Firm Networks and Tax Compliance: Experimental Evidence from Uganda*

Miguel Almunia
CUNEF Universidad

David Henning
Oxford University

Justine Knebelmann
Sciences Po

Dorothy Nakyambadde
Uganda Revenue Authority

Lin Tian
INSEAD

November 2026

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 24% of reporting discrepancies, twenty 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. Nakyambadde: Uganda Revenue Authority, dnakyambadde@ura.go.ug. Tian: INSEAD & CEPR, lin.tian@insead.edu. The authors' names are listed in alphabetical order. Naomi Feldman was coeditor for this article, and Kory Kroft was also consulted. We are grateful to the three anonymous referees for their thoughtful comments. 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. The study received IRB approval from Columbia University (protocol AAAN1410(M01Y04)) and the University of Warwick (reference 104/17-18). 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 a consumption tax collected at each stage of the supply chain: firms charge VAT on their sales and receive credits for the VAT they paid on their purchases. This creates a built-in enforcement mechanism. When a registered firm sells to another registered firm, only the seller issues an invoice, but both parties are required to report the transaction to the tax authority. In principle, this creates opposing incentives that constrain unilateral evasion. A seller who underreports sales risks detection when the buyer claims an input credit for the same transaction, and a buyer who overreports purchases faces the same risk when cross-checked against the seller's report (Pomeranz, 2015). In practice, however, this self-enforcement mechanism may break down when tax authorities lack the capacity to cross-check reports systematically. As a result, widespread misreporting persists even in transactions between registered firms (Almunia, Hjort, Knebelmann and Tian, 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 VAT data in a two-stage experimental design. In the first stage, we construct a sample of seller-buyer links that are disconnected from each other in the transaction network to avoid interference. In the second stage, we randomize which firm in each link receives an enforcement letter from the tax authority—the seller only, the buyer only, or both. The letters inform firms that the revenue authority can now cross-check their reports with trading partners and list specific discrepancies in their past transactions. This 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 detected by our intervention is driven primarily by sellers, who evade by reclassifying business-to-business (B2B) transactions as unverifiable final sales (which include sales to final consumers or to non-registered firms), 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 122,197 unique seller-buyer links. The intervention targets cases where sellers report a smaller amount than buyers, which we label "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. Each letter informed firms that the URA had developed new methods to detect reporting discrepancies, listed up to three months in which their reports did not match their trading partner's, instructed them to file corrections, and reminded them of penalties for evasion.

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, Bohren, McIntosh and Özler, 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, 23.5% of discrepancies listed in the letters are corrected, compared to just 1.2%

[1]SUTVA, 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).

in the control group—an almost twenty-fold increase. Treated links also correct 11.3% of unlisted discrepancies (against 1.5%), a seven-fold increase that shows the response reaches beyond the transactions named in the letter. Looking forward, 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 11.7pp (97%). Together, these results suggest that firms update their beliefs about the revenue authority's monitoring capacity and translate this into sustained changes in reporting. We find no evidence that the intervention affects the probability or volume of subsequent trade in treated links, suggesting that the reduced discrepancies reflect improved compliance rather than collusion.

Our experiment also uncovers sharp asymmetries between sellers and buyers—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, against 2.4% for buyers. The asymmetry is even sharper when we look across treatment arms. When only buyers receive letters, their *seller* partners correct 9.2% of discrepancies (8.2pp above control), while the buyers themselves correct 2.6% (2.6pp above control). The reverse does not hold: when only sellers receive the letters, buyer partners' corrections increase by only 1.3pp. Descriptive evidence from firms' communications with the URA reinforces this pattern: buyers attribute discrepancies to their trading partners in 79% of post-intervention contacts, compared to only 18% for sellers. Together, these findings suggest that sellers are the primary agents of the type of VAT evasion we detect. The cross-arm asymmetry also reveals firm-to-firm communication: sellers respond not only to their own letters but also to letters received by their trading partners. Consistent with this, 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.

We find spillover effects that extend beyond the treated link. Corrections in treated firms' other trading relationships increase by 3.4pp, or elevenfold. 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.

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 \$522 but reduce reported final sales by \$309, resulting in a modest increase in monthly VAT liability of \$14.[2] These firm-level results reveal how sellers evade in this context: they reclassify business-to-business transactions as final sales, which are not subject to cross-checking. In total, our intervention uncovers \$293,670 in previously unreported B2B transactions, 29% of which is due to spillover effects on transactions that were not listed in the letters. The

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

net revenue gain is smaller once reclassification into final sales is taken into account, but the intervention remains highly cost effective: it yields more than six times its administrative 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. 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, Bohren, McIntosh and Özler, 2018; Vazquez-Bare, 2023). Previous 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, Guiso, Lacava and Patacchini, 2022; Boning, Guyton, Hodge and Slemrod, 2020; Brown et al., 2025), across tax types (Lopez-Luzuriaga and Scartascini, 2019), and among geographically proximate entities (Drago, Mengel and Traxler, 2020; Lediga, Riedel and Strohmaier, 2023; Cruces, Tortarolo and Vazquez-Bare, 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, Gurara and Woldeyes, 2017; Holz et al., 2023), growing evidence suggests that firms often offset enforcement by adjusting along untargeted margins (Carrillo, Pomeranz and Singhal, 2017; Slemrod et al., 2017; Almunia and Lopez-Rodriguez, 2018; Mascagni, Mengistu and Woldeyes, 2021; Best, Shah and Waseem, 2021; Hoy, Sinning and McKenzie, 2024). 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 showing that the type of VAT evasion detected through our enforcement instrument is driven primarily by sales underreporting, rather than cost inflation. Unlike previous evidence that focuses on evasion of final sales at the retail stage (Naritomi, 2019), our results demonstrate that firms leverage final sales to evade taxes throughout the supply chain. 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. More broadly, our findings complement recent evidence on the returns to scaling-up automated cross-checking, such as in Pakistan where Shah (2026) shows that a system which denies suspicious input-credit claims in real time when discrepancies are detected increased nationwide VAT revenues by at least 10%.

Finally, our paper is related to the literature on transaction networks. A large literature outside taxation highlights their importance in shaping firm outcomes (Huneeus, Kroft and Lim, 2021; Alfaro-Ureña, Manelici and Vasquez, 2022; Dhyne et al., 2022; Adão et al., 2022; Bernard et al., 2022; Demir, Javorcik, Michalski and Ors, 2022; Cai, Lin and Szeidl, 2024). Within taxation, non-experimental work shows that tax enforcement interventions can also affect firms' trading partners (Lediga, Riedel and Strohmaier, 2023; Carrillo, Donaldson, Pomeranz and Singhal, 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 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.

[3]For 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.

[4]The 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.

[5]In our sample, 59.83% 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.

[6]When 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 low-income countries (Lemgruber, Masters and Cleary, 2015).

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

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 (Uganda Revenue Authority, 2016-2019). 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 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 21,607 unique firms that submit a VAT return and 122,197 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 11.02 (2) unique trading partners over the year, corresponding to 5.66 (0) unique clients and 5.66 (1) unique suppliers.

Final sales. Among VAT-registered firms, 65.8% report at least some final sales and, for 65.9% 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.

[8] For more details on how these transactions are reported, see Almunia, Hjort, Knebelmann and Tian (2024).

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 47.5% 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.

[9]We observe buyer shortfall in 44.9% of the link-month observations, and matching amounts in the remaining 7.6% 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, Hjort, Knebelmann and Tian, 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).

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 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:

Prediction Set 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

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

[11] Seller 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.

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

Prediction Set 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:

Prediction Set 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 122,197 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 10,855 seller-buyer links and 4,452 unique firms. For brevity, we will henceforth refer to the seller-buyer links meeting

[12]We illustrate where the 1 million UGX of accumulated seller shortfall threshold lies in the distribution in Figure B1. Additionally, 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.

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.

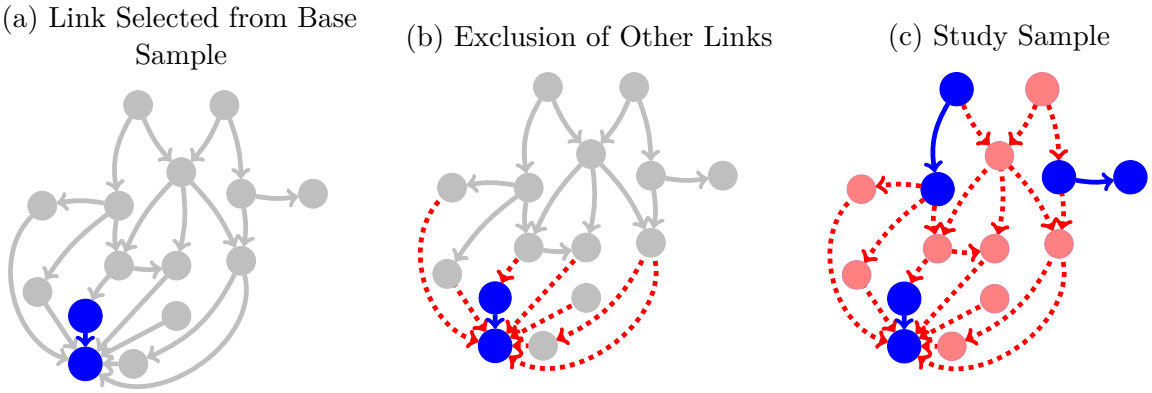
Study Sample. We draw our Study Sample from the Base Sample, applying an iterative sample selection procedure. Figure 1 illustrates the process. Panel (a) represents an example of a firm network of targeted discrepancies: each circle represents a firm and each solid line represents a link between a seller 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 buyer and seller 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 buyer and seller. We repeat the procedure until no eligible links remain. This yields a Study Sample of 1,235 seller-buyer links (and hence 2,470 firms), as shown in Figure 2.[13]

The iterative procedure removes links that share a firm with an already-selected link, but it does not remove the firm itself from the network. A firm with many trading partners may lose one eligible link in a given iteration but retain others through its remaining trading relationships. As a result, more connected firms have a higher probability of entering the Study Sample. We verify this by simulation. Running the sampling algorithm 1,000 times from the Base Sample, we find a clear positive relationship between a firm's number of trading partners and its probability of selection (Figure B2). Figure B3 confirms this in realized samples: the median number of trading partners is 2 in the Full Sample, 13 in the Base Sample, and 26 in the Study Sample. Table B1 in Appendix B.1 provides a full comparison of firm characteristics across the three samples.

The over-representation of more connected firms is a deliberate feature of the design. One of our objectives is to estimate spillover effects on trading partners not listed in the letters, which requires firms with broader networks. The Study Sample is therefore not representative of a randomly selected seller-buyer link in the Ugandan economy; rather, it is designed to be informative about both direct effects and network spillovers. Our estimated treatment effects should be interpreted accordingly.

[13] A data extraction error duplicated all input transactions— transactions where the firm acted as a buyer—for invoices filed between 19/9/2017 and 15/10/2017. This doubled reported purchases, mechanically inflating seller shortfall amounts for those months. As a result, 28% of links in the original Study Sample did not genuinely exceed our 1 million UGX discrepancy threshold. We correct this by excluding those links from the analysis, reducing the sample from 1708 to 1235 links. In Appendix A, we show that the correction is balanced across treatment arms, making bias unlikely.

Figure 1
Sample Selection Strategy



Notes: This figure illustrates the iterative sample selection procedure for the study. Each dot represents a firm and each arrow a "seller-shortfall link". In panel (a), we start from the Base Sample and randomly select one seller-buyer link. In panel (b), we remove all links between other firms and the selected buyer or seller (red dashed lines). The procedure is repeated among the remaining "seller-shortfall" 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 related to the two-stage experimental frameworks used in the literature to study spillover effects (Duflo and Saez, 2003; Miguel and Kremer, 2004; Baird, Bohren, McIntosh and Özler, 2018; Cruces, Tortarolo and Vazquez-Bare, 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.[14]

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.[15] A template is displayed in Appendix Figure C1. The letter instructs the firm to resolve the issue by filing the necessary amendments, and warns them that the URA will carry out

[14]See also Eckles, Karrer and Ugander (2017), which study experimental designs in dense networks. An important difference is that their procedure relies on a data-driven approach to define the relevant clusters, while in our setting, the clusters—seller-buyer pairs—arise naturally from the fact cross-checking of potential tax evasion operates at the link level.

[15]The 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.

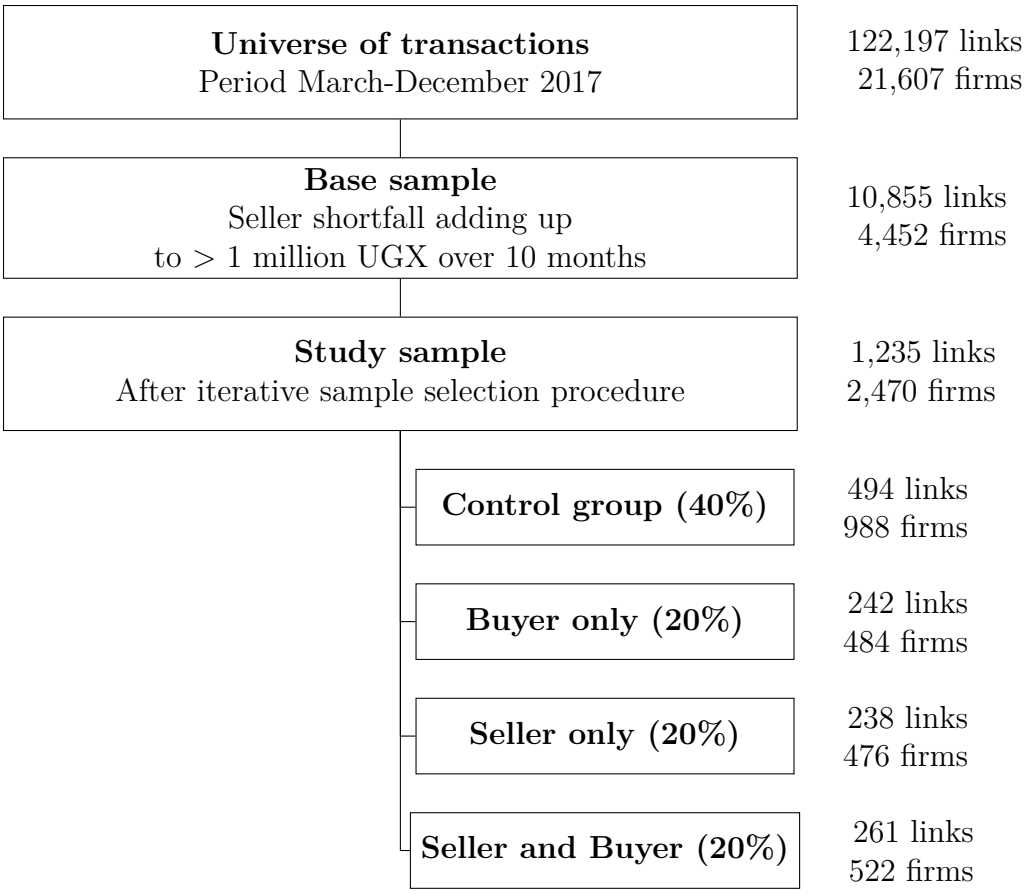


Figure 2
Sample and Experimental Design

Notes: 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 "seller-shortfall links", 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.

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 "Buyer and Seller" 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). We discuss the resulting statistical power in Section 4 and Appendix B.3.

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 "Buyer and Seller" 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 C1 and Figures C2 and C3 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 of the differences is significant at the 5% 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 0.50.[16]

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 C4. 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 links, or indirectly through second-degree links. Section 4.1.4 discusses this in detail and shows that it does not affect our results.

4 Results

This section analyzes the effects of the letter on seller-buyer links, testing our predictions from Section 3.1. We first study how the treatment affects the correction of past discrepancies, then turn to subsequent reporting discrepancies in the ten months after the letters were sent.

[16] In Appendix B.2 Table B2, we replicate the balance tests accounting for multiple hypotheses, using sharpened q-values (Benjamini and Hochberg, 1995; Anderson, 2008). The null hypothesis is never rejected.

Table 1
Study Sample: Descriptive Statistics and Randomization Balance

	Mean	Difference with respect to control			
Link characteristics	Control	Any treatment	Buyer only	Seller only	Buyer and Seller
	(1)	(2)	(3)	(4)	(5)
Monthly transactions	4.82	−0.33*	−0.48	−0.19	−0.32
	(3.27)	[0.19]	[0.26]	[0.26]	[0.25]
Monthly transactions amount	1.10	−0.03	−0.16	−0.02	0.08
	(3.34)	[0.08]	[0.12]	[0.11]	[0.11]
Months with seller shortfall	3.85	−0.10	−0.25	−0.03	−0.03
	(3.09)	[0.18]	[0.25]	[0.25]	[0.23]
Share of extensive margin	0.87	0.03	0.02	0.02	0.05
	(0.31)	[0.02]	[0.03]	[0.03]	[0.03]
Buyer characteristics					
Total output	14.61	−2.62	−1.12	−11.38	3.98
	(57.19)	[5.16]	[4.30]	[8.72]	[3.76]
Total input	11.47	−1.63	−1.62	−6.84	3.10
	(48.39)	[3.18]	[3.68]	[5.18]	[3.10]
Share of final sales	0.51	−0.02	−0.03	−0.01	−0.02
	(0.44)	[0.03]	[0.03]	[0.04]	[0.03]
Audited in 2016	0.06	0.00	0.01	−0.01	0.01
	(0.24)	[0.01]	[0.02]	[0.02]	[0.02]
Unique trading partners	48.85	−1.53	−2.96	−0.36	−1.28
	(81.84)	[4.82]	[6.03]	[6.38]	[6.82]
Unique clients	19.24	0.01	0.83	−1.30	0.45
	(68.03)	[4.03]	[4.78]	[5.27]	[5.94]
Unique suppliers	30.63	−1.77	−4.31	0.76	−1.72
	(34.07)	[2.01]	[2.79]	[2.66]	[2.58]
Seller characteristics					
Total output	27.90	1.57	−0.77	3.06	2.37
	(63.65)	[3.80]	[5.50]	[4.81]	[4.79]
Total input	24.49	2.35	0.56	3.37	3.07
	(56.52)	[3.14]	[4.50]	[4.24]	[4.17]
Share of final sales	0.56	0.00	0.03	−0.05	0.01
	(0.41)	[0.02]	[0.03]	[0.03]	[0.03]
Audited in 2016	0.04	−0.01	−0.03	0.00	−0.01
	(0.21)	[0.01]	[0.02]	[0.02]	[0.02]
Unique trading partners	65.02	−6.71	−5.13	−5.51	−9.28
	(156.90)	[9.72]	[11.22]	[14.36]	[12.35]
Unique clients	45.41	−5.16	−2.12	−3.91	−9.13
	(145.20)	[9.13]	[10.29]	[13.59]	[11.48]
Unique suppliers	21.49	−1.79	−3.27	−1.80	−0.42
	(28.95)	[1.72]	[2.41]	[2.31]	[2.14]
P-value of joint Ftest		0.50	0.55	0.69	0.49
Observations	2380	3812	1281	1191	1340
Links	494	741	242	238	261

Notes: This table reports descriptive statistics for firms in the study sample and assesses 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 differences between each treatment arm and the control group, with standard errors from tests of equality of means in square brackets. The top panel reports characteristics at the seller–buyer link level. The second and third panels report buyer and seller characteristics, respectively, at the firm level. Monthly transaction value is defined as the highest amount reported by either trading partner. Months with seller shortfall is the number of months in which a seller shortfall is observed. Share of extensive margin is the share of seller shortfall instances in which the seller reports no trade with the buyer. Share of final sales is the ratio of final sales plus sales to non-VAT firms to total sales. All monetary values are expressed in thousands of U.S. dollars. *Source*: Data from monthly VAT returns submitted to the URA between March and December 2017.

Specifically, this section focuses on how treatment affects reporting within seller-buyer links. Firm-level outcomes, including the impact on overall tax liabilities, are taken up 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:

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

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 also report the average treatment effect across all three treatment arms, replacing the three treatment-arm dummies from (1) with a single indicator T_i:

$$Y_{it} = \alpha + \beta_T T_i + \delta_t + \xi_{it}, \tag{2}$$

where β_T is the coefficient of interest.

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), 23.5% of discrepancies are corrected in treated links, against 1.2% in the control group—an almost twenty-fold increase, significant at the 1% level. Breaking it down by treatment arm, we find that the effect is largest (30.4pp) for the "Buyer and Seller" treatment, followed by the "Seller only" treatment (24.7pp), and smallest for the "Buyer only" treatment (11.5pp). The estimates are statistically significant at the 1% significance level in

all treatment arms. Appendix Table D1 displays the full regression results and shows that the "Buyer only" treatment is significantly smaller than both the "Buyer and Seller" and "Seller only" treatment, while the latter two do not differ significantly.[17]

In panel (b) of Figure 3, because we focus on discrepancies not listed in the letter, the sample is restricted to the 530 seller-buyer links for which there were more than three discrepancies in the pre-treatment period. Even though these discrepancies are not mentioned, we observe a significant effect on the correction of these discrepancies. The overall effect is a 10.1pp increase in the share of reduced discrepancies (bottom bar), reaching up to 13.4pp in the "Buyer and Seller" group and 13.9pp in the "Seller only" group, compared to 1.5% in the control group.[18] 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.

While the resulting correction rates may seem low, they are higher than those observed by Carrillo, Pomeranz and Singhal (2017) in a similar intervention in Ecuador (between 10% and 20%), a country with higher income per capita than Uganda. Thus, the correction rates we observe are comparatively large. 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. Upon receipt of the letter, firms had three options: (i) correcting the discrepancy, (ii) contacting the URA to inform them why there was a discrepancy, and (iii) doing nothing. 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 C2. 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 treated links respond to the treatment. The reaction rate is highest in the "Buyer and Seller" 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 two arms (48% and 50%, respectively).[19] 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

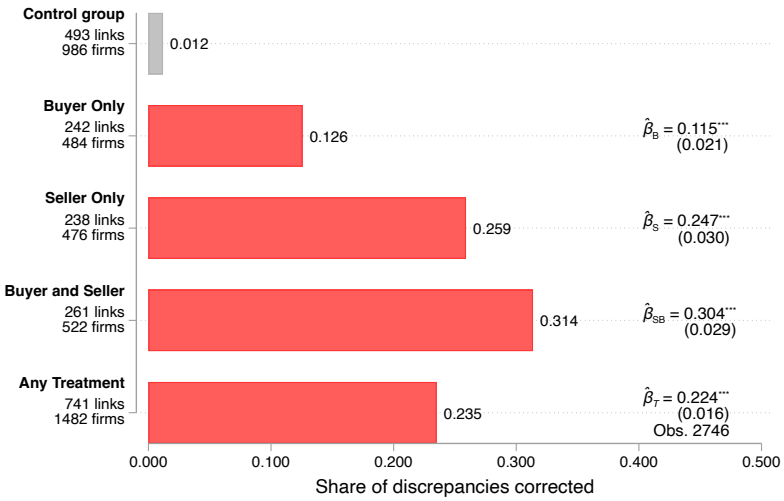
[17] In Appendix Table D1 we also present the results from an alternative specification, where the outcome variable is defined more narrowly and takes value one only if the discrepancy has been fully resolved through the corrections, or "full correction." As shown in columns 5–6, The same pattern of results holds, with slightly muted effects. Figure D1 shows how the distribution of reporting discrepancies shifts to the right after corrections in the treatment group, but not in the control.

[18] Appendix Table D1, columns 7 and 8, shows qualitatively similar results for full corrections, though smaller in magnitude.

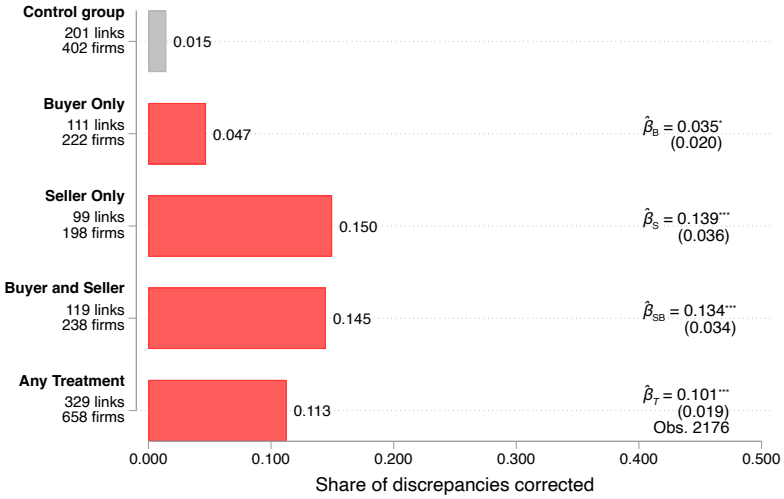
[19] These statistics offer suggestive evidence rather than causal effects, since equivalent figures are unavailable for control firms. Qualitative evidence from discussions with the URA suggests that no follow-up actions were taken with firms beyond those captured in the communication tracker.

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

(a) Listed Discrepancies



(b) Unlisted Discrepancies



Notes: 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 "Buyer and Seller" treatment group. The bottom bar reports the share aggregating all three treatment groups. To the right of each bar, we report the $\hat{\beta}_h$ coefficients from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced. $\hat{\beta}_T$ 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, in parentheses, are clustered at the link level. *p< 0.10; **p< 0.05; ***p< 0.01. All regression results are shown in Appendix Table D1. *Source*: Data from monthly VAT returns submitted to the URA.

characteristics. We return to a discussion of this in Section 5.2. Power calculations in Appendix B.3 confirm that the study is well powered to detect effects of the magnitudes reported above.

4.1.2 Asymmetric effects for sellers and buyers

A core prediction of the conceptual framework is that sellers are more likely to respond than buyers (Predictions 2a and 2b). To test this, we examine which firm in a treated link makes the correction. Results are displayed in Figure 4. We estimate regressions (1) and (2) on listed discrepancies, with the outcome now equal to one if the discrepancy is corrected by the seller (top panel) or the buyer (bottom panel). We focus on listed discrepancies.

The correction rate is systematically higher for sellers. When both the buyer and seller are treated, the correction rate attributable to sellers is 25.1pp larger than the control group (top panel), compared to only 3.2pp for buyers (bottom panel). The same pattern applies when only sellers receive letters: 22.6pp increase for sellers vs. 1.3pp for buyers.

Notably, sellers respond more even when they do not receive a letter. They are 8.2pp more likely to correct when only their buyer receives the letter—a more than sevenfold increase relative to the control group. In contrast, buyer corrections in that treatment arm rise by only 2.6pp.[20]

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.[21] This result also shows that firm's treatment can shift its trading partner's tax reporting even against that partner's own interest: when buyers receive letters, sellers correct, which raises the seller's own tax liabilities.

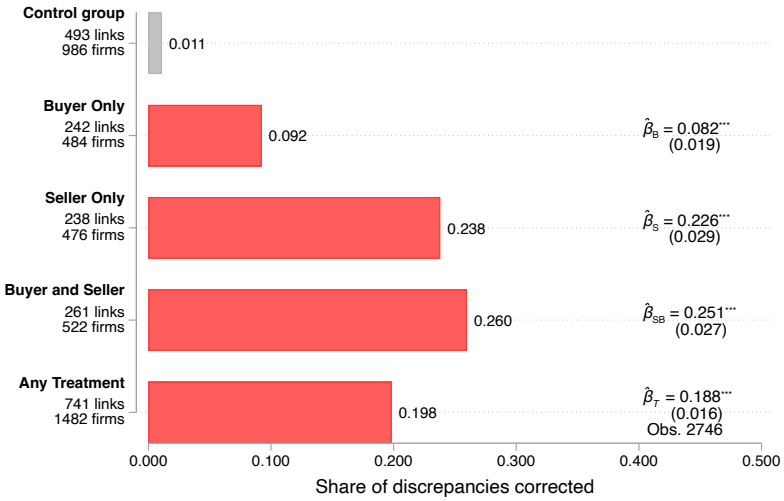
Second, the stark asymmetry in firms' responses suggests that sellers are more likely to be responsible for the discrepancies (Prediction 2a). This is in line with our conceptual framework, which highlights the role played by final sales. Among firms that contact the URA, 79% of buyers in the "Buyer only" treatment say that the discrepancy is due to their trading partner. In contrast, the most common feedback by sellers in the "Seller only" treatment is a stated intention to amend the return (30%), followed by the claim that the transaction had

[20]Full regression results are reported in Appendix Table D2.

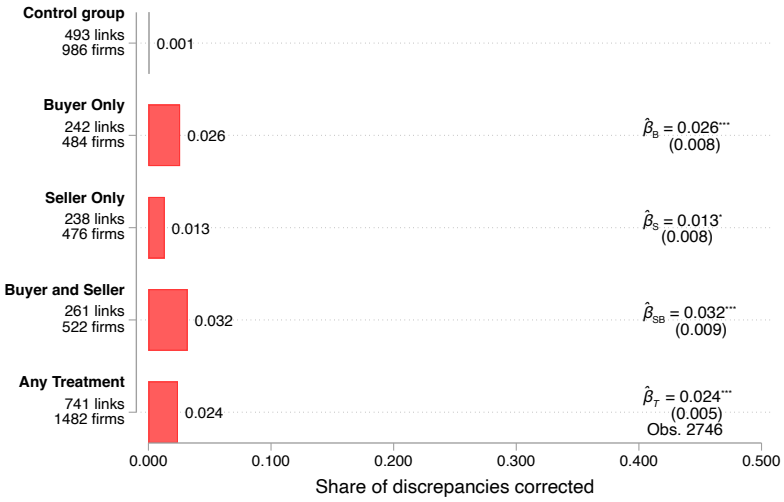
[21]We assume the communication occurs primarily between firms directly. While tax preparers may also facilitate the flow of information, we cannot study that channel in this context because the identity of the individual submitting the tax return is not consistently recorded.

Figure 4
Within-Link Asymmetric Effects: Who Corrects Discrepancies?

(a) Correction by Seller



(b) Correction by Buyer



Notes: This figure reports the effect of the letter on corrections made by each firm (buyer and seller) 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 "Buyer and Seller" treatment group. The bottom bar reports the share aggregating all three treatment groups. To the right of each bar, we report the $\hat{\beta}_h$ coefficients from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced. $\hat{\beta}_T$ 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, in parentheses, are clustered at the link level. *p< 0.10; **p< 0.05; ***p< 0.01. All regression results are shown in Appendix Table D2. *Source*: Data from monthly VAT returns submitted to the URA.

been erroneously classified as final sales (29%).[22]

Finally, for both seller and buyer corrections, the effect is stronger when both firms in the link are treated. However, the difference from the "Seller only" arm is not statistically significant, and the joint effect is smaller than the sum of the "Seller only" and "Buyer only" effects.

4.1.3 Effects on corrections outside treated links

We next test for outside-link effects, as laid out in Prediction 3a. We focus on sellers, since our results from Figure 4 show that corrections are primarily made by sellers. We estimate regressions (1) and (2) with the outcome equal to one if the seller reduces the discrepancy. Results are reported in Figure 5. The sample consists of seller shortfall discrepancies between the sellers from the Study Sample and their other 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 12,554 link-month observations between March and December 2017.[23]

We find a 3.4pp increase in the outside-link correction rate across all treatment arms—an elevenfold increase, significant at the 5% level. The effect is driven by the "Buyer and Seller" arm, which produces a 8.6pp increase (significant at the 10% level); the "Buyer only" and "Seller only" coefficients are both small in magnitude and not statistically significant.[24]

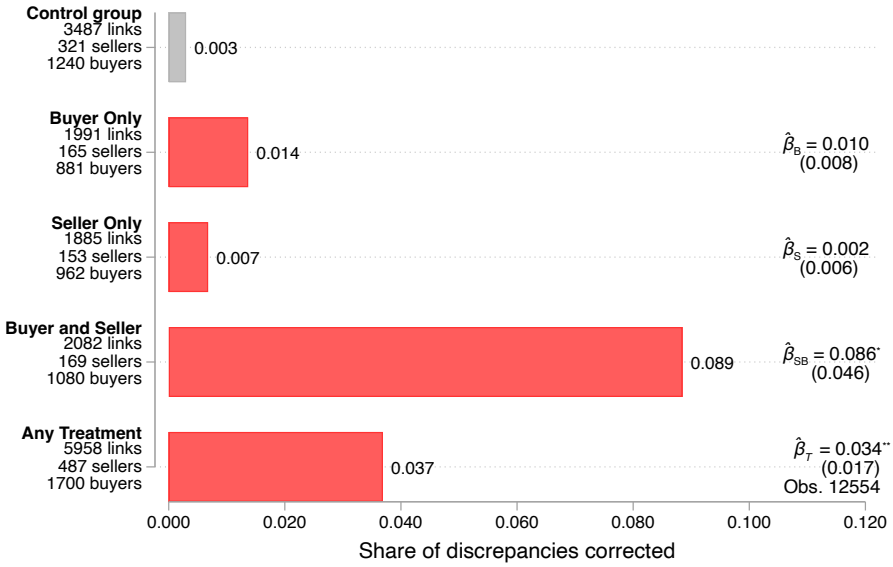
While direct effects do not differ significantly between "Seller only" and "Buyer and Seller" (p-value = 0.163, Appendix Table D1), sending letters to both firms leads to significantly larger spillover effects on untreated links compared to targeting sellers alone (p-value = 0.072, Appendix Table D6). 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 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.

[22]See Appendix Figure C3 for a detailed summary of firms' communications with the URA in response to the treatment.

[23]We show regression results for all months in Appendix Table D5, where the sample includes 45,517 observations. Results are smaller in magnitude but qualitatively similar

[24]The full regression results underlying Figure 5 are shown in Appendix Table D6 (cols. 1-2). Columns 3-4 of that table re-estimate at the seller level, addressing the concern that effects could be driven by highly connected sellers. The results are qualitatively similar. We confirm in Appendix B.3 that the study is well-powered to detect the effects shown in Figure 5. Appendix Table D5 reports specifications using all months and alternative outcomes, again yielding comparable results. These results show sellers extend their expectation of cross-checking beyond the partners named in the letter.

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



Notes: 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 "Buyer and Seller" treatment group. The bottom bar reports the share aggregating all three treatment groups. To the right of each bar, we report the $\hat{\beta}_h$ coefficients from regression (1) where the outcome variable is an indicator set to 1 if a discrepancy is reduced by the seller. $\hat{\beta}_T$ is the coefficient for the aggregate effect of any treatment from regression (2). Standard errors, in parentheses, are clustered at the seller level. *p< 0.10; **p< 0.05; ***p< 0.01. All regression results are shown in Appendix Table D6. *Source*: Data from monthly VAT returns submitted to the URA.

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. Second, that spillovers do not extend beyond one degree of separation in the network. If either fails, but are not controlled for, our estimates could be biased. Thanks to the richness of our network data, we can control for such potential spillovers and assess whether this affects our estimates.

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.[25] 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 E1.

Appendix E 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 E1, columns 1–3), and through second-degree connections (Table E2, columns 1–3). This approach to account for spillovers follows recent advances in the econometric literature (see, e.g., Vazquez-Bare, 2023).[26] In both cases, we find that the estimated treatment effect remains statistically unchanged when these controls are included. Furthermore, the coefficient estimates on the indirect exposure variables are jointly insignificant in both cases.[27]

Together, these checks support the robustness of our link-level estimates: neither spillovers through non-seller-shortfall links nor higher-order network connections meaningfully change the estimated treatment effect. This in turn lends support to the iterative sampling procedure and the link-level cluster definition on which our design rests.[28] Previous studies of enforcement interventions typically assume such indirect effects away rather than testing for them.

These checks concern internal validity. Whether the estimates generalize beyond the Study Sample is a separate question. As discussed in Section 3.2.1, our sample selection procedure leads the Study Sample to contain firms that are more connected than the typical VAT-registered firm in Uganda. We therefore reweigh the Study Sample using entropy balancing (Hainmueller, 2012) so that its degree distribution matches that of the Full Sample, and re-estimate our main link-level specification. As shown in Table D4 of Appendix D.2, the reweighted estimates are close to the unweighted ones—the pooled effect declines from

[25]The sellers (resp., buyers) in our Study Sample have an average of 61.0 (resp., 47.9) unique trading partners. Previous randomized trials involving networks are typically implemented in settings with fewer connections. As an illustration, Vazquez-Bare (2023) considers spillovers for up to five peers in the empirical application.

[26]Our regression is analogous to the correctly-specified interacted reduced form-linear in means regression discussed by Vazquez-Bare (2023). We additionally implement a pooled specification where an indicator takes value one if any trading partner is treated, also following Vazquez-Bare (2023). While this is less appropriate in our setting where the variable is one for the vast majority of links, our results remain qualitatively similar, see column 4 in Table E1 and E2.

[27]In Table E3, we provide similar checks for the outside-link effects. We find that the effects shown in Figure 5 are slightly larger when accounting for indirect exposure, suggesting that our baseline estimates, if anything, slightly underestimate the true effects.

[28]While 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.

22.4 to 20.1 percentage points—and the pattern across treatment arms is unchanged. This suggests that the responses we document generalize to the universe of VAT paying firms.

4.2 Impact on Post-treatment Reporting Behavior

So far, we have shown that firms correct past discrepancies, both within and outside treated links, indicating that they update their expectations about the URA's enforcement capacity. We next examine whether the letter also produces lasting changes in firms' reporting behavior in the months following the treatment. To prevent post-treatment amendments from contaminating the pre-treatment baseline, we restrict the pre-treatment returns to those originally filed before 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:

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

where Y_{it} is the outcome variable of interest for link i in month t, T_i is an indicator equal to 1 if link i is in any treatment arm, 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.[29] The sample includes all monthly trades reported by a given seller-buyer link. We also implement the analogous difference-in-differences (DiD) estimation, which gives the average treatment effect on the treated across all post-treatment months. Our DiD specification is formally written as:

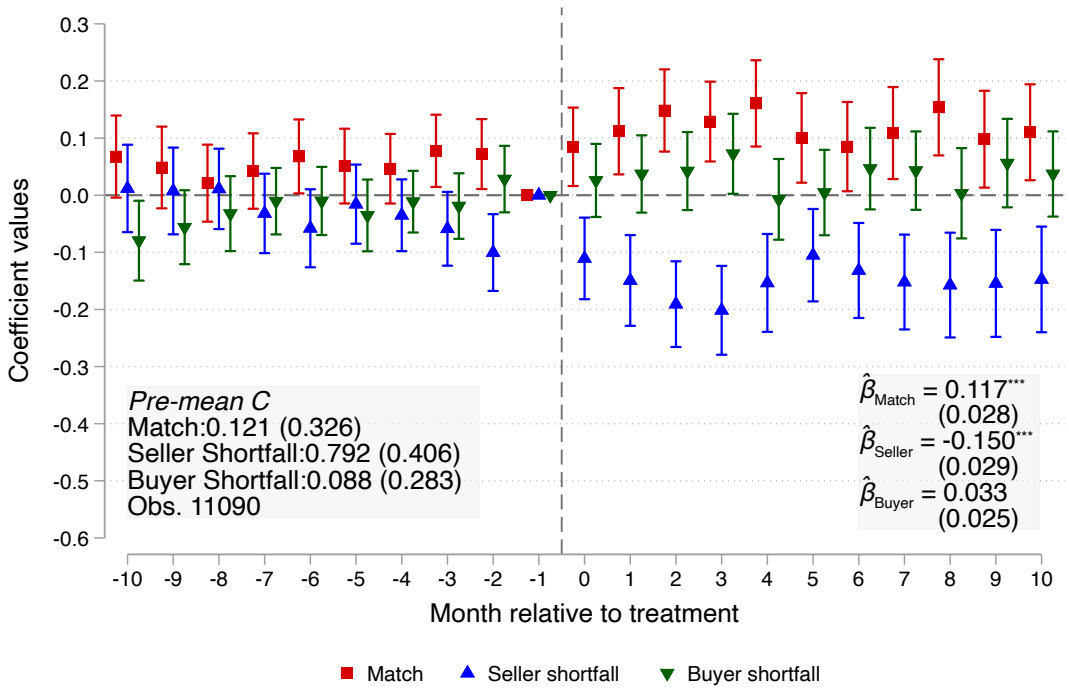
$$Y_{it} = \tau \left(T_i \cdot 1_{j>0}\right) + \delta_t + \gamma_i + \sum_{j=-10}^{-1} \tau_j \left(m_j \cdot T_i\right) + \xi_{it}, \tag{4}$$

where T_i is an indicator taking value one for treated links. This is interacted with $1_{j>0}$, a dummy taking value one for all post-treatment months. The coefficient estimate for τ captures the cumulative treatment effect for the entire post-treatment period between March and December 2018. Following Roth and Sant'Anna (2023), we include separate controls for

[29]The 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.

each pre-treatment period to gain precision.[30]

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



Notes: 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). For pre-treatment months, outcomes use the latest returns filed prior to the intervention. For post-treatment months, the latest returns. This excludes any amendments triggered by the intervention from the pre-treatment outcomes. 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, in parentheses, are clustered at the link level and the bars report 95% confidence intervals. *p< 0.10; **p< 0.05; ***p< 0.01. See Appendix Tables F1 and F2 for the full regression results. *Source*: Data from monthly VAT returns submitted to the URA.

Figure 6 displays the event-study results. We consider three outcome variables in turn: indicator for whether the link displays seller shortfall, buyer shortfall, or matching reports (i.e., no discrepancy). Observations are conditional on trade being reported between the two firms, making the sample an unbalanced panel. We also report the DiD coefficient from (4) at the bottom right of the figure.

[30]We also estimate a specification pooling all pre-treatment months into a single baseline, rather than controlling for each separately. The results are robust to this alternative, both within links (Appendix F.1) and outside them (Appendix F.3). We report the specification with separate pre-treatment controls throughout.

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. The decline is accompanied by a 11.7pp (97%) increase in the probability of matching reports. The buyer shortfall rate rises slightly, but not significantly. These results indicate that firms update their beliefs about the tax authority's monitoring capacity and adjust their reporting within treated links accordingly, consistent with Prediction 1c.

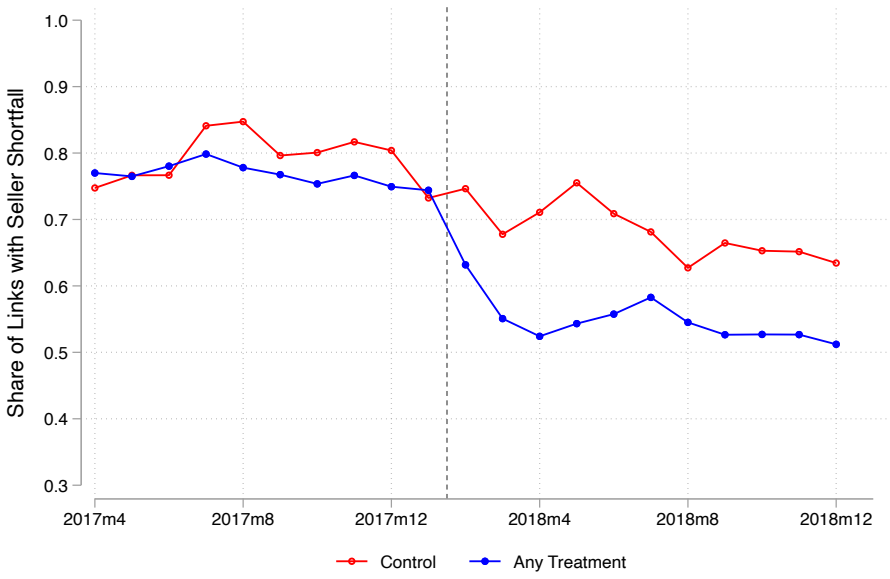
Figure 6 shows a small difference in seller-shortfall and match rates between the two groups in the months before the intervention. We provide evidence that this reflects idiosyncratic variation in the control group unrelated to our intervention rather than a pre-trend. Figure 7 reports the raw proportion of seller shortfall occurrences by month. The control group displays a slightly higher share of seller shortfall in pre-treatment months and then a sharp decrease in the last pre-treatment month (January 2018). Since the intervention was not pre-announced, this drop cannot reflect anticipation effects. The treatment group, by contrast, is stable before treatment and falls sharply in the three months after the letters were delivered. While both the treatment and control groups trend downward in the share of seller shortfall cases in the months following the intervention, only the treatment series features a sharp drop immediately afterward. We confirm the robustness of our estimates using the "honest" difference-in-differences estimator of Rambachan and Roth (2023), reported in Figure F3 of Appendix F.2 (with additional checks in Figures F1 and F2).

Figure 8 provides additional evidence on the mechanisms driving the post-treatment behavior. Panel (a) shows that the seller shortfall amount (conditional on being positive) does not change significantly. The reduction in seller shortfall therefore operates through an extensive margin—sellers reporting more transactions—rather than an intensive margin in reported amounts.

One potential concern is that this improved reporting behavior masks collusion between firms. If the buyer and seller agree to report lower amounts following the intervention, this would show up as a "match" in the data. We provide evidence against this. Panel (b) of Figure 8 shows that, conditional on trading, the intervention does not change the size of reported transactions within links.[31] 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."

[31]Transaction size is defined as the maximum amount reported by the buyer or the seller.

Figure 7
Raw Proportion of Seller Shortfall by Month in Treatment and Control Links

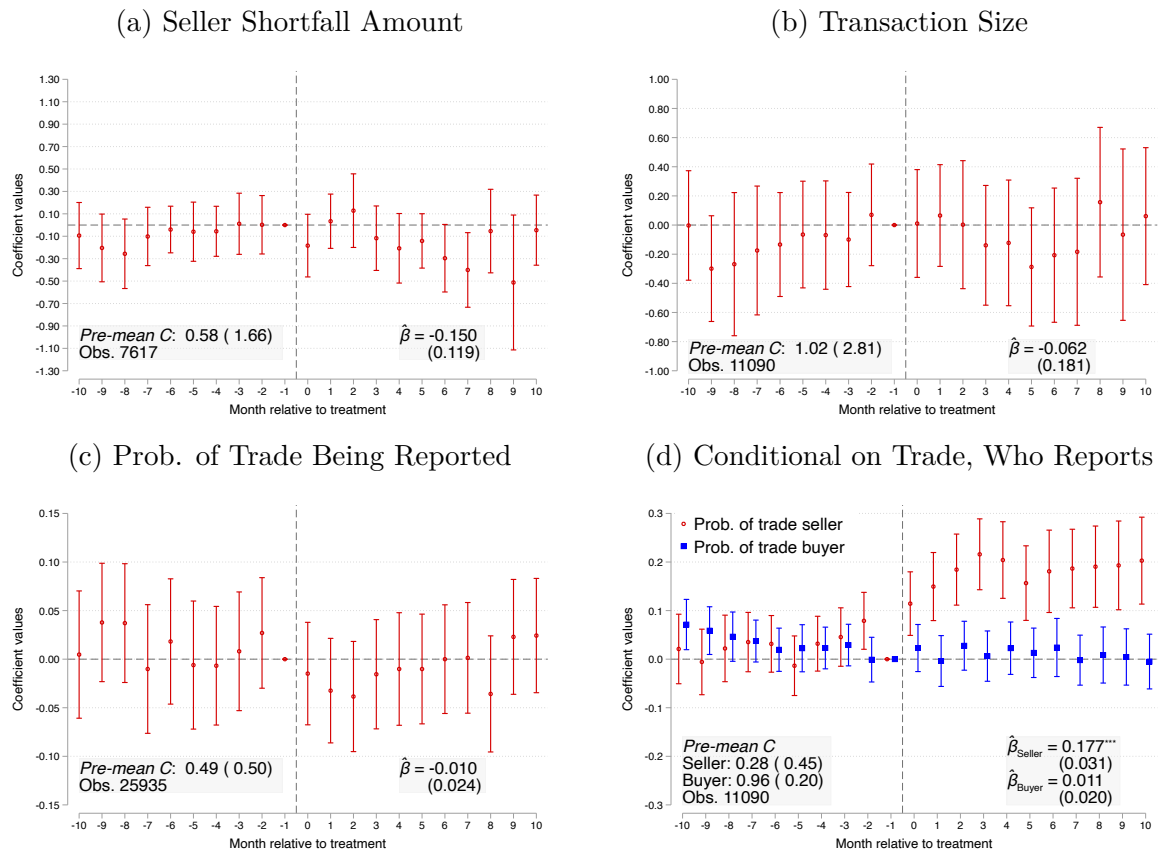


Notes: This figure reports the raw proportion of seller-buyer links displaying seller shortfall in every month, separately for the control group (in red) and the treatment group (in blue). The dashed vertical line represents the time of the treatment. Outcomes are conditional on trade being reported within the link. *Source*: Data from monthly VAT returns submitted to the URA.

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 8 shows that, conditional on trade being reported within the link, there is a 17.7pp (63%) 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 8 with the seller-shortfall finding in Figure 6, recall that "trade being reported" counts if at least one firm in the link reports the transaction. Panel (d) of Figure 8 shows that the decrease in seller shortfall shown in Figure 6 is entirely driven by sellers reporting at higher rates—specifically, sellers starting to report transactions that their buyers were already reporting. Since the buyer's report alone was already sufficient to count the transaction as reported, the probability of trade being reported (panel c) does not change. In other words, consistent with our theoretical framework, buyers report transactions at a stable rate, while sellers increase their reporting rate as a result of the intervention.

Figure 8
Post-Treatment Effect of the Letter on Within-Link Trade



Notes: 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 $\hat{\tau}_j$ 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 $\hat{\tau}$ 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, in parentheses, 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 F1 for full regression results. *Source*: Data from monthly VAT returns submitted to the URA.

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 repeat the event-study analysis on the sample of outside links described in Section 4.1.3: seller shortfall discrepancies between the sellers from the Study Sample and their other buyers, excluding the buyer mentioned in the letter. We find no significant effects on these outside-link reporting discrepancies in the post-treatment period.[32] Hence, while firms update their beliefs about the URA's ability to cross check transactions with the trading

[32]Results are shown in Appendix Figure F4 and Appendix Tables F5, F6, and F7.

partner mentioned in the letter, this does not translate into sustained adjustments of their reporting behavior with other firms.

5 Firm-Level Effects

This section considers firm-level rather than link-level outcomes. In particular, we focus on analyzing the results for sellers, since the previous section shows that they drive the observed responses. Corresponding results for buyers are reported in Appendix G.

5.1 Firm-Level Conceptual Predictions

Our conceptual framework in Section 3.1 extends naturally to firm-level outcomes, since link-level responses aggregate into changes in firms' overall VAT liability. At the firm level, however, firms have an additional margin of adjustment—they can relabel sales across reporting categories. We therefore formulate an additional set of predictions, focusing on firms treated as sellers in the Study Sample.

Prediction Set 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 characteristics.

5.2.1 Effect of treatment on VAT liability

To examine how the intervention affects amounts reported in the monthly VAT returns prior to treatment, we estimate regressions similar to equations (1) and (2), with the unit of

observation now a firm-month rather than a link-month.

Table 2 reports the results for four outcomes: changes in B2B sales, final sales, taxable inputs, and overall VAT liability (all in thousands of US dollars).[33] The sample is restricted to sellers' outcomes in the months listed in the letters. In columns 1 and 2, we look at the change in reported B2B sales. We find that, across all treatment arms, the letter raises monthly B2B sales by \$522, consistent with Prediction 4b. Breaking down by treatment arm, we find that the "Buyer and Seller" treatment has the largest effect (\$701), consistent with the previous link-level results. Columns 3 and 4 examine the reported final sales and find a decrease of \$309 (\$447 in the "Buyer and Seller" group). About 60% of the increase in B2B sales is thus offset by a reduction in reported final sales.[34] Firms are thus exploiting the final sales loophole, in line with Prediction 4a. Columns 5 and 6 show no significant impact on reported taxable inputs. Finally, columns 7 and 8 report the overall effect on VAT liability. The letter raises VAT liability slightly by \$14 for all treatments combined, which is less than 1% of the median Study Sample firm's monthly VAT liability. Correspondingly, the effect is \$20 in the "Seller only" arm and \$18 in the "Buyer and Seller" arm.[35]

To confirm the role of offsetting via final sales, we replicate the analysis on the subsample of sellers with no final sales before the intervention (Table G5). The effect on VAT liability is more than three times larger (\$44) for these firms, consistent with the mechanism: when the final-sales margin is unavailable, the offset cannot occur.

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 net effect on reported VAT liability is modest. Nonetheless, firms are not offsetting all of the increase, and hence the intervention successfully increased VAT liability for the targeted periods, consistent with Prediction 4b. Additionally, firms do not compensate by increasing their reported inputs, suggesting that input overreporting is not their preferred evasion margin in our context, where evasion is detected through B2B sales cross-checking, in contrast to what is documented for corporate income taxes in Ecuador by Carrillo, Pomeranz and Singhal (2017).

[33] We focus on taxable sales and purchases. Monthly VAT liability is the difference between output VAT charged and input VAT paid, excluding tax credits carried forward from past negative liabilities.

[34] The 60% is calculated as $309/522 \approx 447/701 \approx 0.6$.

[35] The effect on VAT liability is above the minimum effect our experimental design can detect, see Appendix B.3.

Table 2
Effect of Treatment on Firm-Level VAT Liability (Sellers)

Dependent variable:	ΔB2B Sales		ΔFinal Sales		ΔTaxable Inputs		ΔVAT Liability	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	0.522***		-0.309***		0.002		0.014***	
	(0.060)		(0.037)		(0.003)		(0.004)	
Buyer only		0.156***		-0.080***		0.005		0.003
		(0.058)		(0.029)		(0.004)		(0.005)
Seller only		0.701***		-0.391***		0.003		0.020**
		(0.119)		(0.074)		(0.003)		(0.008)
Buyer and Seller		0.701***		-0.447***		-0.001		0.018***
		(0.113)		(0.076)		(0.005)		(0.006)
R-squared	0.036	0.056	0.035	0.058	0.007	0.007	0.009	0.012
Observations	2738	2738	2738	2738	2738	2738	2738	2738
No. of Firms	1232	1232	1232	1232	1232	1232	1232	1232
Pre-T Mean of Dep. in Control	60.732	60.732	414.630	414.630	414.475	414.475	10.960	10.960
Mean of Diff. in Control	0.016	0.016	-0.003	-0.003	0.001	0.001	0.005	0.005
Pre-T Median of Dep. in Control	13.005	13.005	40.083	40.083	102.450	102.450	1.813	1.813
P-value of β_S=β_B		0.000		0.000		0.539		0.057
P-value of β_{SB}=β_B		0.000		0.000		0.286		0.035
P-value of β_{SB}=β_S		1.000		0.597		0.538		0.878
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 (2) (odd columns) and (1) (even columns), where the unit of analysis is at the firm-month level. 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 across the entire estimation period. *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< 0.05; ***p< 0.01. Appendix Table G1 reports the equivalent results for buyers. *Source*: Data from monthly VAT returns submitted to the URA.

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, two firms can still be directly connected even when their treated links are not. For example, suppose link ij between firms i and j is treated. The randomization ensures that the first-degree link jk cannot be treated, but firm k may still be treated through another relationship kl (its second-degree connection). In this case, firms k and j are both treated and directly connected, even though their respective treated links, ij and kl, are not. Essentially, what constitutes a second-degree connection at the link level becomes a first-degree connection at the firm level. Appendix Figure E1 provides an illustration of these connections.

To verify that such indirect exposure does not bias our firm-level estimates, we test whether potential spillovers affect treated and control firms equally. This does not imply the absence of spillovers, but if they are symmetric across the two groups they cancel out, yielding consistent estimates of the intervention's direct effect.[36] Following Vazquez-Bare (2023), we re-estimate regression (2) at the firm level including as controls the share of treated partners, separately for treated and untreated firms. Detailed results are presented in Appendix Table E4. The coefficients of interest remain stable and statistically unchanged when including these controls.

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 use specifications analogous to (3) and (4), with observations now at the firm level. Panels (a), (b), and (c) of Figure 9 show no significant effects on reported B2B sales, final sales, and total inputs. Panel (d) reports the effect on monthly VAT liability. Most event-study coefficients are positive but none is individually significant. The difference-in-differences estimate is \$457, which is not statistically significant either. Despite the lack of statistical significance, this coefficient is economically sizable at about 14% of the average monthly VAT liability.[37]

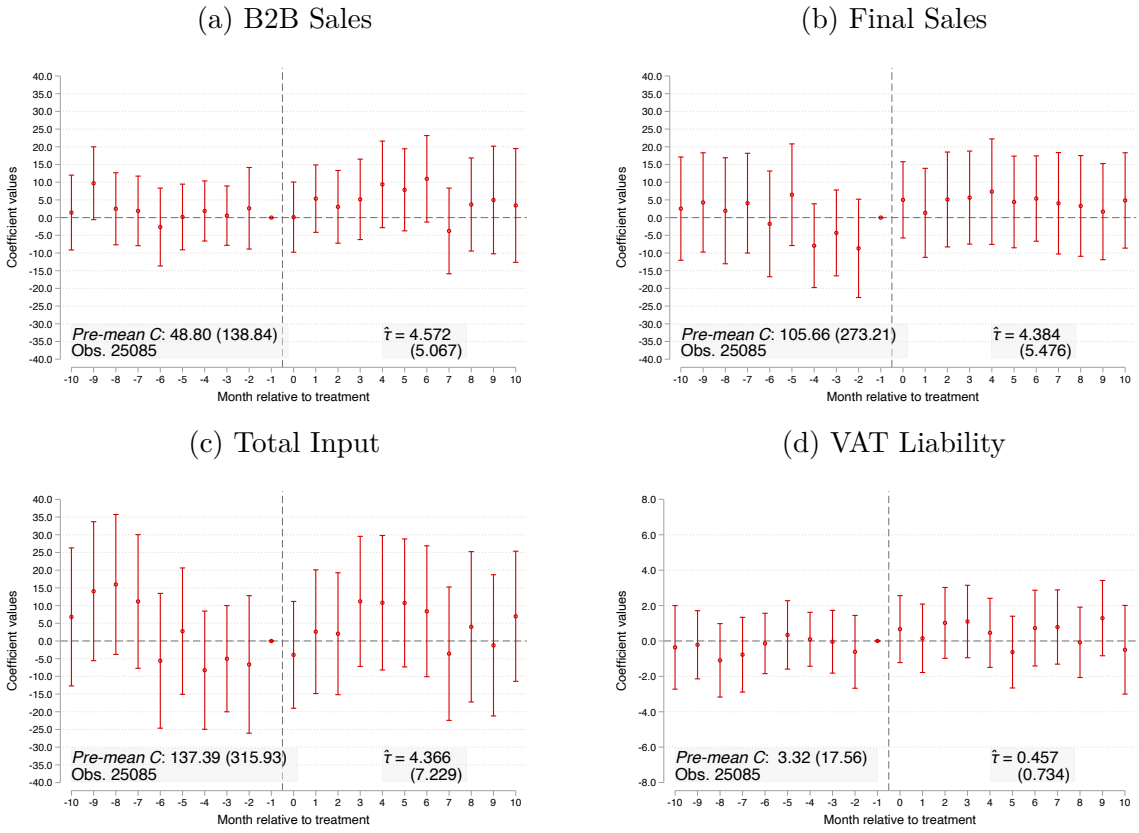
In summary, the intervention leads to higher tax liabilities for amended past returns (Table 2) and triggers a sustained decline in the share of seller-shortfall with the trading partner mentioned in the letter (Figure 6). We do not, however, detect increases in firms' aggregate B2B sales or VAT liabilities in the months after treatment (Figure 9). We interpret

[36] Note 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).

[37] In Table G3, we show that the results do not change when taking indirect exposure into account.

this as the intervention successfully inducing corrections for targeted transactions, but not generating broader increases in subsequent VAT liability. That said, our study lacks statistical power to detect an increase in VAT liability of the size we estimate.[38]

Figure 9
Post-treatment Effects on Firm-Level VAT Liability (Sellers)



Notes: 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 $\hat{\tau}_j$ 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 $\hat{\tau}$ 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, in parentheses, are clustered at the seller level and the bars report 95% confidence intervals. *p< 0.10; **p< 0.05; ***p< 0.01. See Appendix Table G2 for full regression results. *Source*: Data from monthly VAT returns submitted to the URA.

[38]We present detailed power calculations in Appendix B.3. The minimum detectable effect is more than four times larger than the point estimate we find in Figure 9(d), indicating that we do not have sufficient statistical power for this specific test.

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. We estimate a two-stage least squares model with the following first- and second-stage equations:

$$h_i = \lambda T_i + \epsilon_i \tag{5}$$

$$h_i \times k_i = \theta h_i + \nu_i, \tag{6}$$

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 coefficient θ captures the average baseline value of the characteristic of interest for the complying 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 than those in the Study Sample (61.8% vs. 56.5%; p-value of 0.094). This is consistent with our conceptual framework: sellers with a larger share of final sales have more scope to have hidden B2B sales, and therefore more discrepancies to correct. Compliers are also less connected to other firms at baseline (significant at the 10% level). We find no significant differences in total reported sales, the share of monthly tax returns with negative VAT liability, firm age, whether they were audited in the year prior to the intervention, or sector.

Panel B of Table 3 reports analogous results for buyers. Buyers in complying 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 G4, 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

Table 3
Heterogeneity Analysis: Baseline Characteristics of the Compliers

	Average characteristics		Difference	P-value
	Compliers (1)	Sample (2)	(3)	(4)
Panel A: Seller Characteristics				
Sales	135.640	171.486	-35.846	0.142
Final sales ratio	0.618	0.565	0.053	0.094
Share of negative returns	0.407	0.446	-0.039	0.189
Existence length	90.887	97.387	-6.500	0.263
Degree	55.333	69.092	-13.759	0.076
Audited in 2016	0.043	0.053	-0.010	0.479
Sector - Agriculture	0.016	0.012	0.004	0.639
Sector - Construction	0.067	0.065	0.002	0.918
Sector - Manufacturing	0.115	0.125	-0.010	0.680
Sector - Retail	0.193	0.172	0.021	0.484
Sector - Service	0.324	0.275	0.049	0.176
Sector - Wholesale	0.286	0.349	-0.063	0.067
Panel B: Buyer Characteristics				
Sales	150.564	103.806	46.758	0.183
Final sales ratio	0.499	0.529	-0.030	0.380
Share of negative returns	0.407	0.401	0.006	0.831
Existence length	79.491	79.166	0.325	0.954
Degree	62.082	49.759	12.323	0.153
Audited in 2016	0.064	0.056	0.008	0.616
Sector - Agriculture	0.013	0.018	-0.004	0.730
Sector - Construction	0.176	0.130	0.046	0.107
Sector - Manufacturing	0.091	0.079	0.012	0.526
Sector - Retail	0.168	0.163	0.006	0.838
Sector - Service	0.329	0.404	-0.076	0.044
Sector - Wholesale	0.206	0.199	0.007	0.823
Panel C: Pair Characteristics				
High trading volume	0.487	0.450	0.037	0.335
Initial discrepancy	-0.669	-1.105	0.435	0.000

Notes: This table reports the average characteristics of the compliers—correcting at least one discrepancy—(column 1), together with the average characteristics of the sample (column 2), the difference between them (column 3) and the results of a significance test of the difference in means (column 4). Results in column 1 are the θ coefficient derived from estimating equation (5) and (6). We present 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 logarithmic transformation of 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, Manufacturing, Retail, Service, and Wholesale) (viii) whether the number of months traded exceed the median, and (ix) initial value of the discrepancy. *Source*: Data from monthly VAT returns submitted to the URA.

results show that sellers in the wholesale sector and buyers in the service sector are less likely to respond, on average. However, these differences are only marginally significant. Finally, as in the compliers analysis, firms with smaller initial discrepancies are 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 0.06. In column 4, we apply a LASSO estimator to select the baseline characteristics most predictive of correction behavior. None of the covariates from columns 1—3 are selected.

Finally, we examine two dimensions of heterogeneity that have been found to matter in previous studies. Figure G1 plots the correction rate across quintiles of firm size measured by pre-treatment sales of the seller before the intervention. The correction rate is remarkably consistent across the size distribution (0.23 in the lowest quintile, and 0.19 in the highest). This literature has documented size-responsiveness gradients in either direction—larger firms less responsive in certain contexts (Pomeranz, 2015; Brockmeyer, Hernandez, Smith and Kettle, 2019), more responsive in others (Holz et al., 2023). The flat profile we observe instead suggests that the type of evasion our intervention detects is common across the firm-size distribution, implying a low perceived risk to firms of all sizes in the Ugandan context. Figure G2 reports correction rates by quintiles of pre-treatment seller shortfall (our measure of suspected evasion), and finds the rate is similarly stable.

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.

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, Pomeranz and Singhal (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 out of 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 the type of evasion we study. 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

The empirical analysis distinguishes direct from spillover effects. 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), treated firms report an additional $293,670 in B2B trade. Of this, $209,030 (71%) comes from corrections of listed discrepancies (direct effects) and $84,640 (29%) from corrections of unlisted discrepancies (spillover effects). Spillover effects include corrections within the treated link beyond the listed discrepancies, corrections by the firm that did not receive the letter (e.g., the seller in a "Buyer only" link), and corrections involving other trading partners.

To assess the magnitude of the aggregate effect from a causal perspective, we net out comparable corrections made by control firms. Using the same definitions, controls contribute \$9,470 in direct-equivalent corrections and −\$137,430 in spillover-equivalent corrections. The negative sign arises because we count corrections in both directions: amendments can increase or decrease reported B2B sales.[39] The net effect, calculated as corrections by treated firms minus those made by control firms, is an increase of \$421,630 in reported B2B trade, split roughly evenly between direct (47%) and spillover (53%) effects. Thus, spillovers contribute substantially 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 \$96,870 for "Seller only", \$24,460 for "Buyer only", and \$172,330 for "Buyer and Seller". Interestingly, direct effects dominate for "Seller only" and "Buyer and Seller", while spillovers are more important for "Buyer only". In line with the estimates presented in Section 4.1, the Buyer and Seller 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.

Table 4
Aggregate Amounts of Uncovered B2B Sales

	Total	Direct	Spillover
Panel (a): Aggregate Amounts			
All Treated	293.67	209.03	84.64
Control	-127.97	9.47	-137.43
Net (Treated − Control)	421.63	199.56	222.07
Panel (b): Amounts by Treatment Arm			
Seller only	96.87	107.75	-10.87
Buyer only	24.46	-3.38	27.84
Seller & Buyer	172.33	104.66	67.67

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

[39]For 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.

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, across all margins of adjustment, is \$38,332.[40] The total implementation cost—including courier services and URA staff time—was \$5,716 (see Appendix H for details on the calculation). Thus, the intervention generated more than six times its implementation cost in additional revenue. We conclude that, despite the modest increase in VAT liability at the firm level, 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 offset B2B corrections by reducing reported final sales, 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. The large revenue gains documented by Shah (2026) from the more sophisticated automated cross-checking in Pakistan illustrate the potential of such an approach. Our results suggest that even a simpler intervention, relying solely on existing VAT transaction data, can produce substantial compliance gains in settings with limited administrative capacity.

Second, the experiment has informed subsequent URA practices. Following the 2018 experiment, the URA substantially expanded its data cross-checks. In 2021, it 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 out electronic billing machines to record transaction information at the point of sale. These reforms 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 trade with 353 buyers, makes 295 corrections, and uncovers \$5,800 in previously unreported B2B sales, whereas some sellers make no corrections at all. The administration also faces a tradeoff between notifying both firms in a link, which produces the strongest direct response, and notifying only one side which allows to reach twice as many links with the same budget and rely on firm-to-firm communication to reach the untreated partner.

[40] This is calculated by multiplying the average change in tax liability estimated in column 7 of Table 2 times the number of firm-month observations: \$14 $\times$ 2,738 = \$38,332.

Finally, we detect no effects on VAT liability in the months after treatment. We attribute this to the limited scope of the intervention. The intervention only focused on discrepancies with one trading partner, allowing the firm to adjust other margins to offset any increases in VAT liability. In our case the margin was reductions in sales to final consumers. While firms increase their probability of having matched transactions with the indicated trading partner for a sustained period after the intervention, that alone was not sufficient to increase their overall VAT liabilities. Future research could explore whether interventions targeted at more trading partners—similar to what is studied in Shah (2026)—or combining letter interventions with enforcement follow-ups could generate more durable gains in VAT liability.

7 Conclusion

This paper designs and evaluates a randomized trial using transaction-level VAT data from Uganda, 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 the type of VAT evasion targeted by our intervention 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. 2022. "Imports, Exports, and Earnings Inequality: Measures of Exposure and Estimates of Incidence." *The Quarterly Journal of Economics*, 137(3): 1553–1614.

Agha, Ali, and Jonathan Haughton. 1996. "Designing VAT systems: Some efficiency considerations." *Review of Economics and Statistics*, 78: 303–308.

Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P Vasquez. 2022. "The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages." *The Quarterly Journal of Economics*, 137(3): 1495–1552.

Almunia, Miguel, and David Lopez-Rodriguez. 2018. "Under the Radar: The Effects of Monitoring Firms on Tax Compliance." *American Economic Journal: Economic Policy*, 10: 1–38.

Almunia, Miguel, Jonas Hjort, Justine Knebelmann, and Lin Tian. 2018. "Information, Fiscal Capacity and Tax Revenues: An Experimental Evaluation." AEA RCT Registry.

Almunia, Miguel, Jonas Hjort, Justine Knebelmann, and Lin Tian. 2024. "Strategic or Confused Firms? Evidence from Missing Transactions in Uganda." *Review of Economics and Statistics*, 106(1): 256–265.

Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association*, 103(484): 1481–1495.

Baird, Sarah, J. Aislinn Bohren, Craig McIntosh, and Berk Özler. 2018. "Optimal Design of Experiments in the Presence of Interference." *The Review of Economics and Statistics*, 100(5): 844–860.

Battaglini, Marco, Luigi Guiso, Chiara Lacava, and Eleonora Patacchini. 2022. "Tax Professionals and Tax Evasion." mimeo.

Benjamini, Yoav, and Yosef Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1): 289–300.

Bernard, Andrew B., Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes. 2022. "The Origins of Firm Heterogeneity: A Production Network Approach." *Journal of Political Economy*, 130(7): 1765–1804.

Besley, Timothy, and Torsten Persson. 2013. "Chapter 2 - Taxation and Development." In *handbook of public economics, vol. 5.* Vol. 5 of *Handbook of Public Economics*, , ed. Alan J. Auerbach, Raj Chetty, Martin Feldstein and Emmanuel Saez, 51–110. Elsevier.

Besley, Timothy, and Torsten Persson. 2014. "Why Do Developing Countries Tax So Little?" *Journal of Economic Perspectives*, 28(4): 99–120.

Best, Michael, Anne Brockmeyer, Henrik Kleven, Johannes Spinnewijn, and Mazhar Waseem. 2015. "Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan." *Journal of Political Economy*, 123(6): 1311–1355.

Best, Michael, Jawad Shah, and Mazhar Waseem. 2021. "Detection Without Deterrence: Long-Run Effects of Tax Audit on Firm Behavior." mimeo.

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

Brockmeyer, Anne, Giulia Mascagni, Vedanth Nair, Mazhar Waseem, and Miguel Almunia. 2024. "Does the Value-Added Tax Add Value? Lessons Using Administrative Data from a Diverse Set of Countries." *Journal of Economic Perspectives*, 38(1): 107–32.

Brockmeyer, Anne, Marco Hernandez, Spencer Smith, and Stewart Kettle. 2019. "Casting a Wider Tax Net: Experimental Evidence from Costa Rica." *American Economic Journal: Economic Policy*, 11: 55–87.

Brown, Julia, Sebastian Jilke, David Schwegman, Mattie Toma, and Mary Clair Turner. 2025. "Improving Tax Compliance among the Clients of Tax Preparers." Available at SSRN.

Cai, Jing, Wei Lin, and Adam Szeidl. 2024. "Firm-to-Firm Referrals."

Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal. 2023. "Ghosting the Tax Authority: Fake Firms and Tax Fraud in Ecuador." *American Economic Review: Insights*, 5(4): 427–44.

Carrillo, Paul, Dina Pomeranz, and Monica Singhal. 2017. "Dodging the Taxman: Firm Misreporting and Limits to tax enforcement." *American Economic Journal: Applied Economics*, 9(2): 144–164.

Cruces, Guillermo, Dario Tortarolo, and Gonzalo Vazquez-Bare. 2024. "Design of Partial Population Experiments with an Application to Spillovers in Tax Compliance." Working Paper.

Demir, Banu, Beata Javorcik, Tomasz K. Michalski, and Evren Ors. 2022. "Financial Constraints and Propagation of Shocks in Production Networks." *The Review of Economics and Statistics*, 1–46.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot. 2022. "Firm Responses and Wage Effects of Foreign Demand Shocks with Fixed Labor Costs and Monopsony."

Drago, Francesco, Friederike Mengel, and Christian Traxler. 2020. "Compliance Behavior in Networks: Evidence from a Field Experiment." *American Economic Journal: Applied Economics*, 12(2): 96–133.

Duflo, Esther, and Emmanuel Saez. 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *The Quarterly Journal of Economics*, 118(3): 815–842.

Eckles, Dean, Brian Karrer, and Johan Ugander. 2017. "Design and analysis of experiments in networks: Reducing bias from interference." *Journal of Causal Inference*, 5(1): 20150021.

Garriga, Pablo, and Dario Tortarolo. 2024. "Firms as tax collectors." *Journal of Public Economics*, 233: 105092.

Hainmueller, Jens. 2012. "Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies." *Political Analysis*, 20(1): 25–46.

Holz, Justin E., John A. List, Alejandro Zentner, Marvin Cardoza, and Joaquin E. Zentner. 2023. "The $100 million nudge: Increasing tax compliance of firms using a natural field experiment." *Journal of Public Economics*, 218: 104779.

Hoy, Christopher, Mathias Sinning, and Luke McKenzie. 2024. "Improving Tax Compliance without Increasing Revenue: Evidence from Population-Wide Randomized Controlled Trials in Papua New Guinea." *Economic Development and Cultural Change*, 72: 691–723.

Huneeus, Federico, Kory Kroft, and Kevin Lim. 2021. "Earnings Inequality in Production Networks." National Bureau of Economic Research. NBER Working Papers 28424.

IMF. 2014. "Revenue Administration Gap Analysis Program: The Value-Added Tax Gap." Fiscal Affairs Department - International Monetary Fund.

IMF. 2022. "Uganda: Staff Report for the 2021 Article IV Consultation." International Monetary Fund Article IV Report 22/77.

Keen, Michael, and Ben Lockwood. 2010. "The Value Added Tax: Its Causes and Consequences." *Journal of Development Economics*, 92(2): 138–151.

Lediga, Collen, Nadine Riedel, and Kristina Strohmaier. 2023. "Tax Enforcement Spillovers - Evidence from Business Audits in South Africa." mimeo.

Lemgruber, Andrea, Andrew Masters, and Duncan Cleary. 2015. "Understanding Revenue Administration: An Initial Data Analysis Using the Revenue Administration Fiscal Information Tool." International Monetary Fund Fiscal Affairs Department Paper Series.

Lopez-Luzuriaga, Andrea, and Carlos Scartascini. 2019. "Compliance spillovers across taxes: The role of penalties and detection." *Journal of Economic Behavior and Organization*, 164: 518–534.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, 60(3): 531–542.

Mascagni, Giulia, Andualem T. Mengistu, and Firew B. Woldeyes. 2021. "Can ICTs increase tax compliance? Evidence on taxpayer responses to technological innovation in Ethiopia." *Journal of Economic Behavior and Organization*, 189: 172–193.

Miguel, Edward, and Michael Kremer. 2004. "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities." *Econometrica*, 72(1): 159–217.

Naritomi, Joana. 2019. "Consumers as Tax Auditors." *American Economic Review*, 109: 3031–3072.

OECD. 2022. *Revenue Statistics.*

OECD/ATAF. 2022. *Revenue Statistics in Africa.*

Pinotti, Paolo. 2017. "Clicking on Heaven's Door: The Effect of Immigrant Legalization on Crime." *American Economic Review*, 107(1): 138–168.

Pomeranz, Dina. 2015. "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax." *American Economic Review*, 105(8): 2539–69.

Rambachan, Ashesh, and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *The Review of Economic Studies*, 90(5): 2555–2591.

Rosenbaum, Paul R. 2007. "Interference between Units in Randomized Experiments."

Journal of the American Statistical Association, 102(477): 191–200.

Roth, Jonathan, and Pedro H. C. Sant'Anna. 2023. "Efficient Estimation for Staggered Rollout Designs." *Journal of Political Economy Microeconomics*, 1(4): 669–709.

Rubin, Donald B. 1980. "Randomization analysis of experimental data: The Fisher randomization test comment." *Journal of the American Statistical Association*, 75(371): 591–593.

Shah, Jawad. 2026. "Using computerized information to enforce VAT: Evidence from Pakistan." *Journal of Development Economics*, 178: 103595.

Shimeles, Abebe, Daniel Zerfu Gurara, and Firew Woldeyes. 2017. "Taxman's Dilemma: Coercion or Persuasion? Evidence from a Randomized Field Experiment in Ethiopia." *American Economic Review*, 107(5): 420–24.

Slemrod, Joel, Brett Collins, Jeffrey Hoopes, Daniel Reck, and Michael Sebastiani. 2017. "Does Credit-card Information Reporting Improve Small-business Tax Compliance?" *Journal of Public Economics*, 149: 1–19.

Uganda Revenue Authority. 2016-2019. "Administrative Tax Records." Confidential administrative data. Restricted access. Uganda.

Vazquez-Bare, Gonzalo. 2023. "Identification and Estimation of Spillover Effects in Randomized Experiments." *Journal of Econometrics*, 237(1): 105237.

World Bank. 2021. "GNI per capita in PPP (current international dollars)." data retrieved from World Development Indicators, https://data.worldbank.org/indicator/NY.GNP.PCAP.PP.CD?locations=UG.

Supplemental Appendix
For web publication only

Firm Networks and Tax Compliance: Experimental Evidence from Uganda

November 2026

Miguel Almunia (CUNEF Universidad)
David J. Henning (Oxford University)
Justine Knebelmann (Sciences Po)
Dorothy Nakyambadde (Uganda Revenue Authority)
Lin Tian (INSEAD)

A Data correction following extraction problem

In this appendix, we discuss in detail the data extraction error that caused 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. All input transactions—those where a given firm acted as a buyer—for invoices filed between 19/9/2017 and 15/10/2017 were duplicated during the data extraction. This doubled reported purchases, mechanically inflating seller shortfall amounts for those months. As a result, some links that did not genuinely exceed our 1 million UGX threshold were incorrectly included in the original Study Sample. We correct this by excluding all links whose inclusion depended on the inflated amounts, reducing the sample from 1,708 to 1,235 links. In Table A1, we show that the share of excluded links is balanced across treatment arms (27% in control vs. 28% in treatment).

Table A1
Balance Table for Number of Firm-Pairs missing

	N pairs	N missing pairs	Share missing pairs
Control	677	185	0.27
Treatment	1031	291	0.28
By treatment arm			
Buyer and Seller	347	86	0.25
Buyer only	343	101	0.29
Seller only	341	104	0.30

Notes: This table displays the number and share of firm pairs whose seller shortfall fell below the 1 million UGX threshold once the data extraction error was corrected. The error duplicated all input transactions—transactions where the firm acted as a buyer—for invoices filed between 19/9/2017 and 15/10/2017, doubling reported purchases and mechanically inflating seller shortfall amounts.

We also formally test whether the probability of being included in our study because of this extraction error differs across treatment arms. The result of this test is presented in Table A2. In column 1, we estimate a regression using an indicator variable for whether a firm-pair was treated, while in column 2 we estimate the regression using separate indicator variables for each treatment arm. The omitted indicator is therefore always the control group. The individual coefficients are small and never statistically significant. Consistent with this, the p-value of the F-test for the joint significance of all coefficients in column 2 is 0.337.

Table A2
Balance of the Probability of Being Removed from Sample

Dependent variable:	Prob. of missing	
	(1)	(2)
Treatment	0.009	
	(0.022)	
Buyer and Seller		-0.025
		(0.029)
Buyer only		0.021
		(0.030)
Seller only		0.032
		(0.030)
Constant	0.273	0.273
	(0.017)	(0.017)
R-squared	0.000	0.002
Observations	1708	1708
P-value of F-test	0.685	0.337

Notes: This table tests whether the probability of a link being excluded after correcting the data extraction error is balanced across treatment and control. The sample includes all 1,708 links in the original Study Sample, prior to correction. The outcome variable is an indicator taking value 1 if the link's accumulated seller shortfall fell below the 1 million UGX threshold once the duplicated input transactions were removed. Standard errors are clustered at the link level. *p<.10; **p<.05; ***p<.01. *Source*: Data from monthly VAT returns submitted to the URA.

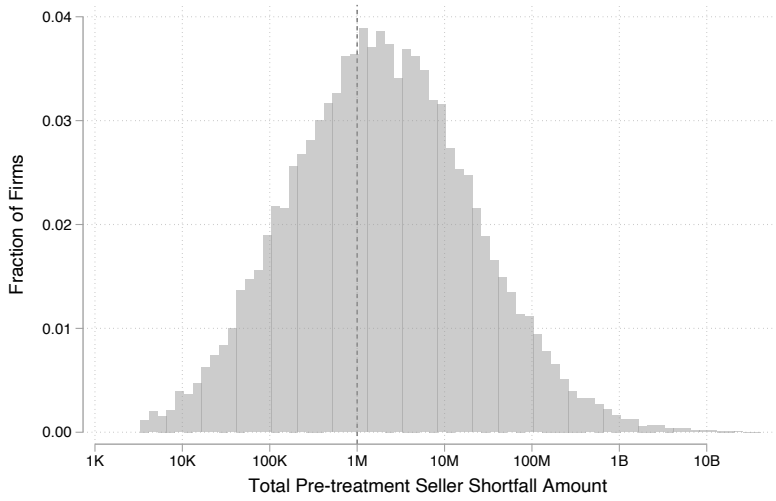
B Experimental Design

In this appendix, we show how the firms of our Study Sample compare with those in the Full Sample of VAT transactions and in the Base Sample, before providing details on the statistical power our experimental design provides.

B.1 Comparing Samples

In Table B1, we report descriptive statistics for links and firms of the Full Sample, Base Sample, and Study Sample, described in Section 3.2.1. Figure B1 shows the distribution of accumulated seller-shortfall amounts in the pre-treatment period for the Full Sample, alongside the 1 million UGX threshold used to select the Base Sample. Figure B2 reports a sampling simulation in which we apply our selection algorithm 1,000 times to the Base Sample. It plots each firm's selection probability against its number of trading partners. The positive relationship confirms that our procedure draws firms with more partners. Figure B3 corroborates this in the realized samples, showing the distribution of trading partners across the Full, Base, and Study Samples.

Figure B1
Selection Cut-off in Seller Shortfall Amounts for Base Sample



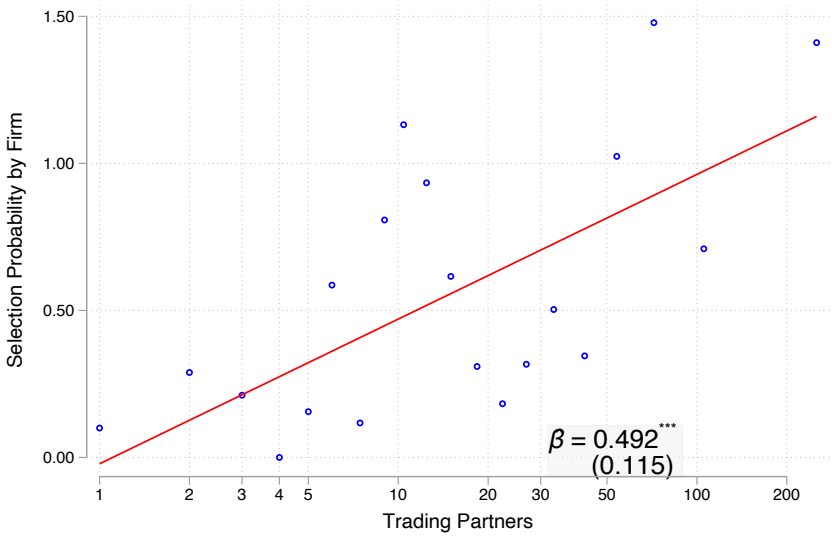
Notes: This histogram shows the distribution of total seller shortfall amounts accumulated between March and December 2017 for seller-buyer links in the Full Sample for which both firms filed a VAT return in December 2017. The dashed line shows the 1 million UGX threshold above which links are selected for the Base Sample. The figure is discussed in Section 3.2. *Source*: Monthly VAT returns submitted to the URA, March–December 2017.

Table B1
Descriptive Statistics for Links and Firms of the Full Sample, Base Sample and Study Sample

Link characteristics	Mean: Full Sample (1)	Mean: Base Sample (2)	Mean: Study Sample (3)	Difference: Base Sample - Full Sample (4)	Difference: Study Sample - Full Sample (5)	Difference: Study Sample Base Sample (6)
Monthly transactions	3.47	6.02	5.01	2.55***	1.55***	−1.00***
	(3.06)	(3.26)	(3.27)	[0.03]	[0.09]	[0.10]
Monthly transactions amount	1.16	1.98	1.12	0.82***	−0.04	−0.86***
	(9.30)	(8.76)	(3.19)	[0.04]	[0.12]	[0.11]
Months with seller shortfall	0.44	0.72	0.80	0.28***	0.36***	0.08***
	(0.49)	(0.41)	(0.32)	[0.00]	[0.01]	[0.01]
Share of extensive margin	0.95	0.87	0.85	−0.08***	−0.10***	−0.02**
	(0.21)	(0.31)	(0.33)	[0.00]	[0.01]	[0.01]
Buyer characteristics						
Audited in 2016	0.03	0.05	0.06	0.02***	0.03***	0.00
	(0.17)	(0.22)	(0.23)	[0.00]	[0.01]	[0.01]
Unique trading partners	18.09	50.17	49.77	32.08***	31.67***	−0.41
	(67.10)	(110.71)	(82.95)	[1.29]	[1.98]	[3.42]
Unique clients	9.05	24.83	19.23	15.78***	10.18***	−5.60*
	(58.89)	(99.69)	(69.42)	[1.14]	[1.74]	[3.05]
Unique suppliers	9.53	26.73	31.69	17.20***	22.16***	4.96***
	(20.47)	(32.74)	(34.51)	[0.39]	[0.62]	[1.08]
Seller characteristics						
Audited in 2016	0.04	0.05	0.05	0.01	0.01	0.00
	(0.20)	(0.21)	(0.22)	[0.00]	[0.01]	[0.01]
Unique trading partners	37.79	57.59	69.05	19.80***	31.26***	11.46**
	(109.95)	(145.97)	(167.30)	[2.29]	[3.52]	[4.99]
Unique clients	24.57	39.65	48.51	15.07***	23.94***	8.87*
	(100.38)	(135.27)	(157.10)	[2.10]	[3.23]	[4.64]
Unique suppliers	14.40	19.76	22.57	5.36***	8.17***	2.81***
	(25.60)	(30.08)	(29.60)	[0.51]	[0.78]	[0.99]
Observations	423629	65327	6192			
Links	122197	10855	1235			
Unique Sellers	7678	2874	1235			
Unique Buyers	20039	2950	1235			

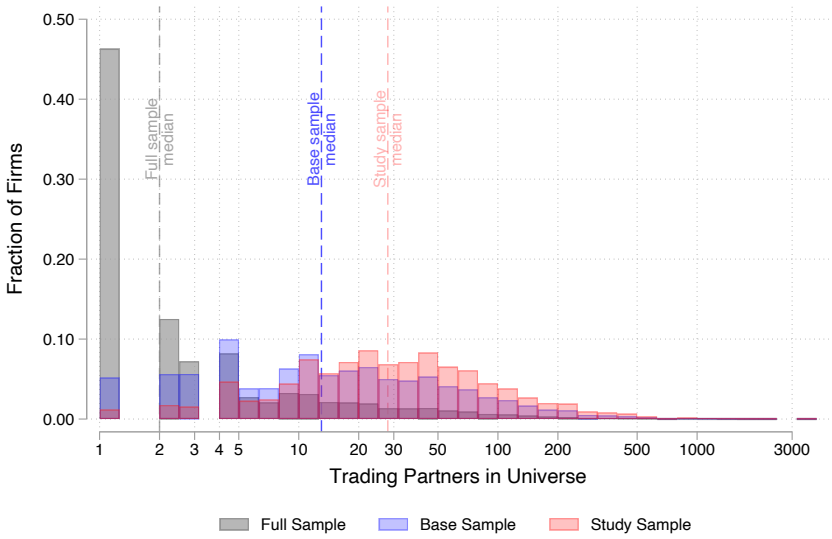
Notes: This table reports descriptive statistics for firms in the Full, Base, and Study Samples, and compares these samples. Columns 1–3 report means, with standard deviations in parentheses. Columns 4–6 report differences in means relative to the comparison sample indicated in the column heading, with standard errors in square brackets. The top panel reports seller-buyer link characteristics. The second and third panels report buyer and seller characteristics at the firm level. Monthly transaction amount is computed as the higher of the amounts reported by the seller and the buyer. Months with seller shortfall is the number of months in which seller shortfall is observed. Share of extensive margin is the share of seller-shortfall instances in which the seller reports no trade with the buyer. Share of final sales is the ratio of sales to final consumers and sales to non-VAT firms to total sales. All monetary values are in thousands of USD. The table is discussed in Section 3.2. *Source*: Monthly VAT returns submitted to the URA, March–December 2017.

Figure B2
Binscatter of Selection Probability and Trading Partners for Firms in the Base Sample



Notes: This figure illustrates how the selection probability varies with the number of trading partners of a firm. The blue circles represent a binscatter plot of a firm's selection probability across its number of trading partners. The trading-partner variable is split into equally sized bins, and each plotted point represents the average selection probability within that bin. The red line represents the linear fit of the binscatter plot. The β coefficient and the standard errors in parentheses are obtained by estimating the regression: $y_i = \alpha + \beta \log(\text{trading partners}_i) + \varepsilon_i$ where y_i denotes the probability that firm i is selected into the study sample, $\text{trading partners}_i$ is the total number of trading partners of firm i and ε_i represent standard errors robust to heteroskedasticity. The selection probability is derived by running the selection algorithm 1,000 times and counting how many times a link that a firm is part of was selected throughout the 1000 simulations. The figure is discussed in Section 3.2. *Source*: Monthly VAT returns submitted to the URA, March–December 2017.

Figure B3
Distribution of the Number of Trading Partners Across Samples



Notes: This figure shows the distributions of the number of trading partners for firms in each of the three samples, Full Sample, Base Sample and Study Sample. In dashed lines we present the median value for each of the samples. The median number of trading partners is 2 in the Full sample, 13 in the base sample, and 26 in the study sample. The figure is discussed in Section 3.2. *Source*: Monthly VAT returns submitted to the URA, March–December 2017.

B.2 Balance Checks Accounting for Multiple Hypothesis Testing

Since our main balance checks (shown in Table 1) test multiple hypotheses (18 coefficients and four comparisons), we verify that the diagnostic holds when adjusting the p-values for Multiple Hypothesis Testing. In Table B2, we display sharpened q-values following Benjamini and Hochberg (1995) and Anderson (2008), adjusting for the testing of 72 hypotheses. The results show that the null hypothesis is never rejected.

B.3 Power Calculations

For each main specification, we report the sample size, unit of observation, level of clustering, and minimum detectable effects (MDE). We compute the MDE using parameters from the design and the control group. We implement power calculations in STATA using the`power` command. Table B3 provides a summary.

Within-Link Effects of the Letter on the Correction of Past Discrepancies (Results shown in Figure 3). There are 2,746 observations at the seller-buyer (link) discrepancy level. Treatment is randomized at the seller-buyer (link) level, and we cluster at the link level. The sample contains 1,234 links, with 493 links assigned to control and 741 to treatment, with a

Table B2
Study Sample: Descriptive Statistics and Randomization Balance adjusting for Multiple Hypothesis Testing

	Mean and SD	Difference with respect to control			
Link characteristics	Control	Any treatment	Buyer only	Seller only	Seller and Buyer
	(1)	(2)	(3)	(4)	(5)
Monthly transactions	4.82 (3.27)	−0.33*	−0.48*	−0.19	−0.32
		(0.09)	(0.07)	(0.47)	(0.20)
		[1.00]	[1.00]	[1.00]	[1.00]
Monthly transactions amount	1.10 (3.34)	−0.03	−0.16	−0.02	0.08
		(0.72)	(0.17)	(0.89)	(0.46)
		[1.00]	[1.00]	[1.00]	[1.00]
Months with seller shortfall	0.83 (0.29)	0.05**	0.03	0.04	0.07***
		(0.01)	(0.20)	(0.13)	(0.00)
		[0.32]	[1.00]	[1.00]	[0.31]
Share of extensive margin	0.87 (0.31)	0.03	0.02	0.02	0.05*
		(0.14)	(0.47)	(0.48)	(0.07)
		[1.00]	[1.00]	[1.00]	[1.00]
Buyer characteristics					
Buy: Total output	14.61 (57.19)	−2.62	−1.12	−11.38	3.98
		(0.61)	(0.79)	(0.19)	(0.29)
		[1.00]	[1.00]	[1.00]	[1.00]
Buy: Total input	11.47 (48.39)	−1.63	−1.62	−6.84	3.10
		(0.61)	(0.66)	(0.19)	(0.32)
		[1.00]	[1.00]	[1.00]	[1.00]
Buy: Share of final sales	0.51 (0.44)	−0.02	−0.03	−0.01	−0.02
		(0.50)	(0.39)	(0.83)	(0.66)
		[1.00]	[1.00]	[1.00]	[1.00]
Buy: Audited in 2016	0.06 (0.24)	0.00	0.01	−0.01	0.01
		(0.72)	(0.46)	(0.51)	(0.46)
		[1.00]	[1.00]	[1.00]	[1.00]
Buy: Unique trading partners	48.85 (81.84)	−1.53	−2.96	−0.36	−1.28
		(0.75)	(0.62)	(0.96)	(0.85)
		[1.00]	[1.00]	[1.00]	[1.00]
Buy: Unique clients	19.24 (68.03)	0.01	0.83	−1.30	0.45
		(1.00)	(0.86)	(0.81)	(0.94)
		[1.00]	[1.00]	[1.00]	[1.00]
Buy: Unique suppliers	30.63 (34.07)	−1.77	−4.31	0.76	−1.72
		(0.38)	(0.12)	(0.78)	(0.50)
		[1.00]	[1.00]	[1.00]	[1.00]
Seller characteristics					
Sell: Total input	27.90 (63.65)	1.57	−0.77	3.06	2.37
		(0.68)	(0.89)	(0.53)	(0.62)
		[1.00]	[1.00]	[1.00]	[1.00]
Sell: Total output	24.49 (56.52)	2.35	0.56	3.37	3.07
		(0.46)	(0.90)	(0.43)	(0.46)
		[1.00]	[1.00]	[1.00]	[1.00]
Sell: Share of final sales	0.56 (0.41)	0.00	0.03	−0.05	0.01
		(0.94)	(0.36)	(0.10)	(0.65)
		[1.00]	[1.00]	[1.00]	[1.00]
Sell: Audited in 2016	0.04 (0.21)	−0.01	−0.03*	0.00	−0.01
		(0.30)	(0.09)	(0.88)	(0.43)
		[1.00]	[1.00]	[1.00]	[1.00]
Sell: Unique trading partners	65.02 (156.91)	−6.71	−5.13	−5.51	−9.28
		(0.49)	(0.65)	(0.70)	(0.45)
		[1.00]	[1.00]	[1.00]	[1.00]
Sell: Unique clients	45.41 (145.19)	−5.16	−2.12	−3.91	−9.13
		(0.57)	(0.84)	(0.77)	(0.43)
		[1.00]	[1.00]	[1.00]	[1.00]
Sell: Unique suppliers	21.49 (28.95)	−1.79	−3.27	−1.80	−0.42
		(0.30)	(0.17)	(0.44)	(0.85)
		[1.00]	[1.00]	[1.00]	[1.00]
P-value of joint Ftest		0.50	0.49	0.70	0.49
Observations	2380	3812	1281	1191	1340
Links	494	741	242	238	261

Notes: This table displays descriptive statistics for firms of the Study Sample and their balance across treatment arms and the control group. Column 1 reports the mean and standard deviation (in parentheses) for the control group. Columns 2–5 report the differences between treatment arms and the control group, with the p-values in parenthesis, and the sharpened False Discovery Rate-adjusted q-values for Multiple Hypothesis Testing in square brackets computed following Benjamini and Hochberg (1995) and Anderson (2008). The table is discussed in Section 3.3. *Source*: Data from monthly VAT returns submitted to the URA between March and December 2017.

mean of 2.2 observations (discrepancies) per link.[41] The harmonic mean number of clusters per study arm is 592.[42] The outcome is a binary indicator, *Any Correction*, which has a mean of 0.012 in the control group and an intraclass correlation coefficient of 0.82. Based on these inputs, the MDE is a 2.3 percentage point increase in the *Any Correction* rate. In the results, *Any Correction* increases by 22.4, from 1.2% to 23.6%—an effect that is well above the MDE. We verify that the study is very well powered to detect this effect, with a 100% probability of avoiding a false negative.

Outside-Link Effects of the Letter on the Correction of Past Discrepancies (Results shown in Figure 5). Observations are at the seller-buyer link-month level, as in the within-link analysis. The sample now includes all discrepancies that sellers in our Study Sample have with partners outside the Study Sample (12,554 observations). Standard errors are clustered at the seller level, with 311 sellers in the control group and 469 in the treatment group (harmonic mean: 374 sellers per arm). *Any Correction* has a mean of 0.003 in the control group, yielding an MDE of 1.6 percentage point. The estimated treatment effect is 3.4, well above the MDE, and the implied power is approximately 99%.

Effect of Treatment on Firm-Level VAT Liability (Sellers) (Results shown in Table 2). This analysis is restricted to sellers, with observations defined at the seller–tax return level, yielding 2,738 observations in total. Standard errors are clustered at the seller level. The sample includes 1,232 sellers: 492 in the control group and 740 in the treatment group. The outcome variable is Δ(VAT Liability), defined as the change in VAT liability between the initial and revised return (measured in thousands of US dollars). In the control group, Δ(VAT Liability) has a mean of 0.0055, a standard deviation of 0.0719, and intraclass correlation coefficient of 0.33. Given these inputs, the MDE is 0.0093 in Δ(VAT Liability) in thousands of \$, assuming 80% power. The estimated effect is 0.014 thousand \$, which is slightly above the MDE. Hence we are also powered to draw conclusions from this estimate.

Post-treatment Effects on Firm-Level VAT Liability (Sellers) (Results shown in Figure 9, β coefficient in bottom right of each panel). This analysis is restricted to sellers. The unit of observations is at the seller–tax return level, yielding 25,085 observations. Standard errors are clustered at the seller level. The sample includes 1,235 distinct sellers: 494 in the control group and 741 in the treatment group. On average, each seller contributes 20 observations. The outcome is VAT liability, measured in thousands of USD and winsorized at the 0.5% level. In the control group, VAT liability has a mean of 3.321, a standard deviation of 17.557, and an intraclass correlation coefficient of 0.53. Given these inputs, the MDE is an increase of 2.116 thousand \$ in VAT liability. The estimated coefficient is 0.457 (and is not

[41] While our study sample includes 1,235 links, the sample size for the analysis of within-link effects on the correction of past discrepancies is 1,234, and the sample size for the analysis for sellers of effects on VAT liability is 1,232. The reason for the former is that discrepancies for one link were included in our initial treatment assignment (which is why the link remains in the study sample), but dropped following a correction in the data cleaning procedure. The reason for the latter is that two sellers of the study sample have a discrepancy for a given month, but did not file VAT in the same month.

[42] We use the harmonic mean, as recommended when cluster counts differ across arms.

statistically significant). Consistent with this, the study does not have sufficient power to detect an effect of this size; the implied power is approximately 9%.

Table B3
Power Calculations Summary

Specification	Cluster	MDE	Point Estimate	Implied Power
Within-Link corrections	Link	2.3 pp	22.4 pp	$\approx 100\%$
Outside-Link corrections	Seller	1.6 pp	3.4 pp	99%
Effect on VAT Liability	Seller	9	14	99%
Post-treatment effect on VAT Liability	Seller	2,116	457	9%

Notes: This table presents a summary of the power calculations described in detail in Appendix B.3. Here we present the level of clustering, the Minimum Detectable Effect size (MDE), our point estimate, and the implied power for that point estimate across each of our main specifications.

C Implementation Details

This appendix provides further details on the experimental implementation and follow-ups. Figure C1 shows a template of the notification letter. Table C1 reports the shares of the treated firms that provided feedback, and Figure C2 the overall response rates counting both amendments and feedback. Figure C3 displays the types of feedback submitted to the URA by treated firms of each arm. Figure C4 shows the absence of differential attrition between treatment and control groups.

Figure C1
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.

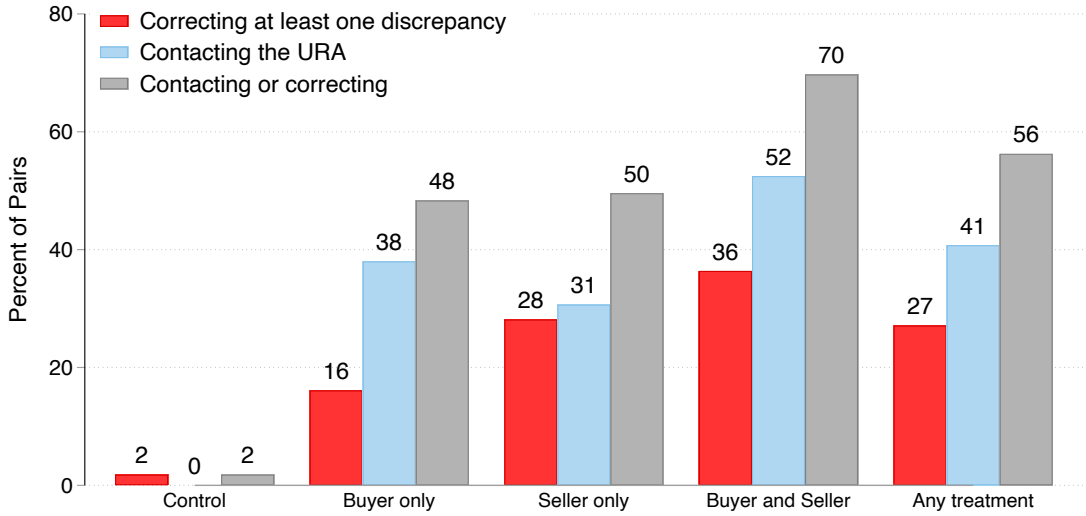
We thank you for your usual cooperation with the URA as we Develop Uganda Together.

«Assistant Commissioner Compliance Management»
For: Commissioner General

Copy: «Jurisdiction»

Notes: This figure displays the template of the letter used by the Ugandan Revenue Authority to notify firms about discrepancies in their tax returns during the intervention. The figure is mentioned in Section 3.2.

Figure C2
Share of Links Reacting to the Letter or Contacting the URA



Notes: This figure displays the types of feedback submitted by firms in treated links to the URA. If any firm in the link provided feedback corresponding to multiple categories, all are counted. This figure is mentioned in Sections 3.3 and 4.1. *Source*: Data from the communication tracker designed in collaboration with the URA, and monthly VAT returns submitted to the URA.

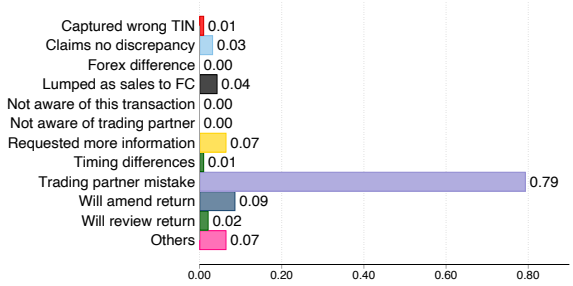
Table C1
Feedback from Treated Firms

	LTO	MTO	STO	Overall
Number of firms	80	168	743	991
Frac receiving letters or giving feedback	0.98	0.93	0.91	0.92
Frac contacting URA	0.55	0.44	0.34	0.37

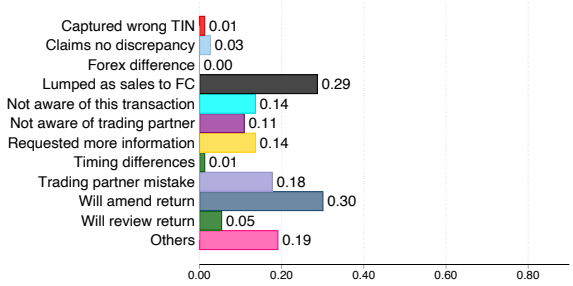
Notes: This table reports summary statistics on firms' interaction with the URA following the intervention by types of firms. LTO, MTO, and STO refer to large taxpayers' office, medium taxpayers' office, and small taxpayers' office, respectively. We consider a firm received the letter or gave feedback 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 C3
Responses from Firms that Gave Feedback by Treatment Arm

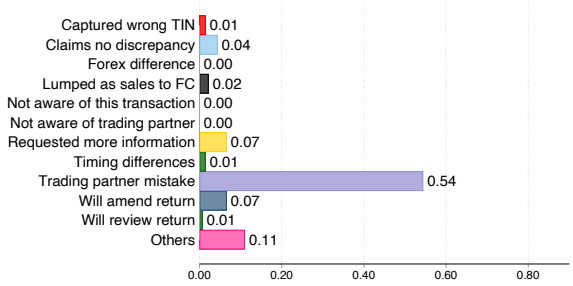
(a) Buyers in Buyer only treatment



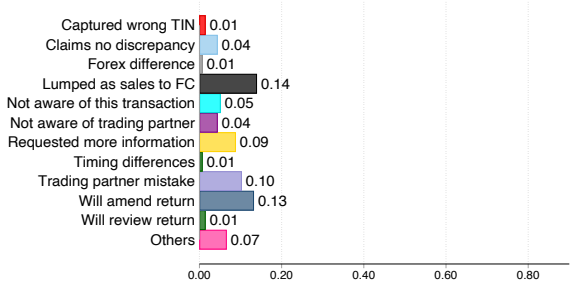
(b) Sellers in Seller only treatment



(c) Buyers in Buyer and Seller treatment

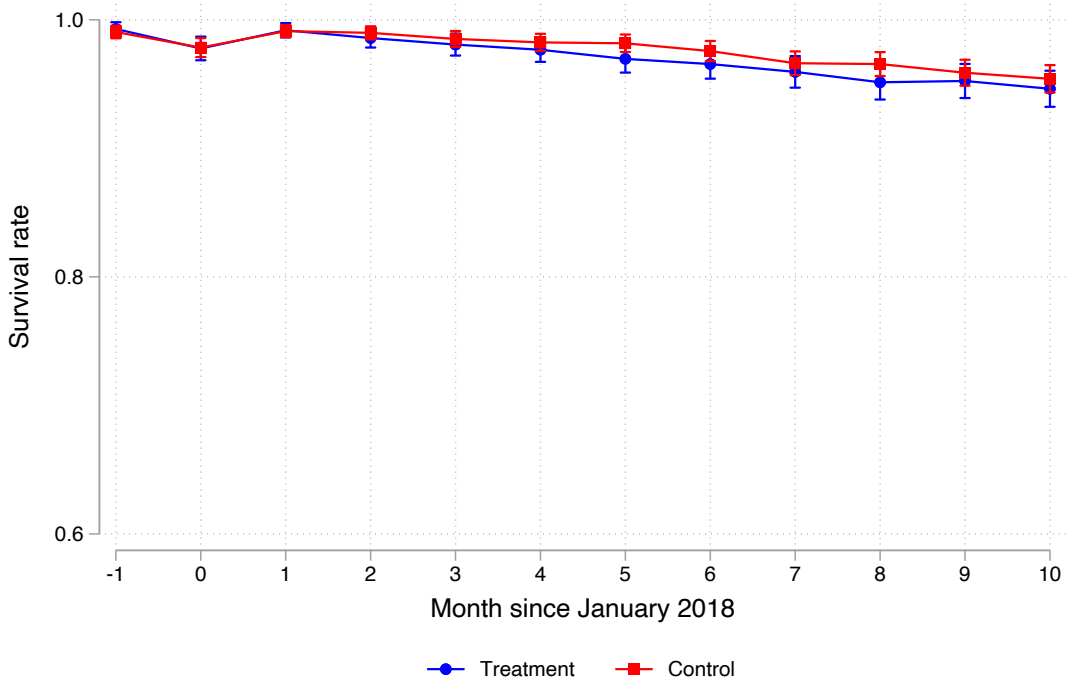


(d) Sellers in Buyer and Seller treatment



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. This figure is discussed in Section 3.3 and 4.1. *Source*: Data from the communication tracker designed in collaboration with the URA

Figure C4
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 firms from treated links is plotted in blue, while the share for 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.

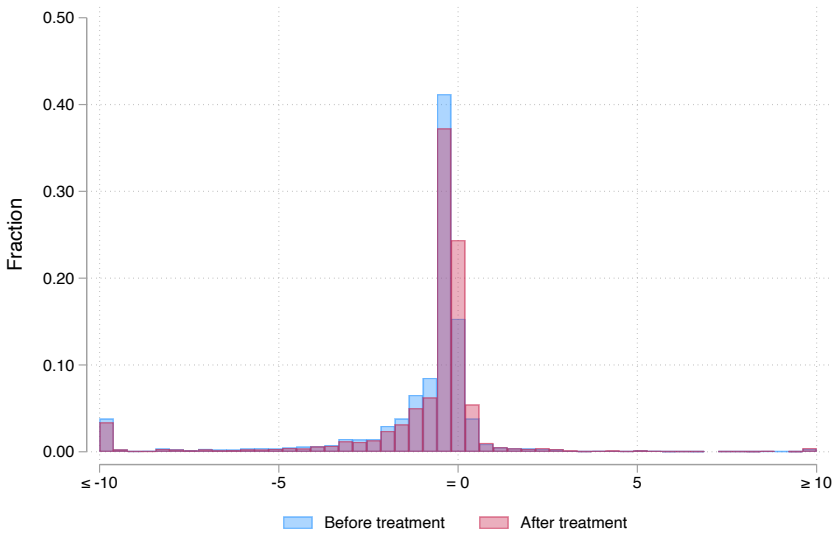
D Additional Results on Corrections of Past VAT Returns

This appendix provides additional results on corrections of past VAT returns, supplementing those shown in Section 4.1. In Section D.1 we present additional results for the within-link estimations. Figure D1 compares the distributions of discrepancies before and after the treatment for the treatment and control groups. Tables D1–D2 show the regression results that correspond to Figures 3-4 discussed in the main text. Table D1 also includes an alternative outcome variable: an indicator taking value one only if the discrepancy has been fully resolved through the corrections, i.e., *Full correction*. In Section D.2 we assess to what extent our results can be extrapolated to the universe of VAT paying firms by reweighting the sample. Table D3 compares the Study Sample to the Full Sample before and after reweighting, while Table D4 presents the results. In Section D.3 we provide additional results for our outside-link estimations. Table D5 presents results using all months and the alternative outcome *Full correction*, and Table D6 presents the regression results underlying Figure 5 together with the seller level share of corrected discrepancies.

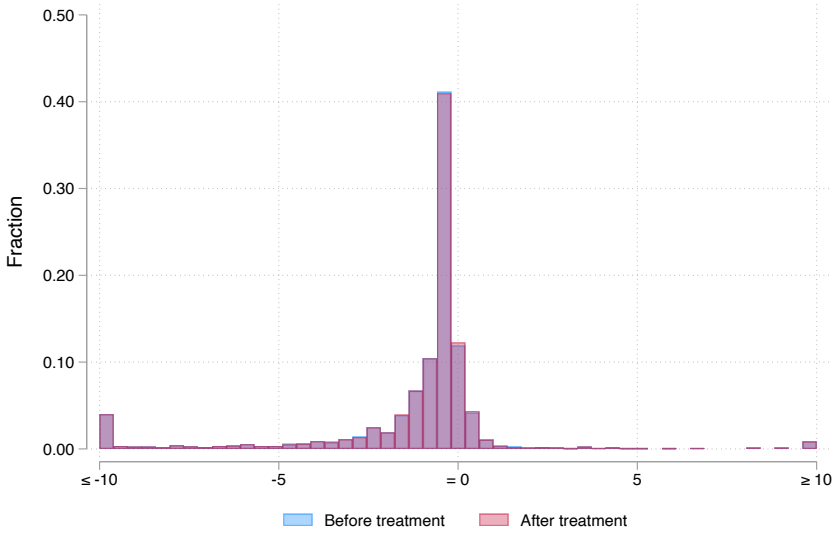
D.1 Additional Within-Link Results

Figure D1
Distribution of Discrepancies Before and After Treatment

(a) Treatment



(b) Control



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

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

Dependent variable:	Any correction				Full correction			
	Listed		Unlisted		Listed		Unlisted	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	0.224***		0.101***		0.143***		0.065***	
	(0.016)		(0.019)		(0.013)		(0.014)	
Buyer only		0.115***		0.035*		0.077***		0.025*
		(0.021)		(0.020)		(0.017)		(0.015)
Seller only		0.247***		0.139***		0.158***		0.077***
		(0.030)		(0.036)		(0.024)		(0.026)
Buyer and Seller		0.304***		0.134***		0.189***		0.095***
		(0.029)		(0.034)		(0.024)		(0.026)
R-squared	0.095	0.125	0.031	0.051	0.058	0.074	0.019	0.031
Observations	2746	2746	2176	2176	2733	2733	2173	2173
No. of Unique Links	1234	1234	530	530	1234	1234	529	529
Mean of Dep. in Control	0.012	0.012	0.015	0.015	0.006	0.006	0.006	0.006
P-value of β_S=β_B		0.000		0.009		0.005		0.078
P-value of β_{SB}=β_B		0.000		0.011		0.000		0.017
P-value of β_{SB}=β_S		0.163		0.909		0.357		0.627
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 estimating equation (2) (odd columns), and equation (1) (even columns). In column 1 to 4 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< 0.05; ***p< 0.01. Results for Any Correction are shown in Figure 3 in Section 4.1. *Source*: Data from monthly VAT returns submitted to the URA.

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

Correction by:	Listed				Unlisted			
	Seller		Buyer		Seller		Buyer	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	0.188***		0.024***		0.089***		0.005	
	(0.016)		(0.005)		(0.018)		(0.003)	
Buyer only		0.082***		0.026***		0.034*		0.003
		(0.019)		(0.008)		(0.020)		(0.002)
Seller only		0.226***		0.013*		0.123***		0.012
		(0.029)		(0.008)		(0.036)		(0.009)
Buyer and Seller		0.251***		0.032***		0.113***		0.000
		(0.027)		(0.009)		(0.032)		(0.001)
R-squared	0.080	0.111	0.013	0.015	0.030	0.046	0.006	0.012
Observations	2746	2746	2746	2746	2176	2176	2176	2176
No. of Unique Links	1234	1234	1234	1234	530	530	530	530
Mean of Dep. in Control	0.011	0.011	0.001	0.001	0.013	0.013	0.000	0.000
P-value of β_S=β_B		0.000		0.227		0.024		0.264
P-value of β_{SB}=β_B		0.000		0.623		0.032		0.313
P-value of β_{SB}=β_S		0.543		0.109		0.823		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 (buyer and seller) 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 estimating equation (2) (odd columns), and equation (1) (even columns) 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< 0.05; ***p< 0.01. Results for Listed discrepancies are shown in Figure 4 in Section 4.1. *Source*: Data from monthly VAT returns submitted to the URA.

D.2 Reweighting the Study Sample

Firms in the Study Sample trade with more partners than firms in the Full Sample. This gap is introduced at both stages of sample construction. Restricting the Full Sample to links with accumulated seller shortfall above the threshold yields the Base Sample, where the median firm has 13 trading partners against 2 in the Full Sample. Applying the iterative selection algorithm then raises the median to 26 (Figure B3; Table B1 reports the full comparison across samples).

This appendix asks what effect the intervention would have had in a population whose firms are connected like the typical VAT-registered firm in Uganda. We reweight the Study Sample to match the connectivity of the Full Sample and re-estimate our main link-level specifications using these weights.

We construct weights using entropy balancing (Hainmueller, 2012). The procedure directly targets covariate balance, rather than relying on a propensity score estimated from an auxiliary

model. It solves for observation-specific weights that satisfy pre-specified moment conditions, while deviating as little as possible from uniform weights in terms of information loss:

$$\min_{w_i} \sum_i w_i \log(w_i) \quad \text{s.t.} \quad \sum_i w_i X_i = \bar{X}^{Full}, \quad \sum_i w_i = 1, \quad w_i \geq 0,$$

where $\bar{X}^{Full}$ collects the corresponding moments in the Full Sample. We set

$$X_i = \left(d_i^b,\ d_i^s,\ (d_i^b)^2,\ (d_i^s)^2,\ d_i^b \cdot d_i^s \right),$$

so that the reweighted sample reproduces the mean, variance, and covariance of buyer-degree d_i^b and seller-degree d_i^s, where degree is the number of unique partners. We match the joint moments because the two dimensions are not separable in a network setting. We balance at the link level, using the degrees of the two firms forming each link, since this is the level at which our specifications are estimated. Table D3 reports the moments of buyer- and seller-degree in the Full and Study Samples, before and after reweighting. Before reweighting, links in the Study Sample involve considerably less connected firms: the average buyer has 49.8 trading partners against 87 in the Full Sample, and the average seller 69 against 334.9. After reweighting, both means match the Full Sample almost exactly (87 and 334.9), as do the variances of the two degree distributions, since these are the moments the weights target.

Table D4 reports unweighted (odd columns) and reweighted (even columns) estimates of the effect on whether a link corrects at least one listed discrepancy. Columns 1 and 2 report the pooled effect from equation (2); columns 3 and 4 report the three treatment arms separately, following equation (1). Reweighting leaves the estimates close to their unweighted counterparts: the pooled effect declines from 22.4 to 20.1 percentage points. The pattern across arms also holds.

Table D3
Within-Link Effects of the Letter on the Correction of Past Discrepancies - Reweighted Sample

Variable:	Full Sample			Study Sample		
	Mean (1)	S.D. (2)	Skewness (3)	Mean (4)	S.D. (5)	Skewness (6)
Before reweighting						
d_i^b	87.07	198.62	8.96	49.77	82.95	8.52
d_i^s	334.91	600.69	3.28	69.05	167.30	10.66
$(d_i^b)^2$	47031.66	456079.34	19.72	9352.14	80546.00	21.48
$(d_i^s)^2$	472997.42	1694645.25	5.06	32734.30	378467.65	23.41
$d_i^b \cdot d_i^s$	16210.06	75861.53	32.71	4735.09	30640.52	24.02
After reweighting						
d_i^b	87.07	198.62	8.96	87.06	198.68	5.38
d_i^s	334.91	600.69	3.28	334.85	600.88	2.71
$(d_i^b)^2$	47031.66	456079.34	19.72	47021.89	260694.99	7.17
$(d_i^s)^2$	472997.42	1694645.25	5.06	472892.84	1510906.57	5.37
$d_i^b \cdot d_i^s$	16210.06	75861.53	32.71	16207.97	37158.27	6.17

Notes: This table reports the moments of buyer-degree d_i^b and seller-degree d_i^s in the Full and Study Samples, before and after reweighting. *Source*: Data from monthly VAT returns submitted to the URA.

Table D4
Within-Link Effects of the Letter on the Correction of Past Discrepancies - Reweighted Sample

Dependent variable:	Any correction			
	(1)	(2)	(3)	(4)
Any treatment	0.224***	0.201***		
	(0.016)	(0.023)		
Buyer only			0.115***	0.109***
			(0.021)	(0.022)
Seller only			0.247***	0.253***
			(0.030)	(0.034)
Buyer and Seller			0.304***	0.232***
			(0.029)	(0.049)
R-squared	0.095	0.085	0.125	0.104
Observations	2746	2746	2746	2746
No. of Unique Links	1234	1234	1234	1234
Mean of Dep. in Control	0.012	0.012	0.012	0.012
P-value of β_S=β_B			0.000	0.000
P-value of β_{SB}=β_B			0.000	0.020
P-value of β_{SB}=β_S			0.163	0.724
Month-Year FE	Yes	Yes	Yes	Yes
Reweighted Sample	No	Yes	No	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) are analyzed at the seller-buyer link-month level. The table reports results from estimating equation (2). The outcome variable is an indicator equal to one if a discrepancy is reduced, and the analysis focuses on listed discrepancies. Odd-numbered columns report the baseline estimates using the original Study Sample, while even-numbered columns report estimates using the reweighted Study Sample, where firm connectivity matches that of the Full Sample. Standard errors are clustered at the link level. *p< 0.10; **p< 0.05; ***p< 0.01. *Source*: Data from monthly VAT returns submitted to the URA.

D.3 Additional Outside-Link Results

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

	Any correction		Full correction	
Correction by:	Seller (1)	Seller (2)	Seller (3)	Seller (4)
Treatment	0.025*		0.013*	
	(0.014)		(0.007)	
Buyer only		0.008		0.002
		(0.007)		(0.003)
Seller only		-0.001		0.000
		(0.003)		(0.002)
Buyer and Seller		0.067*		0.035*
		(0.039)		(0.019)
R-squared	0.009	0.038	0.005	0.020
Observations	45493	45493	45493	45493
No. of Unique Links	14184	14184	14184	14184
No. of Study Firms	992	992	992	992
Mean of Dep. in Control	0.004	0.004	0.002	0.002
P-value of β_S=β_B		0.176		0.590
P-value of β_{SB}=β_B		0.142		0.087
P-value of β_{SB}=β_S		0.088		0.069
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. The table reports results from estimating equation (2) (odd columns), and equation (1) (even columns). In columns 1–2 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< 0.05; ***p< 0.01. This table is mentioned in Section 4.1. *Source*: Data from monthly VAT returns submitted to the URA.

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

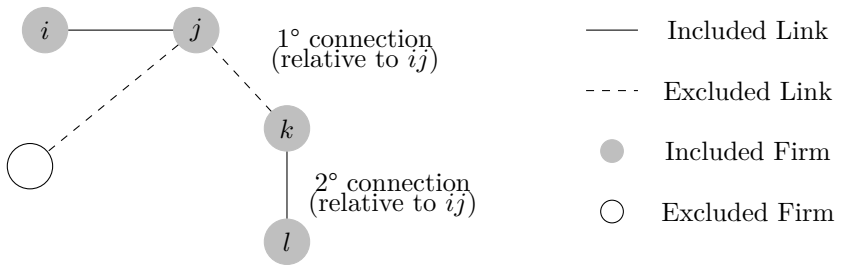
	Any Correction (Links)		Share of Corrections (Sellers)	
Correction by:	Seller (1)	Seller (2)	Seller (3)	Seller (4)
Treatment	0.034**		0.018**	
	(0.017)		(0.008)	
Buyer only		0.010		0.004
		(0.008)		(0.008)
Seller only		0.002		0.002
		(0.006)		(0.008)
Buyer and Seller		0.086*		0.048***
		(0.046)		(0.017)
R-squared	0.024	0.062	0.006	0.027
Observations	12554	12554	780	780
No. of Study Firms	780	780	780	780
Mean of Dep. in Control	0.003	0.003	0.008	0.008
P-value of β_S=β_B		0.374		0.822
P-value of β_{SB}=β_B		0.109		0.009
P-value of β_{SB}=β_S		0.072		0.007
Month-Year FE	Yes	Yes	Yes	Yes

Notes: This table reports the effect of the letter on corrections outside the seller-buyer links in the study sample. VAT returns for pre-treatment months, March–December 2017, are 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 for the months listed in the letter. We focus on corrections made by sellers. The table reports results from estimating equation (2) (odd columns), and equation (1) (even columns). In columns 1-2, the outcome variable is an indicator set to 1 if a discrepancy is reduced by the seller. Observations are at the link level and standard errors are clustered at the seller level. In columns 3-4, the outcome is the share of corrections made by a given seller and observations are at the seller level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. The link-level results are shown in Figure 5. This table is discussed in Section 4.1. *Source*: Monthly VAT returns submitted to the URA.

E 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 E1
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.

E.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:

$$Y_{it} = \alpha + \beta T_i + \left[\gamma_s^0 \bar{T}_{si} + \gamma_b^0 \bar{T}_{bi}\right](1 - T_i) + \left[\gamma_s^1 \bar{T}_{si} + \gamma_b^1 \bar{T}_{bi}\right] T_i + \delta_t + \epsilon_{it}, \tag{7}$$

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 E1. 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 (7), 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 (0.196 vs 0.224), highly significant, and the indirect exposure controls are jointly insignificant. Vazquez-Bare (2023) also proposes a "pooled" estimand for settings where the two key assumptions underlying regression (7)—exchangeability and linearity—would not hold. In that specification, the shares are replaced with an indicator S_i for whether any trading partner is treated:

$$Y_{it} = \alpha + \beta T_i + \left[\gamma_s^0 1(S_{si} > 0) + \gamma_b^0 1(S_{bi} > 0)\right](1 - T_i) + \left[\gamma_s^1 1(S_{si} > 0) + \gamma_b^1 1(S_{bi} > 0)\right] T_i + \delta_t + \epsilon_{it}, \tag{8}$$

where $1(S_{si} > 0)$ (respectively $1(S_{bi} > 0)$) indicates that at least one partner of the seller (resp., buyer) is treated. We show results in column 4 of Table E1. The coefficient decreases, but is still significantly different from zero. This approach is less relevant 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). A direct test confirms the coefficients from column 1 and column 3 (respectively, column 4) are not statistically different (p-value = 0.357 and p-value = 0.203). We interpret this as evidence that exposure through non-seller shortfall links exerts limited bias on our estimates.

We next extend the analysis to second-degree links. Figure E1 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 (7) 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, and the results are in columns 1–3 of Table E2. In column 4, we rely on the pooled specification controlling for whether any partner is treated. The results remain consistent: the estimated coefficients are similar (0.216 and 0.189 vs 0.224), 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.

Finally, we conduct similar robustness checks for the outside-link results shown in Figure 5. In Table E3, column 1 reproduces our main estimate, columns 2–3 control for the share of non-seller-shortfall partners of both the buyer and the seller that are treated, and columns 4–5 control for the share of second-degree partners of the buyer and seller that are treated.

Table E1
Robustness to Non-Seller-Shortfall Spillovers

	(1) Any correction	(2) Any correction	(3) Any correction	(4) Any correction
T_i	0.224***	0.224***	0.196***	0.111**
	(0.016)	(0.016)	(0.034)	(0.054)
$\bar{T}_{si}$		-0.000		
		(0.061)		
$\bar{T}_{bi}$		0.039		
		(0.071)		
$\bar{T}_{si} \times (1 - T_i)$			-0.019	
			(0.022)	
$\bar{T}_{si} \times T_i$			0.009	
			(0.099)	
$\bar{T}_{bi} \times (1 - T_i)$			-0.028	
			(0.019)	
$\bar{T}_{bi} \times T_i$			0.088	
			(0.124)	
$1(S_{si}) \times (1 - T_i)$				-0.000
				(0.012)
$1(S_{bi}) \times (1 - T_i)$				0.050
				(0.048)
$1(S_{si}) \times T_i$				0.016***
				(0.006)
$1(S_{bi}) \times T_i$				0.088*
				(0.047)
R-squared	0.098	0.098	0.099	0.102
Observations	2746	2746	2746	2746
No. of Firms	1234	1234	1234	1234
Mean of Dep. in Control	0.012	0.012	0.012	0.012
P-Value testing for coefficient equality		0.670	0.357	0.203
P-value (joint F-test of shares)		0.861	0.583	0.009
Month-Year FE	Yes	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 buyer and seller of the links, while in column 3 we report results from estimating equation (7). In column 4, we control for an indicator taking value one if any partner is treated, following equation (8). In the second to last line, we present the p-value of a test for whether the coefficient in column 2 (respectively 3 and 4) 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 table focuses on listed discrepancies, mentioned in the letter. Standard errors are clustered at the link level. *p< 0.10; **p< 0.05; ***p< 0.01. This Table is mentioned in Section 4.1. *Source*: Data from monthly VAT returns submitted to the URA.

The difference between these coefficients and our baseline coefficient in column 1 is significant. The interaction terms in columns 3 and 5 show that the probability of correcting is lower for firms that have more treated partners. This is consistent with the idea that treated partners might submit corrections themselves, lowering the additional effect of a firm being treated. Overall, these results suggest that in our main estimation that does not control for indirect exposure, we are slightly *underestimating* treatment effects.

E.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 E1, 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.[43]

To test whether such indirect exposure could bias our firm-level results from Section 5, we re-estimate (2), now at the firm-level, controlling for the share of a firm's trading partners in the Study Sample that are treated. Formally, we estimate the following specification:

$$Y_{jt} = \alpha + \beta_h T_j + \gamma^0 \bar{T}_j(1 - T_j) + \gamma^1 \bar{T}_j T_j + \delta_t + \epsilon_{jt}, \tag{9}$$

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 E4 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 (9). 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 significant at 5%, 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.

[43] The 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 E2
Robustness to Second-Degree Link Spillovers

	(1) Any correction	(2) Any correction	(3) Any correction	(4) Any correction
T_i	0.224***	0.224***	0.216***	0.189***
	(0.016)	(0.016)	(0.027)	(0.033)
$\bar{T}_{si}$		0.013		
		(0.030)		
$\bar{T}_{bi}$		0.003		
		(0.029)		
$\bar{T}_{si} \times (1 - T_i)$			-0.009	
			(0.023)	
$\bar{T}_{si} \times T_i$			0.028	
			(0.048)	
$\bar{T}_{bi} \times (1 - T_i)$			0.013	
			(0.022)	
$\bar{T}_{bi} \times T_i$			-0.004	
			(0.047)	
$1(S_{si} > 0) \times (1 - T_i)$				-0.015
				(0.014)
$1(S_{si} > 0) \times T_i$				0.018
				(0.033)
$1(S_{bi} > 0) \times (1 - T_i)$				0.002
				(0.012)
$1(S_{bi} > 0) \times T_i$				0.021
				(0.033)
R-squared	0.098	0.098	0.098	0.099
Observations	2746	2746	2746	2746
No. of Firms	1234	1234	1234	1234
Mean of Dep. in Control	0.012	0.012	0.012	0.012
P-Value testing for coefficient equality		0.896	0.725	.
P-value (joint F-test of shares)		0.893	0.949	0.513
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 second-degree links that area treated for the buyer and seller of the links, while in column 3 we report results from estimating equation (7). In column 4, we control for an indicator taking value one if any second-degree partner is treated, following equation (8). In the second to last line, we present the p-value of a test for whether the coefficient in column 2 (respectively 3 and 4) 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 table focuses on listed discrepancies, mentioned in the letter. Standard errors are clustered at the link level. *p< 0.10; **p< 0.05; ***p< 0.01. This Table is mentioned in Section 4.1. *Source*: Data from monthly VAT returns submitted to the URA.

Table E3
Effect on Outside-Link Corrections: Robustness to Non-Seller-Shortfall and Second-Degree Link Spillovers

		Non-Seller-Shortfall		Second Degree	
	(1) Any correction	(2) Any correction	(3) Any correction	(4) Any correction	(5) Any correction
T_i	0.034***	0.035***	0.078***	0.035***	0.071***
	(0.003)	(0.003)	(0.011)	(0.003)	(0.010)
$\bar{T}_{si}$		-0.015**		-0.017**	
		(0.007)		(0.007)	
$\bar{T}_{bi}$		-0.082***		-0.027***	
		(0.019)		(0.008)	
$\bar{T}_{si} \times (1 - T_i)$			0.006		0.006
			(0.006)		(0.006)
$\bar{T}_{si} \times T_i$			-0.027**		-0.031***
			(0.011)		(0.011)
$\bar{T}_{bi} \times (1 - T_i)$			-0.005		-0.003
			(0.013)		(0.004)
$\bar{T}_{bi} \times T_i$			-0.116***		-0.041***
			(0.027)		(0.012)
R-squared	0.024	0.028	0.030	0.027	0.028
Observations	12554	12554	12554	12554	12554
No. of Firms	7515	7515	7515	7515	7515
Mean of Dep. in Control	0.005	0.005	0.005	0.005	0.005
P-Value testing for coefficient equality		0.023	0.000	0.007	0.000
P-value (joint F-test of shares)		0.000	0.000	0.000	0.000
Month-Year FE	Yes	Yes	Yes	Yes	

Notes: This table reports the effects of the letters on the correction of past discrepancies for links outside the seller-buyer links of the Study Sample, when controlling for the share of treated trading partners in first-degree non-Seller Shortfall links (column 2-3) and the share of treated trading partners in second-degree Seller Shortfall links (column 4-5). 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 columns 2 and 4, we control for the share of treated partners, where the definition of partner is indicated in the column titles. In columns 3 and 5 we do so using estimating equation (7). In the second to last line, we present the p-value of a test for whether the coefficient in column 2 (respectively 3, 4 and 5) 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 Table focuses on the months mentioned in the letter. Standard errors are clustered at the link level. *p< 0.10; **p< 0.05; ***p< 0.01. This Table is mentioned in Section 4.1. *Source*: Data from monthly VAT returns submitted to the URA.

Table E4
Effect of Treatment on Firm-Level VAT Liability controlling for Treated Partners (Sellers)

Dependent variable:	ΔB2B Sales		ΔFinal Sales		ΔTaxable Inputs		ΔVAT Liability	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T_i	0.522***	0.534***	-0.309***	-0.274***	0.002	0.003	0.014***	0.022***
	(0.060)	(0.088)	(0.037)	(0.049)	(0.003)	(0.004)	(0.004)	(0.006)
$\bar{T}_i \times (1 - T_i)$		-0.025		0.001		-0.009**		0.015**
		(0.025)		(0.006)		(0.004)		(0.006)
$\bar{T}_i \times T_i$		-0.058		-0.092		-0.009*		-0.007
		(0.158)		(0.099)		(0.005)		(0.012)
R-squared	0.036	0.036	0.035	0.036	0.007	0.008	0.009	0.010
Observations	2738	2738	2738	2738	2738	2738	2738	2738
No. of Firms	1232	1232	1232	1232	1232	1232	1232	1232
Mean of Dep. in Control	42.102	42.102	134.186	134.186	157.784	157.784	3.331	3.331
Mean of Diff. in Control	0.016	0.016	-0.003	-0.003	0.001	0.001	0.005	0.005
Median of Dep. in Control	1.730	1.730	11.413	11.413	25.452	25.452	0.237	0.237
P-value of test for eq. of coefficients		0.839		0.343		0.928		0.088
P-value (joint F-test)		0.557		0.632		0.021		0.035
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 (2) at the firm-level, while even columns present results from estimating equation (9). 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< 0.05; ***p< 0.01. This Table is mentioned in Section 5.2. *Source*: Data from monthly VAT returns submitted to the URA.

F Analysis of Subsequent Behavior: Additional Results

This appendix provides additional results for the analysis of post-treatment reporting behavior discussed in Section 4.2. In Section F.1 we provide additional descriptive and regression results supporting our analysis, while in Section F.2 we implement the "Honest Difference-in-Difference" estimator following Rambachan and Roth (2023).

F.1 Within-Link Results

We provide additional results supporting the findings shown in Figures 6 and 7. Columns 1 to 3 of Table F1 provide the regression results underlying Figure 6. Columns 4 to 6 of Table F1 provide the regression results underlying Figure 8. Tables F2 and F3 run the same event study analysis as in Table F1, but pool together all pre- and post-treatment months.

Next, we extend the event-study analysis to include 22 (instead of 10) months before the intervention. Figure F1 reports the raw proportions of seller shortfall in the treated and control groups, and Figure F2 reproduces the event-study regression results for this longer time horizon. Both figures provide reassurance that the treated and control links were not on different structural trends before the intervention.

Table F1
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	0.012	-0.080**	0.068*	-0.003	-0.094	0.005
	(0.039)	(0.036)	(0.037)	(0.191)	(0.150)	(0.033)
-9	0.007	-0.056*	0.049	-0.299	-0.204	0.038
	(0.039)	(0.033)	(0.036)	(0.185)	(0.154)	(0.031)
-8	0.011	-0.032	0.021	-0.268	-0.256	0.037
	(0.036)	(0.033)	(0.034)	(0.250)	(0.158)	(0.031)
-7	-0.032	-0.010	0.042	-0.175	-0.102	-0.010
	(0.035)	(0.030)	(0.034)	(0.225)	(0.133)	(0.034)
-6	-0.058*	-0.010	0.068**	-0.134	-0.040	0.018
	(0.035)	(0.030)	(0.033)	(0.182)	(0.106)	(0.033)
-5	-0.016	-0.035	0.051	-0.065	-0.060	-0.006
	(0.035)	(0.032)	(0.033)	(0.187)	(0.134)	(0.034)
-4	-0.035	-0.011	0.046	-0.069	-0.056	-0.007
	(0.032)	(0.028)	(0.031)	(0.189)	(0.114)	(0.031)
-3	-0.059*	-0.019	0.078**	-0.099	0.011	0.008
	(0.033)	(0.029)	(0.032)	(0.165)	(0.139)	(0.031)
-2	-0.100***	0.028	0.072**	0.070	0.002	0.027
	(0.034)	(0.030)	(0.031)	(0.178)	(0.133)	(0.029)
-1	0.000	0.000	0.000	0.000	0.000	0.000
	(.)	(.)	(.)	(.)	(.)	(.)
0	-0.111***	0.026	0.085**	0.011	-0.184	-0.015
	(0.036)	(0.033)	(0.035)	(0.188)	(0.142)	(0.027)
1	-0.149***	0.037	0.112***	0.065	0.034	-0.032
	(0.040)	(0.034)	(0.038)	(0.178)	(0.123)	(0.027)
2	-0.191***	0.042	0.148***	0.002	0.128	-0.038
	(0.038)	(0.035)	(0.037)	(0.224)	(0.167)	(0.029)
3	-0.202***	0.073**	0.129***	-0.139	-0.117	-0.016
	(0.040)	(0.036)	(0.036)	(0.209)	(0.147)	(0.029)
4	-0.154***	-0.007	0.161***	-0.122	-0.207	-0.010
	(0.044)	(0.036)	(0.039)	(0.220)	(0.158)	(0.030)
5	-0.105**	0.005	0.100**	-0.287	-0.142	-0.010
	(0.041)	(0.038)	(0.040)	(0.207)	(0.124)	(0.029)
6	-0.132***	0.047	0.085**	-0.207	-0.296*	0.000
	(0.042)	(0.036)	(0.040)	(0.235)	(0.153)	(0.029)
7	-0.152***	0.043	0.109***	-0.184	-0.401**	0.001
	(0.042)	(0.035)	(0.041)	(0.257)	(0.170)	(0.029)
8	-0.157***	0.003	0.154***	0.157	-0.054	-0.036
	(0.047)	(0.040)	(0.043)	(0.262)	(0.190)	(0.030)
9	-0.154***	0.056	0.098**	-0.066	-0.512*	0.023
	(0.048)	(0.039)	(0.043)	(0.300)	(0.307)	(0.030)
10	-0.147***	0.037	0.110**	0.061	-0.046	0.024
	(0.047)	(0.038)	(0.043)	(0.239)	(0.159)	(0.030)
$\sum_{j=-10}^{-2} \tau_j$	-0.269	-0.226	0.495**	-1.042	-0.798	0.110
	(0.230)	(0.215)	(0.231)	(1.390)	(1.002)	(0.210)
R-squared	0.535	0.334	0.469	0.748	0.736	0.468
Observations	11090	11090	11090	11090	7617	25935
No. of Firms	1079	1079	1079	1079	937	1235
Mean of Dep. in Control	0.743	0.105	0.152	1.023	0.540	0.423
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 $\hat{\tau}_j$ 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 8. This table is mentioned in Section 4.2. *Source*: Data from monthly VAT returns submitted to the URA.

Table F2
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***	0.033	0.117***	-0.062	-0.150	-0.010
	(0.029)	(0.025)	(0.028)	(0.181)	(0.119)	(0.024)
R-squared	0.534	0.333	0.468	0.748	0.735	0.467
Observations	11090	11090	11090	11090	7617	25935
No. of Links	1079	1079	1079	1079	937	1235
Mean of Dep. in Control	0.792	0.088	0.121	1.020	0.583	0.487
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

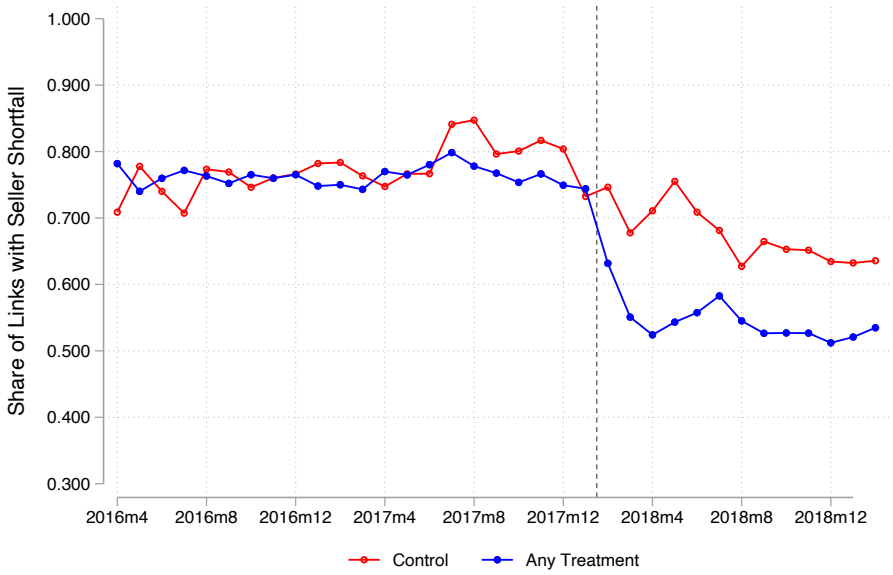
Notes: This table reports the effect of the letter on subsequent reporting discrepancies within the link analysed in the Study Sample. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). The table reports the $\hat{\tau}$ 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< 0.05; ***p< 0.01. 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 F3
Post-Treatment Effect of the Letter on Within-Link Trade: Diff-in-Diff Specification

	Transaction size		Seller shortfall amount		Continuing trade		Seller Reports		Buyer Reports	
DiD specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Any Treatment	-0.062		-0.150		-0.010		0.177***		0.011	
	(0.181)		(0.119)		(0.024)		(0.031)		(0.020)	
Buyer only		-0.256		-0.212		-0.035		0.124***		0.047**
		(0.198)		(0.141)		(0.028)		(0.038)		(0.021)
Seller only		0.088		-0.159		-0.014		0.186***		0.004
		(0.235)		(0.145)		(0.028)		(0.044)		(0.022)
Buyer and Seller		-0.026		-0.076		0.016		0.214***		-0.013
		(0.221)		(0.123)		(0.028)		(0.042)		(0.023)
R-squared	0.748	0.748	0.735	0.735	0.467	0.468	0.680	0.681	0.356	0.358
Observations	11090	11090	7617	7617	25935	25935	11090	11090	11090	11090
No. of Firms	1079	1079	937	937	1235	1235	1079	1079	1079	1079
Mean of Dep. in Control	1.020	1.020	0.583	0.583	0.487	0.487	0.280	0.280	0.959	0.959
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

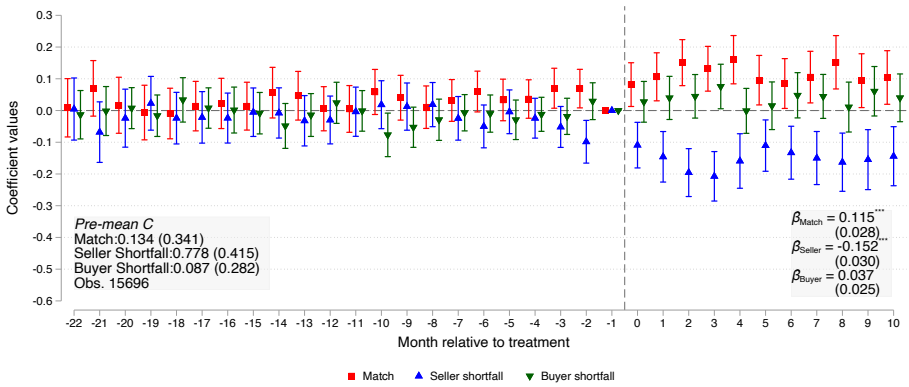
Notes: This table reports the effect of the letter on subsequent reporting discrepancies. VAT returns for 10 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level (April 2017 to December 2018). The table reports the $\hat{\tau}$ coefficients from our difference-in-differences 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< 0.05; ***p< 0.01. The results are displayed in Figure 8. This table is mentioned in Section 4.2. *Source*: Data from monthly VAT returns submitted to the URA.

Figure F1
Raw Proportion of Seller Shortfall by Month in Treatment and Control Links with 12 Additional Months



Notes: This figure is an extension of Figure 7 to which we add 12 pre-treatment months. It reports the raw proportion 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. Outcomes are conditional on trade being reported within the link. This figure is mentioned in Section 4.2. *Source*: Data from monthly VAT returns submitted to the URA.

Figure F2
Post-Treatment Effect of the Letter on Within-Link Discrepancies with 12 Additional Months



Notes: This figure is an extension of Figure 6. It reports the effect of the letter on subsequent reporting discrepancies. VAT returns for 22 months before and 10 months after the treatment were analyzed at the seller-buyer link-month level. 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, in parentheses, are clustered at the link level and the bars report 95% confidence intervals. *p< 0.10; **p< 0.05; ***p< 0.01. This figure is mentioned in Section 4.2. *Source*: Data from monthly VAT returns submitted to the URA.

Table F4
Post-Treatment Effect of the Letter on Within-Link Discrepancies:
Diff-in-Diff Specification - Reweighted Sample

	Seller shortfall		Buyer shortfall		No discrepancy	
	(1)	(2)	(3)	(4)	(5)	(6)
Any Treatment	-0.150***	-0.120***	0.033	0.053***	0.117***	0.067***
	(0.029)	(0.023)	(0.025)	(0.016)	(0.028)	(0.019)
R-squared	0.534	0.534	0.333	0.333	0.468	0.468
Observations	11090	11090	11090	11090	11090	11090
No. of Links	1079	1079	1079	1079	1079	1079
Mean of Dep. in Control	0.792	0.792	0.088	0.088	0.121	0.121
Link FE	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Pre-treatment dummies	Yes	No	Yes	No	Yes	No

Notes: This table reports the effect of the letter on subsequent reporting discrepancies within the link analysed in 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). The table reports the $\hat{\tau}$ coefficients from our difference-in-differences regression (4) pooling together all post-treatment months, with the following outcomes: probability of seller shortfall in columns 1-2; probability of buyer shortfall in column 3-4; and probability of there being no discrepancy ("Match") in column 5-6. Odd-numbered columns report the baseline estimates using the original Study Sample, while even-numbered columns report estimates using the reweighted Study Sample, where firm connectivity matches that of the Full Sample. *p< 0.10; **p< 0.05; ***p< 0.01. Standard errors are clustered at the link level. *Source*: Data from monthly VAT returns submitted to the URA.

F.2 Honest Difference in Differences

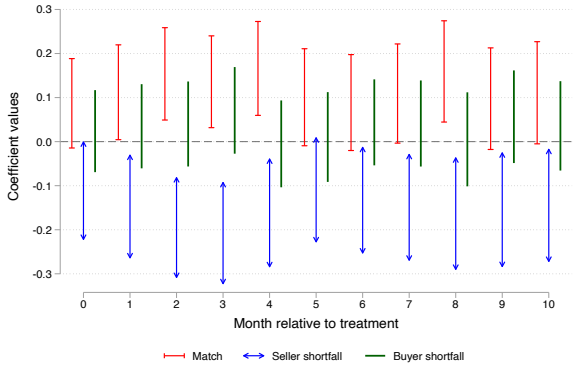
As noted in Section 4.2, the coefficients in Figure 6 reveal some pre-treatment differences in the rate of seller shortfall between the treatment and control groups. Figure 7 plots the raw share of seller shortfall instances for both groups. In the control group, we observe a higher share of seller shortfall between July and December 2017, followed by a pronounced drop in the last pre-treatment month (January 2018). Because the intervention was not announced beforehand, we attribute this fluctuation to idiosyncratic variation rather than to any anticipation of the intervention.

To address possible violations of the parallel trends assumption, we conduct a robustness check following Rambachan and Roth (2023). This method recomputes confidence intervals for each post-treatment coefficient while allowing for bounded deviations from parallel trends. In particular, we posit that the unobserved forces driving any non-parallel trends after treatment are not substantially larger than those evident before treatment. We examine three values of the sensitivity parameter M, which limits the maximum post-treatment deviation relative to the largest pre-treatment gap: (i) $M = 0.5$, (ii) $M = 0.75$, and (iii) $M = 1$. These correspond to allowing post-treatment deviations equal to one-half, three-quarters, and the full size of the largest pre-treatment discrepancy, respectively.

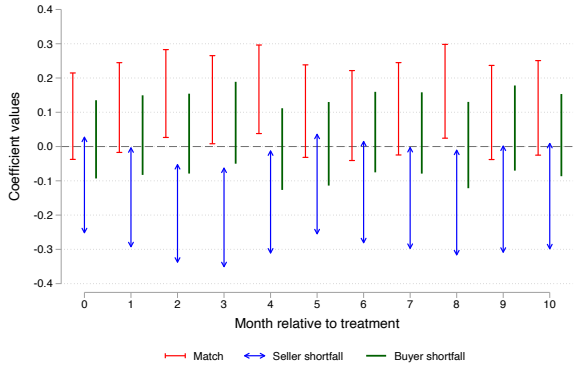
Figure F3 summarizes the results. Under the baseline setting ($M = 1$), which permits post-treatment deviations up to the maximum absolute pre-treatment difference, the estimated impacts on seller shortfall are negative and statistically significant in the second and third months following treatment. With tighter constraints ($M = 0.75$ and $M = 0.5$), the effects remain statistically significant in 4 and 10 out of the 11 post-treatment months, respectively. Taken together, these findings suggest that our results are robust to moderate violations of the parallel trends assumption, on the order of those observed during the pre-treatment period.

Figure F3
Robustness Check: Honest Difference-in-Differences Estimator

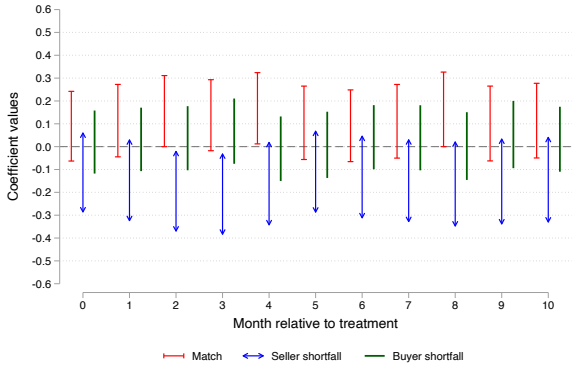
(a) $M = 0.5$



(b) $M = 0.75$



(c) $M = 1$

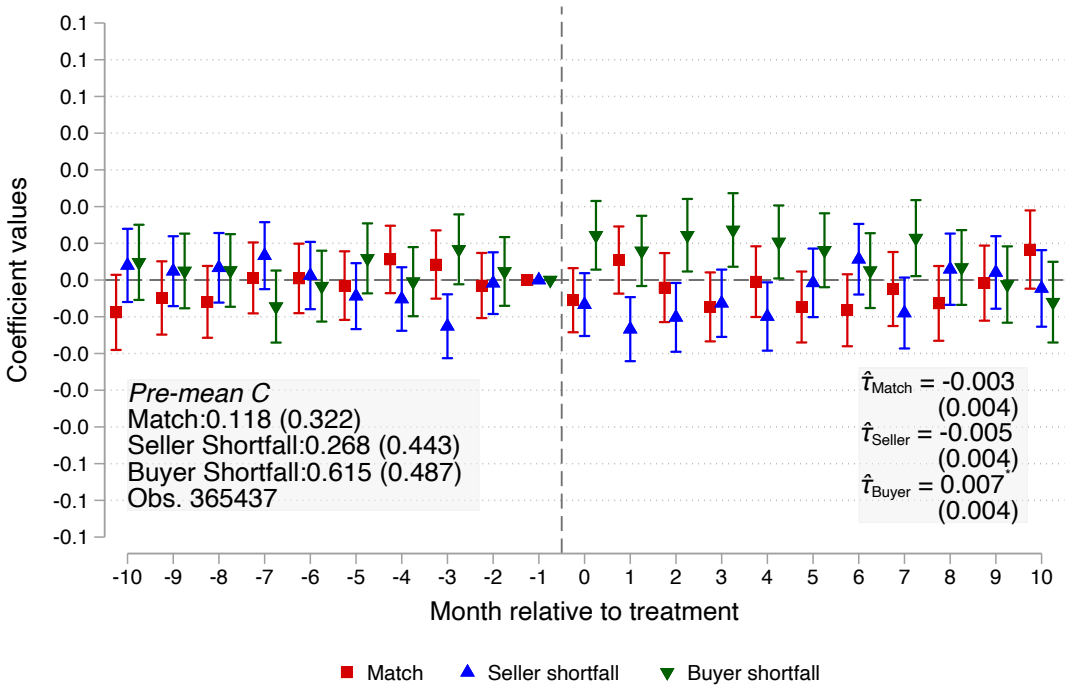


Notes: This figure reports sensitivity-adjusted confidence intervals for the event-study coefficients following Rambachan and Roth (2023). The parameter M governs the allowable deviation from parallel trends, expressed as a fraction of the maximum absolute pre-treatment deviation. Specifically, $M \in 0.5, 0.75, 1$ correspond to allowing post-treatment violations equal to one half (panel a), three quarters (panel b), and the full (panel c) magnitude of the largest pre-treatment discrepancy, respectively. Confidence intervals are constructed to be robust to such deviations, with larger values of M implying more conservative inference. This figure is mentioned in Section 4.2. *Source*: Data from monthly VAT returns submitted to the URA.

F.3 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 F4 presents event-study estimates from estimating equation (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 F5 provides the regression results underlying Figure F4. Tables F6 and F7 run the same event study analysis as in the one in Table F5, but pool together all pre- and post-treatment months.

Figure F4
Post-Treatment Effect of the Letter on Outside-Link Discrepancies



Notes: 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 $\hat{\tau}_j$ 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 (4) pooling together all post-treatment months. Standard errors, in parentheses, are clustered at the link level and the bars report 95% confidence intervals. *p< 0.10; **p< 0.05; ***p< 0.01. See Appendix Table F5 for full regression results. This figure is mentioned in Section III.B. *Source*: Data from monthly VAT returns submitted to the URA.

Table F5
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	0.004	0.005	-0.009*	0.024	-0.037	-0.016***	0.002	0.003
	(0.005)	(0.005)	(0.005)	(0.025)	(0.033)	(0.004)	(0.004)	(0.004)
-9	0.002	0.002	-0.005	0.030	-0.017	0.006	0.014***	0.001
	(0.005)	(0.005)	(0.005)	(0.026)	(0.027)	(0.004)	(0.004)	(0.004)
-8	0.003	0.003	-0.006	-0.022	-0.038	-0.007*	0.007*	0.001
	(0.005)	(0.005)	(0.005)	(0.026)	(0.031)	(0.004)	(0.004)	(0.004)
-7	0.007	-0.007	0.001	0.007	0.014	-0.018***	0.006	0.006
	(0.005)	(0.005)	(0.005)	(0.025)	(0.028)	(0.004)	(0.004)	(0.004)
-6	0.001	-0.002	0.000	-0.029	-0.001	-0.004	0.005	0.000
	(0.005)	(0.005)	(0.005)	(0.026)	(0.026)	(0.004)	(0.004)	(0.004)
-5	-0.004	0.006	-0.002	-0.002	0.001	-0.009**	0.010***	-0.002
	(0.005)	(0.005)	(0.005)	(0.023)	(0.026)	(0.004)	(0.004)	(0.003)
-4	-0.005	-0.000	0.006	0.026	0.005	-0.008*	0.014***	0.003
	(0.004)	(0.005)	(0.005)	(0.024)	(0.026)	(0.004)	(0.004)	(0.003)
-3	-0.013***	0.008*	0.004	-0.002	-0.037	-0.002	0.013***	-0.003
	(0.004)	(0.005)	(0.005)	(0.023)	(0.028)	(0.004)	(0.004)	(0.003)
-2	-0.001	0.002	-0.001	-0.013	-0.007	-0.007*	0.006*	-0.005
	(0.004)	(0.005)	(0.005)	(0.024)	(0.028)	(0.004)	(0.004)	(0.003)
-1	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(.)	(.)	(.)	(.)	(.)	(.)	(.)	(.)
0	-0.007	0.012**	-0.005	0.022	0.021	0.005	0.010***	0.001
	(0.004)	(0.005)	(0.004)	(0.021)	(0.025)	(0.004)	(0.004)	(0.003)
1	-0.013***	0.008	0.005	0.011	-0.001	0.010***	0.017***	0.001
	(0.004)	(0.005)	(0.005)	(0.023)	(0.027)	(0.004)	(0.004)	(0.003)
2	-0.010**	0.012**	-0.002	0.004	-0.045	-0.005	0.010**	-0.000
	(0.005)	(0.005)	(0.005)	(0.023)	(0.031)	(0.004)	(0.004)	(0.004)
3	-0.006	0.014***	-0.007	0.027	0.001	-0.010***	0.008**	0.003
	(0.005)	(0.005)	(0.005)	(0.023)	(0.027)	(0.004)	(0.004)	(0.004)
4	-0.010**	0.010**	-0.000	0.030	0.006	-0.006	0.006	-0.001
	(0.005)	(0.005)	(0.005)	(0.025)	(0.029)	(0.004)	(0.004)	(0.004)
5	-0.001	0.008	-0.007	0.014	0.018	0.001	0.010**	0.002
	(0.005)	(0.005)	(0.005)	(0.025)	(0.027)	(0.004)	(0.004)	(0.004)
6	0.006	0.003	-0.008*	0.009	-0.003	0.004	0.014***	-0.002
	(0.005)	(0.005)	(0.005)	(0.024)	(0.026)	(0.004)	(0.004)	(0.004)
7	-0.009*	0.011**	-0.002	-0.021	-0.034	-0.012***	0.015***	-0.000
	(0.005)	(0.005)	(0.005)	(0.026)	(0.031)	(0.004)	(0.004)	(0.004)
8	0.003	0.003	-0.006	-0.006	0.018	-0.002	0.007	0.000
	(0.005)	(0.005)	(0.005)	(0.026)	(0.029)	(0.004)	(0.004)	(0.004)
9	0.002	-0.001	-0.001	-0.040	-0.012	-0.005	0.002	0.001
	(0.005)	(0.005)	(0.005)	(0.027)	(0.028)	(0.004)	(0.004)	(0.004)
10	-0.002	-0.006	0.008	-0.045	0.006	-0.001	-0.003	-0.002
	(0.005)	(0.006)	(0.005)	(0.030)	(0.029)	(0.004)	(0.005)	(0.004)
$\sum_{j=-10}^{-2} \tau_j$	-0.005	0.017	-0.012	0.021	-0.116	-0.065**	0.077***	0.004
	(0.031)	(0.034)	(0.033)	(0.165)	(0.204)	(0.028)	(0.027)	(0.024)
R-squared	0.785	0.817	0.642	0.776	0.726	0.408	0.844	0.913
Observations	365437	365437	365437	365437	79280	1465191	365437	365437
No. of Firms	47303	47303	47303	47303	11118	69771	47303	47303
Mean of Dep. in Control	0.249	0.626	0.124	0.637	0.293	0.252	0.764	0.379
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 τ_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. Columns 1 to 4 and 7 to 8 contain reported transactions, column 5 includes only observations on seller shortfalls, and column 6 analyses all monthly observations as a balanced panel. 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 F4. This table is mentioned in Section 4.2. *Source*: Data from monthly VAT returns submitted to the URA.

Table F6
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	-0.004		0.005**		-0.002	
	(0.002)		(0.002)		(0.002)	
Buyer only		-0.004		0.005		-0.001
		(0.004)		(0.004)		(0.004)
Seller only		0.000		0.008*		-0.008*
		(0.004)		(0.004)		(0.004)
Buyer and Seller		-0.010**		0.009**		0.001
		(0.004)		(0.004)		(0.004)
R-squared	0.785	0.785	0.817	0.817	0.642	0.642
Observations	365437	365437	365437	365437	365437	365437
No. of Firms	47303	47303	47303	47303	47303	47303
Mean of Dep. in Control	0.268	0.268	0.615	0.615	0.118	0.118
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 $\hat{\tau}$ 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< 0.05; ***p< 0.01. 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 F7
Post-Treatment Effect of the Letter on Outside-Link Trade:
Diff-in-Diff Specification

	Transaction size		Seller shortfall amount		Continuing trade		Seller Reports		Buyer Reports	
DiD specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Any Treatment	0.002		-0.002		-0.002		0.009***		0.000	
	(0.018)		(0.022)		(0.003)		(0.003)		(0.003)	
Buyer only		0.012		-0.015		-0.012***		0.006		0.002
		(0.020)		(0.023)		(0.003)		(0.004)		(0.003)
Seller only		-0.008		0.024		-0.000		0.003		0.001
		(0.020)		(0.023)		(0.004)		(0.004)		(0.003)
Buyer and Seller		0.004		-0.012		0.004		0.017***		-0.002
		(0.020)		(0.023)		(0.003)		(0.004)		(0.003)
R-squared	0.776	0.776	0.726	0.726	0.408	0.408	0.844	0.844	0.913	0.913
Observations	365437	365437	79280	79280	1465191	1465191	365437	365437	365437	365437
No. of Firms	47303	47303	11118	11118	69771	69771	47303	47303	47303	47303
Mean of Dep. in Control	0.629	0.629	0.301	0.301	0.264	0.264	0.742	0.742	0.394	0.394
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 $\hat{\tau}$ 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< 0.05; ***p< 0.01. This table is mentioned in Section 4.2. *Source*: Data from monthly VAT returns submitted to the URA.

G 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 G1, we show the corresponding results for buyers. Table G2 shows the regression results that correspond to Figure 9. In Table G3, we report effects on subsequent VAT liability controlling for the share of treated partners, following equation (9), discussed in detail in Appendix E. In Section G.3, we report heterogeneity results by seller characteristics (Table G4), by quintiles of pre-treatment sales (Figure G1) or discrepancy size (Figure G2), and the revenue effects for firms with no final sales (Table G5).

G.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 G1
Effect of Corrections on Firm-Level VAT Liability (Buyers)

Dependent variable:	ΔB2B Sales		ΔFinal Sales		ΔTaxable Inputs		ΔVAT Liability	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	-0.008		0.013		-0.013***		0.008**	
	(0.035)		(0.013)		(0.004)		(0.004)	
Buyer only		-0.015		0.003		-0.012**		0.006
		(0.039)		(0.019)		(0.006)		(0.005)
Seller only		-0.033		0.018		-0.010**		0.003
		(0.038)		(0.012)		(0.004)		(0.005)
Buyer and Seller		0.020		0.018		-0.016**		0.014**
		(0.045)		(0.012)		(0.007)		(0.007)
R-squared	0.006	0.007	0.004	0.005	0.008	0.009	0.005	0.007
Observations	2737	2737	2737	2737	2737	2737	2737	2737
No. of Firms	1231	1231	1231	1231	1231	1231	1231	1231
Pre-T Mean of Dep. in Control	46.992	46.992	21.050	21.050	57.809	57.809	1.842	1.842
Mean of Diff. in Control	0.055	0.055	-0.020	-0.020	0.002	0.002	0.005	0.005
Pre-T Median of Dep. in Control	0.866	0.866	2.871	2.871	12.185	12.185	0.119	0.119
P-value of β_S=β_B		0.535		0.336		0.744		0.574
P-value of β_{SB}=β_B		0.396		0.325		0.655		0.261
P-value of β_{SB}=β_S		0.150		0.980		0.418		0.130
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 (2) (odd columns) and (1) (even columns) at the firm-level. 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< 0.05; ***p< 0.01. This table is mentioned in Section 5.2. *Source*: Data from VAT monthly returns submitted to the URA.

G.2 Subsequent VAT Liability

Table G2
Post-treatment Effects on Firm-Level VAT Liability (Sellers)

	(1) B2B Sales	(2) Final Sales	(3) Taxable Inputs	(4) VAT Liability
-10	1.420	2.536	6.797	-0.362
	(5.388)	(7.434)	(9.938)	(1.204)
-9	9.703*	4.285	14.061	-0.215
	(5.243)	(7.147)	(10.005)	(0.982)
-8	2.494	1.928	15.964	-1.096
	(5.188)	(7.632)	(10.069)	(1.058)
-7	1.898	4.077	11.154	-0.776
	(5.009)	(7.181)	(9.622)	(1.076)
-6	-2.668	-1.771	-5.603	-0.141
	(5.605)	(7.609)	(9.703)	(0.869)
-5	0.178	6.463	2.775	0.342
	(4.726)	(7.319)	(9.100)	(0.984)
-4	1.865	-7.936	-8.254	0.094
	(4.325)	(6.030)	(8.512)	(0.776)
-3	0.554	-4.328	-5.007	-0.046
	(4.274)	(6.182)	(7.646)	(0.903)
-2	2.629	-8.687	-6.630	-0.614
	(5.870)	(7.074)	(9.902)	(1.049)
-1	0.000	0.000	0.000	0.000
	(.)	(.)	(.)	(.)
0	0.129	5.015	-3.919	0.668
	(5.051)	(5.485)	(7.692)	(0.964)
1	5.356	1.340	2.608	0.151
	(4.851)	(6.409)	(8.905)	(0.986)
2	3.044	5.115	2.044	1.020
	(5.235)	(6.831)	(8.775)	(1.020)
3	5.153	5.648	11.176	1.099
	(5.792)	(6.699)	(9.369)	(1.043)
4	9.382	7.324	10.790	0.459
	(6.229)	(7.596)	(9.681)	(0.995)
5	7.850	4.417	10.746	-0.630
	(5.909)	(6.599)	(9.214)	(1.033)
6	10.953*	5.377	8.389	0.727
	(6.227)	(6.141)	(9.423)	(1.091)
7	-3.775	4.046	-3.594	0.787
	(6.171)	(7.309)	(9.602)	(1.069)
8	3.702	3.292	4.001	-0.079
	(6.702)	(7.247)	(10.825)	(1.013)
9	4.973	1.687	-1.239	1.290
	(7.746)	(6.922)	(10.169)	(1.084)
10	3.427	4.845	6.968	-0.497
	(8.189)	(6.867)	(9.368)	(1.277)
$\sum_{j=-10}^{-2} \hat{\tau}_j$	18.073	-3.433	25.256	-2.813
	(35.421)	(53.702)	(66.450)	(6.325)
R-squared	0.846	0.901	0.876	0.550
Observations	25085	25085	25085	25085
No. of Firms	1235	1235	1235	1235
Mean of Dep. in Control	53.514	106.029	142.991	3.337
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 $\hat{\tau}_j$ 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 9. This table is mentioned in Section 5.2. *Source*: Data from monthly VAT returns submitted to the URA.

Table G3
Effect of Treatment on Subsequent Firm-Level VAT Liability for Sellers controlling for Treated Partners (Difference-in-Differences)

	(1) VAT Liability	(2) VAT Liability	(3) VAT Liability
T_i	1.278	1.272	0.981
	(0.829)	(0.829)	(0.765)
$\bar{T}_i$		2.213***	
		(0.679)	
$\bar{T}_i \times (1 - T_i)$			1.723*
			(0.909)
$\bar{T}_i \times T_i$			2.540***
			(0.957)
R-squared	0.002	0.002	0.002
Observations	25085	25085	25085
No. of Firms	1235	1235	1235
Pre-T Mean of Dep. in Control	3.321	3.321	3.321
P-Value testing for coefficient equality		0.650	0.923
P-value (joint F-test of shares)		0.001	0.005
Month-Year FE	Yes	Yes	Yes

Notes: This table reports the effects of the letters on subsequent VAT liability when controlling for the share of trading partners that are treated. In column 1, T_i is an indicator taking value one for treated sellers in the months post-intervention. In column 2, $\bar{T}$ is the share of partners of the seller that are treated. VAT returns for 10 months before and 10 months after the treatment are analyzed at the seller-month level (April 2017 to December 2018). *Mean of Dep. in Control* reports the average value of the dependent variable for firms in the control group. All outcomes are reported in thousands of US dollars and are winsorized at the 0.5% level. Standard errors are clustered at the seller level. *p<0.10; **p<0.05; ***p<0.01. This table is mentioned in Section 5.2. *Source*: Data from monthly VAT returns submitted to the URA.

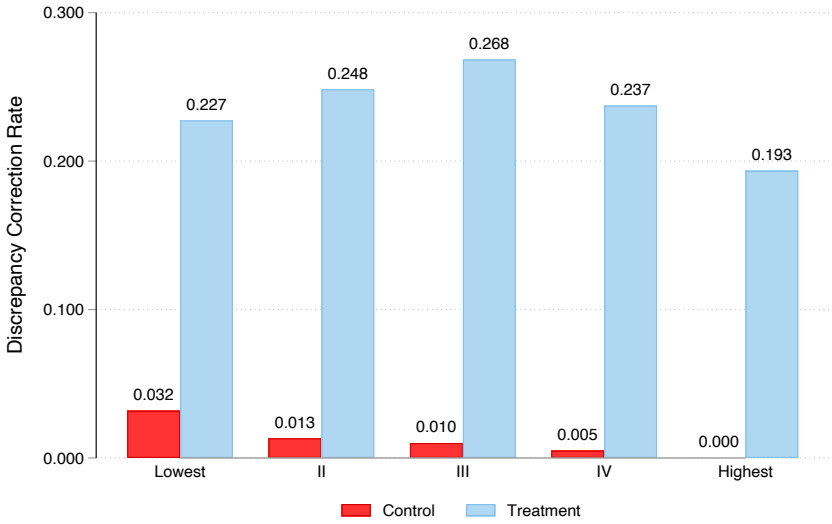
G.3 Heterogeneity

Table G4
Heterogeneity Analysis: Any Correction by the Seller

OLS						Lasso	
Seller Characteristics (1)		**Buyer Characteristics** (2)		**Pair Characteristics** (3)		**Relevant Characteristics** (4)	
Degree	0.000	Degree	0.000	High trading volume	0.051		
	(0.000)		(0.000)		(0.036)		
Sector: Agriculture	0.132	Sector: Agriculture	0.109	Initial discrepancy	−0.016**		
	(0.151)		(0.148)		(0.007)		
Sector: Construction	0.065	Sector: Construction	0.104*				
	(0.072)		(0.060)				
Sector: Manufacturing	0.030	Sector: Manufacturing	0.057				
	(0.055)		(0.071)				
Sector: Mining	−0.253	Sector: Mining	0.335				
	(0.316)		(0.203)				
Sector: Retail	0.085*	Sector: Retail	0.010				
	(0.048)		(0.055)				
Sector: Service	0.130***	Sector: Service	−0.040				
	(0.044)		(0.046)				
Ln Sales	0.000	Ln(Sales)	0.000				
	(0.000)		(0.000)				
Share of negative returns	−0.054	Share of negative returns	0.043				
	(0.045)		(0.046)				
Existence length	0.000**	Existence length	0.000				
	(0.000)		(0.000)				
FC ratio	0.059	FC ratio	−0.072				
	(0.046)		(0.044)				
Audited in 2016	−0.051	Audited in 2016	−0.025				
	(0.072)		(0.076)				
R-squared	0.065						
Observations	741					Observations	741

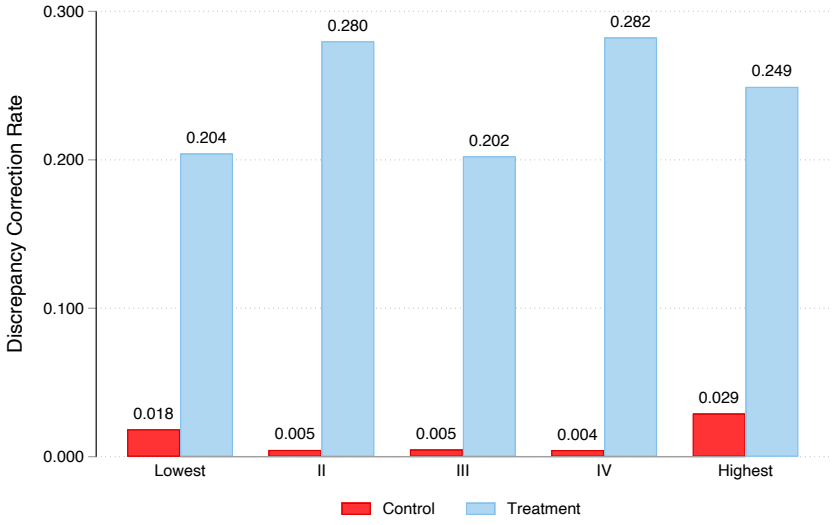
Notes: This table reports results on the baseline characteristics of firms that are correlated with whether a firm responds to the letter. We estimate an OLS regression (columns 1-3) and a Robust Lasso (column 4). The outcome variable is an indicator taking value 1 if at least one discrepancy is corrected by either firm in the link. We regress the outcome variable 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 logarithmic transformation of the amount of sales (in thousands of USD), (v) the share of monthly VAT returns with negative Value-Added, (vi) number of months since firm registered for VAT, (vii) the ratio of final sales to total sales, 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 5.2. *Source*: Data from monthly VAT returns submitted to the URA.

Figure G1
Discrepancy Correction Rate by Sales Quintile in Pre-treatment Periods



Notes: This figure shows the share of corrected discrepancies for treated and control links, by quintile of sales of the seller in the pre-treatment period. This figure is mentioned in Section 5.2. *Source*: Data from monthly VAT returns submitted to the URA.

Figure G2
Discrepancy Correction Rate by Seller Shortfall Amount in Pre-treatment Periods



Notes: This figure shows the share of corrected discrepancies for treated and control links, by quintile of the amount of seller shortfall in the pre-treatment period. This figure is mentioned in Section 5.2. *Source*: Data from monthly VAT returns submitted to the URA.

Table G5
Effect of Treatment on Firm-Level VAT Liability (Sellers with no Final Sales)

Dependent variable:	ΔB2B Sales		ΔFinal Sales		ΔTaxable Inputs		ΔVAT Liability	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any treatment	0.340**		0.000		0.005		0.044***	
	(0.148)		(0.000)		(0.009)		(0.015)	
Buyer only		0.197		0.000		0.009		0.022
		(0.198)		(0.000)		(0.013)		(0.018)
Seller only		0.585*		0.000		0.005		0.062*
		(0.350)		(0.000)		(0.009)		(0.035)
Buyer and Seller		0.271		0.000		0.001		0.048*
		(0.175)		(0.000)		(0.007)		(0.025)
R-squared	0.044	0.051	0.000	0.000	0.028	0.029	0.053	0.059
Observations	377	377	377	377	377	377	377	377
No. of Firms	197	197	197	197	197	197	197	197
Mean of Dep. in Control	81.496	81.496	0.000	0.000	77.946	77.946	0.639	0.639
Mean of Diff. in Control	0.057	0.057	0.000	0.000	-0.002	-0.002	0.002	0.002
Median of Dep. in Control	3.224	12.108	0.000	0.000	0.905	2.035	0.000	0.597
P-value of β_S=β_B		0.319		0.000		0.630		0.303
P-value of β_{SB}=β_B		0.772		0.000		0.392		0.401
P-value of β_{SB}=β_S		0.409		0.000		0.149		0.750
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 (2) (odd columns) and (1) (even columns) at the firm-level. We restrict the sample to sellers that have no final sales prior to the intervention. 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. Standard errors are clustered at the seller level. *p< 0.10; **p< 0.05; ***p< 0.01. This table is mentioned in Section 5.2. *Source*: Data from monthly VAT returns submitted to the URA.

H Implementation 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 1708 letters for send-off.[44] Taxpayers got in touch with the URA upon receiving the letter. The time it took to respond to each taxpayer depended on the type of communication. Specifically, when a firm visited the URA it took around 45 minutes, when a firm sent an email or physical letter it took around 25 minutes to respond (since the information provided by the taxpayer in the email often needed to be reconciled with information in the URA's database), and if a firm called it took around 5 minutes. Finally, after firms had responded, 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 multiplying the number of hours with the average hourly salary, before deductions, for a junior URA officer. We calculate that the monetary cost associated with the number of hours worked is $2,415.

The total cost incurred by the URA for the intervention is thus $5,716.[45]

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

[45] In 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.