unreported B2B sales, while some sellers make no corrections at all. An avenue for future research could be to explore how accounting for such heterogeneity could enhance cost-effectiveness.

However, the intervention we study here was less successful in changing the subsequent tax-reporting behavior of firms. Although we observe a decrease in the share of seller-shortfall discrepancies *within* treated links, this effect does not extend to treated firms' behavior with *other* trading partners. Hence, it is unsurprising that the overall effect on subsequent reported VAT liability is insignificant. One possible explanation for the limited longer-term impact is the lack of follow-up intervention to reinforce firms' perception of the new cross-checking system. Future research could examine the effectiveness of such follow-up measures in sustaining compliance gains in different contexts.

Policy implications

Taken together, our results suggest that interventions targeting sellers' reporting practices and exploiting network-based strategies have the potential to improve tax compliance.

The contrast between treating both the buyer and the seller in a link versus treating only a single firm provides a nuanced perspective on how information on firm networks can be leveraged to improve both compliance and the cost-effectiveness of enforcement strategies. Although the combined effects are stronger when both firms in a link are treated, the magnitude of the treatment effects is smaller than the sum of the effects observed in the seller-only and buyer-only treatments. This outcome is consistent with our results on communication-driven spillovers: letters sent to one firm often induce changes in the behavior of their trading partners, even when those partners do not directly receive a letter. These findings offer valuable guidance for future enforcement strategies. Revenue authorities can maximize cost effectiveness by leveraging firm networks to design interventions that strategically target firms located in different clusters in the transaction network, ensuring a

greater impact on compliance with fewer resources.

Finally, it is worth highlighting that this study has had a direct impact on the actual policies implemented. Following the 2018 experimental study described in this paper, the data cross-checks by the URA were substantially expanded. In 2021, the URA established a new program called electronic fiscal receipting and invoicing (EFRIS), whereby all input claims must be validated electronically with the corresponding invoice. Furthermore, in 2022, the URA started rolling out electronic billing machines to record transaction information at the point of sale. While these measures represent substantial advances in tax enforcement, their impact has not yet been rigorously evaluated, providing an avenue for future research.

6 Conclusion

This paper uses a near universal dataset of firm-to-firm transaction records from VAT returns in Uganda to design a randomized trial aimed at improving tax compliance and reporting behavior among VAT-registered firms. By leveraging the data and carefully implemented sampling and randomization procedures, we rigorously estimate both the direct and spillover effects of the intervention. To our knowledge, this is the first experimental study to causally measure the impact of a tax enforcement intervention within a firm network at such a fine-grained level.

The intervention was implemented successfully, with 92% of the treated firms confirming receipt of the letter. The letters elicited strong responses, significantly increasing corrections of both listed and unlisted discrepancies. Spillover effects extended to other trading partners, both within treated links and with partners outside the Study Sample, and across time periods, affecting months not listed in the letters. In particular, some treated buyers induced their suppliers to correct reported sales, offering rare insights into firm-to-firm transmission of enforcement responses. Additionally, the asymmetric responses observed suggest that VAT evasion in this setting is primarily driven by sellers. Overall, the intervention slightly increased VAT liability for the past months and proved to be cost-effective.

These findings underscore the potential of leveraging firm networks to strengthen tax enforcement, particularly in addressing past discrepancies and improving reporting behavior. However, the limited impact on overall VAT revenue suggests opportunities for refinement of interventions using detailed network data. Future research could investigate how such interventions influence compliance over longer horizons and how other government policies propagate through firm networks.

References

- Adão, Rodrigo, Paul Carrillo, Arnaud Costinot, Dave Donaldson, and Dina Pomeranz**, “Imports, Exports, and Earnings Inequality: Measures of Exposure and Estimates of Incidence*,” *The Quarterly Journal of Economics*, August 2022, *137* (3), 1553–1614.
- Agha, Ali and Jonathan Haughton**, “Designing VAT systems: Some efficiency considerations,” *Review of Economics and Statistics*, 1996, *78*, 303–308.
- Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P Vasquez**, “The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages,” *The Quarterly Journal of Economics*, 01 2022, *137* (3), 1495–1552.
- Almunia, Miguel and David Lopez-Rodriguez**, “Under the Radar: The Effects of Monitoring Firms on Tax Compliance,” *American Economic Journal: Economic Policy*, 2018, *10*, 1–38.
- , **Jonas Hjort, Justine Knebelmann, and Lin Tian**, “Information, Fiscal Capacity and Tax Revenues: An Experimental Evaluation,” 2018. AEA RCT Registry, <https://doi.org/10.1257/rct.2958-1.0>.
- , – , – , and – , “Strategic or Confused Firms? Evidence from Missing Transactions in Uganda,” *Review of Economics and Statistics*, 2024, *106* (1), 256–265.
- Atkin, David, Amit K. Khandelwal, and Adam Osman**, “Exporting and Firm Performance: Evidence from a Randomized Experiment,” *The Quarterly Journal of Economics*, 2017, *132* (2), 551–615.
- Baird, Sarah, J. Aislinn Bohren, Craig McIntosh, and Berk Özler**, “Optimal Design of Experiments in the Presence of Interference,” *The Review of Economics and Statistics*, December 2018, *100* (5), 844–860.
- Basri, M. Chatib, Mayara Felix, Rema Hanna, and Benjamin A. Olken**, “Tax Administration versus Tax Rates: Evidence from Corporate Taxation in Indonesia,” *American Economic Review*, December 2021, *111* (12), 3827–71.
- Battaglini, Marco, Luigi Guiso, Chiara Lacava, and Eleonora Patacchini**, “Tax Professionals and Tax Evasion,” 2019. NBER working paper 25745.
- Bernard, Andrew B., Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes**, “The Origins of Firm Heterogeneity: A Production Network Approach,” *Journal of Political Economy*, 2022, *130* (7), 1765–1804.
- Besley, Timothy and Torsten Persson**, “Chapter 2 - Taxation and Development,” in Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, eds., *handbook of public economics, vol. 5*, Vol. 5 of *Handbook of Public Economics*, Elsevier, 2013, pp. 51–110.
- and – , “Why Do Developing Countries Tax So Little?,” *Journal of Economic Perspectives*, November 2014, *28* (4), 99–120.
- Best, Michael, Anne Brockmeyer, Henrik Kleven, Johannes Spinnewijn, and Mazhar Waseem**, “Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan,” *Journal of Political Economy*, 2015, *123* (6), 1311–1355.
- , **Jawad Shah, and Mazhar Waseem**, “Detection Without Deterrence: Long-Run

- Effects of Tax Audit on Firm Behavior,” 2021. mimeo.
- 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.
- 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*, February 2024, *38* (1), 107–32.
- Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal**, “Ghosting the Tax Authority: Fake Firms and Tax Fraud in Ecuador,” *American Economic Review: Insights*, December 2023, *5* (4), 427–44.
- , **Dina Pomeranz, and Monica Singhal**, “Dodging the Taxman: Firm Misreporting and Limits to tax enforcement,” *American Economic Journal: Applied Economics*, 2017, *9* (2), 144–164.
- Cruces, Guillermo, Dario Tortarolo, and Gonzalo Vazquez-Bare**, “Design of Partial Population Experiments with an Application to Spillovers in Tax Compliance,” 2024. Working Paper.
- Demir, Banu, Beata Javorcik, Tomasz K. Michalski, and Evren Ors**, “Financial Constraints and Propagation of Shocks in Production Networks,” *The Review of Economics and Statistics*, 01 2022, pp. 1–46.
- Deserranno, Erika, Stefano Caria, Philipp Kastrau, and Gianmarco León-Ciliotta**, “The Allocation of Incentives in Multi-Layered Organizations,” 2022. Working Paper.
- Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot**, “Firm Responses and Wage Effects of Foreign Demand Shocks with Fixed Labor Costs and Monopsony,” September 2022.
- Drago, Francesco, Friederike Mengel, and Christian Traxler**, “Compliance Behavior in Networks: Evidence from a Field Experiment,” *American Economic Journal: Applied Economics*, April 2020, *12* (2), 96–133.
- Duflo, Esther and Emmanuel Saez**, “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment*,” *The Quarterly Journal of Economics*, August 2003, *118* (3), 815–842.
- Gadenne, Lucie, Tushar Nandi, and Roland Rathelot**, “Taxation and Supplier Networks: Evidence from India,” 2024. Working Paper.
- Garriga, Pablo and Dario Tortarolo**, “Firms as tax collectors,” *Journal of Public Economics*, 2024, *233*, 105092.
- Hjort, Jonas, Vinayak Iyer, and Golvine de Rochambeau**, “Informational Barriers to Market Access: Experimental Evidence from Liberian Firms,” CEPR Discussion Papers 15219, C.E.P.R. Discussion Papers December 2021.
- Holz, Justin E., John A. List, Alejandro Zentner, Marvin Cardoza, and Joaquin E. Zentner**, “The \$100 million nudge: Increasing tax compliance of firms using a natural field experiment,” *Journal of Public Economics*, 2023, *218*, 104779.
- Hoy, Christopher, Mathias Sinning, and Luke McKenzie**, “Improving Tax Compliance

- without Increasing Revenue: Evidence from Population-Wide Randomized Controlled Trials in Papua New Guinea,” *Economic Development and Cultural Change*, 2022, *forthcoming*.
- Hudgens, Michael G. and M. Elizabeth Halloran**, “Toward Causal Inference With Interference,” *Journal of the American Statistical Association*, June 2008, *103* (482), 832–842.
- IMF**, “Revenue Administration Gap Analysis Program: The Value-Added Tax Gap,” 2014. Fiscal Affairs Department - International Monetary Fund.
- IMF**, “Uganda: Staff Report for the 2021 Article IV Consultation,” Article IV Report 22/77, International Monetary Fund March 2022.
- Keen, Michael and Ben Lockwood**, “The Value Added Tax: Its Causes and Consequences,” *Journal of Development Economics*, 2010, *92* (2), 138–151.
- Lediga, Collen, Nadine Riedel, and Kristina Strohmaier**, “Tax Enforcement Spillovers - Evidence from Business Audits in South Africa,” 2022. mimeo.
- Lemgruber, Andrea, Andrew Masters, and Duncan Cleary**, “Understanding Revenue Administration: An Initial Data Analysis Using the Revenue Administration Fiscal Information Tool,” 2015. International Monetary Fund Fiscal Affairs Department Paper Series.
- Lopez-Luzuriaga, Andrea and Carlos Scartascini**, “Compliance spillovers across taxes: The role of penalties and detection,” *Journal of Economic Behavior and Organization*, 2019, *164*, 518–534.
- Manski, Charles F.**, “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 1993, *60* (3), 531–542.
- Mascagni, Giulia, Andualem T. Mengistu, and Firew B. Woldeyes**, “Can ICTs increase tax compliance? Evidence on taxpayer responses to technological innovation in Ethiopia,” *Journal of Economic Behavior and Organization*, 2021, *189*, 172–193.
- , **Fabrizio Santoro, Denis Mukama, John Karangwa, and Napthal Hakizimana**, “Active Ghosts: Nil-filing in Rwanda,” *World Development*, 2022, *152*, 105806.
- McMillan, John and Christopher Woodruff**, “Interfirm Relationships and Informal Credit in Vietnam*,” *The Quarterly Journal of Economics*, 11 1999, *114* (4), 1285–1320.
- OECD**, *Revenue Statistics 2022* 2022.
- OECD/ATAF**, *Revenue Statistics in Africa 2022* 2022.
- Pinotti, Paolo**, “Clicking on Heaven’s Door: The Effect of Immigrant Legalization on Crime,” *American Economic Review*, January 2017, *107* (1), 138–168.
- Pomeranz, Dina**, “No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax,” *American Economic Review*, 2015, *105* (8), 2539–69.
- Rosenbaum, Paul R.**, “Interference between Units in Randomized Experiments,” *Journal of the American Statistical Association*, 2007, *102* (477), 191–200.
- Roth, Jonathan and Pedro H. C. Sant’Anna**, “Efficient Estimation for Staggered Rollout Designs,” *Journal of Political Economy Microeconomics*, 2023, *1* (4), 669–709.
- Shimeles, Abebe, Daniel Zerfu Gurara, and Firew Woldeyes**, “Taxman’s Dilemma: Coercion or Persuasion? Evidence from a Randomized Field Experiment in Ethiopia,” *American Economic Review*, May 2017, *107* (5), 420–24.

- 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.
- , **Brett Collins, Jeffrey Hoopes, Daniel Reck, and Michael Sebastiani**, “Does Credit-card Information Reporting Improve Small-business Tax Compliance?,” *Journal of Public Economics*, 2017, 149, 1–19.
- Vazquez-Bare, Gonzalo**, “Identification and Estimation of Spillover Effects in Randomized Experiments,” *Journal of Econometrics*, November 2023, 237 (1), 105237.
- Waseem, Mazhar**, “Overclaimed refunds, undeclared sales, and invoice mills: Nature and extent of noncompliance in a value-added tax,” *Journal of Public Economics*, 2023, 218, 104783.
- World Bank**, “GNI per capita in PPP (current international dollars),” 2021. data retrieved from World Development Indicators, <https://data.worldbank.org/indicator/NY.GNP.PCAP.PP.CD?locations=UG>.

Online Appendix

For web publication only

A Assumptions

In this appendix, we evaluate the identification assumptions outlined in Section 3.2, which are critical for establishing and interpreting the causal effects of the intervention within trading networks. This section presents empirical tests to assess the validity of these assumptions.

Assumption 1:

Our first assumption is that the letter intervention only impacts firm links that meet the Base Sample selection criteria, with an important condition being that the seller-buyer links exhibit seller shortfall. We test this assumption by showing that the letter induces corrections of seller shortfall instances, but not other types of seller-buyer month observations (buyer shortfall or match cases). To test whether this is the case, we run the following regression:

$$Y_{it} = \alpha + \sum_{h=1}^3 \beta_h T_{ih} + \delta_t + \epsilon_{it}, \quad (6)$$

where Y_{it} is a dummy variable indicating whether the discrepancy for month t of seller-buyer link i has been changed at any point in the ten months after treatment. Modifications in discrepancies reported by both sellers and buyers are included. T_{ih} denotes a set of dummy variables capturing the three mutually exclusive treatment arms, δ_t is a month fixed effect, and ϵ_{it} is an error term. Observations are at the link-month level, and standard errors are clustered at the link level. The sample includes all non-seller shortfall observations for 419 unique seller-buyer links.

As shown in Table A1, we find that, whether it be any treatment or treatment arms taken separately, there is no effect of receiving the letter on the probability of making corrections to non-seller seller shortfall instances.

Assumption 2:

The second assumption is that spillover effects across links do not differentially affect the control and treated groups. Since all links in the Study Sample are separated by at least one degree in the network, a sufficient condition to validate this assumption is that spillovers do not extend beyond direct trading relationships. In other words, firms connected to treated links should remain unaffected. To test this, we conduct an exercise examining which firms are responsible for making corrections in untreated links. As shown in Table A2, nearly all corrections are made by treated firms themselves, i.e., those directly involved in treated links, rather than by their untreated partners outside these links.

Furthermore, even if second-degree spillovers were present, the assumption is still valid if the second-degree exposure to treatment—defined as treatment received by a firm’s direct trading partners—does not affect the reporting behaviors of treated and control firms differentially. We show, in Table A3, that the effect of the treatment does not vary with the strength of exposure to the treatment.

Table A1
Test for Assumption 1: Any Correction of Non-Seller Shortfall Discrepancies

Dependent variable:	Any correction		Match		Buyer shortfall	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.021 (0.013)		0.006 (0.007)		0.050 (0.031)	
Buyer only		0.037** (0.018)		0.019 (0.013)		0.077* (0.041)
Seller only		0.005 (0.014)		-0.002 (0.007)		0.033 (0.037)
Buyer and Seller		0.019 (0.019)		0.002 (0.009)		0.038 (0.045)
R-squared	0.016	0.019	0.009	0.014	0.045	0.048
Observations	1376	1376	862	862	514	514
No. of Unique Links	421	421	309	309	246	246
Mean of Dep. in Control	0.008	0.008	0.007	0.007	0.010	0.010
P-value of $\beta_S = \beta_B$		0.078		0.105		0.324
P-value of $\beta_{SB} = \beta_B$		0.434		0.223		0.448
P-value of $\beta_{SB} = \beta_S$		0.453		0.618		0.914
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table tests whether there are effects of the letters on the corrections of non-seller shortfall discrepancies (Assumption 1), using equation (6). VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The outcome variable is an indicator set to 1 if a discrepancy is reduced. Columns 1-2 includes all non-seller shortfall discrepancies, columns 3-4 focuses on transactions with matched amounts, and columns 5-6 focuses on buyer shortfall discrepancies. Standard errors are clustered at the link level. *p < 0.10; **p < .05; ***p < .01. This table is mentioned in Section 3.2. *Source:* Data from monthly VAT returns submitted to the URA.

Table A2
 Test for Assumption 2: Any Correction in Untreated Links with a Treated Firm

Dependent variable:	Any Correction		Correction by Study Firm		Correction by Partner Firm	
	(1)	(2)	(3)	(4)	(5)	(6)
Any Treatment	0.006*** (0.001)		0.006*** (0.000)		-0.000 (0.001)	
Buyer only		0.003* (0.002)		0.002** (0.001)		0.001 (0.001)
Seller only		0.002 (0.002)		0.001 (0.001)		0.001 (0.001)
Buyer and Seller		0.001 (0.001)		0.001 (0.001)		0.000 (0.001)
R-squared	0.002	0.001	0.002	0.001	0.000	0.000
Observations	200924	200924	200924	200924	200924	200924
No. of Unique Links	47390	47390	47390	47390	47390	47390
No. of Study Firms	2391	2391	2391	2391	2391	2391
No. of Partners Firms	7550	7550	7550	7550	7550	7550
Mean of Dep. in Control	0.009	0.009	0.005	0.005	0.004	0.004
P-value of $\beta_S = \beta_B$		0.656		0.491		0.807
P-value of $\beta_{SB} = \beta_B$		0.348		0.240		0.512
P-value of $\beta_{SB} = \beta_S$		0.653		0.652		0.700
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table tests whether there are effects of the letters in untreated links and whether they are driven by firms of a treated link (Assumption 2). VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The sample includes all untreated links which include a treated firm. The outcome variable is an indicator set to 1 if a discrepancy is reduced. In columns 1-2 the outcome is defined as any correction, columns 3-4 focuses on corrections made by the firm which is in the study sample, and columns 5-6 focuses on correction made by the partner firm from the untreated link outside the study sample. Standard errors are clustered at the link level. *p< 0.10; **p< .05; ***p< .01. This table is mentioned in Section 3.2. *Source:* Data from monthly VAT returns submitted to the URA.

Table A3
 Test for Assumption 2: Effect of Exposure to Treatment on Any
 Correction in Untreated Links with a Treated Firm

Dependent variable:	Any correction			
	Exposure (Continuous)		Exposure (High-Low)	
	(1)	(2)	(3)	(4)
Treatment exposure	0.009 (0.006)	0.009 (0.005)		
Any Treatment		0.013** (0.006)		0.006*** (0.002)
Any Treatment × Treatment exposure		-0.011 (0.009)		
High exposure			-0.000 (0.003)	-0.002 (0.003)
Any Treatment × High exposure				0.001 (0.002)
R-squared	0.001	0.002	0.001	0.002
Observations	200924	200924	200924	200924
No. of Unique Links	47390	47390	47390	47390
No. of Study Firms	2391	2391	2391	2391
No. of Partners Firms	7550	7550	7550	7550
Mean of Dep. in Control	0.009	0.009	0.009	0.009
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table tests whether there are effects of the letters in untreated links depending on the exposure to the treatment (Assumption 2). VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The sample includes all untreated links which include a treated firm. The outcome variable is an indicator set to 1 if a discrepancy is reduced. Exposure is defined as the share of a firms' partners that are treated, taking values 0 to 1 in columns 1 and 2, and an indicator taking value 1 if exposure is above the mean in columns 3 and 4. Standard errors are clustered at the link level. *p < 0.10; **p < .05; ***p < .01. This table is mentioned in Section 3.2. *Source:* Data from monthly VAT returns submitted to the URA.

B Data correction following extraction problem

In this appendix, we discuss in detail the data extraction error that cause us to inflate the occurrences and amounts of seller shortfall between firm-pairs. We show that it is balanced across treatment arm, indicating that correcting for it did not bias our results.

The error occurs as a result of a mistake in the script that extracted the administrative data from the URA database. As one can observe in Table B1, the share of firm-pairs this affected was balanced across treatment arms indicating that this did not bias our sample.

Table B1
Balance table for number of firm-pairs missing

	N pairs	N missing pairs	Share missing pairs
Control	676	182	0.27
Treatment	1029	288	0.28
By treatment arm			
Buyer and Seller	346	85	0.25
Buyer only	343	101	0.29
Seller only	340	102	0.30

Notes: This table displays the number and share of firm pairs that were notified of a discrepancy purely because of the data extraction error that duplicated the reported inputs by firms for the period 19/9/2017 to 15/10/2017.

We formally test whether being part of a treatment arm is correlated with whether the firm was included because of the data correction error in Table B2. In column (1) we run the regression using an indicator variable for whether the firm-pair was treated and in column (2) we run the regression with separate indicator variables for each treatment arm. The omitted indicator is therefore always the control group. The individual coefficient are small and never statistically significant. Furthermore, the F-test for the joint significance of all coefficients is clearly rejecting that they are different from the constant.

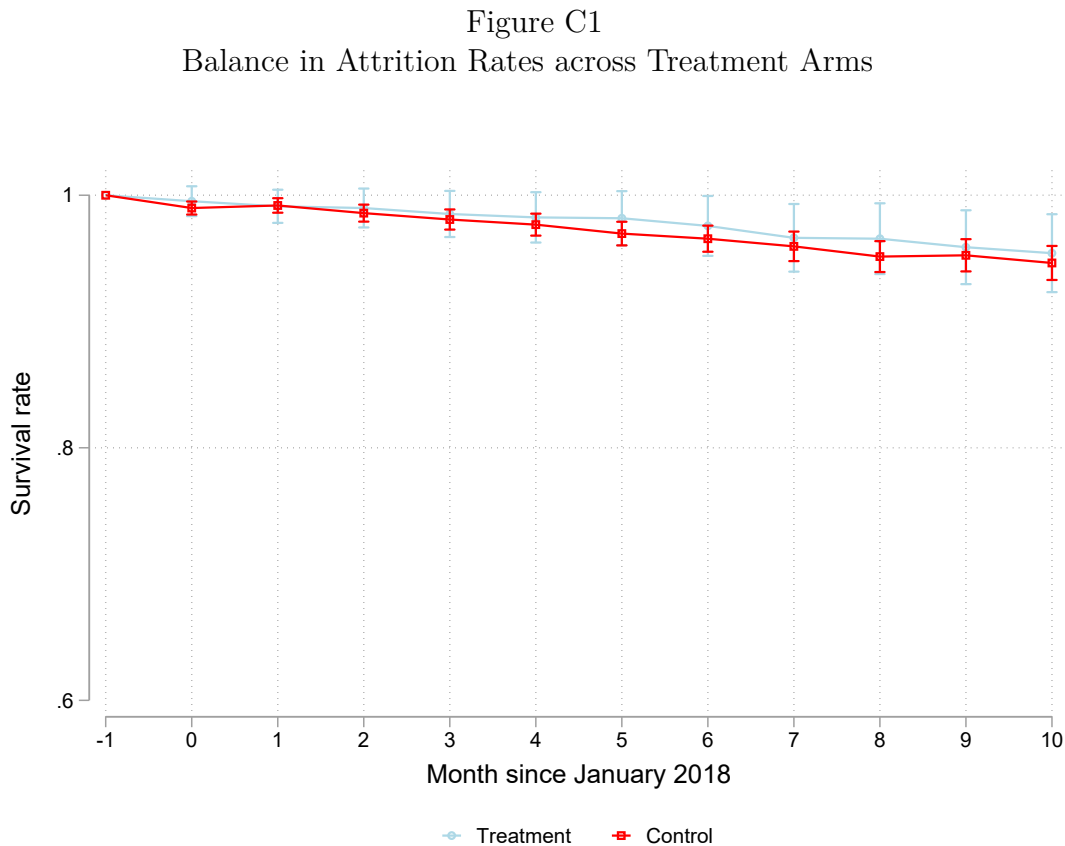
Table B2
Balance of the Probability of Being
Removed from Sample

Dependent variable:	Prob. of missing	
	(1)	(2)
Treatment	0.011 (0.022)	
Buyer and Seller		-0.024 (0.029)
Buyer only		0.025 (0.030)
Seller only		0.031 (0.030)
Constant	0.269 (0.017)	0.269 (0.017)
R-squared	0.000	0.002
Observations	1705	1705
P-value of F-test	0.630	0.341

Notes: This table verifies whether the probability of a link being affected by the data extraction error is balanced across treatment and control. The sample includes all links that were originally included in our sample. The outcome variable is an indicator taking value 1 if the link only featured a discrepancy because of the data extraction error and is hence excluded from our Base Sample. Standard errors are clustered at the link level. * $p < .10$; ** $p < .05$; *** $p < .01$. *Source:* Data from monthly VAT returns submitted to the URA.

C Evidence on Implementation

In this appendix, we provide more information on the experimental implementation and follow-up details. We first show that there is no differential attrition rate in treatment and control groups in Figure C1. We then show a template of the notification letters sent to firms in Figure C2. Table C1 shows shares of treated firms that provided feedback on the letters and Figure C3 displays the types of feedback submitted by the treated firms.



Notes: This figure displays the share of firms observed just before the intervention starts that keep filing every month afterwards. The share for the 1482 firms from treated links is plotted in blue, while the share for the 988 firms from control links is plotted in red. The figure is mentioned in Section 3.4. *Source:* Data from monthly VAT returns submitted to the URA.

Figure C2
Template of the Notification Letter

Ref: URA/DTD/CMHQ/RISK/«TIN_of_the_buyer»

28th February, 2018

«Name_Of_The_Buyer»
«Physical_Location»
«District», Uganda
Tel: «Tel_no»

Dear Sir/Madam,

RE: NOTIFICATION OF INCREASED VAT COMPLIANCE MONITORING

The Uganda Revenue Authority has developed a new system of monitoring value-added tax (VAT) compliance through reviewing VAT declarations. Therefore, from now on, your VAT declarations are being **closely monitored** to determine your compliance status.

This communication is to draw your attention to a discrepancy **«Sum_of_Des» UGX** that has been detected arising from mismatches between the input VAT claimed by your company and the output VAT declared by **your trading partner «Nameof_The_Seller» (TIN: «TIN_of_the_seller»)** on their sales to you for transactions reported from March 2017 to December 2017. For clarification, some illustrative cases are reported in **annexure**, attached.

Given these inconsistencies, this is to request you to check your VAT declarations and amend your returns accordingly. You are advised to comply with the above requirement by the 16th of April, 2018. Failure to comply will not only result in additional enforcement measures raised against you in accordance with Section 23 of the Tax Procedures Code (TPC) Act, but could also lead to prosecution in accordance with Section 58 of the TPC. Both sections are quoted in **annexure**, attached.

If you require any clarification, please contact **XXXX** (email: xxx or Tel: xxx) or at 3rd Floor, Tall tower-Crested Towers or the under signed.

We thank you for your usual cooperation with the URA as we Develop Uganda Together.

«Assistant Commissioner Compliance Management»

For: Commissioner General


Copy: «Jurisdiction»

Table C1
Feedback from Treated Firms

	LTO	MTO	STO	Overall
Number of firms	80	168	754	1002
Share receiving letters or giving feedback	0.98	0.93	0.89	0.91
Share contacting URA	0.55	0.44	0.34	0.37

Notes: This table reports summary statistics on firms' interaction with the URA following the intervention. We consider a firm received the letter or gave feedback if it received the physical letter and signed off, and/or got in touch with the URA either in person, by call, or by email. Shares are calculated using the number of firms receiving the letter as a denominator. This figure is mentioned in Section 3.4. *Source:* Data from monthly VAT returns submitted to the URA.

Figure C3
Responses from Firms that Gave Feedback




Notes: This figure displays the types of feedback submitted by treated firms to the URA. If firms provided feedback corresponding to multiple categories, all are counted. This figure is mentioned in Section 3.4. *Source:* Data from monthly VAT returns submitted to the URA.

D Additional Results on on Corrections of Past VAT Returns


This appendix provides additional results on corrections of past VAT returns, supplementing those shown in Section 4.1. Figure D1 compares the distributions of discrepancies before and after the treatment for the treatment and control groups. Tables D1-D3 show the regression results that correspond to Figures 3-5 discussed in the main texts and include additionally the alternative outcomes that take value one only if the discrepancy has been fully resolved through the corrections, i.e., “Full correction.” Table D4 runs a similar exercise as in the one in Table D3 but distinguishes between months with listed and unlisted discrepancies. Table D5 also runs the same test as the one in Table D3 but focuses on the sample of outside links in which the sellers in the Study Sample act as buyers in these outside links. Finally, Tables D6 and D7 display results of the outside-link effects as in Tables D3 and D5, but focus on buyers in the Study Sample.

Figure D1
Distribution of Discrepancies Before and After Treatment

(a) Treatment



(b) Control



Notes: This figure shows the distribution of discrepancy amounts before and after treatment. In Panel (a), we show the distribution for treated links, while Panel (b) shows the distribution for control links. Discrepancies are computed as seller amount minus buyer amount, hence values below zero imply seller shortfall, while values above zero imply buyer shortfall. This figure is mentioned in Section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table D1
Within-Link Effects of the Letter on the Correction of Past Discrepancies

Dependent variable:	Any correction				Full correction			
	Listed		Unlisted		Listed		Unlisted	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	0.223*** (0.017)		0.106*** (0.019)		0.140*** (0.013)		0.069*** (0.013)	
Buyer only		0.116*** (0.021)		0.039** (0.019)		0.074*** (0.017)		0.028* (0.015)
Seller only		0.243*** (0.030)		0.144*** (0.036)		0.155*** (0.024)		0.081*** (0.026)
Buyer and Seller		0.303*** (0.029)		0.138*** (0.034)		0.186*** (0.024)		0.099*** (0.025)
R-squared	0.092	0.121	0.034	0.054	0.055	0.071	0.022	0.035
Observations	2744	2744	2153	2153	2731	2731	2150	2150
No. of Unique Links	1235	1235	528	528	1235	1235	527	527
Mean of Dep. in Control	0.016	0.016	0.011	0.011	0.008	0.008	0.003	0.003
P-value of $\beta_S = \beta_B$		0.000		0.009		0.005		0.076
P-value of $\beta_{SB} = \beta_B$		0.000		0.010		0.000		0.017
P-value of $\beta_{SB} = \beta_S$		0.149		0.902		0.353		0.635
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on the correction of past discrepancies within links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The table reports results from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced. In columns 5 to 8, the outcome variable is defined more narrowly, taking value 1 only in cases where a discrepancy is fully resolved. Columns 1-2 and 5-6 focus in listed discrepancies, mentioned in the letter, while columns 3-4 and 7-8 focus on unlisted discrepancies. The sample size in this case is smaller since we drop links for which all discrepancies observed in the pre-treatment period were listed on the letter. Standard errors are clustered at the link level. *p < 0.10; **p < .05; ***p < .01. Results for Any Correction are shown in Figure 3 in section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table D2
Within-Link Direct and Spillover Effects: Who Reduces Discrepancies?

Correction by:	Listed				Unlisted			
	Seller		Buyer		Seller		Buyer	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	0.185*** (0.016)		0.026*** (0.005)		0.094*** (0.018)		0.005 (0.003)	
Buyer only		0.079*** (0.019)		0.029*** (0.009)		0.038** (0.019)		0.003 (0.002)
Seller only		0.223*** (0.029)		0.014* (0.008)		0.128*** (0.035)		0.012 (0.009)
Buyer and Seller		0.249*** (0.027)		0.033*** (0.009)		0.119*** (0.032)		0.001 (0.001)
R-squared	0.077	0.107	0.013	0.016	0.033	0.050	0.006	0.012
Observations	2744	2744	2744	2744	2153	2153	2153	2153
No. of Unique Links	1235	1235	1235	1235	528	528	528	528
Mean of Dep. in Control	0.014	0.014	0.001	0.001	0.010	0.010	0.000	0.000
P-value of $\beta_S=\beta_B$		0.000		0.181		0.023		0.263
P-value of $\beta_{SB}=\beta_B$		0.000		0.754		0.027		0.318
P-value of $\beta_{SB}=\beta_S$		0.509		0.118		0.850		0.161
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on corrections made by each firm (seller and buyer) within links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. The table reports results from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced. We distinguish whether the correction is made by the seller (columns 1-2 and 5-6) or the buyer (columns 3-4 and 7-8). Columns 1-4 focus on listed discrepancies, mentioned in the letter, while columns 5-8 focus on unlisted discrepancies. The sample size in this case is smaller since we drop links for which all discrepancies observed in the pre-treatment period were listed on the letter. Standard errors are clustered at the link level. *p < 0.10; **p < .05; ***p < .01. Results for Listed discrepancies are shown in Figure 4 in section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table D3
 Outside-Link Effects of the Letter on the Correction of Past
 Discrepancies (Sellers in Study Sample as Sellers)

Correction by:	Any correction		Full correction	
	Seller (1)	Seller (2)	Seller (3)	Seller (4)
Treatment	0.020 (0.012)		0.010* (0.006)	
Buyer only		0.007 (0.007)		0.001 (0.003)
Seller only		-0.001 (0.003)		0.001 (0.002)
Buyer and Seller		0.058 (0.036)		0.029* (0.016)
R-squared	0.007	0.031	0.004	0.016
Observations	41366	41366	41366	41366
No. of Unique Links	12783	12783	12783	12783
No. of Study Firms	968	968	968	968
Mean of Dep. in Control	0.005	0.005	0.003	0.003
P-value of $\beta_S = \beta_B$		0.245		0.829
P-value of $\beta_{SB} = \beta_B$		0.158		0.086
P-value of $\beta_{SB} = \beta_S$		0.098		0.075
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on corrections outside the seller-buyer links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. To identify outside-link effects, the sample includes all seller shortfall discrepancies of the sellers from the study sample with all their buyers, excluding the buyer from the study sample link. We focus on corrections made by sellers. We report results from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced by the seller. In columns 3-4, the outcome variable is defined more narrowly, taking value 1 only in cases where a discrepancy is fully resolved. Standard errors are clustered at the seller level. *p < 0.10; **p < .05; ***p < .01. This table is mentioned in Section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table D4
 Outside-Link Effects of the Letter on the Correction of Past
 Discrepancies (Sellers in Study Sample as Sellers)

Correction by:	Listed Months		Unlisted Months	
	Seller (1)	Seller (2)	Seller (3)	Seller (4)
Treatment	0.027** (0.013)		0.032* (0.019)	
Buyer only		0.010 (0.009)		0.013 (0.010)
Seller only		0.001 (0.006)		0.000 (0.002)
Buyer and Seller		0.072* (0.038)		0.098 (0.061)
R-squared	0.019	0.047	0.016	0.066
Observations	11530	11530	14525	14525
No. of Unique Links	6866	6866	5365	5365
No. of Study Firms	744	744	416	416
Mean of Dep. in Control	0.004	0.004	0.001	0.001
P-value of $\beta_S = \beta_B$		0.407		0.193
P-value of $\beta_{SB} = \beta_B$		0.122		0.178
P-value of $\beta_{SB} = \beta_S$		0.078		0.111
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on corrections outside the seller-buyer links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. To identify outside-link effects, the sample includes all seller shortfall discrepancies of the sellers from the study sample with all their buyers, excluding the buyer from the study sample link. We focus on corrections made by sellers. We report results from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced. Columns 1-2 focus on listed discrepancies, mentioned in the letter, while columns 3-4 focus on unlisted discrepancies. *p < 0.10; **p < .05; ***p < .01. Standard errors are clustered at the seller level. Results for Listed Months are shown in Figure 5. This table is mentioned in Section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table D5
 Outside-Link Effects of the Letter on the Correction of Past
 Discrepancies (Sellers in Study Sample as Buyers)

Correction by:	Any correction		Full correction	
	Buyer (1)	Buyer (2)	Buyer (3)	Buyer (4)
Treatment	-0.001* (0.001)		0.000 (0.000)	
Buyer only		-0.001 (0.001)		0.000 (0.000)
Seller only		-0.001 (0.001)		0.000 (0.000)
Buyer and Seller		-0.002** (0.001)		-0.000** (0.000)
R-squared	0.001	0.001	0.000	0.000
Observations	50489	50489	50489	50489
No. of Unique Links	16431	16431	16431	16431
No. of Study Firms	1108	1108	1108	1108
Mean of Dep. in Control	0.003	0.003	0.000	0.000
P-value of $\beta_S = \beta_B$		0.767		0.995
P-value of $\beta_{SB} = \beta_B$		0.402		0.322
P-value of $\beta_{SB} = \beta_S$		0.251		0.324
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on corrections outside the seller-buyer links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. To identify outside-link effects, the sample includes all seller shortfall discrepancies of the sellers from the study sample, in instances where they are buyers (sellers-as-buyers) with other sellers. We focus on corrections made by these sellers-as-buyers. We report results from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is corrected by the seller-as-buyer. In columns 3-4, the outcome variable is defined more narrowly, taking value 1 only in cases where a discrepancy is fully resolved. *p < 0.10; **p < .05; ***p < .01. Standard errors are clustered at the seller level. This table is mentioned in Section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table D6
 Outside-Link Effects of the Letter on the Correction of Past
 Discrepancies (Buyers in Study Sample as Sellers)

Correction by:	Any correction		Full correction	
	Seller (1)	Seller (2)	Seller (3)	Seller (4)
Treatment	-0.000 (0.001)		-0.001 (0.001)	
Buyer only		0.001 (0.002)		-0.000 (0.002)
Seller only		0.000 (0.002)		0.000 (0.002)
Buyer and Seller		-0.002 (0.001)		-0.002* (0.001)
R-squared	0.007	0.007	0.004	0.004
Observations	8320	8320	8320	8320
No. of Unique Links	3385	3385	3385	3385
No. of Study Firms	567	567	567	567
Mean of Dep. in Control	0.005	0.005	0.003	0.003
P-value of $\beta_S = \beta_B$		0.853		0.707
P-value of $\beta_{SB} = \beta_B$		0.228		0.111
P-value of $\beta_{SB} = \beta_S$		0.332		0.088
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on corrections outside the seller-buyer links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. To identify outside-link effects, the sample includes all seller shortfall discrepancies of the buyers from the study sample, in instances where they are sellers (buyers-as-sellers) with other buyers. We focus on corrections made by these buyers-as-sellers. We report results from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced by the buyer-as-seller. In columns 3-4, the outcome variable is defined more narrowly, taking value 1 only in cases where a discrepancy is fully resolved. *p < 0.10; **p < .05; ***p < .01. Standard errors are clustered at the buyer-as-seller level. This table is mentioned in Section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table D7
 Outside-Link Effects of the Letter on the Correction of Past
 Discrepancies (Buyers in Study Sample as Sellers)

Correction by:	Listed months		Unlisted months	
	Seller	Seller	Seller	Seller
Treatment	-0.002 (0.001)		-0.002 (0.002)	
Buyer only		-0.000 (0.002)		-0.003 (0.002)
Seller only		-0.003 (0.002)		-0.000 (0.002)
Buyer and Seller		-0.003 (0.002)		-0.003 (0.002)
R-squared	0.035	0.036	0.007	0.008
Observations	2041	2041	1813	1813
No. of Unique Links	1441	1441	924	924
No. of Study Firms	366	366	186	186
Mean of Dep. in Control	0.006	0.006	0.003	0.003
P-value of $\beta_S = \beta_B$		0.323		0.276
P-value of $\beta_{SB} = \beta_B$		0.273		0.615
P-value of $\beta_{SB} = \beta_S$		0.961		0.291
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on corrections outside the seller-buyer links of the study sample. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. To identify outside-link effects, the sample includes all seller shortfall discrepancies of the buyers from the study sample, in instances where they are sellers (buyers-as-sellers) with other buyers. We focus on corrections made by these buyers-as-sellers. We report results from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced by the buyer-as-seller. Columns 1-2 focus on listed discrepancies, mentioned in the letter, while columns 3-4 focus on unlisted discrepancies. * $p < 0.10$; ** $p < .05$; *** $p < .01$. Standard errors are clustered at the buyer-as-seller level. This table is mentioned in Section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.


E Analysis of Subsequent Behavior: Additional Results

This appendix provides additional results for the analysis of post-treatment reporting behavior discussed in Section 4.2.

E.1 Within-Link Results

Figure E1 presents a simple comparison of seller shortfall instances observed over time in treatment and control links. Columns 1 and 2 of Table E1 provide the regression results underlying Figure 6. Columns 3-6 of Table E1 provide the regression results underlying Figure 7. Tables E2 and E3 run the same event study analysis as in the one in Table E1, but pool together all pre- and post-treatment months.

Figure E1
Probability of Seller Shortfall by Month in Treatment vs. Control Links



Note: This figure reports the share of seller-buyer links displaying seller shortfall in every month, separately for the control group (in red) and the treatment group (in blue). The dashed vertical line represents the time of the treatment. This share is around 80% for both groups before the intervention, with broadly similar trends. Then, the seller shortfall rate drops to about 50% in the treatment group and 65% in the control group. Outcomes are conditional on trade being reported within the link. This figure is consistent with the event-study results presented in Figure 6 and is mentioned in Section 4.1. *Source:* Data from monthly VAT returns submitted to the URA.

Table E1
 Post-Treatment Effect of the Letter on Within-Link Discrepancies and Trade: Event-Study
 Coefficients

	(1) Seller shortfall	(2) Buyer shortfall	(3) No discrepancy	(4) Transaction size	(5) Seller shortfall amount	(6) Continuing trade
-10	0.024 (0.039)	-0.096** (0.038)	0.072** (0.037)	0.013 (0.192)	-0.053 (0.180)	0.015 (0.033)
-9	-0.008 (0.039)	-0.066* (0.035)	0.074** (0.036)	-0.290 (0.186)	-0.126 (0.182)	0.040 (0.031)
-8	0.003 (0.036)	-0.054 (0.035)	0.051 (0.034)	-0.265 (0.251)	-0.245 (0.188)	0.038 (0.031)
-7	-0.034 (0.036)	-0.020 (0.031)	0.054 (0.033)	-0.169 (0.228)	-0.092 (0.167)	-0.003 (0.034)
-6	-0.041 (0.035)	-0.022 (0.032)	0.063* (0.032)	-0.143 (0.184)	-0.072 (0.151)	0.027 (0.033)
-5	-0.008 (0.036)	-0.051 (0.033)	0.058* (0.033)	-0.064 (0.188)	-0.021 (0.188)	-0.004 (0.034)
-4	-0.043 (0.031)	-0.029 (0.029)	0.072** (0.030)	-0.059 (0.191)	-0.050 (0.173)	-0.006 (0.031)
-3	-0.060* (0.033)	-0.030 (0.030)	0.091*** (0.031)	-0.095 (0.166)	0.040 (0.180)	0.014 (0.031)
-2	-0.110*** (0.035)	0.021 (0.031)	0.089*** (0.031)	0.066 (0.179)	0.048 (0.191)	0.032 (0.029)
-1	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)
0	-0.116*** (0.036)	0.013 (0.034)	0.103*** (0.034)	0.014 (0.191)	-0.162 (0.175)	-0.010 (0.027)
1	-0.160*** (0.040)	0.002 (0.036)	0.158*** (0.036)	0.066 (0.180)	0.045 (0.163)	-0.027 (0.027)
2	-0.197*** (0.039)	0.027 (0.038)	0.170*** (0.036)	0.014 (0.228)	0.116 (0.191)	-0.034 (0.029)
3	-0.195*** (0.040)	0.051 (0.038)	0.144*** (0.035)	-0.137 (0.212)	-0.093 (0.182)	-0.008 (0.029)
4	-0.132*** (0.044)	-0.006 (0.038)	0.139*** (0.038)	-0.137 (0.223)	-0.193 (0.209)	-0.001 (0.030)
5	-0.103** (0.042)	0.011 (0.040)	0.091** (0.039)	-0.255 (0.204)	-0.178 (0.168)	-0.005 (0.029)
6	-0.133*** (0.044)	0.022 (0.039)	0.111*** (0.040)	-0.210 (0.238)	-0.291 (0.184)	0.004 (0.029)
7	-0.149*** (0.043)	0.036 (0.038)	0.113*** (0.041)	-0.186 (0.263)	-0.366* (0.199)	0.008 (0.029)
8	-0.157*** (0.047)	0.014 (0.042)	0.143*** (0.041)	0.164 (0.267)	-0.014 (0.219)	-0.028 (0.031)
9	-0.156*** (0.048)	0.021 (0.043)	0.135*** (0.042)	-0.061 (0.305)	-0.540 (0.332)	0.026 (0.030)
10	-0.156*** (0.048)	0.038 (0.041)	0.118*** (0.042)	0.054 (0.245)	-0.113 (0.198)	0.026 (0.030)
R-squared	0.533	0.354	0.463	0.749	0.720	0.464
Observations	10943	10943	10943	10943	7529	25935
No. of Firms	1077	1077	1077	1077	933	1235
Mean of Dep. in Control	0.741	0.122	0.138	1.034	0.555	0.417
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on subsequent reporting discrepancies and trade. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). The table reports the β_j coefficients estimated in the event-study laid out in equation (2), with the following outcomes: probability of seller shortfall in column 1; probability of buyer shortfall in column 2; probability of there being no discrepancy ("Match") in column 3; transaction size in column 4; seller shortfall amount in column 5; and the probability of continuing trade in column 6. In columns 1-5 the outcome is conditional on trade occurring within the link. All amounts are in thousands of USD and winsorized at the 0.5% level. Standard errors are clustered at the link level. The results are displayed in Figures 6 and 7. This table is mentioned in Section 4.2. *Source:* Data from monthly VAT returns submitted to the URA.

Table E2
 Post-Treatment Effect of the Letter on Within-Link Discrepancies:
 Diff-in-Diff Specification

DiD specification	Seller shortfall		Buyer shortfall		No discrepancy	
	(1)	(2)	(3)	(4)	(5)	(6)
Any Treatment	-0.150*** (0.030)	0.020 (0.027)	0.130*** (0.027)	-0.059 (0.183)	-0.147 (0.159)	-0.005 (0.024)
R-squared	0.533	0.354	0.463	0.749	0.720	0.464
Observations	10943	10943	10943	10943	7529	25935
No. of Links	1077	1077	1077	1077	933	1235
Mean of Dep. in Control	0.792	0.095	0.113	1.026	0.602	0.483
Link FE	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Pre-treatment dummies	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on subsequent reporting discrepancies. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). The table reports the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months, for three outcomes: probability of seller shortfall in columns 1-2; probability of buyer shortfall in columns 3-4; probability of there being no discrepancy ("Match") in columns 5-6. Outcomes are conditional on trade occurring within the link. * $p < 0.10$; ** $p < .05$; *** $p < .01$. Standard errors are clustered at the link level. The results are displayed in Figure 6. This table is mentioned in Section 4.2. *Source:* Data from monthly VAT returns submitted to the URA.

Table E3
 Post-Treatment Effect of the Letter on Within-Link Trade:
 Diff-in-Diff Specification


DiD specification	Transaction size		Seller shortfall amount		Continuing trade		Seller Reports		Buyer Reports	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Any Treatment	-0.059 (0.183)		-0.147 (0.159)		-0.005 (0.024)		0.176*** (0.031)		0.011 (0.021)	
Seller only		-0.246 (0.201)		-0.201 (0.174)		-0.030 (0.028)		0.123*** (0.038)		0.054** (0.022)
Buyer and Seller		0.088 (0.237)		-0.169 (0.187)		-0.006 (0.028)		0.186*** (0.044)		0.010 (0.023)
Buyer only		-0.025 (0.225)		-0.065 (0.157)		0.020 (0.028)		0.214*** (0.042)		-0.026 (0.026)
R-squared	0.749	0.749	0.720	0.720	0.464	0.464	0.686	0.687	0.397	0.399
Observations	10943	10943	7529	7529	25935	25935	10943	10943	10943	10943
No. of Firms	1077	1077	933	933	1235	1235	1077	1077	1077	1077
Mean of Dep. in Control	1.026	1.026	0.602	0.602	0.483	0.483	0.283	0.283	0.953	0.953
Link FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pre-treatment dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on subsequent reporting discrepancies. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). The table reports the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months, for five outcomes: transaction size conditional on trade in columns 1-2, defined as the maximum amount reported by the buyer or seller; seller shortfall amount conditional on there being seller shortfall in columns 3-4; probability of any trade being reported columns 5-6; conditional on trade, whether it is reported by the seller in columns 7-8, and/or the buyer in columns 9-10. All amounts are in thousands of USD and winsorized at the 0.5% level. Standard errors are clustered at the link level. *p< 0.10; **p< .05; ***p< .01. The results are displayed in Figure 7. This table is mentioned in Section 4.2. *Source:* Data from monthly VAT returns submitted to the URA.

E.2 Outside-Link Results

This subsection runs a parallel set of analyses as in the previous subsection, but focusing on the indirect effects of the treatment outside the treated links. Figure E2 presents event-study estimates from (2), but focusing on outside-link effects by selecting a subsample of links formed by a seller from the Study Sample with a buyer outside of the Study Sample. Table E4 provides the regression results underlying Figure E2. Tables E2 and E3 run the same event study analysis as in the one in Table E4, but pool together all pre- and post-treatment months.

Figure E2
Post-Treatment Effect of the Letter on Outside-Link Discrepancies



Note: This figure reports the effect of the letter on subsequent reporting discrepancies outside the seller-buyer links of the study sample. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). To identify outside-link effects, the sample includes all links formed by the sellers of the study sample with all their buyers, excluding the buyer from the study sample link. The figure plots $\hat{\beta}_j$ coefficients estimated in the event-study laid out in equation (2), with three different outcomes: probability of there being no discrepancy (“Match”) in red; probability of seller shortfall in blue; and probability of buyer shortfall in green. Outcomes are conditional on trade occurring within the link. In the bottom right corner, we report the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months. Standard errors are clustered at the link level and the bars report 95% confidence intervals. * $p < 0.10$; ** $p < .05$; *** $p < .01$. See Appendix Table E4 for full regression results. This figure is mentioned in Section 4.2. *Source:* Data from monthly VAT returns submitted to the URA.

Table E4
 Post-Treatment Effect of the Letter on Outside-Link Discrepancies and Trade: Event-Study
 Coefficients

	Seller shortfall (1)	Buyer shortfall (2)	No discrepancy (3)	Transaction size (4)	Seller shortfall amount (5)	Report trade (6)	Seller Reports (7)	Buyer Reports (8)
-10	0.003 (0.006)	0.003 (0.007)	-0.000 (0.007)	0.027 (0.021)	0.031 (0.085)	-0.013 (0.011)	-0.002 (0.006)	-0.002 (0.006)
-9	0.009 (0.006)	-0.007 (0.007)	-0.000 (0.006)	0.018 (0.024)	0.043 (0.064)	0.002 (0.009)	-0.003 (0.005)	-0.003 (0.005)
-8	0.007 (0.006)	-0.006 (0.006)	0.002 (0.005)	-0.007 (0.027)	-0.000 (0.070)	-0.006 (0.010)	0.000 (0.004)	0.000 (0.004)
-7	0.002 (0.006)	0.001 (0.006)	-0.003 (0.006)	0.004 (0.025)	-0.014 (0.062)	-0.014 (0.011)	0.004 (0.005)	0.004 (0.005)
-6	0.001 (0.006)	0.003 (0.006)	0.004 (0.005)	-0.012 (0.023)	0.018 (0.098)	-0.010 (0.010)	0.002 (0.005)	0.002 (0.005)
-5	0.001 (0.007)	0.005 (0.007)	0.001 (0.005)	-0.001 (0.024)	0.070 (0.079)	-0.015* (0.008)	-0.003 (0.005)	-0.003 (0.005)
-4	0.008 (0.006)	-0.006 (0.006)	0.000 (0.005)	0.027 (0.022)	0.046 (0.086)	-0.010 (0.008)	-0.002 (0.004)	-0.002 (0.004)
-3	0.001 (0.005)	-0.006 (0.007)	0.006 (0.006)	-0.004 (0.024)	-0.056 (0.073)	-0.008 (0.009)	-0.006 (0.006)	-0.006 (0.006)
-2	0.007 (0.006)	-0.005 (0.005)	0.001 (0.004)	-0.010 (0.025)	0.086 (0.067)	-0.010 (0.007)	-0.008** (0.004)	-0.008** (0.004)
-1	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)
0	-0.007 (0.007)	0.011 (0.009)	-0.001 (0.004)	0.037 (0.022)	0.135* (0.076)	-0.001 (0.006)	-0.001 (0.004)	-0.001 (0.004)
1	-0.005 (0.005)	0.005 (0.007)	0.002 (0.005)	0.026 (0.022)	0.074 (0.071)	0.017** (0.008)	-0.004 (0.004)	-0.004 (0.004)
2	-0.006 (0.008)	0.008 (0.010)	-0.001 (0.005)	0.005 (0.021)	-0.004 (0.060)	0.007 (0.011)	-0.001 (0.005)	-0.001 (0.005)
3	-0.003 (0.011)	0.013 (0.012)	-0.006 (0.005)	0.035* (0.021)	0.149** (0.063)	0.000 (0.012)	0.002 (0.005)	0.002 (0.005)
4	0.001 (0.008)	0.005 (0.008)	-0.002 (0.005)	0.009 (0.023)	0.006 (0.087)	-0.000 (0.010)	-0.003 (0.005)	-0.003 (0.005)
5	0.003 (0.006)	0.006 (0.007)	-0.005 (0.006)	0.012 (0.018)	0.195** (0.089)	-0.002 (0.010)	-0.003 (0.006)	-0.003 (0.006)
6	0.025 (0.021)	-0.010 (0.015)	-0.011 (0.007)	-0.001 (0.022)	0.110* (0.057)	0.018 (0.018)	0.000 (0.007)	0.000 (0.007)
7	0.003 (0.007)	0.004 (0.006)	-0.004 (0.005)	-0.004 (0.022)	-0.006 (0.076)	0.005 (0.010)	0.003 (0.005)	0.003 (0.005)
8	0.011 (0.010)	-0.001 (0.008)	-0.007 (0.005)	-0.006 (0.020)	0.044 (0.060)	0.012 (0.010)	-0.004 (0.006)	-0.004 (0.006)
9	0.015 (0.012)	-0.012 (0.009)	0.001 (0.006)	-0.032 (0.031)	0.056 (0.082)	0.006 (0.010)	-0.009 (0.006)	-0.009 (0.006)
10	0.002 (0.006)	-0.016 (0.016)	0.015 (0.015)	-0.029 (0.029)	0.112 (0.093)	0.010 (0.012)	-0.020 (0.015)	-0.020 (0.015)
R-squared	0.789	0.819	0.690	0.768	0.812	0.389	0.862	0.862
Observations	235540	235540	235540	235540	23748	1634850	235540	235540
No. of Firms	35856	35856	35856	35856	3551	77850	35856	35856
Mean of Dep. in Control	0.137	0.788	0.070	0.528	0.464	0.142	0.861	0.861
Link FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on subsequent reporting discrepancies and trade outside the seller-buyer links of the study sample. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). To identify outside-link effects, the sample includes all links formed by the sellers of the study sample with all their buyers, excluding the buyer from the study sample link. The table reports the β_j coefficients estimated in the event-study laid out in equation (2), with the following outcomes: probability of seller shortfall in column 1; probability of buyer shortfall in column 2; probability of there being no discrepancy ("Match") in column 3; transaction size in column 4; seller shortfall amount in column 5; and the probability of continuing trade in column 6. In columns 1-5 the outcome is conditional on trade occurring within the link. All amounts are in thousands of USD and winsorized at the 0.5% level. Standard errors are clustered at the link level. The results are displayed in Figure E2. This table is mentioned in Section 4.2. *Source:* Data from monthly VAT returns submitted to the URA.

Table E5
 Post-Treatment Effect of the Letter on Outside-Link Discrepancies: Diff-in-Diff Specification

DiD specification	Seller shortfall		Buyer shortfall		No discrepancy	
	(1)	(2)	(3)	(4)	(5)	(6)
Any Treatment	-0.003 (0.002)		0.006** (0.002)		-0.003 (0.002)	
Seller only		-0.003 (0.004)		0.005 (0.004)		-0.002 (0.004)
Buyer and Seller		-0.007* (0.004)		0.001 (0.004)		0.006* (0.004)
Buyer only		-0.006 (0.004)		0.004 (0.004)		0.002 (0.004)
R-squared	0.796	0.796	0.819	0.819	0.640	0.640
Observations	357665	357665	357665	357665	357665	357665
No. of Firms	46512	46512	46512	46512	46512	46512
Mean of Dep. in Control	0.258	0.258	0.655	0.655	0.087	0.087
Link FE	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Pre-treatment dummies	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on subsequent reporting discrepancies outside the seller-buyer links of the study sample. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). To identify outside-link effects, the sample includes all links formed by the sellers of the study sample with all their buyers, excluding the buyer from the study sample link. The table reports the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months, for three outcomes: probability of seller shortfall in columns 1-2; probability of buyer shortfall in columns 3-4; probability of there being no discrepancy ("Match") in columns 5-6. Outcomes are conditional on trade occurring within the link. * $p < 0.10$; ** $p < .05$; *** $p < .01$. Standard errors are clustered at the link level. This table is mentioned in Section 4.2. *Source:* Data from monthly VAT returns submitted to the URA.

Table E6
 Post-Treatment Effect of the Letter on Outside-Link Trade:
 Diff-in-Diff Specification

DiD specification	Transaction size		Seller shortfall amount		Continuing trade		Seller Reports		Buyer Reports	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Any Treatment	0.000 (0.018)		0.001 (0.021)		-0.002 (0.003)		0.008*** (0.003)		0.001 (0.003)	
Seller only		-0.007 (0.021)		0.032 (0.023)		-0.001 (0.004)		0.005 (0.003)		0.001 (0.003)
Buyer and Seller		0.001 (0.020)		-0.013 (0.023)		0.006* (0.003)		0.013*** (0.004)		0.002 (0.003)
Buyer only		0.009 (0.021)		-0.014 (0.023)		-0.013*** (0.004)		0.006* (0.004)		0.002 (0.003)
R-squared	0.776	0.776	0.731	0.731	0.406	0.406	0.860	0.860	0.922	0.922
Observations	357665	357665	73406	73406	1444674	1444674	357665	357665	357665	357665
No. of Firms	46512	46512	10189	10189	68794	68794	46512	46512	46512	46512
Mean of Dep. in Control	0.638	0.638	0.306	0.306	0.263	0.263	0.755	0.755	0.358	0.358
Link FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pre-treatment dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on subsequent reporting discrepancies. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). To identify outside-link effects, the sample includes all links formed by the sellers of the study sample with all their buyers, excluding the buyer from the study sample link. The table reports the $\hat{\beta}$ coefficients from our difference-in-differences regression (footnote 31) pooling together all post-treatment months, for five outcomes: transaction size conditional on trade in columns 1-2, defined as the maximum amount reported by the buyer or seller; seller shortfall amount conditional on there being seller shortfall in columns 3-4; probability of any trade being reported columns 5-6; conditional on trade, whether it is reported by the seller in columns 7-8, and/or the buyer in columns 9-10. All amounts are in thousands of USD and winsorized at the 0.5% level. Standard errors are clustered at the link level. *p< 0.10; **p< .05; ***p< .01. This table is mentioned in Section 4.2. *Source:* Data from monthly VAT returns submitted to the URA.

F Firm-Level Analysis: Additional Results

This appendix shows additional results for the firm-level analysis discussed in Section 4.3. In the main text, we focus on discussing the results for sellers (Table 3). In Table F1, we show the corresponding results for buyers. Table F2 shows the regression results that correspond to Figure 8. Finally, Tables F3, F5 and F6 show additional results on firm heterogeneity analysis.

F.1 Past VAT Liability (Buyers)

The table below shows that there is no significant effect of the treatment on buyers' VAT liability, potentially because they submit very few corrections.

Table F1
Effect of Corrections on Firm-Level VAT Liability (Buyers)

Dependent variable:	$\Delta B2B$ Sales		Δ Final Sales		Δ Taxable Inputs		Δ VAT Liability	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	-0.016 (0.035)		0.016 (0.013)		-0.013*** (0.004)		0.008* (0.004)	
Buyer only		-0.023 (0.038)		0.006 (0.020)		-0.012** (0.006)		0.006 (0.005)
Seller only		-0.040 (0.037)		0.020 (0.013)		-0.010** (0.004)		0.003 (0.005)
Buyer and Seller		0.011 (0.044)		0.021* (0.012)		-0.017** (0.007)		0.014** (0.007)
R-squared	0.008	0.009	0.004	0.004	0.009	0.009	0.005	0.007
Observations	2736	2736	2736	2736	2736	2736	2736	2736
No. of Firms	1232	1232	1232	1232	1232	1232	1232	1232
Mean of Dep. in Control	45.913	45.913	20.285	20.285	55.434	55.434	1.938	1.938
Mean of Diff. in Control	0.060	0.060	-0.023	-0.023	0.002	0.002	0.005	0.005
Median of Dep. in Control	0.551	0.551	2.602	2.602	11.773	11.773	0.111	0.111
P-value of $\beta_S = \beta_B$		0.565		0.369		0.719		0.570
P-value of $\beta_{SB} = \beta_B$		0.379		0.350		0.649		0.262
P-value of $\beta_{SB} = \beta_S$		0.150		0.976		0.395		0.129
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the effects of the letters on past VAT liability for buyers, estimated using equation (3). The sample includes VAT returns of the buyers for all months mentioned in the letters. The outcome variables are defined as the change in a given entry of the VAT return before and after the treatment, where the changes occur through amendments. Columns 1-2 report results for *B2B Sales* (sales to other VAT firms), columns 3-4 for *Final Sales* (sales to final consumers or non-VAT firms), columns 5-6 for *Taxable Inputs* (purchases from VAT firms) and columns 7-8 for *VAT Liability* (total output tax minus total input tax). *Mean* (resp. *Median*) of *Dep. in Control* reports the average (resp. median) value of each entry for the firms in the control group. *Mean of Diff. in Control* reports the average differences in the values before and after treatment for the firms in the control group. Standard errors are clustered at the seller level. *p < 0.10; **p < .05; ***p < .01. This table is mentioned in Section 4.3. *Source:* Data from VAT monthly returns submitted to the URA.

F.2 Subsequent VAT Liability

Table F2
Post-treatment Effects on Firm-Level VAT Liability (Sellers)

	(1) B2B Sales	(2) Final Sales	(3) Taxable Inputs	(4) VAT Liability
-10	1.430 (5.388)	2.538 (7.434)	5.397 (9.846)	-0.264 (1.203)
-9	9.712* (5.245)	4.286 (7.147)	12.183 (9.929)	-0.031 (0.985)
-8	2.570 (5.187)	1.912 (7.632)	14.134 (9.951)	-0.912 (1.054)
-7	1.861 (5.010)	4.045 (7.181)	9.188 (9.554)	-0.591 (1.080)
-6	-2.721 (5.606)	-1.765 (7.609)	-7.553 (9.578)	0.045 (0.867)
-5	-1.365 (4.489)	6.197 (7.312)	-0.728 (8.861)	0.491 (0.983)
-4	0.902 (4.217)	-8.182 (6.024)	-11.361 (8.269)	0.280 (0.756)
-3	-0.367 (4.163)	-4.645 (6.173)	-8.175 (7.344)	0.145 (0.887)
-2	2.625 (5.869)	-8.688 (7.074)	-8.516 (9.796)	-0.431 (1.037)
-1	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)
0	0.103 (5.050)	4.984 (5.485)	-5.880 (7.804)	0.854 (0.964)
1	5.427 (4.850)	1.331 (6.409)	0.786 (9.040)	0.334 (0.990)
2	3.041 (5.235)	5.115 (6.831)	0.262 (8.904)	1.184 (1.023)
3	3.352 (5.511)	5.500 (6.696)	8.467 (9.410)	1.080 (1.023)
4	7.835 (6.050)	7.022 (7.589)	5.533 (9.853)	0.932 (1.060)
5	7.847 (5.908)	4.417 (6.599)	9.150 (9.292)	-0.499 (1.031)
6	10.957* (6.228)	5.323 (6.141)	6.448 (9.512)	0.911 (1.092)
7	-1.788 (6.500)	4.268 (7.312)	-2.328 (9.862)	0.801 (1.052)
8	3.709 (6.703)	3.295 (7.247)	2.124 (10.948)	0.104 (1.015)
9	5.034 (7.748)	1.787 (6.923)	-2.998 (10.299)	1.479 (1.083)
10	3.510 (8.189)	4.883 (6.867)	5.140 (9.507)	-0.302 (1.276)
R-squared	0.844	0.901	0.875	0.552
Observations	25085	25085	25085	25085
No. of Firms	1235	1235	1235	1235
Mean of Dep. in Control	53.516	106.028	142.974	3.346
Firm FE	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on subsequent VAT liability of sellers. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-month level (April 2017 to December 2018). The table reports the β_j coefficients estimated in an event-study similar to equation (2), but at the firm level, with four different outcomes: amount of B2B sales (sales to other VAT firms) in column 1, final sales (sales to final consumers or non-VAT firms) in column 2, total inputs in column 3, and VAT liability (total output tax minus total input tax) in column 4. All amounts are in thousands of USD and winsorized at the 0.5% level. Standard errors are clustered at the firm level. *p<.10; **p<.05; ***p<.01. The results are shown in Figure 8. This table is mentioned in Section 4.3. *Source:* Data from monthly VAT returns submitted to the URA.

F.3 Heterogeneity

Table F3
Heterogeneity Analysis: Any Correction by the Seller

OLS				Lasso			
Seller Characteristics		Buyer Characteristics		Pair Characteristics		Relevant Characteristics	
Received letter	0.195*** (0.041)	Received letter	0.071* (0.041)	High trading volume	0.053 (0.036)	Seller received letter	0.020
Degree	0.000 (0.000)	Degree	0.001* (0.000)	Initial discrepancy	0.018* (0.009)		
Sector: Agriculture	0.048 (0.159)	Sector: Agriculture	0.077 (0.153)				
Sector: Construction	0.051 (0.077)	Sector: Construction	0.082 (0.063)				
Sector: Manufacturing	0.051 (0.056)	Sector: Manufacturing	0.050 (0.073)				
Sector: Mining	-0.252 (0.313)	Sector: Mining	0.377* (0.202)				
Sector: Retail	0.075 (0.049)	Sector: Retail	0.016 (0.055)				
Sector: Service	0.098** (0.047)	Sector: Service	-0.058 (0.047)				
Ln(Sales)	0.002 (0.010)	Ln(Sales)	-0.008 (0.012)				
Share of negative returns	-0.093** (0.046)	Share of negative returns	0.051 (0.047)				
Existence length	0.000** (0.000)	Existence length	0.000 (0.000)				
FC ratio	0.067 (0.048)	FC ratio	-0.073 (0.045)				
Audited in 2016	-0.066 (0.072)	Audited in 2016	0.040 (0.080)				
R-squared	0.094						
Observations	694					Observations	694

Notes: This table reports results on the heterogeneity in sellers' response to the letter. We estimate an OLS regression (columns 1-3) and a Robust Lasso (column 4) where the outcome variable is an indicator taking value 1 if the seller corrects at least one of the discrepancies of pre-treatment months, regressed on a set of seller, buyer and pair characteristics. The characteristics are computed using VAT returns from 10-months prior to the intervention. The firm-level characteristics are: (i) whether the firm was the recipient of the letter, (ii) degree of connection within the network computed as the number of unique partners, (iii) a categorical variable indicating the sector (with Wholesale as the reference category), (iv) the natural logarithm of sales, (v) the share of monthly VAT returns with negative Value-Added, (vi) the ratio of final sales to total sales, (vii) number of months since firm exists, and (viii) whether the firm was audited in 2016. The pair-level characteristics are: (ix) whether the transaction volume exceeds the mean, and (x) initial value of discrepancy. The Lasso specification only retains one seller characteristic. *p<.10; **p<.05; ***p<.01. This table is mentioned in Section 4.3. *Source:* Data from monthly VAT returns submitted to the URA.

Table F4
Heterogeneity Analysis: Any Correction by the Seller for All Discrepancies

Ind. Var.	Dep. Var: Any correction made by Seller			
	Trading Volume (1)	Degree (2)	FC ratio (3)	Audited in 2016 (4)
Any treatment	0.250*** (0.023)	0.240*** (0.026)	0.234*** (0.026)	0.264*** (0.018)
Above median	0.024 (0.015)	-0.016 (0.014)	-0.008 (0.014)	
Any treatment x Above median	0.019 (0.037)	0.042 (0.036)	0.048 (0.036)	
Audited in 2016				-0.004*** (0.001)
Any treatment x Audited in 2016				-0.012 (0.012)
R-squared	0.112	0.111	0.112	0.112
Observations	1235	1235	1235	1235
No. of Unique Pairs	1235	1235	1235	1235
Mean of Dep. in Control	0.024	0.024	0.024	0.024

Notes: This table reports results on the heterogeneity in sellers' response to the letter. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level, and this table focuses on all discrepancies. The outcome variable is an indicator set to 1 if a discrepancy is reduced by the seller. It is regressed on the treatment dummy, an indicator for whether the link has an above median value for trade volume (column 1), degree (column 2), ratio of final sales (column 3), and for whether the seller was audited in 2016 (column 4), and their interaction. Standard errors clustered at link level. *p<.10; **p<.05; ***p<.01. This table is mentioned in Section 4.3. *Source:* Data from monthly VAT returns submitted to the URA.

Table F5
Heterogeneity Analysis: Any Correction by the Seller for Listed Discrepancies

Ind. Var.	Dep. Var: Any correction made by Seller			
	Trading Volume (1)	Degree (2)	FC ratio (3)	Audited in 2016 (4)
Any treatment	0.250*** (0.023)	0.275*** (0.025)	0.236*** (0.025)	0.266*** (0.019)
Above median	0.024 (0.015)	0.007 (0.014)	-0.024* (0.013)	
Any treatment x Above median	0.019 (0.037)	-0.032 (0.036)	0.056 (0.036)	
Audited in 2016				-0.010*** (0.003)
Any treatment x Audited in 2016				-0.042* (0.024)
R-squared	0.112	0.111	0.111	0.113
Observations	1235	1235	1235	1235
No. of Unique Pairs	1235	1235	1235	1235
Mean of Dep. in Control	0.024	0.024	0.024	0.024

Notes: This table reports results on the heterogeneity in sellers' response to the letter. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level, and this table focuses on listed discrepancies. The outcome variable is an indicator set to 1 if a discrepancy is reduced by the seller. It is regressed on the treatment dummy, an indicator for whether the link has an above median value for trade volume (column 1), degree (column 2), ratio of final sales (column 3), and for whether the seller was audited in 2016 (column 4), and their interaction. Standard errors clustered at link level. *p<.10; **p<.05; ***p<.01. This table is mentioned in Section 4.3. *Source:* Data from monthly VAT returns submitted to the URA.

Table F6
Heterogeneity Analysis: Any Correction by the Seller for Unlisted Discrepancies

Ind. Var.	Dep. Var: Any correction made by Seller			
	Trading Volume (1)	Degree (2)	FC ratio (3)	Audited in 2016 (4)
Any treatment	0.165** (0.066)	0.255*** (0.050)	0.174*** (0.054)	0.272*** (0.030)
Above median	0.012 (0.034)	-0.033 (0.030)	-0.077** (0.038)	
Any treatment x Above median	0.120 (0.073)	0.025 (0.061)	0.144** (0.064)	
Audited in 2016				-0.008*** (0.003)
Any treatment x Audited in 2016				-0.012 (0.021)
R-squared	0.111	0.103	0.109	0.105
Observations	528	528	528	528
No. of Unique Pairs	528	528	528	528
Mean of Dep. in Control	0.040	0.040	0.040	0.040

Notes: This table reports results on the heterogeneity in sellers' response to the letter. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level, and this table focuses on unlisted discrepancies. The outcome variable is an indicator set to 1 if a discrepancy is reduced by the seller. It is regressed on the treatment dummy, an indicator for whether the link has an above median value for trade volume (column 1), degree (column 2), ratio of final sales (column 3), and for whether the seller was audited in 2016 (column 4), and their interaction. Standard errors clustered at link level. *p<.10; **p<.05; ***p<.01. This table is mentioned in Section 4.3. *Source:* Data from monthly VAT returns submitted to the URA.

G Cost Calculations

This appendix discusses details of the cost calculations mentioned in Section 5. There are two major types of costs associated with this experiment: 1) delivery costs of the physical letters 2) the time of the URA staff.

Letter delivery: The first part is calculated by multiplying the price of sending a letter to a given region with the number of letters sent to that region. The total cost of sending all letters amounts to approximately \$3,301.

Staff time: The second is calculated based on lengthy discussion with officers at the URA. We calculate that it took 8 full working days for 2 officers to prepare the 1325 letters for send-off.⁴⁴ Taxpayers got in touch with the URA upon receiving the letter. The time it took to respond to each taxpayer depended on the type of communication. Specifically, when a firm visited the URA it took around 45 minutes, when a firm sent an email or physical letter it took around 25 minutes to respond (since the information provided by the taxpayer in the email often needed to be reconciled with information in the URA's database), and if a firm called it took around 5 minutes. Finally, after firms had responded, DT compared the firms' responses to the information in their database. All in all, responding and reconciling the information from taxpayers took approximately 5 full working days for 2 officers. Assuming a full working day is 8 hours, the total numbers of hours it took to undertake this study was approximately 450. We finally convert this to monetary costs by multiply the number of hours with the average hourly salary, before deductions, for a junior URA officer. We calculate that the monetary cost associated with the number of hours worked is \$2,415.

The total cost incurred by the URA for the intervention is thus \$5,716.⁴⁵

⁴⁴Note that this is larger than the number quoted throughout the paper, which is due to the data error discussed in detail in Appendix B.

⁴⁵In this calculation we have excluded the cost of the time used by the researchers and staff in the Research and Planning Department to identify the discrepancies and seller-buyer links. We choose to do so because this can be automated.