

Miguel Almunia
David Henning
Justine Knebelmann
Dorothy Nakyambadde
Lin Tian

25/51

Working paper

Firm networks and tax compliance: experimental evidence from Uganda

Firm Networks and Tax Compliance: Experimental Evidence from Uganda*

Miguel Almunia
CUNEF Universidad

David Henning
Oxford University

Justine Knebelmann
Sciences Po

Dorothy Nakymbadde
Uganda Revenue Authority

Lin Tian
INSEAD

October 2025

Abstract

How do policy interventions diffuse through firm transaction networks? We design a novel two-stage randomization strategy that assigns a tax enforcement treatment at the seller–buyer link level and ensures separation within the network to identify direct and spillover effects. Using Ugandan transaction-level VAT data, we find that treated links correct 23.8% of reporting discrepancies, fourteen times the control rate. Corrections are driven by sellers—even when only buyers receive letters—providing evidence of communication between firms. Spillovers extend to other transactions, with persistent improvements in post-treatment reporting. Sellers evade by reclassifying firm-to-firm transactions as unverifiable final sales, weakening the VAT’s self-enforcing property.

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

JEL codes: H25, H26, L14, O12.

*Almunia: CUNEF Universidad, CEPR, IGC, IFS, miguel.almunia@cunef.edu. Henning: Oxford & CSAE, david.henning@economics.ox.ac.uk. 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 thank Guillermo Cruces, Antoine Ferey, Pablo Garriga, Jonas Hjort, Isabelle Mejean, Dina Pomeranz, Luca Salvadori, Timothy Van Zandt, Gonzalo Vázquez-Bare, Mazhar Waseem and many seminar participants for useful comments. 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 grants RYC2021-031858-I and PID2023-150638NB-I00. 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). Because firms are embedded in complex transaction networks, tax enforcement measures may affect not only targeted firms but also spread to their trading partners through information flows. Quantifying these spillover effects is essential for designing effective enforcement policies. Yet we know little about how enforcement effects propagate through firm networks, the role of firm-to-firm communication, and what this reveals about firms’ evasion strategies. Answering these questions is difficult because firms’ transaction networks are highly dense, making it hard to disentangle direct from spillover responses.

We study these questions in the context of the Value-Added Tax (VAT), which accounts for about a third of total tax revenue in low-income countries (Keen and Lockwood, 2010; Brockmeyer et al., 2024). The VAT is designed to be self-enforcing: each transaction between firms generates two invoices, giving sellers incentives to underreport sales and buyers incentives to overreport purchases. While there is evidence that this mechanism can improve compliance in higher-capacity settings (Pomeranz, 2015), widespread misreporting persists in low-capacity environments, even in transactions between registered firms (Almunia et al., 2024). In this context, communication between trading partners becomes central to enforcement: information exchanged within networks can either facilitate evasion or magnify the impact of enforcement efforts.

In this paper, we present experimental evidence on how tax enforcement policies spread through firm networks. We leverage transaction-level data from VAT returns to construct a subsample of seller-buyer links that are disconnected from each other, and then randomize treatment across these selected links. This experimental design allows us to disentangle direct from spillover effects, provide evidence of firm-to-firm communication, and identify asymmetric responses between sellers and buyers. In combination, our results show that VAT evasion in this context is driven primarily by sellers, who evade by reclassifying business-to-business (B2B) transactions as unverifiable final sales, thereby undermining the VAT’s self-enforcing property.

We partner with the Uganda Revenue Authority (URA) to implement the randomized intervention. Using transaction-level VAT data from March to December 2017, we map the full network of VAT-registered firms in Uganda, comprising 115,856 unique seller-buyer links. The intervention targets cases where sellers report a smaller amount than buyers, which we label as “seller shortfall.” These reporting discrepancies lower tax liabilities and are often indicative of evasion. Such cases are central to VAT compliance gaps in low-income countries,

where firms often believe revenue authorities cannot effectively cross-check tax returns. To test how changing this belief affects behavior, the URA sent official letters to a randomly selected subset of seller–buyer links. The letters notified firms of the new detection method, listed up to three discrepancies between the firm pair, and asked them to amend past returns.

The experimental design incorporates three innovations to explicitly account for the network structure of the data. First, the unit of randomization is the seller-buyer link rather than the individual firm. Randomizing at the firm level would not allow causal identification in the presence of spillovers, since highly-connected firms are more likely to be exposed to treatment. This would violate SUTVA and conflate direct and spillover effects (Rosenbaum, 2007; Baird et al., 2018).¹ Second, we implement an iterative sampling procedure: once a link is included in the Study Sample, all directly connected eligible links are excluded. Together, these two features ensure that no firm is exposed to multiple treatments, which allows for clean identification of direct and spillover effects of the intervention. The sampling procedure produces a Study Sample with 1,235 seller–buyer links that are separated by at least one degree in the network. Third, we vary treatment intensity across arms: letter sent to seller only (20%), buyer only (20%), or both (20%), with the remaining 40% forming the control group (no letter). With this variation, we can test whether targeting pairs of firms is more effective than treating individual firms and, by comparing who corrects discrepancies, also directly observe how information is shared between partners.

The intervention substantially increases corrections of past reporting discrepancies. In treated links, the correction rate for discrepancies listed in the letters rises by 22.3 percentage points (pp), a fourteen-fold increase over the 1.6% in the control group. Corrections of unlisted discrepancies also increase by 10.6pp. Moreover, the intervention leads to changes in reporting behavior: over the following ten months, seller shortfall drops by 15.0pp (19%) and matching reports increase by 13.0pp (115%). These results suggest that firms update their beliefs about the revenue authority’s monitoring capacity, generalize the enforcement threat beyond the transactions named in the letters, and also internalize it into a sustained improvement in reporting accuracy.

Our experiment also uncovers sharp asymmetries between sellers and buyers, as well as clear evidence of firm-to-firm communication—patterns that are visible thanks to the link-level design. Sellers drive nearly all corrections: in the treatment group, they correct 19.8% of discrepancies, while buyers only correct 2.6%. Strikingly, when only buyers receive letters, corrections by their *seller* partners rise by 7.9pp, while their own corrections only increase by

¹SUTVA, the stable unit-treatment value assumption, requires that each unit’s potential outcomes depend only on its own treatment status and that there is a single, well-defined version of each treatment. The assumption was first formalized in Rubin (1980).

2.9pp. In contrast, when only sellers get the letter, corrections by buyers increase by 1.4pp. These findings suggest that sellers are the primary agents of the type of VAT evasion we study. Supporting this interpretation, descriptive evidence from firms’ communications with the URA shows that buyers attribute discrepancies to their trading partners in 79% of post-intervention contacts with the URA, compared to only 19% for sellers. Moreover, the results highlight the effectiveness of communication to trigger changes in partners’ behavior. Finally, effects are strongest when both sides of a link receive the letter, although the combined effect is no greater than the sum of individual treatments—suggesting that firm-to-firm communication amplifies single-firm interventions rather than diminishing their impact.

We find spillover effects that extend beyond the treated link. Corrections in treated firms’ other trading relationships increase by 2.7pp, or sixfold. These effects are again driven by sellers, and they are strongest when both firms in a link receive a letter, indicating 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, suggesting that the reduced discrepancy reflects improved compliance rather than collusion.

Because tax liabilities are ultimately determined at the firm level, we also examine how link-level responses aggregate. Sellers increase monthly reported sales to other VAT firms by \$512 but reduce reported final sales by \$309, resulting in a modest increase in monthly VAT liability of \$13.² These firm-level findings further clarify the underlying evasion mechanism: tax evasion in this setting is driven primarily by underreporting of sales rather than overreporting of costs. In total, our intervention uncovers \$300,230 in previously unreported B2B transactions, of which 30% is due to spillover effects on transactions that were not listed in the letters. The net revenue gain is smaller once reclassification into final sales is taken into account. Nonetheless, because of the low administrative cost of sending the letters, the intervention is highly cost effective, yielding more than six times its cost in additional VAT revenue.

We contribute to three strands of the literature. Our first contribution lies in the design of a randomized experiment within a firm network, enabling us to distinguish between direct and spillover effects of an enforcement intervention. Being able to make this distinction is essential for designing optimal policies (Miguel and Kremer, 2004), but is increasingly challenging as the density of the network increases (Manski, 1993; Rosenbaum, 2007). While two-stage network experimental designs have been explored, they often focus on social or geographic settings where the network is less dense (Baird et al., 2018; Vazquez-Bare, 2023). Previous

²Consistent with this mechanism, when sellers contacted the URA, their most common explanation was that the transactions in question had been classified as final sales.

experiments have shown the influence of spillovers in firms’ tax compliance behavior, but focusing on other forms of spillovers, such as through tax preparers (Battaglini et al., 2022; Boning et al., 2020; Brown et al., 2025), across tax types (Lopez-Luzuriaga and Scartascini, 2019), and among geographically proximate entities (Drago et al., 2020; Lediga et al., 2023; Cruces et al., 2024). The study closest to ours is Pomeranz (2015), which shows that audit threats to Chilean VAT firms affected their suppliers. Because treatment in that study was assigned at the firm level, however, it is difficult to separate direct effects from indirect exposure. Mapping networks *ex ante*, randomizing at the seller–buyer link level, and varying which trading partner is treated allows us to both disentangle direct from spillover effects and identify the extent of firm-to-firm communication.

Second, we contribute to the literature on tax enforcement and firms’ evasion behavior by highlighting the potential of leveraging transaction networks to design more effective enforcement measures. While some interventions have successfully deterred tax evasion (e.g., Shimeles et al., 2017; Holz et al., 2023), growing evidence suggests that firms often offset enforcement by adjusting along untargeted margins (Carrillo et al., 2017; Slemrod et al., 2017; Mascagni et al., 2021; Best et al., 2021; Hoy et al., 2022). Our results align with this mixed picture: the intervention only modestly increases reported tax liabilities for past returns (reflecting firms’ ability to reclassify sales) but has no significant impact on subsequent VAT revenues. We extend this literature by pointing to sales underreporting, rather than cost inflation, as a key evasion driver, and by showing that final sales play a central role in dampening enforcement gains. This mechanism is related but distinct from the “last mile” problem (e.g., Naritomi, 2019): rather than simply underreporting final sales, firms reclassify B2B transactions as final sales, exploiting the unverifiable nature of this margin. This helps explain both the asymmetric, seller-driven responses we observe and reveals how network structure interacts with the VAT’s self-enforcing design. Together, our results allow us to pin down how firms evade VAT in a low-income economy like Uganda. The URA’s subsequent adoption of cross-checking and enforcement messages from our design underscores its policy relevance.

Finally, our paper is related to the literature on transaction networks. A large literature outside taxation highlights their importance in shaping firm outcomes (Huneus et al., 2021; Alfaro-Ureña et al., 2022; Dhyne et al., 2022; Adão et al., 2022; Bernard et al., 2022; Demir et al., 2022; Cai et al., 2024). Within taxation, non-experimental work shows that tax enforcement interventions can also affect firms’ trading partners (Lediga et al., 2023; Carrillo et al., 2023; Garriga and Tortarolo, 2024) but, without experimental variation, they cannot identify direct and spillover effects as rigorously. We advance this literature by implementing, to the best of our knowledge, the first two-stage randomized experiment studying spillovers

through firm-to-firm networks. Our results provide experimental evidence that transaction networks play an important role in tax enforcement, as firms communicate and transmit information through these transaction links.

The rest of the paper is organized as follows. Section 2 provides the context and describes the data. Section 3 presents a conceptual framework to guide the empirical analysis and outlines the experimental design. Sections 4 and 5 present the link-level and firm-level results, respectively, followed by a discussion of their implications and relation to existing evidence in Section 6. Section 7 concludes.

2 Background

In this section, we first describe the institutional context of the VAT in Uganda, and then present the data used to map firm networks and detect reporting discrepancies.

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.³ Firms with annual turnover above 150 million UGX (\$45,000)⁴ are required to register for VAT, while smaller firms can choose to pay a simplified turnover tax.⁵

VAT-registered firms must 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

³For 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.

⁴The average exchange rate between Ugandan Shillings (UGX) and US dollars (USD) was 3,300UGX per USD in 2018. We use this exchange rate throughout the paper.

⁵In 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.

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.⁶ 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 sales and purchases with other VAT firms. This system is designed such that the URA receives two reports for each B2B transaction, one from the seller and one from the buyer.⁷ This dual-reporting mechanism enables us to map firm-to-firm trading networks and detect potential misreporting. In contrast, sales to non VAT-registered firms or to final consumers—henceforth referred to as “final sales”—are reported only as aggregate totals, with no corresponding third-party verification. This asymmetry in reporting and verifiability between B2B and final sales plays a central role in our analysis, as discussed below.

2.2 Data

Our data covers the universe of monthly VAT returns filed by Ugandan firms in the period between March 2017 and December 2018. A firm’s 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, final sales, 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.⁸ For each entry, firms report the TIN of the trading partner, the amount and the date of the transaction, whereas final sales 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 final sales, while input VAT matches in 99% of declarations. The internal consistency of the returns confirms that the transaction-level records provide a reliable account of firms’ reported VAT declarations, though not necessarily

⁶When 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).

⁷This reporting requirement is stricter than what is commonly observed in advanced countries, where transaction-level information is usually only requested during tax audits.

⁸For more details on how these transactions are reported, see Almunia et al. (2024).

their true liabilities.

Seller-buyer links and networks. We define two firms as forming a link in a given month if both firms are VAT-registered, and 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 transactions to the link-month level. During our pre-intervention period, March to December 2017, there are 21,548 unique firms that submit a VAT return and 115,856 distinct seller-buyer links. We construct each firm’s network based on its directly reported clients and suppliers. These VAT-registered firms have an average (respectively, median) of 10.51 (2) unique trading partners over the year, corresponding to 5.38 (0) unique clients and 5.38 (1) unique suppliers.

Final sales. Among VAT-registered firms, 75.8% report at least some final sales and, for 53.7% of these, final sales account for over half of their total reported sales during the pre-intervention period. This suggests that the majority of firms in the economy have the scope to engage in VAT evasion by misclassifying B2B sales as final sales, for which no third-party verification is available.

Amendments of past returns. To submit an amendment, a firm must re-file the VAT return it had previously submitted for that reporting period, marking the new filing as an amendment by ticking the designated box. 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. Amendments are relatively rare: of all monthly VAT returns filed in 2016, about 10% were amended in the subsequent 12 months.

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. We document important limitations to VAT performance, potentially attributed in part to lack of technological capabilities and qualified staff.

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 pre-intervention period, we find widespread discrepancies. Specifically, we observe seller shortfall in 41.4% of the link-month observations.⁹ More than 92% of these are extensive-margin discrepancies, meaning that only one firm in the link reported trading with the other firm.

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

Against this backdrop, the objective of the letter intervention, designed in collaboration with the URA, is to reduce VAT evasion, by targeting seller shortfall discrepancies and leveraging the information on the transaction network.

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 us to identify the direct and spillover effects of the intervention and disentangle the asymmetry in the VAT reporting, before describing the implementation of the experiment.

3.1 A Conceptual Framework

We define a link as a seller-buyer pair, in which at least one party reports trading with the other. Let $\hat{s}_{jk}$ denote the sales reported by firm j to the URA on transactions with their client k , $\hat{c}_{kj}$ denote the purchases reported by firm k on transactions 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., $\hat{s}_{jk} < \hat{c}_{kj}$.¹⁰ While seller shortfall reduces the total tax remitted to the URA, it is not clear ex ante which of the two parties of

⁹We observe buyer shortfall in 51.6% 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](#)).

¹⁰Conversely, 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 $\hat{s}_{jk} = \hat{c}_{kj}$.

the transaction is liable for the discrepancy.¹¹

The prevalence of seller shortfall suggests that firms perceive the risk of detection to be low. Our letter intervention aims to increase the perceived risk by notifying firms of specific discrepancies (the *listed* discrepancies) identified by the administration and warning of potential penalties if tax evasion is confirmed. Revealing the URA’s detection capacity is therefore expected to change firms’ beliefs and motivate corrective 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 seller shortfall discrepancies *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 after treatment.

Predictions 1a and 1b focus on the contemporary effects of the letter. Prediction 1a reflects the direct impact on listed discrepancies, while Prediction 1b captures firms’ learning about the URA’s enhanced detection capacity, leading them to correct unlisted discrepancies as well. Prediction 1c tests if this shift in beliefs is persistent, resulting in improved compliance in subsequent reporting.

Asymmetry. The first set of predictions treats sellers and buyers symmetrically, but this breaks down once final sales are considered. Suppose firm j sells both to VAT-registered firms and final consumers, with total sales being $s_j \equiv \sum_k s_{jk} + s_j^F$, where s_j^F denotes final sales. If the firm underreports sales ($\hat{s}_j < s_j$), and the tax authority cross-checks only a subset of its transactions, the firm can plausibly deny evasion by claiming the missing transactions had been reported as final sales. By relabeling some final sales as sales to VAT-registered firms, it can keep overall sales, and thereby tax liability, constant. The tax authority can unambiguously detect evasion only if total purchases declared by firm j ’s trading partners exceed the firm’s declared total sales, i.e., $\sum_i \hat{c}_{ij} > \hat{s}_j^F + \sum_k \hat{s}_{jk}$. This level of verification requires substantial administrative and computational resources, which are often scarce in low-income settings like Uganda. By contrast, on the input side, no such plausible deniability

¹¹Seller shortfall can occur if the seller underreports its sales ($\hat{s}_{jk} < s_{jk} = c_{kj}$), the buyer overreports its purchases ($\hat{c}_{kj} > c_{kj} = s_{jk}$), or if both happen simultaneously.

exists: if declared purchases exceed reported sales for a single link, i.e., $\hat{c}_{kj} > \hat{s}_{jk}$, this immediately flags potential evasion attempt.

This difference in the risk of detection gives sellers stronger incentives to engage in seller shortfall than buyers. Consequently, the letter intervention may prompt asymmetric responses between buyers and sellers, leading to a second 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. Buyers receiving letters are more likely than sellers to communicate with trading partners to induce corrections.

Untreated Links. 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.

Our predictions thus far focus on responses at the seller-buyer link level. Therefore, they do not address potential effects on the overall VAT liability of individual firms. We return to the discussion of firm-level VAT liability in Section 5.

3.2 Two-Stage Experimental Design

In order to empirically test the predictions from the previous section, we leverage the structure of the firm network to design a novel network experiment. Our experiment involves two stages: sample selection and randomization.

3.2.1 Sample Selection

Our first innovation is to define the unit of analysis at the link rather than the firm level. We do this for two reasons. First, randomizing at the firm level would potentially violate

SUTVA: in the presence of spillover effects, more connected firms would be more exposed to a randomly-assigned treatment (Rosenbaum, 2007). Second, seller shortfall discrepancies are observed at link level, making it the natural unit of study.

However, a “naive” randomization even at the link level still creates the risk that the same firm, whether as seller or buyer, is assigned to multiple treatment arms through different trading relationships. 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.

Base Sample. In consultation with the URA, we select a time window of interest, March to December 2017. During that time window VAT-registered firms formed a total of 115,856 seller-buyer links. We define a Base Sample of links with seller shortfall discrepancies that are suggestive of potential tax evasion. Links in the Base Sample meet the following criteria within our time window: both firms in the link are registered for VAT, filed a VAT return in December 2017, and the link had an accumulated seller shortfall greater than 1 million UGX (\$303).¹² With these filters, the Base Sample consists of 11,036 seller-buyer links and 4,514 unique firms. For brevity, we will henceforth refer to the seller-buyer links meeting these criteria as “seller-shortfall links,” and to all excluded links—including buyer shortfalls, matched cases, and seller shortfalls below the threshold—as “non-seller-shortfall” links.


Study Sample. We draw our Study Sample from the Base Sample, applying an iterative sample selection procedure. Figure 1 illustrates the iterative process. Panel (a) represents an example of a firm network of targeted discrepancies: each circle represents a firm and 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 seller-shortfall links of that seller and buyer with other firms, and remove them from the pool of eligible links. These are the red dotted lines in panel (b). Third, we randomly select another link from the remaining sample of eligible links, as shown in panel (c). Similarly, we remove all remaining seller-shortfall 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 2.¹³ Firms in the Study Sample are more connected than the typical VAT-registered firm

¹²Additionally, 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.

¹³A data extraction error duplicated all input transactions filed between 19/9/2017 and 15/10/2017, mechanically inflating seller shortfall cases. As a result, 28% of links in the original Study Sample did not actually meet our discrepancy criteria. We correct this by excluding those links from the analysis, as detailed in Appendix A, and show that the correction is balanced across treatment arms, making bias unlikely.

in Uganda, with an average of 55.8 distinct trading partners, compared to 10.5 in the full sample.

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. 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 procedure is repeated among the remaining eligible links, until no links are left. Panel (c) displays the resulting Study Sample.


3.2.2 Treatment Assignment

Using the Study Sample, we implement a stratified randomization to allocate seller-buyer links into three treatment arms and a control group. Our link-level design is consistent with the two-stage experimental frameworks used in the literature to study spillover effects (Duflo and Saez, 2003; Miguel and Kremer, 2004; Baird et al., 2018; Cruces et al., 2024). In those settings, researchers randomly assign clusters to different treatment intensities, and then vary which individuals are treated within clusters. In our context, the cluster corresponds to a link with seller-shortfall.

Treatment. The treatment letters notify firms that the URA has developed new analytical methods to detect discrepancies in VAT declarations. Each letter lists the three largest instances of monthly seller shortfall discrepancies between the recipient firm and its partner.¹⁴ A template is displayed in Appendix Figure B1. The letter instructs the firm to resolve

¹⁴The final number of discrepancies listed on the letter might be smaller than three if there are only one or two occurrences of seller shortfall for that link in the pre-intervention period. We use the same criteria to define counterfactual listed discrepancies for the control group.

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, while the Study Sample refers to the links included in the study after applying the iterative sample selection procedure describe in Section 3.2.1. 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.

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 credibility.

Treatment Arms. We divide our treatment group into three arms, as shown in Figure 2. In the “Buyer only” arm, a letter is sent exclusively to the buyer, in the “Seller only” arm exclusively to the seller, and in the “Seller and Buyer” to both. Out of the 1,235 links in the Study Sample, 741 links are divided up equally across these three treatment arms (20% each). The control group consists of the 494 remaining links (40%), which receive no letter

at all. The randomization is stratified by three seller characteristics: 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. outside).

Our iterative sampling procedure ensures that all links in the Study Sample are at least one degree apart within the network, and the subsequent randomization assignment varies treatment intensity at the link level. Together, this two-stage experimental design allows us to isolate direct effects from spillover effects while addressing potential SUTVA concerns. A key advantage of this design—and of having full transaction-level data—is that we can empirically assess, in Section 4.1.4, identification assumptions that are typically taken as given in studies of network spillovers.

Additionally, our design allows us to examine distinct dimensions of firms’ reporting responses. We observe which firm in a link corrects the discrepancy in the VAT data. By comparing corrections between the two sides of the transaction, we can examine whether there is asymmetry in the responses from buyers and sellers (Prediction 2a). Moreover, we are able to detect firm-to-firm communication: any adjustment by the buyers in the “Seller only” arm, or by the sellers in the “Buyer only” arm, reveals information sharing between firms (Prediction 2b). The “Seller and Buyer” arm captures effects of contacting both firms. While joint treatment could encourage partners to collude by moving transactions off the books, it could also amplify compliance by prompting corrections from both sides, and/or reducing scope for putting the blame on the other firm.

3.3 Experimental Implementation

This section provides details on the implementation and discusses potential issues associated with the randomized experiment, including balance and attrition.

3.3.1 Timeline

The timeline of the experiment is 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 B1 and Figures B2 and B3 in the Appendix provide qualitative information on what firms communicated to the URA about 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. We discuss results from the communication tracker in Section 4.1.

3.3.2 Balance

Table 1 reports averages of baseline observable characteristics for the control group and the coefficients from regressions of those characteristics on treatment indicators, both for the whole treatment group and each treatment arm separately. For the majority of baseline characteristics, there is no statistically significant difference between treatment arms and the control group at the 10% level (and none are significantly different at the 5% significance level). Because we estimate 72 regressions—18 for each of the three treatment arms plus 18 for treatment group as a whole—some marginally significant differences are expected by chance. Consistent with this, the p-value for an F-test of joint significance of all baseline covariates comparing treated firms against control firms is 0.67.

3.3.3 Attrition

Attrition is a common concern in randomized experiments, particularly for enforcement interventions that might induce firms to exit the tax system. Reassuringly, in our case, the attrition rates in treatment and control groups are very similar, as shown in Appendix Figure B4. After 10 months, 95.9% of treatment firms and 95.2% of control firms were still filing VAT returns.

3.3.4 Additional links

Links in the Study Sample are separated by at least one degree, meaning they are not directly connected through seller-shortfall relationships. They may, however, still be connected to each other via non-seller-shortfall instances, or indirectly through second-degree links. Section 4.1.4 discusses this in detail and shows that it does not affect our results.

Table 1
Study Sample: 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	4.77 (3.25)	-0.34 [0.19]	-0.51 [0.26]	-0.20 [0.26]	-0.32 [0.24]
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)	1.58 [3.79]	-1.06 [5.51]	3.04 [4.81]	2.70 [4.77]
Total Output	24.48 (56.52)	2.41 [3.13]	0.45 [4.51]	3.32 [4.24]	3.38 [4.15]
Share of final sales	0.56 (0.41)	0.00 [0.02]	0.03 [0.03]	-0.06* [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
P-value of joint F-test		0.67	0.49	0.71	0.55

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 final sales plus sales to non-VAT firms over total sales. All monetary values are in thousands of USD. *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 seller-buyer links in two parts, testing our predictions from Section 3.1. First, we study how the treatment affects the probability that past discrepancies are corrected. Second, we analyze the impact of the treatment on subsequent reporting discrepancies in the post-treatment period, up until ten months after the letters were sent. Specifically, this section focuses on how treatment affects reporting within seller-buyer links. We turn to firm-level outcomes, including the impact on overall tax liabilities, in Section 5.

4.1 Impact on the Correction of Past Reporting Discrepancies

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

$$Y_{it} = \alpha + \sum_{h \in \{S, B, SB\}} \beta_h T_{ih} + \delta_t + \epsilon_{it}, \quad (1)$$

where Y_{it} is the outcome of interest in month t of seller-buyer link i . 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 of each treatment arm. Observations are at the link-month level, and standard errors are clustered at the link level.

To estimate the overall effect of the intervention, we report the average treatment effect across all three treatment arms, obtained from the following regression:

$$Y_{it} = \alpha + \beta_T T_i + \delta_t + \xi_{it}, \quad (2)$$

where β_T is the coefficient of interest. T_i indicates any treatment, and standard errors ξ_{it} are clustered at the link level.

4.1.1 Effect on corrections in treated links

We start by analyzing corrections induced by the intervention to test Predictions 1a and 1b. In panel (a) of Figure 3, the outcome variable is a dummy variable indicating whether the discrepancy listed in the letter 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. In panel (b), the outcome variable is defined analogously, but for discrepancies *not* listed in the letter.

In each panel of Figure 3, the bars indicate the share of discrepancies corrected, in

the control group (top bar), in each treatment arm (second to fourth bar), and across all treatment arms (bottom bar). The β_h coefficients and their standard errors, shown to the right of each bar, are estimated from regression equations (1) and (2). As shown in the bottom bar of panel (a), the share of corrected discrepancies for listed months is 22.3pp higher (significant at the 1% level) for any treated link. Given that the average in the control group is 1.6%, this corresponds to a fourteen-fold increase in the correction rate. Breaking it down by treatment arm, we find that the effect is largest (30.3pp) for the “Seller and Buyer” treatment, followed by the “Seller only” treatment (24.3pp), and smallest for the “Buyer only” treatment (11.6pp). The estimates are statistically significant at the 1% significance level in all treatment arms. Appendix Table C1 displays the full regression results and shows that the “Buyer only” treatment is significantly smaller than both the “Seller and Buyer” and “Seller only” treatment (1% level), while the latter two do not differ significantly.¹⁵


In panel (b) of Figure 3, because we focus on discrepancies not listed in the letter, 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, 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.¹⁶ This shows that firms extend what they learned about the URA’s detection capacity to discrepancies within the link that were not listed in the letter.

These regression results provide a lower bound on the extent to which firms react to the letter, as revealed by additional evidence collected during the intervention. In collaboration with URA officers, we designed a tool to keep track of all communications between treated firms and the administration. Results are displayed in Appendix Figure B2. When taking into account all forms of interactions (submitting a correction, contacting the URA whether by email, in person or phone), we find that 56% of firms respond to the treatment. The reaction rate is highest in the “Seller and Buyer” arm (70%). While firms in the “Buyer only” arm are more likely to contact the URA than those in the “Seller only” arm (38% vs 31%), they are less likely to make corrections, thus the overall reaction rates are similar across these

¹⁵In Appendix Table C1 we also present the results from an alternative specification, where 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). 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. Figure C1 shows how the distribution of reporting discrepancies shifts to the right after corrections in the treatment group, but not in the control.

¹⁶Columns 7 and 8 of Appendix Table C1 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 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. τ_T is the coefficient for the aggregate effect of any treatment from regression (2). 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 C1. *Source:* Data from monthly VAT returns submitted to the URA.

two arms (50% and 49%, respectively).¹⁷ This also hints at asymmetric reactions, which we turn to in the next section.

An extensive heterogeneity analysis does not reveal strong differences across most baseline characteristics. We return to a discussion of heterogeneous effects in Section 5.2.

4.1.2 Asymmetric effects for sellers and buyers

A core prediction highlighted by the conceptual framework is that sellers are more likely to respond than buyers (Predictions 2a and 2b). To test for this, we examine whether it is the seller or the buyer in a treated link who corrects the discrepancies. Results are displayed in Figure 4. We estimate regressions (1) and (2), but the outcome variable is now a dummy taking value one if the discrepancy is corrected by the seller (top panel) or by the buyer (bottom panel). We focus on listed discrepancies.

The correction rate is systematically higher for sellers. When both the seller and buyer are treated, the correction rate attributable to sellers is 24.9pp larger than the control group (top panel), compared to only 3.3pp for buyers (bottom panel). The same pattern applies when only sellers receive letters: 22.3pp increase for sellers vs. 1.4pp for buyers.

Notably, sellers respond more even when they do not receive a letter. They are 7.9pp more likely to correct when only their buyer receives the letter, a sixfold increase relative to the control group. In contrast, the buyer correction rate in that treatment arm increases by only 2.9pp.¹⁸


Two novel findings emerge from this set of results. First, because links were isolated in the transaction network, the strong effect of the “Buyer only” treatment on sellers’ behavior can only be attributed to communication between trading partners (Prediction 2b). This provides unique evidence of information flows between firms, which has been challenging to capture in empirical research. This result also illustrates that firms can influence their trading partners’ tax reporting even in ways that may not align with the partners’ immediate interests.

Second, in line with Prediction 2a, the stark asymmetry in firms’ responses suggests that sellers are more likely to be responsible for the discrepancies. This is in line with our conceptual framework, which highlights the role played by final sales. Among firms that contact the URA, 79% of buyers in the “Buyer only” treatment say that the discrepancy is due to their trading partner. In contrast, the most common feedback by sellers of the “Seller

¹⁷These statistics offer suggestive evidence rather than causal effects, since equivalent figures are unavailable for control firms. However, because such interactions with the URA are rare for this subset of firms, they still provide meaningful insights into the extent of firms’ reactions to the treatment.

¹⁸Full regression results are reported in Appendix Table C2.

Figure 4
Within-Link Asymmetric 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. τ_T is the coefficient for the aggregate effect of any treatment from regression (2). 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. * $p < .10$; ** $p < .05$; *** $p < .01$. All regression results are shown in Appendix Table C2. *Source:* Data from monthly VAT returns submitted to the URA.

only” treatment is that the transaction had been erroneously classified as final sales (29%) and that they will amend the return (30%).¹⁹

Finally, whether it be for corrections by the seller or by the buyer, the effect is stronger when both firms in the link are treated. The difference with effect for the “Seller only” arm is not statistically significant in either case. Moreover, the magnitude of the effect is smaller than the sum of the “Seller only” and “Buyer only” effects.

4.1.3 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. Again, we estimate regressions (1) and (2), with the outcome variable defined as a dummy variable taking value one when the discrepancy is reduced by the seller. We focus on sellers, since our results from Figure 4 show that corrections are primarily made by sellers. The sample consists of seller shortfall discrepancies of the sellers from the Study Sample with all their buyers, excluding the buyer mentioned in the letter. To facilitate comparison to the within-link effects, we further restrict the sample to the months of the discrepancies listed in the letters. This leads to 11,530 link-month observations between March and December 2017.²⁰

We find a 2.7pp (sixfold) increase in the outside-link correction rate across all treatment arms, significant at the 5% level. 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 both small in magnitude and not statistically significant.²¹


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 C1), 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 C4). In other words, it is crucial that both trading partners receive the letter for the spillover effects on untreated links to materialize. One possible explanation is that, when firms realize that their trading partner also received a similar letter, they infer that the new cross-checking system is more

¹⁹See Appendix Figure B3 for a detailed summary of firms’ communications with the URA in response to the treatment.

²⁰We show regression results for all months in Appendix Table C3, where the sample includes 12,783 observations.

²¹The full regression results are shown in Appendix Table C4. In columns 3 and 4, we show results for unlisted months. The effects are of similar magnitudes (a 3.2pp increase), but only significant at the 10% level. Including all months leads to less precise estimates (Appendix Table C3); the 1.0pp increase is significant at the 10% level only when full corrections are used as the outcome.

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. τ_T is the coefficient for the aggregate effect of any treatment from regression (2). Standard errors are clustered at the seller level. * $p < 0.10$; ** $p < .05$; *** $p < .01$. All regression results are shown in Appendix Table C4. *Source:* Data from monthly VAT returns submitted to the URA.

widespread than when only one of the firms in the link receives it. Accounting for these spillover effects to transactions outside the link may change our assessment of the most effective treatment. We return to this in Section 6.

Taken together, the results on outside-link effects show that the treatment leads some sellers to expect that the revenue authority will cross-check discrepancies beyond the particular partners indicated in the letter.

4.1.4 Robustness checks

Our identification of the link-level results relies on two assumptions: first, that spillover effects do not operate through non-seller-shortfall links; and second, that spillovers do not extend beyond one degree of separation in the network. If spillovers beyond those presented in Figure 5 meaningfully affected link-level responses, but are not controlled for, this could

bias our estimates. Thanks to the richness of our network data, we can control for such potential spillovers and assess whether this affects our estimates.

More precisely, we conduct two robustness checks. First, we test whether spillovers operate through non-seller-shortfall links. While our sampling design ensures that no two links in the Study Sample are directly connected via seller shortfall instances, some links may still be connected through other types of transaction relationships. Second, we examine whether links respond to treatment through second-degree connections. Our two-stage randomization strategy ensures that if a given link ij is treated, any directly connected link jk (its first-degree connection) is not treated. However, higher-order connections are not removed, as doing so would make it infeasible to construct a sufficiently large sample in the dense network.²² We therefore test whether a treated link ij reacts indirectly to treatment received by another link kl , which is connected only through an intermediary firm k , making kl a second-degree connection relative to ij . An illustration of this structure is provided in Appendix Figure D1.

Appendix D provides details on the specifications used for these robustness checks and presents the results. We re-estimate the average treatment effect following (2), including controls for treated and untreated links' exposure to the treatment through non-seller-shortfall links (Table D1), and through second-degree connections (Table D2). This approach to account for spillovers follows recent advances in the econometric literature (see, e.g., [Vazquez-Bare, 2023](#)).²³ In both cases, we find that the estimated treatment effect remains stable and statistically unchanged when these controls are included. Furthermore, the coefficient estimates on the indirect exposure variables are jointly insignificant in both cases.

This provides empirical support for the robustness of our estimates identified through the iterative sampling procedure and suggests that our definition of a firm-pair cluster is appropriate.²⁴ It is also worth noting that previous studies on the causal effects of enforcement interventions typically assume that such indirect effects are absent, but are rarely able to empirically test this assumption.

²²The sellers (resp., buyers) in our Study Sample have an average of 60.5 (resp., 46.4) unique trading partners, while previous randomized trials involving networks are typically implemented in settings where fewer connections are considered. As an illustration [Vazquez-Bare \(2023\)](#) considers spillovers for up to five peers in the empirical application.

²³Our regression is analogous to the correctly-specified interacted reduced form-linear in means regression discussed by [Vazquez-Bare \(2023\)](#).

²⁴While our network construction uses all available VAT transaction data between registered firms, some trading relationships may be entirely absent from the data if neither party reports the transaction. For these links, a direct test is, by definition, impossible.

4.2 Impact on Post-treatment Reporting Behavior

So far, we have shown that firms correct past discrepancies, both with the trading partner mentioned in the letter, and with other trading partners. While this indicates that firms are updating their expectations about the URA’s enforcement capacity, it is not clear whether this is a lasting effect. We next examine whether the letter leads to changes in firms’ reporting behavior in the months following the treatment.

4.2.1 Effect on post-treatment reporting in treated links

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

$$Y_{it} = \sum_{\substack{j=-10 \\ j \neq -1}}^{10} \tau_j (m_j \cdot D_i) + \delta_t + \gamma_i + \epsilon_{it}, \quad (3)$$

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 τ_j ’s. The analysis period runs between April 2017 and December 2018, with January 2018 taken as the reference month.²⁵ The sample includes all monthly trades reported by a given seller-buyer link. We also implement the analogous difference-in-differences (DiD) estimation to discuss the average treatment effect on the treated across all post-treatment months. Our DiD specification is formally written as:


$$Y_{it} = \tau (D_i \cdot \mathbb{1}_{j>0}) + \delta_t + \gamma_i + \sum_{j=-10}^{-1} \tau_j (m_j \cdot D_i) + \epsilon_{it}, \quad (4)$$

where D_i is an indicator taking value one for treated links. This is interacted with $\mathbb{1}_{j>0}$, a dummy taking value one for all post-treatment months. The coefficient estimate for τ reveals the cumulative treatment effect for the entire post-treatment period between March and December 2018. Since we are conducting an event-study analysis for a randomized intervention, we follow [Roth and Sant’Anna \(2023\)](#) and include separate controls for each of the pre-treatment periods to gain precision.

Figure 6 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

²⁵The 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.

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 $\hat{\tau}_j$ coefficients estimated in the event-study laid out in equation (3), 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 $\hat{\tau}$ coefficients from our difference-in-differences regression (equation 4) pooling together all post-treatment months. In the bottom left corner, we report the means in the control group for the pre-treatment period and the number of observations. 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.


reports (i.e., no discrepancy). Observations are conditional on trade being reported between the two firms, so the underlying sample is an unbalanced panel. In addition to the event-study coefficients from (3), we also report the DiD coefficient from (4) at the bottom right 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.²⁶ This is in line with Prediction 1c. The decline is accompanied by a 13.0pp (115%) increase in the probability of matching reports. The buyer

²⁶While the last pre-treatment coefficient on seller shortfall is statistically different from zero, Appendix Figure E1 shows that this reflects a change in the control group rather than the treatment group, ruling out anticipation effects.

shortfall rate rises slightly, but not significantly. Taken together, these results indicate that firms update their beliefs about the tax authority's monitoring capacity and adjust their reporting within treated links accordingly.

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 (3), 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 (4) pooling together all post-treatment months. In the bottom left corner, we report the means in the control group for the pre-treatment period and the number of observations. Standard errors are clustered at the link level and the bars report 95% confidence intervals. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. 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.

Figure 7 provides additional evidence on the mechanisms driving post-treatment behavior. Panel (a) shows that the seller shortfall amount (conditional on being positive) does not change significantly, suggesting that the effect operates mainly through a higher probability of reporting transactions with the trading partner rather than adjustments in reported amounts.

One potential concern is that this improved reporting behavior masks collusion between firms. If the seller and buyer agree to report lower amounts following the intervention, this

would show up as a “match” in the data. We provide evidence that this is not the case. Panel (b) of Figure 7 shows that, conditional on trading, the intervention does not change the size of reported transactions within links.²⁷ Finally, panel (c) shows that there is no change in the probability of reporting transactions. These results indicate that the intervention did not disrupt trading relationships in a significant way, nor push firms to start trading “under the radar.”

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. Panel (d) of Figure 7 shows that, conditional on trade being reported within the link, there is a 17.6pp (62%) higher probability of the seller reporting a transaction, with no significant change for the buyer. This again underscores the asymmetry underlying VAT evasion, outlined in earlier sections. It also points to a more optimistic reading: a simple letter intervention led to sustained improvements in reporting behavior within the link that was mentioned in the letter.

To reconcile the effects reported in panels (c) and (d) of Figure 7 and the corresponding panels in Figure 6, recall that we consider “trade being reported” when at least one firm in the link reports a given transaction. Panel (d) of Figure 7 shows that the decrease in seller shortfall shown in Figure 6 is entirely driven by a higher reporting rate among sellers. The null effects on the probability of trade being reported are thus driven by the lack of extensive-margin response on the buyer side. In other words, consistently with our theoretical framework buyers report transactions at a stable rate, while sellers increase their reporting rate as a result of the intervention.

4.2.2 Effect on post-treatment 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. We do not find significant effects on these outside-link reporting discrepancies subsequent to the treatment.²⁸ Hence, while firms update their beliefs about the URA’s ability to cross check transactions with the trading partner mentioned in the letter, this does not translate into sustained adjustments of their reporting behavior with other firms.

²⁷Transaction size is defined as the maximum amount reported by the buyer or the seller.

²⁸Results are shown in Appendix Figure E2 and Appendix Tables E4 and E5.

5 Firm-Level Effects

In this section, we consider the effects of the intervention on outcomes at the firm, instead of link, level. Particularly, we focus on analyzing 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.

5.1 Firm-Level Conceptual Predictions

Our conceptual framework in Section 3.1 aggregates naturally to the firm level, since the effects identified at the link level may also impact a firm’s overall reported VAT liability. Additionally, once we shift from links to firms, additional firm-level adjustments may come into play, as firms can relabel their reporting across types of sales. We therefore formulate an additional set of predictions, focusing on firms treated as sellers in the Study Sample.

Set of Predictions 4. Firm-level Effects.

- 4a.** Since the tax administration cannot cross-check final sales (s_j^F), sellers may relabel part of s_j^F as B2B sales s_{jk} , dampening the increase in total reported sales s_j .
- 4b.** The reductions in seller shortfall through the correction of past returns (Predictions 1a and 1b) aggregate to higher total B2B sales and, as a result, higher VAT liabilities for treated sellers, in spite of some relabeling (Prediction 4a).
- 4c.** The reductions in seller shortfall in months following the treatment (Prediction 1c) aggregate to higher total B2B sales and, as a result, higher VAT liabilities for treated sellers.

5.2 Firm-Level Results

With these predictions, we now turn to the firm-level results. We first examine the effects of the treatment on firms’ past VAT liabilities, then on their subsequent liabilities, and finally explore heterogeneity across firms to assess which characteristics drive these effects.

5.2.1 Effect of treatment on VAT liability

To analyze how the information treatment changes the amounts reported in the monthly VAT returns prior to treatment, we estimate the following regression:

$$Y_{jt} = \alpha + \sum_{h=S,B,SB}^3 \beta_h T_{jh} + \delta_t + \epsilon_{jt}, \quad (5)$$

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 treatment arms, δ_t is a month fixed effect, and ϵ_{jt} is the error term. Note that this is analogous to regression (1), but in this case the unit of observation is a firm-month instead of a link-month.

Similar to Section 4, to estimate the overall effect of the intervention, we report the average treatment effect across all three treatment arms, obtained from the following regression:

$$Y_{jt} = \alpha + \beta T_j + \delta_t + \xi_{jt}, \quad (6)$$

where β is the coefficient of interest and standard errors ξ_{it} are clustered at the firm level.

Table 2 reports the results for four outcomes: changes in B2B sales, in final sales, in taxable inputs, and in the overall VAT liability.²⁹ The sample is restricted to sellers’ outcomes in the months listed in the letters. All amounts are in thousands of US dollars. In columns 1 and 2, we look at the change in reported B2B sales. We find that, across all treatment arms, the letter led to an average increase of \$512 in reported B2B sales per month, consistent with Prediction 4b. Breaking the result down by treatment arm, we find that the “Seller and Buyer” treatment has the largest effect (\$689), consistent with the previous results on corrections. In columns 3 and 4, we examine the impact on reported final sales and find a decrease of \$309 (\$446 in the “Seller and Buyer” group). 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.³⁰ Firms are thus exploiting the final sales loophole, in line with Prediction 4a. In columns 5 and 6, we find no significant impact on reported taxable inputs. Finally in columns 7 and 8, we estimate the overall effect on the VAT liability reported in 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 effect is \$19 for the “Seller only” group and \$16 for the “Seller and Buyer” group.

Taken together, 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. As a result, the overall rise in reported tax liability for the months mentioned in the letter is relatively modest. Nonetheless, firms are not offsetting all of the increase, and hence the intervention successfully increased VAT liability for the targeted periods, in line with Prediction 4b. Additionally, they do not compensate by increasing their reported inputs, suggesting that

²⁹We focus on taxable sales and purchases. To calculate monthly VAT liability, we calculate the difference between output VAT charged and input VAT paid, excluding tax credits carried forward from past negative liabilities.

³⁰The 60% is calculated as follows: $309/512 \approx 446/689 \approx 0.6$.

Table 2
Effect of Treatment 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 (5). 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. All outcomes are reported in thousands of US dollars and winsorized at the 0.5% level (results are similar when we winsorize at 1%). 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.

input overreporting is not their preferred evasion margin in this context, in contrast to what is documented for corporate income taxes in Ecuador by Carrillo et al. (2017).

5.2.2 Robustness Checks

Having established the firm-level effects of the intervention, we next assess their robustness to potential spillovers within the network. Our randomization design isolates links but not firms. Because firms participate in multiple seller-buyer relationships, isolated links can remain connected through shared firms. In other words, if link ij between firms i and j is treated, while the connected link jk (its first-degree connection) cannot be treated, firm k could be treated through another relationship, such as kl (its second-degree connection). In this case, treated firm j is directly connected to another treated firm k , even though their respective treated links, ij and kl , are not directly connected. Essentially, what constitutes a second-degree connection at the link level becomes a first-degree connection at the firm level. Appendix Figure D1 provides an illustration of these connections.

To verify that such indirect exposure does not bias our firm-level estimates, we test whether potential spillover effects are identical for treated and control firms. While this

does not imply the absence of spillovers, it ensures that any spillover effects for treated and control firms cancel out, yielding consistent estimates of the intervention’s direct effect.³¹ Following [Vazquez-Bare \(2023\)](#), we re-estimate regression (5) including as controls the share of treated partners for treated and untreated firms separately. The corresponding table and detailed results are presented in Appendix D. The coefficients of interest remain stable and statistically unchanged when including these controls. We cannot reject the null hypothesis that all the variables capturing the share of treated partners are jointly insignificant.

5.2.3 Effect on subsequent VAT liability

Next, we analyze the impact of the letters on reported VAT liability in the months after treatment to test Prediction 4c. We rely on specifications similar to (3) and (4), except that here, observations are at the firm level. As shown in panels (a), (b), and (c) of Figure 8, there are no significant effects in the reported B2B sales, final sales, and total inputs after the treatment. Panel (d) 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 difference-in-differences specification, we find 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 leads to higher tax liabilities for amended past returns (Table 3) and triggers a sustained decline in the share of seller-shortfall (Figure 6), our study lacks sufficient statistical power to detect an increase in overall reported VAT liability among treated firms in the months following the intervention.

5.2.4 Heterogeneity analysis

To better understand which firms are driving our results, we estimate the average characteristics of compliers—i.e., seller-buyer links that are induced into reducing at least one discrepancy—following the methodology of [Pinotti \(2017\)](#). The outcome is defined at the link level, with pre-treatment baseline characteristics drawn from either the seller or the buyer within a link or from the link itself. To estimate the average characteristics of the compliers we estimate a two-stage least squares model with the following first and second

³¹Note that this is the same condition—often implicitly maintained—that underlies much of the existing literature estimating direct effects in the presence of potential spillovers (e.g., [Pomeranz, 2015](#)).

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 (3), 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 (4) pooling together all post-treatment months. In the bottom left corner, we report the means in the control group for the pre-treatment period and the number of observations. 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.

stage equations:

$$h_i = \lambda T_i + \epsilon_i \quad (7)$$

$$h_i \times k_i = \theta T_i + \nu_i, \quad (8)$$

where h_i is an indicator variable equal to one if at least one discrepancy of link i is corrected, k_i is the characteristic of interest—either a seller, buyer, or link characteristics—for link i , and T_i indicates whether link i is assigned to treatment. The characteristics of the complying links are given by θ . Intuitively, θ captures the average baseline value of the characteristics

of interest for the complier links.

Table 3 reports average baseline characteristics for compliers and all firms in the Study Sample, along with the difference between the two and the corresponding p-value. Panel A reports results for seller characteristics. Sellers in complying links have a higher share of final sales (61.8% vs. 56.4% in the Study Sample), with the difference significant at the 10% level. This is consistent with our conceptual framework: sellers that correct discrepancies in response to the treatment tend to have a larger share of final sales that can be relabeled as B2B sales. Compliers also exhibit a lower degree of connectedness, i.e. they are less connected to other firms at baseline, with this difference significant at the 10% level. We further examine whether compliers differ in total reported sales, the share of monthly tax returns with negative VAT liability, firm age, or whether they were audited in the year prior to the intervention. None of these differences is statistically significant. Finally, we find no systematic differences across sectors.

Panel B of Table 3 reports analogous results for buyers. Buyers in complier links are significantly less likely to be in the service sector, we do not find significant differences for other characteristics. Panel C extends the analysis to link-level characteristics—specifically, trading volume and the size of the initial discrepancy. Complying links have significantly smaller initial discrepancies than the rest of the Study Sample, consistent with the idea that the cost of correcting is lower when the discrepancy is small.

We next investigate whether specific firm characteristics predict the likelihood of correcting discrepancies in response to the letter. In columns 1-3 of Appendix Table F3, we estimate an OLS regression in which the outcome variable—whether at least one discrepancy has been corrected—is regressed on a large set of baseline seller, buyer, and link characteristics. The results show that sellers in the service sector, sellers that received the letter, and younger firms are more likely to respond. Among buyer characteristics, only one variable—whether the firm is in the mining sector—is significant at the 10% level. Finally, similar to what we find in the compliers analysis, firms with a lower initial discrepancies are also more likely to correct. Despite including a rich set of covariates, the explanatory power of these regressions remains limited, with an R-squared of only 0.08. In column 4, we apply a LASSO estimator to select the baseline characteristics most predictive of correction behavior. Despite including all covariates from columns 1-3, none of the variables is selected by the LASSO estimator.

Overall, these findings suggest limited heterogeneity in firms’ responses to the letter based on observable characteristics. While this makes it somewhat difficult to pinpoint which types of firms drive the observed effects, the evidence suggests that the behavioral response to the intervention is broad-based across VAT-registered firms in Uganda.

Table 3
Heterogeneity Analysis: Baseline Characteristics of the Compliers

	Average characteristics		Difference	P-value
	Compliers (1)	Sample (2)		
<i>Panel A: Seller Characteristics</i>				
Sales	144.058	174.197	-30.139	0.221
Final sales ratio	0.618	0.564	0.054	0.088
Share of negative returns	0.404	0.434	-0.031	0.299
Existence length	90.865	97.315	-6.450	0.254
Degree	44.106	53.883	-9.777	0.070
Audited in 2016	0.041	0.053	-0.011	0.410
Sector - Agriculture	0.016	0.012	0.003	0.672
Sector - Construction	0.070	0.065	0.005	0.789
Sector - Manufacturing	0.111	0.125	-0.013	0.565
Sector - Retail	0.187	0.172	0.014	0.624
Sector - Service	0.319	0.274	0.044	0.210
Sector - Wholesale	0.298	0.349	-0.051	0.134
<i>Panel B: Buyer Characteristics</i>				
Sales	145.165	105.113	40.051	0.229
Final sales ratio	0.511	0.527	-0.016	0.634
Share of negative returns	0.398	0.385	0.013	0.644
Existence length	80.091	79.197	0.894	0.871
Degree	50.078	41.435	8.643	0.117
Audited in 2016	0.067	0.056	0.011	0.486
Sector - Agriculture	0.013	0.018	-0.005	0.699
Sector - Construction	0.171	0.131	0.040	0.150
Sector - Manufacturing	0.088	0.079	0.010	0.612
Sector - Retail	0.174	0.163	0.011	0.687
Sector - Service	0.329	0.404	-0.075	0.041
Sector - Wholesale	0.210	0.199	0.011	0.712
<i>Panel C: Pair Characteristics</i>				
High trading volume	0.606	0.567	0.039	0.270
Initial discrepancy	-0.689	-1.106	0.417	0.000

Notes: This table reports heterogeneity results estimated in equation (8). We run regressions where the outcome is an indicator equal to 1 if there is a correction for at least one discrepancy in a link, interacted in turn with characteristics of the seller in the link (panel A), the buyer in the link (panel B), or the link itself (panel C). The firm-level characteristics are computed using VAT returns from 10-months prior to the intervention. The characteristics are (i) the amount of sales in thousands of USD, winsorized at the 0.5% level, (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 the firm registered for VAT, (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) (viii) whether the number of months traded exceed the median, and (ix) initial value of the discrepancy. Column (1) reports the θ coefficient on the indicator for receiving the letter, from equation (8). Column (2) reports the average of the characteristic in the Study Sample. Column (3) reports the difference between the two, while column (4) reports results of a significance test of the difference in means. *Source:* Data from monthly VAT returns submitted to the URA.

6 Discussion

This section interprets our findings, situating them within the broader literature and discussing their policy relevance.

6.1 Mechanisms

The results in Sections 4 and 5 demonstrate that a simple letter intervention exploiting firm network data can improve tax compliance, both through correction of past returns and improvements in post-treatment reporting behavior. Treated sellers correct discrepancies by increasing reported B2B sales but partly offset these adjustments by reducing reported final sales, resulting in only a modest increase in overall tax liability. This mechanism contrasts with the one documented in Carrillo et al. (2017), where firms compensate for additional reported sales by inflating reported input costs to leave tax liability unchanged. In our setting, reported inputs do not increase.

Instead, evasion operates through misclassification across reporting categories. Before the intervention, firms report some B2B transactions as final sales to move activity off the cross-checkable margin. After the intervention, this misclassification is partly reversed, with reported B2B increasing and reported final sales declining. This behavior is related to, but distinct from, the well-known “last mile” problem of the VAT (e.g., Naritomi, 2019). The key margin here is relabeling between B2B and final sales, not simply underreporting final sales. Even though seller-buyer cross-checks are effective for B2B transactions, the lack of paper trail for final sales creates a persistent verification gap, particularly when firms report on both margins. This gap propagates upstream and weakens the VAT’s self-enforcing properties throughout the production chain.

The evidence also points to sellers, rather than buyers, as the primary agents of VAT evasion in this context. Because they can reclassify sales across reporting categories, sellers have greater scope to respond to enforcement, and the resulting adjustments would not substantially increase their tax liability. This pattern is consistent with the findings of Pomeranz (2015), which shows that VAT enforcement interventions tend to affect upstream trading partners (suppliers/sellers) more strongly than downstream ones (clients/buyers). However, the mechanism in our setting differs: while Pomeranz’s conceptual framework predicts spillovers only in the presence of collusive evasion—when firms misreport matching amounts that evade simple cross-checks—we find spillovers in cases of seller shortfall (what Pomeranz terms “unilateral evasion”), where a single firm misreports and discrepancies remain detectable through cross-checking. Notably, we find no evidence suggesting that the intervention leads to increased collusive evasion, even when both firms in a treated link

receive the letter.

6.2 Direct and Spillover Effects

In the empirical analysis, we differentiate between direct and spillover effects of the tax enforcement intervention. A natural question is which of the two is more important in aggregate revenue terms. Table 4 reports an accounting exercise that decomposes corrections accordingly. As shown in panel (a), across all the corrections, treated firms report an additional \$300,230 in B2B trade, of which \$210,150 (70%) reflect corrections of discrepancies listed in the letters (direct effects) and \$90,070 (30%) reflect corrections of unlisted discrepancies (spillover effects), including corrections within the treated link that were not listed, corrections by the firm that did not receive the letter (e.g., the seller in a “Buyer only” link), or corrections involving other trading partners.

To assess magnitude of the aggregate effect from a causal perspective, we net out the corrections made by control firms. Among controls, corrections that would correspond to direct effects total \$8,120, while corrections equivalent to spillover effects amount to $-\$124,880$. The negative sign arises because we consider corrections in either direction, increasing or decreasing reported B2B sales.³² The net effect, calculated as corrections by treated firms minus those made by control firms, is an increase of \$416,990 in aggregate B2B trade, split roughly evenly between direct (48%) and spillover (52%) effects. Thus, spillovers contribute largely to the aggregate increase in reported trade.

Comparing across treatment arms in panel (b) of Table 4, the aggregate increase in reported B2B sales (including direct and spillover effects) is \$98,010 for “Seller only”, \$92,210 for “Buyer only”, and \$110,010 for “Seller and Buyer”. Interestingly, direct effects dominate for “Seller only” and “Seller and Buyer”, while spillovers are more important for “Buyer only”. In line with the estimates presented in Section 4.1, the Seller and Buyer treatment arm yields the largest revenue gain, though differences across arms are small. Overall, these patterns indicate that targeting sellers is the most effective way to induce large aggregate effects on reported B2B sales.

³²For example, when a firm submits an amendment with a larger seller shortfall discrepancy than the original report, the implied change in reported B2B sales is negative.

Table 4
Aggregate Amounts of Uncovered B2B Sales

	Total	Direct	Spillover
<i>Panel (a): Aggregate Amounts</i>			
All Treated	300.23	210.15	90.07
Control	-116.76	8.12	-124.88
Net (Treated – Control)	416.99	202.04	214.95
<i>Panel (b): Amounts by Treatment Arm</i>			
Seller only	98.01	107.75	-9.74
Buyer only	92.21	-3.20	95.41
Seller & Buyer	110.01	105.61	4.40

Notes: This figure reports the changes in aggregate amounts of B2B sales reported by firms of the Study Sample for returns corresponding to the months listed in the letters. There are a total of 1,772 of these changes in the treated group and 303 in the control group. All amounts are in thousands of USD. Negative amounts mean that the corrections decreased reported B2B trade. One outlier is dropped from the sample for spillover effects in the control group, with a massive decrease in reported sales – its inclusion would further increase the difference between the treated and control group.

6.3 Cost Effectiveness and Policy Implications

After accounting for the partial reclassification of final sales as B2B sales, the estimated revenue gain from the amendments, accounting for all margins of adjustment, is \$35,555.³³ The total implementation cost—including courier services and the URA staff time—was \$5,716 (see Appendix G for details on the calculation). Thus, the intervention generated over six times its cost in additional revenue. We conclude that, despite the modest increase in VAT liability from amendments, the intervention is highly cost effective.

Taken together, our results suggest that interventions targeting firms’ reporting practices, and explicitly leveraging network information, can meaningfully improve tax compliance. Several implications for policy design emerge. First, while sellers can reclassify final sales as B2B sales to limit the impact on their tax liability, this strategy is inherently limited: final sales cannot be reduced indefinitely. A scaled-up intervention in which the revenue authority systematically notifies firms of all detected discrepancies could further constrain this margin of evasion and increase truthful reporting.

Second, the experiment has informed subsequent URA practices. Following the 2018 experiment, the data cross-checks were substantially expanded. In 2021, the URA launched a new program called electronic fiscal receipting and invoicing (EFRIS), which requires that all input claims be validated with the corresponding VAT invoices. In 2022, it began rolling

³³This is calculated by multiplying the change in tax liability estimated in column 7 of Table 2 times the number of firm-month observations: $\$13 \times 2,735 = \$35,555$.

out electronic billing machines to record transaction information at the point of sale. These measures represent substantial advances in tax enforcement capacity, though their impact has yet to be rigorously evaluated.

Third, the effectiveness of network-based interventions depends crucially on which firms are targeted. Because spillovers scale with the number of connected buyers, targeting sellers with many trading partners may yield larger aggregate gains. A granular investigation of firms' response in our study reveals strong heterogeneity: the seller making the most corrections traded with 353 buyers, makes 295 corrections, and uncovers \$5,800 in previously unreported B2B sales, whereas some sellers make no corrections. Understanding how such heterogeneity maps into targeting rules is a natural direction for future research. While letters sent to both firms in a link elicit the strongest responses, we also observe that treated firms induce non-treated partners to make corrections via firm-to-firm communication. The administration hence faces a tradeoff between reinforcing credibility for a smaller number of pairs and maximizing reach across the network.

Finally, we detect no effects on VAT liability in the months after treatment. One interpretation is that, absent follow-up measures, firms' updated beliefs about the administration's improved monitoring capacity do not persist. Future research could explore whether reinforcement strategies—such as repeated notifications or complementary monitoring tools—can generate more durable compliance gains.

7 Conclusion

This paper uses transaction-level VAT data from Uganda to design a randomized trial aimed at improving tax compliance and reporting behavior among VAT-registered firms. We identify seller-buyer links with seller shortfalls—reporting discrepancies that reduce tax liability and suggest tax evasion—and randomly assign them to receive enforcement letters from the URA. Methodologically, we extend the network-experiment literature by implementing a novel two-stage experimental design. This design together with the variation in treatment intensity across links allow us to rigorously estimate the direct and spillover effects of the intervention, and document firm-to-firm communication in response to the treatment.

The letters elicit strong responses, significantly increasing corrections of both listed and unlisted discrepancies, and improving reporting in subsequent months. Spillover effects extend to other trading partners, both within treated links and beyond them. These patterns provide rare evidence of firm-to-firm communication of enforcement responses, as some treated buyers induce their seller to correct reported sales. The asymmetric responses observed suggest that VAT evasion in this setting is primarily driven by sellers, who exploit the ability to reclassify

B2B sales as unverifiable final sales. Unlike previous evidence that restricts this “last mile” problem to retailers, our results show that such behavior extends throughout the supply chain. These new insights suggest that evasion mechanisms may differ systematically in low-capacity settings like Uganda. Overall, the intervention modestly increases VAT liabilities but remains highly cost effective.

These findings demonstrate the potential of leveraging firm networks to strengthen tax enforcement. At the same time, 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, whether they would benefit from being targeted to specific firms within the network, 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, Jonas Hjort, Justine Knebelmann, and Lin Tian**, “Information, Fiscal Capacity and Tax Revenues: An Experimental Evaluation,” 2018. AEA RCT Registry.
- , – , – , and – , “Strategic or Confused Firms? Evidence from Missing Transactions in Uganda,” *Review of Economics and Statistics*, 2024, *106* (1), 256–265.
- 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.
- Battaglini, Marco, Luigi Guiso, Chiara Lacava, and Eleonora Patacchini**, “Tax Professionals and Tax Evasion,” 2022. mimeo.
- 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.
- Brown, Julia, Sebastian Jilke, David Schwegman, Mattie Toma, and Mary Clair**

- Turner**, “Improving Tax Compliance among the Clients of Tax Preparers,” 2025. Available at SSRN.
- Cai, Jing, Wei Lin, and Adam Szeidl**, “Firm-to-Firm Referrals,” 2024.
- 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.
- 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.
- Garriga, Pablo and Dario Tortarolo**, “Firms as tax collectors,” *Journal of Public Economics*, 2024, 233, 105092.
- 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*.
- Huneus, Federico, Kory Kroft, and Kevin Lim**, “Earnings Inequality in Production Networks,” NBER Working Papers 28424, National Bureau of Economic Research. Feb 2021.
- 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,” 2023. 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.
- Miguel, Edward and Michael Kremer**, “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities,” *Econometrica*, 2004, 72 (1), 159–217.
- Naritomi, Joana**, “Consumers as Tax Auditors,” *American Economic Review*, 2019, 109, 3031–3072.
- OECD**, *Revenue Statistics* 2022.
- OECD/ATAF**, *Revenue Statistics in Africa* 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.
- Rubin, Donald B.**, “Randomization analysis of experimental data: The Fisher randomization test comment,” *Journal of the American Statistical Association*, 1980, 75 (371), 591–593.
- 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, 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.
- 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 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 A1, the share of firm-pairs this affected was balanced across treatment arms indicating that this did not bias our sample.

Table A1
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 A2. 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 A2
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.

B Evidence on Implementation

In this appendix, we provide more information on the experimental implementation and follow-up details. We first show a template of the notification letters sent to firms in Figure B1. We then show that there is no differential attrition rate in treatment and control groups in Figure B4. Table B1 shows shares of treated firms that provided feedback on the letters and Figure B3 displays the types of feedback submitted by the treated firms.

Figure B1
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»

Figure B2
Share of Firms Reacting to the Letter or Contacting the URA



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 3.3. *Source:* Data from monthly VAT returns submitted to the URA.

Table B1
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 table is mentioned in Section 3.3. *Source:* Data from monthly VAT returns submitted to the URA.

Figure B3
Responses from Firms that Gave Feedback by Treatment Arm



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 3.3. *Source:* Data from monthly VAT returns submitted to the URA.

Figure B4
Balance in Attrition Rates across Treatment Arms




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 3.3. *Source:* Data from monthly VAT returns submitted to the URA.

C Additional Results 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 C1 compares the distributions of discrepancies before and after the treatment for the treatment and control groups. Tables C1-C3 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 C4 runs a similar exercise as in the one in Table C3 but distinguishes between months with listed and unlisted discrepancies.

Figure C1
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 C1
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 on 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 C2
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 C3
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 C4
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.

D Robustness Checks Details

In this appendix, we discuss the specifications used to conduct the robustness checks outlined in Sections 4.1 and 5.2. We discuss the specification separately for our link- and firm-level results.

Figure D1
First and Second Degree Connections in the Randomized Network



Notes: This figure illustrates the structure generated by the two-stage randomization, discussed in Sections 4.1 and 5.2. At the *link* level, a link ij included in the Study Sample connects firms i and j , while any directly adjacent link—such as jk , its first-degree connection—is excluded from the Study Sample. Higher-order links, like kl , a second-degree connection relative to ij , may still be included. At the firm level, the randomization isolates links but not firms: a firm such as j on an included link ij may be connected to another included firm k through kl . Thus, what is a second-degree connection at the link level becomes a first-degree connection at the firm level for firms j and k . Links included in the Study Sample are then randomized into control and treatment groups.

D.1 Link-level Checks

Our identification strategy for the link-level results (presented in Section 4.1) relies on two assumptions: (i) spillovers do not operate through non-seller-shortfall links, and (ii) they do not extend beyond one degree of separation in the network. In other words, excluding first-degree seller-shortfall links from the Study Sample during randomization is sufficient to account for potential spillover effects that could bias our estimation.

Because our network data include all transactions between VAT registered firms, we can empirically test whether potential spillovers beyond those accounted for could bias our estimates. To do so, we follow the approach of Vazquez-Bare (2023) and estimate the following reduced-form linear-in-means regression, which explicitly controls for the presence of spillover effects:

$$Y_{it} = \alpha + \beta T_i + \left[\gamma_s^0 \bar{T}_{si} + \gamma_b^0 \bar{T}_{bi} \right] (1 - T_i) + \left[\gamma_s^1 \bar{T}_{si} + \gamma_b^1 \bar{T}_{bi} \right] T_i + \delta_t + \epsilon_{it}, \quad (9)$$

where Y_{it} is the outcome of interest for link i in month t , T_i is an indicator for whether the link is treated, δ_t are month fixed effects, and ϵ_{it} is the error term. The variables $\bar{T}_{si}$ and $\bar{T}_{bi}$ are the share of the seller's and buyer's partners that are treated, with their exact definition varying across the robustness exercises below. This specification is identical to (2), except that it includes the partner-treatment shares as additional controls.

In the first exercise, $\bar{T}_{si}$ and $\bar{T}_{bi}$ denote the shares of the seller’s and buyer’s first-degree trading partners that are treated. These are, by construction, non-seller-shortfall links, because seller shortfall links involving links in the Study Sample are automatically excluded through the iterative procedure explained in Section 3.2. The results are presented in Table D1. Column (1) replicates the baseline specification (as in Figure 3, panel a), while column (2) adds the shares of the seller’s and buyer’s trading partners that are treated. The coefficient estimate on T_i and its standard errors remain virtually unchanged. To further allow for the possibility that spillovers differ between treated and untreated firms, column (3) interacts the shares with treatment status as indicated in equation (9), which is equivalent to the “correctly specified interacted reduced form linear in means regression” proposed by Vazquez-Bare (2023). The coefficient on T_i is qualitatively similar (0.189 vs 0.223), highly significant, and the indirect exposure controls are jointly insignificant. A direct test confirms the coefficients from column (1) and column (3) are not statistically different (p-value = 0.272). We interpret this as evidence that exposure through non-seller shortfall links exerts limited bias on our estimates.³⁴

We next extend the analysis to second-degree links. Figure D1 illustrates our network structure underlying our randomization: when a link ij is included in the Study Sample, any directly adjacent link—such as jk , its first-degree connection relative to ij —is excluded, while higher-order links—such as kl , a second-degree connection relative to ij —remain eligible. To assess whether such indirect exposures affect our main treatment effects, we re-estimate (9) using exposure shares defined at the second-degree level. Specifically, $\bar{T}_{si}$ and $\bar{T}_{bi}$ denote the shares of a seller’s and buyer’s trading partners’ seller shortfall links with other firms that are treated. As shown in columns (1)–(3) of Table D2, the results remain consistent: the estimated coefficients are similar (0.202 vs 0.223), highly significant, and the interacted shares controls are jointly insignificant. Tests for coefficient equality again fail to reject the null. Together, these results reassure that first or second-degree exposures are unlikely to bias our link-level estimates.

D.2 Firm-level Checks

As discussed in Section 5.2, our randomization isolates seller-buyer links but not individual firms in the network. Because firms participate in multiple trading relationships, some firms in the Study Sample remain connected to other sampled firms through seller-shortfall links. As illustrated in Figure D1, what constitutes a second-degree connection at the link level (e.g., links ij and kl connected only through firm k) can therefore correspond to a first-degree connection at the firm level for firms j and k .³⁵

³⁴Vazquez-Bare (2023) also proposes a “pooled” estimand for settings where the two key assumptions underlying regression (9)—exchangeability and linearity—do not hold. In that specification, the shares are replaced with an indicator for whether any trading partner is treated. However, this approach is not appropriate for our context since there is extremely little variation in such an indicator: 91% of buyers and 90% of sellers have at least one trading partner that is treated.

³⁵The high interconnectedness of firms makes it infeasible to construct a sufficiently large Study Sample in which all firms are fully disconnected from each other.

Table D1
Robustness to Non-Seller-Shortfall Spillovers

	(1) Any correction	(2) Any correction	(3) Any correction
T_i	0.223*** (0.017)	0.223*** (0.017)	0.189*** (0.034)
$\bar{T}_{si}$		-0.018 (0.057)	
$\bar{T}_{bi}$		0.053 (0.072)	
$\bar{T}_{si} \times (1 - T_i)$			-0.033 (0.024)
$\bar{T}_{si} \times T_i$			-0.014 (0.092)
$\bar{T}_{bi} \times (1 - T_i)$			-0.030 (0.018)
$\bar{T}_{bi} \times T_i$			0.119 (0.128)
R-squared	0.095	0.096	0.097
Observations	2744	2744	2744
No. of Firms	1235	1235	1235
Mean of Dep. in Control	0.016	0.016	0.016
P-Value testing for coefficient equality		0.775	0.272
P-value (joint F-test of shares)		0.740	0.422
Month-Year FE	Yes	Yes	Yes

Notes: This table reports the effect of the treatment on the correction of past discrepancies within links of the Study Sample when controlling for the share of treated trading partners among those connected via non-seller-shortfall links. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. In column (1) we present results from estimating equation (2), in column (2) we add the share of treated trading partners for the seller and buyer of the links, while in column (3) we report results from estimating equation (9). In the second to last line, we present the p-value of a test for whether the coefficient in column (2) (and (3)) is equal to the coefficient in column (1). In the last line we present the p-value for a test of whether the interacted variables are jointly statistically significant. This tables focuses on listed discrepancies, mentioned in the letter. Standard errors are clustered at the link level. *p< 0.10; **p< .05; ***p< .01. *Source:* Data from monthly VAT returns submitted to the URA.

Table D2
Robustness to Second-Degree Link Spillovers

	(1) Any correction	(2) Any correction	(3) Any correction
T_i	0.223*** (0.017)	0.222*** (0.017)	0.202*** (0.026)
$\bar{T}_{si}$		0.039 (0.041)	
$\bar{T}_{bi}$		-0.023 (0.037)	
$\bar{T}_{si} \times (1 - T_i)$			-0.053** (0.025)
$\bar{T}_{si} \times T_i$			0.070 (0.054)
$\bar{T}_{bi} \times (1 - T_i)$			0.007 (0.027)
$\bar{T}_{bi} \times T_i$			-0.033 (0.048)
R-squared	0.095	0.096	0.098
Observations	2744	2744	2744
No. of Firms	1235	1235	1235
Mean of Dep. in Control	0.016	0.016	0.016
P-Value testing for coefficient equality		0.875	0.300
P-value (joint F-test of shares)		0.552	0.106
Month-Year FE	Yes	Yes	Yes

Notes: This table reports the effect of the treatment on the correction of past discrepancies within links of the Study Sample when controlling for the share of treated trading partners among those connected by second-degree seller shortfall links. VAT returns for pre-treatment months (March 2017 to December 2017) were analyzed at the seller-buyer link-month level. In column (1) we present results from estimating equation (2), in column (2) we add the share of trading partners non-seller-shortfall links that area treated for the seller and buyer of the links, while in column (3) we report results from estimating equation (9). In the second to last line, we present the p-value of a test for whether the coefficient in column (2) (and (3)) is equal to the coefficient in column (1). In the last line we present the p-value for a test of whether the interacted variables are jointly statistically significant. This tables focuses on listed discrepancies, mentioned in the letter. Standard errors are clustered at the link level. *p< 0.10; **p< .05; ***p< .01. *Source:* Data from monthly VAT returns submitted to the URA.

To test whether such indirect exposure could bias our firm-level results from Section 5, we re-estimate (6), controlling for the share of a firm’s trading partners in the Study Sample that are treated. Formally, we estimate the following specification:

$$Y_{jt} = \alpha + \beta_h T_j + \gamma^0 \bar{T}_j (1 - T_j) + \gamma^1 \bar{T}_j T_j + \delta_t + \epsilon_{jt}, \quad (10)$$

where Y_{jt} denotes firm-level outcomes for seller j in month t , T_j denotes a dummy variables for whether a seller is treated, $\bar{T}_j$ is the share of the firm’s trading partners that are treated, δ_t are month fixed effects, and ϵ_{jt} is the error term.

Table D3 reports the results. Odd-numbered columns replicate the baseline estimates from Table 2, while even-numbered columns incorporate the additional controls for partner treatment exposure from (10). In most cases, the coefficients on T_j remain stable and quantitatively similar from the baseline and the indirect exposure controls are jointly insignificant. Nonetheless, as expected, indirect exposure plays a somewhat larger role at the firm-level. In column (6), the indirect exposure variables are jointly marginally significant at 10%, and in column (8), the coefficient on T_j is marginally different from column (7). These differences, however, are small in magnitude and do not change the overall interpretation. Overall, the results confirm that indirect exposure exerts only minor influence even on firm-level estimates.

Table D3
Effect of Treatment on Firm-Level VAT Liability (Sellers)

Dependent variable:	$\Delta B2B$ Sales		Δ Final Sales		Δ Taxable Inputs		Δ VAT Liability	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T_i	0.512*** (0.059)	0.539*** (0.078)	-0.309*** (0.037)	-0.294*** (0.046)	0.002 (0.003)	0.002 (0.004)	0.013*** (0.004)	0.020*** (0.006)
$\bar{T}_i \times (1 - T_i)$		-0.018 (0.037)		-0.005 (0.012)		-0.012* (0.007)		0.020* (0.011)
$\bar{T}_i \times T_i$		-0.108 (0.145)		-0.057 (0.092)		-0.009 (0.005)		-0.010 (0.012)
R-squared	0.036	0.037	0.035	0.035	0.007	0.008	0.010	0.011
Observations	2747	2747	2747	2747	2747	2747	2747	2747
No. of Firms	1235	1235	1235	1235	1235	1235	1235	1235
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	0.017	0.017	-0.003	-0.003	0.001	0.001	0.005	0.005
Median of Dep. in Control	1.661	1.661	11.656	11.656	27.837	27.837	0.255	0.255
P-value of test for eq. of coefficients		0.524		0.556		0.990		0.093
P-value (joint F-test)		0.671		0.772		0.073		0.137
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 when controlling for the share of trading partners that are treated. Odd columns present results from estimating equation (5), while even columns present results from estimating equation (10). 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. In the second to last line, we present the p-value of a test for whether the coefficient in column (2) (and (3)) is equal to the coefficient in column (1). In the last line we present the p-value for a test of whether the interacted variables are jointly statistically significant. All outcomes are reported in thousands of US dollars and winsorized at the 0.5% level (results are similar when we winsorize at 1%). 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.


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 4.2. *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)
$\sum_{j=-10}^{-2} \beta_j$	-0.276 (0.230)	-0.348 (0.225)	0.624*** (0.230)	-1.004 (1.399)	-0.571 (1.397)	0.153 (0.209)
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 (3), 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

	(1) Seller shortfall	(2) Buyer shortfall	(3) No discrepancy	(4) Transaction size	(5) Seller shortfall amount	(6) Continuing trade
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 (4) pooling together all post-treatment months, 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. *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 (1) (2)		Seller shortfall amount (3) (4)		Continuing trade (5) (6)		Seller Reports (7) (8)		Buyer Reports (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 (4) 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 (3), 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 E5 and E6 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 (3), 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 (4) 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 (3), 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

	Seller shortfall		Buyer shortfall		No discrepancy	
DiD specification	(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 (4) 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 (4) 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 5. In the main text, we focus on discussing the results for sellers (Table 2). In Table F1, we show the corresponding results for buyers. Table F2 shows the regression results that correspond to Figure 8.

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		$\Delta Final$ Sales		$\Delta Taxable$ Inputs		ΔVAT 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 (5). 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*, 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 5. *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)
$\sum_{j=-10}^{-2} \beta_j$	14.646 (34.657)	-4.300 (53.682)	4.570 (64.623)	-1.270 (6.262)
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 $\hat{\beta}_j$ coefficients estimated in an event-study similar to equation (3), but at the firm level, with four different outcomes: amount of B2B sales (sales to other VAT firms) in column 1, final sales 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 5. *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 (1)		Buyer Characteristics (2)		Pair Characteristics (3)		Relevant Characteristics (4)	
Degree	0.000 (0.000)	Degree	0.000 (0.000)	High trading volume	0.054 (0.036)		
Sector: Agriculture	0.118 (0.153)	Sector: Agriculture	0.088 (0.149)	Initial discrepancy	-0.014* (0.007)		
Sector: Construction	0.075 (0.073)	Sector: Construction	0.095 (0.061)				
Sector: Manufacturing	0.015 (0.056)	Sector: Manufacturing	0.042 (0.072)				
Sector: Mining	-0.270 (0.320)	Sector: Mining	0.315 (0.205)				
Sector: Retail	0.072 (0.049)	Sector: Retail	0.014 (0.055)				
Sector: Service	0.119** (0.045)	Sector: Service	-0.040 (0.046)				
Ln Sales	0.000 (0.000)	Ln(Sales)	0.000 (0.000)				
Share of negative returns	-0.047 (0.045)	Share of negative returns	0.063 (0.046)				
Existence length	0.000** (0.000)	Existence length	0.000 (0.000)				
FC ratio	0.053 (0.046)	FC ratio	-0.052 (0.044)				
Audited in 2016	-0.056 (0.073)	Audited in 2016	-0.006 (0.077)				
R-squared	0.058						
Observations	741					Observations	741

Notes: This table reports results on the heterogeneity in responses to the letter. We estimate an OLS regression (columns 1-3) and a Robust Lasso (column 4). The outcome variable is an indicator taking value 1 if at least one discrepancy is corrected by either firm in the link. We regress the outcome variable 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 amount of sales (in thousands of USD), (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 registered for VAT, and (viii) whether the firm was audited in 2016. The pair-level characteristics are: (ix) whether the number of months traded exceed the median, and (x) the initial value of discrepancy. Interestingly, the Lasso specification only retains one seller characteristic. *p<.10; **p<.05; ***p<.01. This table is mentioned in Section 5. *Source:* Data from monthly VAT returns submitted to the URA.

G Cost Calculations

This appendix discusses details of the cost calculations mentioned in Section 6. 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.³⁶ 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, URA staff 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 \$5,716.³⁷

³⁶Note that this is larger than the number quoted throughout the paper, which is due to the data error discussed in detail in Appendix A.

³⁷In 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.