Institute for Fiscal Studies

Stephan Dietrich
Yannick Markhof
Rose Camille Vincent

The cost of bureaucratic fragmentation: business tax evasion and revenue mobilization in a low-income country

Working paper

26/39

Economic and Social Research Council

The cost of bureaucratic fragmentation: Business tax evasion and revenue mobilization in a low-income country *

Stephan Dietrich
Maastricht University
UNU-MERIT

Yannick Markhof
ETH Zurich
UNU-MERIT

Rose Camille Vincent
Utrecht University
IFS-TaxDev

May 27, 2026

Abstract

We provide novel evidence on bureaucratic fragmentation and weak tax administrations as central enablers of low revenue mobilization in low-income countries. In collaboration with the municipal and national tax authorities in Kampala, Uganda, we cross-link previously siloed tax records for 155,000 firms and conduct a large-scale experiment with 60,000 firms. We document pervasive and selective tax evasion: only 14% of verifiably active firms comply with both government tiers. Cross-record linkage almost triples detectable non-compliance while offering increased enforcement efficiency. This coordination dividend is left untapped. Firms exploit the resulting loopholes through partial informality, re-registering under new identities, and strategic late payments. In a cross-authority field experiment, deterrence nudges, including messages signaling inter-authority coordination, fail to offer a light-touch alternative to addressing fragmentation directly. Our findings establish bureaucratic fragmentation as a distinct and costly source of passive waste in tax administration that existing approaches to revenue mobilization rarely address.

Keywords: Taxation; Tax evasion; Tax administration; Low-income countries; Nudges

JEL Codes: H26; H20; H71; C93; O12

*We gratefully acknowledge the financial support of ETH4D at ETH Zürich (Research Challenges Grant 2023), the International Centre for Tax and Development (ICTD), and UNU-WIDER. We also appreciate the long-standing partnership with the Kampala Capital City Authority (KCCA) and the Uganda Revenue Authority (URA). In particular, we extend our gratitude to Ezra Ssebuwufu and Nicholas Musoke for their dedication, engagement, and strategic follow-through in implementing this research. We are particularly grateful to Firminus Mugumya for his role in shaping the broader research program underlying this study. We also thank our collaborators at PADRI in Kampala for their excellent work in survey data collection, and Fred Kasalirwe in particular for excellent research assistance. Finally, we acknowledge the valuable contributions of all individuals and institutions involved in this project. Usual disclaimers apply. The experiment was pre-registered as *AEARCTR-0012747*.
Contact: s.dietrich@maastrichtuniversity.nl (SD), ymarkhof@ethz.ch (YM), r.c.vincent@uu.nl (RCV).

1 Introduction

Tight fiscal space constraining investments into economic and social development is among the principal challenges faced by low- and middle-income country (LMIC) governments (Jensen et al., 2024). A large literature in public and development economics is therefore concerned with policy interventions that increase revenue collection in these contexts (see for e.g. Pomeranz and Vila-Belda, 2019; Slemrod, 2019; Okunogbe and Tourek, 2024; Jensen and Weigel, 2025).

In many low-income countries, the fragmentation of available capacity, structural gaps, and narrow tax bases create a conundrum: Policy options such as increased enforcement via audits, third-party reporting, or tax rate reforms have proven effective in middle- and high-income countries. Yet, these options require effective administrative structures that are rarely well established or investments that are out of reach in low-income contexts (Jensen and Weigel, 2025). How tax administration itself shapes fiscal space and policy levers to raise revenue is therefore increasingly gaining economists' attention (Besley and Persson, 2014; Pomeranz and Vila-Belda, 2019; Jensen and Weigel, 2025), but the specific role of fragmentation and firms' behavioral responses to it remain scarcely studied and difficult to quantify.[1]

In this paper, we provide novel empirical evidence on how bureaucratic fragmentation between tax authorities creates "passive waste" (Bandiera et al., 2009) and constrains revenue mobilization in low-income countries. We compile a unique dataset consisting of cross-linked, longitudinal tax records for over 150,000 firms in Kampala, Uganda, across the previously siloed records of two independently operating government authorities (municipal and national). This allows us to document the cost of weak inter-authority coordination, identify structural gaps and loopholes in tax administration, and show how firms strategically exploit these deficiencies. We then investigate how tax administrations should respond to these gaps and rule out a deceptively easy 'fix'. Given capacity constraints, tax nudges—short deterrence or reminder messages sent to taxpayers—have become a popular tool for raising tax compliance in LMICs (Antinyan and Asatryan, 2025; Slemrod, 2019; Mascagni, 2018) precisely because they require no structural reform to implement. In a field experiment involving both tax authorities and 60,317 firms in Kampala, we provide causal evidence that such light-touch approaches fail to provide a shortcut to the revenue potential from improved inter-agency cooperation.

[1]For broader overviews of the structural and institutional challenges to tax administration and revenue mobilization in low-income countries, see Besley and Persson (2014) for a conceptual framework on limited fiscal capacity in developing economies, Moore et al. (2018) for political and institutional perspectives, Okunogbe and Tourek (2024) for the role of technology for revenue collection, and Jensen and Weigel (2025) for challenges within and faced by tax administrations.

Firms in our data are liable to the municipal (Kampala Capital City Authority, KCCA) and national (Uganda Revenue Authority, URA) tax authorities. At a minimum, these firms have to file either presumptive (smaller firms) or corporate (larger firms) business income tax to the URA every fiscal year and hold an active trade license with the KCCA that is payable annually. Despite efforts to increase collaboration, both authorities had previously been working independently of each other, with limited collaboration in data integration, tax evasion detection, or enforcement. We partnered with both tax authorities on what, to our knowledge, was the first large-scale effort to integrate their tax records and conduct a cross-government tier experiment in coordinated detection of tax evasion.

This set-up allows us to address a key challenge in detecting tax evasion in developing countries: in an environment in which half of small firms are predicted to cease operations within six years (McKenzie and Paffhausen, 2019) and firms often do not formally de-register (Mascagni et al., 2022), discerning tax evasion from mere inactivity or permanent firm closure is key. We observe the same firms over the same period in two independent tax records and can therefore ascertain a firm's activity status when it does not file with one authority in the records of the other. Based on this setup, we empirically document several new insights about how businesses' tax-filing behavior responds to incentives set by bureaucratic fragmentation, administrative loopholes, and structural gaps in tax enforcement.

First, tax non-compliance is pervasive and selective, making cross-record linkage highly valuable. Among firms that were verifiably active in a given fiscal year between 2018 and 2023, on average 43% did not file national business tax to URA and 49% did not pay for a municipal trade license to KCCA. Virtually every firm in our data misses at least one tax declaration or payment over this period despite being active. Firms have a high likelihood of being serial offenders: 72% never file taxes with one of the two authorities. This makes tax records from both authorities highly complementary and allows for the detection of over 260,000 *additional* cases of tax non-compliance that could not be detected when relying on a single authority's records. Put differently, cross-record linkage almost triples the total number of detectable tax non-compliance cases in Kampala between 2018 and 2023. We estimate the associated fiscal shortfall to sum to 42% of annual business tax revenues for URA and to 102% of current trade license revenues for KCCA.

Crucially, in a context with limited enforcement resources, inter-authority collaboration not only dramatically increases the number of *detectable* non-filing cases but also reduces the marginal cost of following up on them through audits. The ability to discern non-compliance from mere inactivity or business closure ex-ante—that is, before audit resources are deployed—increases the rate at which an audit deployed should lead to a confirmed case of non-compliance. In a physical tracking exercise we conduct with 1,270 firms, cross-checking KCCA's records for recent activity increases the rate at which URA non-filers can be successfully tracked down one year after missing the filing deadline by 18–21%.

Second, the enforcement capacity of any single authority is severely constrained and leaves a gap exactly where the other authority focuses its efforts. Only 2.5% of registered taxpayers are ever assigned an audit by the national tax authority over a ten-year period. Further, capacity constraints mean that for every assigned audit that is effectively conducted, 15 other audits assigned to suspicious cases by the URA are not followed through with. As URA's audits are predominantly targeted at larger taxpayers, small- to medium-sized firms and unincorporated businesses are unlikely to ever receive enforcement action. Over ten years, fewer than one in a thousand firms filing small business tax ever receive a completed audit, whereas 7.4% of firms that file corporate income tax do. Conversely, the municipal authority uses its boots-on-the-ground capacity in Kampala to conduct geographically targeted trade license inspections. In 2025 alone, these inspections resulted in the temporary closure ("sealing") of 3,870 businesses in Kampala and predominantly covered precisely the small- and medium-sized population of firms that audits by the national tax authority almost entirely miss. This makes the enforcement efforts of both authorities highly complementary in principle. In practice, however, this complementarity is left entirely unexploited.

Third, we document several strategies exploiting fragmentation and loopholes that firms adopt to evade taxation. For example, firms escape enforcement by opting for (partial) informality. We document thousands of cases in which a previously compliant firm stopped filing national business tax but continued to pay for a municipal trade license. Among a sample of 17,040 firms that stopped filing national business taxes, a third (32%) still actively renewed their trade license three years later. Over two-thirds of firms that stopped filing national business taxes, but remained active, still had unresolved tax debt in 2023, compared to 46% of other businesses. We also find firms electing to restart their business under new identities. We document several hundred cases in which a business owner stopped filing national business tax for their firm, only to register an observably near identical business under a new tax identification number (TIN) soon after. As we are limited to identifying only the most glaring cases, this number likely represents a vast underestimate of the true prevalence of the issue. In fact, in a survey of small- to medium-sized firms we conducted in Kampala, business owners estimate that about 30% of their peers would prefer to reopen as a new business rather than pay old tax debts when they are enforced. Finally, we find that even those firms that are compliant at the extensive margin exploit administrative loopholes to lower their effective tax payments. As trade license renewal is required once every twelve months, but trade license validity is not backdated when renewed late, we find that those businesses that renew their trade license in a given year systematically do so late. On average, this affords businesses approximately half a month (15 days) of 'trade license holidays' per renewed trade license and costs the KCCA an estimated 3.2% of total trade license revenue every year.

How should tax administrations respond to these gaps? Our results suggest high returns to reducing fragmentation in tax administration via coordinated data sharing and firm identification, non-compliance detection, and enforcement. However, experience with a previous policy initiative to improve collaboration between tax authorities in Kampala (described in Jouste et al. (2021)) suggests that coordination requires administrative reforms and lasting commitment, without which the long-term effectiveness of such initiatives will remain limited. This increases the outward appeal of solutions that eschew institutional reforms in favor of simpler policy responses.

One such option are tax nudges, short messages with a reminder, deterrence, or moral appeal character that are sent to (prospective) taxpayers ahead of the filing deadline (Mascagni, 2018; Antinyan and Asatryan, 2025). Due to their simple implementation, tax nudges have proliferated from high-income to low- and middle-income countries, including Uganda, and have been found to provide a cost-effective tool for revenue recovery.[2] These properties make the approach promised by them—merely *signaling* increased enforcement and coordination to taxpayers—a seemingly appealing 'quick fix'. In a large-scale field experiment, we offer causal evidence on the effectiveness of such a light-touch approach to reducing (perceived) fragmentation of tax administration.

[2]In a meta-analysis, Antinyan and Asatryan (2025) list 71 randomized controlled trials across a range of schemes, taxpayers, and regions. Deterrence messages, the most effective type of tax nudge, are estimated to increase the probability of compliance by 5.9 percentage points (pp) on average, or approximately 30% relative to a baseline without messages. As a result of the low cost of sending messages, several papers have reported impressively high rates of return estimates. A few selected examples from the literature: In the Dominican Republic, Holz et al. (2023) find that their treatment messages increased the amount of income taxes paid by firms by $184 million, equivalent to about 0.22% of the country's 2018 GDP, compared to a control message, at trivial marginal costs. In Costa Rica, Brockmeyer et al. (2019) finds a 3.4 p.p. return (equivalent of US$ 15) among previously non-filing firms. In a message experiment in Belgium, de Neve et al. (2021) estimates that simplification of messages alone could have increased personal income tax collection by €17.5 million, compared to the costs of the nudge intervention of about €80,000. Hallsworth et al. (2017) estimate that a message experiment in the UK raised more than £9 million at negligibly small costs. Evidence from Sub-Saharan Africa is scarcer but points in a similar direction. An RCT in Rwanda led to additional tax revenue of about $6 million at a cost of about $4000 Mascagni and Nell (2022), in Eswatini, Santoro (2024) finds extra revenue gains of $0.2 million and a cost-benefit ratio of about 1:11, and in Tanzania Collin et al. (2025) find a mean benefit-cost ratio of 36:1 for a property tax intervention. Closest to our study, Cohen (2024) finds a marginal rate of return of 8.8 for a deterrence message in an income tax message experiment in Uganda.

We partner with the national and municipal tax authorities in a cross-government field experiment involving 60,317 firms in Kampala and send deterrence messages via text to a random subsample of businesses before the tax filing deadline. Our messages threaten enforcement if the national business tax is not filed and/or a trade license is not renewed, randomly vary which tax authority sends the message, and whether the message is jointly signed by both tax authorities. However, in line with the light-touch premise of tax nudges, we hold other structural factors of collaboration and enforcement actions constant. Tax nudges, including those signaling increased collaboration, have robust and precise null effects. This finding holds across both tax authorities and businesses of different size or formality status. Using our cross-linked dataset and a deterrence message intervention conducted by the URA four years prior (Cohen, 2024), we also demonstrate that the null effects of simple nudges are replicable over time and hold when considering dynamic effects over multiple years. Taken together, this strongly supports the notion that tax nudges fail to overcome the systematic gaps that weak and fragmented tax administration leaves.

Our results extend our knowledge on the role of state capacity for public revenue generation in low-income countries, firms' behavioral responses to fragmented tax administrations, and the scope to broaden the tax base in this context. Existing evidence has shown that low state capacity constrains policy effectiveness in low-income countries (Bergeron et al., 2024; Cohen, 2024; Henning and Okello Ayo, 2025) and that increasing administrative capacity can raise significant amounts of additional revenue in an LMIC-context (Weigel, 2024; Basri et al., 2021; Balán et al., 2022; Khan et al., 2016). A further strand of literature has demonstrated the value of improving information flows within tax administrations, for example via third-party reporting, VAT paper trails, or electronic filing and invoicing systems (Brockmeyer et al., 2024; Almunia et al., 2024; Mascagni et al., 2023; Okunogbe and Tourek, 2024; Carrillo et al., 2023, 2017; Naritomi, 2019; Pomeranz, 2015). However, evidence on the role of fragmentation in tax administration across different authorities and government tiers, a characteristic shared by many low- and middle-income countries, is exceedingly scarce (Vincent, 2023).

We provide some of the first empirical evidence on how a lack of coordination across government tiers creates exploitable gaps in the tax net, the approximate cost of those gaps, and firms' responses to them. This is made possible by a unique collaboration with two independently acting tax authorities that are responsible for the same group of taxpayers in Kampala and that keep separate, but linkable, records. In contrast to policies that build on institutional investments and capacity expansions, the revenue potential we identify lies in better institutional collaboration and more efficient use of *existing* capacity.

To this end, we document what amounts to “passive waste” but on the revenue side of fiscal policy (Bandiera et al., 2009): revenue potential that falls through the cracks of a fragmented tax system and that may be recovered by fixing administrative loopholes, between and within tax authorities, that firms currently exploit.

Our results have two main implications in an environment in which tax administrations in low-income countries are under immense pressure to deliver higher levels of revenue mobilization under tight resource constraints.

First, they highlight reductions in bureaucratic fragmentation as a source of significant revenue potential and an opportunity to broaden the tax base. The policy levers we identify tackle the reduction of inefficiencies given state capacity rather than requiring high(er) state capacity as a precondition (Jensen and Weigel, 2025). In contrast to interventions that have proven successful in high-state-capacity environments but are ill-equipped for the structural challenges faced by low-income countries, reducing fragmentation and closing loopholes can involve solutions that are comparatively within reach. For example, we demonstrate that integrating tax records and enforcement efforts between authorities responsible for the same set of taxpayers allows to cast a much wider tax net. This facilitates identifying, finding, and confirming tax evasion among tens of thousands of additional businesses every year that currently face close to zero enforcement risk—while improving the efficiency of enforcement for both parties. Other opportunities we identify relate to inefficiencies within the administration of a single authority. For instance, the simple practice of backdating trade license renewals to the date of the previous license’s expiration has a revenue potential that compares favorably to the gains that can be expected from policy interventions that are much more complicated or reliant on state capacity.

Second, closing the gaps and loopholes that siloed record-keeping, insufficient harmonization of taxpayer identification, and uncoordinated enforcement leave takes improved inter-authority coordination, not mere signals to taxpayers. In an environment in which taxpayers have come to expect structural enforcement gaps and learned to actively exploit them, new communication strategies are unlikely to sufficiently shift taxpayers’ priors and offer no shortcut to the revenue potential from improved institutional collaboration. These results add to growing evidence that, despite their popularity and appeal as ‘quick and cheap’, tax nudges in isolation are unlikely to deliver the sustained shifts in revenue mobilization that low-income countries need (Antinyan and Asatryan, 2025).[3]

[3]Our results align with those of Antinyan and Asatryan (2025) who found that tax nudges are, on average, less effective in low-income countries and that published evidence appears to overstate their

In sum, our paper's core message is still an encouraging one for low-income countries. There is substantial revenue potential in making more efficient use of the capacity tax administrations already possess. Tapping into this potential will require sustained commitment to reduce bureaucratic fragmentation and non-trivial institutional changes that are plausibly within reach. Not least, strengthening tax administration in this way—by expanding the tax base, establishing efficient detection and enforcement protocols, and closing loopholes—creates an enabling environment in which other ambitious policy interventions and structural reforms can make strides toward fiscal sustainability in low-income countries.

Our paper is organized as follows. Section 2 provides background on the business tax environment in Kampala and collaboration challenges between authorities prior to our study. Section 3 introduces our data sources and describes the cross-record linking process. We cover the design of the tax nudge experiment in Section 4. In Section 5 we present the empirical results of the descriptive and experimental analysis, and we discuss our findings in Section 6.

2 Institutional context

2.1 Business taxation in Kampala

Businesses in Kampala are taxed at the municipal and national levels by the KCCA and the URA, respectively. At the municipal level, businesses must obtain a yearly-renewable trade license in order to operate (KCCA, 2026b). A firm may need multiple trade licenses if it has multiple branches (such as several restaurants) or engages in multiple business activities on the same premises (for example, a beauty salon located within a hotel). A predetermined schedule issued by the Ministry of Trade and Industry and the KCCA prescribes fees per business type for the different 'grading' or location areas (KCCA, 2026a). The trade license costs are therefore a function of a firm's location and the nature of its business. The latter is assigned upon registration from a classification list. The rate assessment is done by a revenue officer, who issues a KCCA payment advice form, and the corresponding payment is executed using any of the KCCA-designated mobile payment platforms or banks. Trade licenses need to be renewed every 12 months,

average effectiveness through the systematic omission of null results. Our results are consistent with evidence that finds tax nudge interventions to deliver larger gains when paired with changes in actual enforcement practices (see for e.g. Holz et al., 2023; Gil et al., 2024; Brockmeyer et al., 2019) and align with evidence from Cohen (2024), who finds tax nudges in Uganda to be effective only in areas in the bottom quintile of state reach, where they can still shift taxpayers' priors about state capacity and enforcement.

but because renewals used to be based on calendar year, more than 70% of trade licenses are due in January to March.

All firms in Kampala must also file a national business income tax return with the URA. Businesses file either corporate income tax (CIT) or presumptive income tax (PT). The latter has been devised to facilitate the formalization and integration of small and medium-sized enterprises into the tax system. Since 2022, PT applies to firms with a turnover between UGX 10 million and UGX 150 million and follows a progressive tax schedule (Jouste et al., 2024).[4] Marginal tax rates range from 0.4% (for turnover between UGX 10 million and UGX 30 million) to 0.7% (for turnover between UGX 80 million and UGX 150 million). Firms with a turnover above UGX 150 million must file CIT at a rate of 30% of taxable income.[5] URA's business income taxes follow the default financial year, which runs from July 1 to June 30. Tax returns must be filed within six months of the end of the tax year. In practice, most firms file their returns at the end of the calendar year, i.e., near the end of the statutory filing period.

Since the inception of the Taxpayer Registration Expansion Program (TREP) in 2013, closer collaboration between URA and KCCA has been an agreed-upon goal. As part of TREP, all firms in Kampala are mandated to obtain a Taxpayer Identification Number (TIN), which is the universal firm identifier across authorities in Kampala. This is facilitated via one-stop shops for business registration, where new firms are issued a TIN and registered with both URA and KCCA, thereby ensuring linkable identifiers across their records. The TIN allows us to link firms between the records of KCCA and URA.[6]

2.2 Tax enforcement challenges

Prior to our research, collaboration between national and municipal authorities in the detection and enforcement of tax non-compliance had been limited. Both operated largely independently in identifying non-compliers and in enforcing tax obligations on businesses. From extensive conversations with public officials at both authorities, we identified several reasons for this.

First, coordination between the URA and the KCCA is hampered by the fact that both are independent authorities, each with its own mandate, and keep siloed records and

[4] Firms with an annual turnover below UGX 10 million (about USD 2,800 at 2025 exchange rates) are not taxed on their turnover. However, they are still legally obliged to file a return under the Ugandan Income Tax Act (Parliament of Uganda, 1997). The PT tax schedule has been adjusted over the years, with the current schedule in place since 2022.

[5] That is, CIT applies to taxable income after eligible deductions, whereas no deductions exist for PT. Architectural, engineering, accounting, legal, or other professional services must file for corporate income tax (CIT) regardless of their turnover (see for e.g. Waiswa et al., 2021).

[6] See Jouste et al. (2021) for a detailed description of the TREP initiative.

enforcement strategies for the taxes they levy on local businesses in Kampala. Further, the KCCA continues to issue its own City Operator Identification Number (COIN) to identify trade license holders internally while only imperfectly enforcing and recording TINs in its records. This is in spite of TREP, which pledged closer collaboration between the two authorities and the universal availability of a common, unique identifier across both authorities' records. In practice, better coordination is hindered by bureaucratic and political hurdles, as each authority operates under its own mandate and incentives.[7]

Second, coordination is complicated by different administrative processes between URA and KCCA. URA collects business income tax (CIT/PT) on a fiscal-year basis through self-declared returns, while KCCA issues 12-monthly trade license fees, traditionally aligned with the calendar year. Payments are made via separate payment slips and channels, and there is no systematic platform for reconciling transactions across the two systems.[8] This administrative mismatch promotes the retention of two parallel records (as opposed to a unified one), even where records are, in principle, linkable.

Third, both authorities have either lacked or not allocated sufficient capacity to reliably and regularly integrate data in a way that makes it usable for non-compliance detection. Data sharing remains a manual process: one authority transmits a raw extract to the other, which must then be checked for completeness, cleaned, and harmonized before integration is possible. Even when the requisite data are transmitted, it took us considerable time and effort to make them usable at a scale that supports enforcement. Previous attempts by tax authority staff to combine records for non-compliance detection were thwarted by the technical and bureaucratic burden of cross-agency data integration.[9] The lack of a systematic, permanent, and interoperable data infrastructure, even after the inception of the TREP, undermines the efficiency of data integration efforts when attempted and leaves structural loopholes that enable firms to comply with one authority while evading the other with limited risk of detection.

Fourth, the national and municipal authorities predominantly focus their enforcement on different groups of taxpayers. The municipal authority has greater on-the-ground capacity due to its narrower focus and a large group of inspectors who conduct sweeping

[7]For example, anecdotal evidence from conversations with KCCA and URA staff suggests that one-stop shops were sometimes not manned by representatives from both authorities, which limited their effectiveness; and data were not shared or lacked sufficient detail, as interoperability was not treated as a shared priority.

[8]For example, URA requires an annual electronic income tax return filed six months after the end of the fiscal year, while KCCA issues a paper-based payment slip tied to the annual renewal of a trade license, usually concentrated in the first quarter of the calendar year.

[9]In conversation with staff from both authorities, we heard anecdotes of several integration projects that were abandoned because of a lack of sustained capacity for data processing and close collaboration across authorities.

inspections of all businesses in a targeted geographic area. These inspections are typically unannounced and broad in scope, allowing the municipal authority to quickly detect non-compliance among small and semi-informal businesses within a small geographic area (e.g., a street, block, or a mall). Conversely, the national authority conducts targeted, data-driven audits that focus on larger firms, for which there is greater potential for revenue recovery (Henning and Okello Ayo, 2025). Both authorities also differ in their legal mandate, that of the municipal authority being more limited. While the municipal authority temporarily closes businesses ("seals them", in its own words) that operate without a valid trade license, it is only authorized to do so with a court order. This limits its effectiveness among larger firms that can afford legal representation and, therefore, affects its enforcement focus. By contrast, the national authority has a stronger legal mandate and a greater analytical capacity to potentially utilize data sharing. Yet much of the complementary information in the municipal records pertains to small- and medium-sized businesses, which does not align with the national authority's current enforcement efforts that concentrate on large taxpayers.

In sum, these aspects left ample scope for data integration and collaboration between the two authorities, which remained largely untapped before our study.

3 Data sources and description

3.1 Tax records and cross-linkage

We collaborate with the KCCA and URA to harmonize and integrate the tax records of both authorities for Kampala. To do so, we draw on three data sources: (i) KCCA's records of trade license payments; (ii) URA's records of business income tax filings; (iii) the registry of all firms in Kampala with a taxpayer identification number (TIN), the mandatory, unique firm identifier across tax authorities in Kampala (see Section 2).

Linking firms across URA's and KCCA's records enables us to verify firms' activity status across authorities, identify selective non-compliance, and construct the sampling frame for the tax experiment. We describe the key elements of the linking process here. Appendix C provides complete technical details, including the full TIN-recovery pipeline, illustrative examples, and robustness checks.

3.1.1 Cross-authority record linkage

The goal of our exercise is to link taxpayers across the records of KCCA and URA. We do so via their unique Taxpayer Identification Number (TIN). Since one firm can have multiple trade licenses, but the entity assigned a TIN is the parent firm, we first aggregate the KCCA data to the parent-firm level. All firms in the URA's records have a documented TIN that was assigned to them upon registration with the URA. However, only 40% of trade license records in the KCCA data come with a reported TIN. This reflects daily practice deviating from what is agreed under TREP. As the KCCA internally continues to use its own firm identifier, it does not universally demand and record a TIN when a firm pays for a trade license. The TREP mandate is therefore not always enforced (see Section 2.2).

Restricting the analysis to firms with TINs readily documented in both authorities' records would hence exclude active firms without a reported TIN in the KCCA data. As part of this gap may reflect incomplete record keeping—rather than a firm never having registered for a TIN—we attempt to retrieve the TINs of TIN-less firms from the universal TIN registry.

Together with our partners at URA and KCCA, we use a similarity-based matching procedure and contact information from KCCA records (name, phone number, and email address) to recover missing TINs via the TIN registry. All identifiers are first standardized prior to matching. We then impose conservative matching requirements: candidate matches must match exactly on one uniquely identified contact detail (phone number or email) and match closely on business name.[10] In this way, we are able to recover TINs for a third (33%) of firms that were missing this information in KCCA records. The remaining firms either never registered for a TIN or registered under different details than what is recorded in the KCCA data. We therefore cannot link them across the records of URA and KCCA and exclude them from our analysis.

In total, our cross-authority dataset comprises 155,260 unique firms (see also Appendix C.3). Each of these firms is identifiable across both authorities' records via its TIN. Whether a firm actually appears in just one or both authorities' records depends on which of them it has ever filed or paid tax to. Figure 1 visualizes the linking process and resulting cross-record dataset. Appendix C.2 provides the full technical detail of the TIN-recovery procedure, and Table C2 provides illustrative examples. Appendix C.4 conducts a range of sensitivity checks.

[10]Candidates that meet a similarity threshold of 80% but lack an exact contact match are classified as low-confidence and excluded, as manual inspection revealed a high risk of false linkage in this group. This is because two businesses or individuals can have similar names even if they refer to different entities. For phone numbers and emails, the risk of duplication is much less severe. Requiring names to be similar *and* at least one of phone number or email to match exactly, therefore results in confident matches.

Figure 1: Cross-record linking of businesses in Kampala

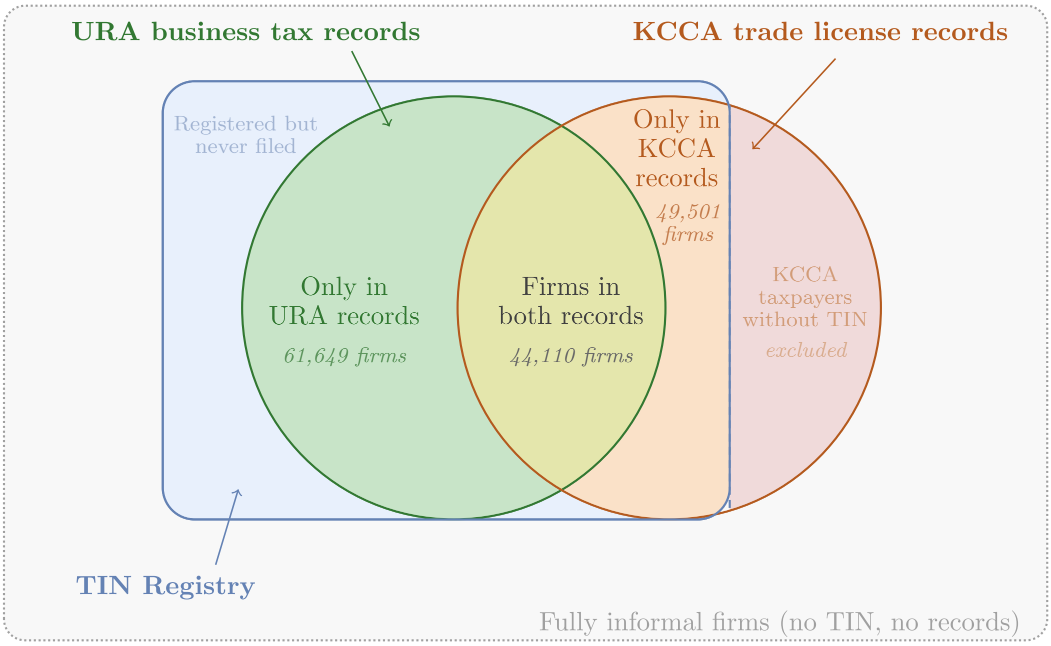

Notes: This figure visualizes the composition of our cross-linked dataset of tax records. Our sample comprises 155,260 firms that are either in URA's tax records, in KCCA's tax records, or in both. URA business tax records also include a small group of firms (2% of the total sample), only filing VAT or PAYE. The sample excludes firms without a Taxpayer Identification Number (TIN) in the universal TIN registry, who are therefore not identifiable across records. This concerns firms in KCCA's records without a documented TIN (either directly in KCCA's records or recoverable via the TIN registry), as well as fully informal firms, which we do not observe at all in any dataset. We also exclude firms that are registered for a TIN but never file any tax returns with any tax authority (and therefore do not appear in any authority's tax records).

3.1.2 Harmonizing filing periods

To link records at the firm-year level across authorities, it is also necessary to harmonize filing periods. In the administrative tax records, KCCA and URA follow different reporting schedules. We harmonize records to align with the URA's standard CIT/PT reporting period, which is by fiscal year (FY), running from July 1 to June 30.[11] As KCCA

[11]For example, FY 2023 would denote the period between July 1, 2022, and June 30, 2023. About 6% of firms in the URA data follow a different accounting period from the standard July 1 to June 30 reporting period. We assign these firms to the July 1 to June 30 fiscal year, covering the majority of the return period using a majority-months rule. For instance, a firm may follow the British fiscal year, running from April 1 to March 31. In this case, we would assign a return filed for April 1, 2021, to March 31, 2022, to FY 2022 because nine of the twelve months fall within (standard) fiscal year 2022. This coding scheme follows the URA's own practice in its records.

licenses are due in advance for the next twelve months and do not follow a unified renewal date, we assign trade license payments to the fiscal year they are mainly valid for.[12] Additionally, we make use of Value Added Tax (VAT) and Pay-As-You-Earn (PAYE) records filed to URA, both of which are due monthly for liable businesses. We assign these records to the fiscal year they belong to.[13] In this paper, we use VAT and PAYE records solely as additional information to verify a firm's activity status, as our focus is on business income taxes. Appendix C.5 provides additional examples illustrating how we harmonize records across different reporting periods.

3.1.3 Tax (non-)compliance

For our main analysis, we define tax (non-)compliance on the extensive margin and focus on national (CIT/PT) and municipal business taxes (trade licenses), for which firms in our sample are universally liable. A firm is compliant with national business income tax requirements at the extensive margin if it files corporate income tax (CIT) or presumptive tax (PT) to the URA for a given year. Similarly, a firm is considered compliant with municipal business taxes if it holds at least one active trade license from KCCA for a given year. This captures basic engagement with both tax authorities, rather than tax payments at the intensive margin.

In turn, we consider a firm non-compliant if it is verifiably active but does not file national business tax with the URA or pay for an active trade license with the KCCA. We determine whether a firm was verifiably active in a given fiscal year by cross-checking between the different tax records we harmonize: corporate income tax (CIT) returns, presumptive tax (PT) returns, VAT returns, PAYE returns (all to URA), and trade license purchases (from KCCA). For example, we can identify firms that held an active trade license with the KCCA but did not file CIT or PT to the URA for the same year, or that filed any tax to the URA in a given year without holding an active trade license with the KCCA. This approach allows us to discern non-compliant firms from those that may have been merely inactive during a given fiscal year, thereby addressing a key empirical challenge in detecting tax evasion in developing countries where rates of firm death are high (McKenzie and Paffhausen, 2019), and firms often do not formally de-register (Mascagni et al., 2022). In addition to extensive margin compliance, our experimental analysis investigates intensive margin compliance by considering the URA

[12]For example, a trade license purchased or renewed on May 1, 2021, is valid from May 1, 2021, until April 30, 2022. It would therefore be assigned to FY 2022, as ten of its twelve months fall within the July 1, 2021, to June 30, 2022 period. This is consistent with how we and the URA treat firms following non-standard reporting periods in the URA data (see footnote 11).

[13]For example, a VAT return filed for June 2022 would fall into FY 2022 while a return filed for July 2022 would fall into FY 2023.

tax liability and the amount paid to the KCCA standardized using the mean amount of 2019.

3.1.4 Business size

A key dimension of heterogeneity for our analysis is business size. Given the absence of a unified definition of business size across the records of URA and KCCA, we use two proxies.

First, we distinguish smaller from larger businesses according to the type of income tax a business is filing for at the URA: presumptive tax (PT), which is applicable to businesses with a turnover below UGX 150 million (smaller businesses), and corporate income tax (CIT), applicable to businesses with a turnover above UGX 150 million (larger businesses). In consultation with our local partners, we classify businesses that *never* filed PT or CIT as smaller firms.[14] This choice is substantiated by evidence from Table A1 which indicates that this group is, if anything, comprised of even smaller firms than the PT-paying group according to two proxies of business size in the KCCA data: the likelihood to be classified as Grades 2–4 and the firm's trade license cost.[15]

Second, the TIN registry contains information on the type of entity that was registered, distinguishing between non-individuals (predominantly corporations) and individuals (predominantly unincorporated businesses in the name of the entrepreneur).[16] We take corporations to typically represent larger businesses.

3.1.5 Robustness to TIN retrieval errors

Because our main sample includes firms whose TINs we had to recover from the TIN registry (see Section 3.1.1), our results could be sensitive to errors in TIN-retrieval. We

[14]Businesses in our dataset who never filed national business income tax are those only found in KCCA's records. Because trade license rates are lump-sum fees determined by industry and location, the KCCA does not collect information on business turnover.

[15]The KCCA assigns Grade 1 to businesses who operate in prime, "high-value" business locations of Kampala. Conversely, Grades 2–4 are assigned to firms in more peripheral business areas. Within a Grade (i.e., within similarly economically central locations), the KCCA charges higher trade license fees to businesses who deal in valuable goods or services with generally higher turnover and higher amounts of capital (such as banks, hotels, or supermarkets) and lower fees to businesses who are typically of smaller nature (such as beauty salons, mobile money agents, or service repair workshops). Firms that never filed business income tax to URA are more than twice as likely (14 percentage points likelier) to be of Grades 2–4 than firms that file corporate income tax (those that we know to be larger firms). Conditional on Grade (that is, conditional on economic zone within Kampala), the trade licenses of firms that never filed business income tax to URA are, on average, 42% cheaper (0.354 log points) than those of CIT-filing firms. This strongly supports the notion that firms that never filed business income tax to the URA are, on average, of smaller size.

[16]The TIN registry has some missing values on this variable. Following instructions from staff at the URA data laboratory, these firms are unincorporated businesses that registered more recently.

address this concern with a range of robustness and sensitivity checks (see Appendix C.4 for details).

First, there could be cases where we retrieve an incorrect TIN from the TIN registry. Our criteria for TIN retrieval are deliberately conservative to mitigate this issue. Still, we replicate our results excluding firms whose TINs we had to retrieve ourselves (Appendix C.4.1).

Second, there could also be cases where we fail to recover the TIN of a firm even though the firm in reality has a TIN (false negatives in TIN-retrieval).[17] We address this concern by replicating our analysis using only firms found in both authorities' records (the intersection in Figure 1). This eliminates the potential for false negatives in TIN-retrieval (Appendix C.4.2).

Finally, we conduct a sensitivity exercise that bounds the effects of TIN-retrieval errors on our main estimates of cross-record complementarity (Appendix C.4.3). As such, it estimates the effect of closing all TIN-documentation gaps in KCCA's records on tax non-compliance detection via cross-record linkage.

3.2 Survey data and tracking exercise

We complement the cross-linked tax records data with survey data that we collected from 1,923 firms in November 2024. We sampled firms from the KCCA's records, which contain location data at the street or village level and contact information to help track these businesses down. We limited the sample to firms that had ever held an active KCCA trade license since 2020 to maintain fieldwork efficiency in light of high rates of firm closure. Furthermore, we focused on firms with an annual turnover between UGX 20 million (approximately USD 5,600) and UGX 300 million (USD 84,000), i.e., small- to medium-sized businesses. Appendix B and Table B1 contain more details on the survey sample.

The survey serves two purposes. First, it facilitates a physical tracking exercise in which we simulate the extent to which recent activity data from cross-record linkage improves the efficiency of in-person enforcement action (such as via an audit). We compare successful tracking rates for a sample of 1,270 URA non-filers between businesses who recently had renewed their trade license (since January 2022 or later) to those whose last renewal was slightly less recent (since January 2020). Given tax authorities' limited enforcement capacity, this provides us with an estimate of the extent to which recent activity data

[17] In other words, we would conclude the firm is in the set of KCCA taxpayers without a TIN in Figure 1, whereas it is actually in both records—just without being linked.

from cross-record linkage allows to target audits to businesses that an auditor would be more likely to successfully track down.[18] The same tracking exercise also allows us to verify the activity status of a sub-sample of businesses that did not file any tax during the 2023 fiscal year. This provides insights into the extent of fully non-compliant businesses whose activity status we could otherwise not verify in tax records.

Second, the survey data provide descriptive insights into the business tax climate in Kampala and relevant background information on the tax nudge experiment.

3.3 Audit and inspection data

To explore bureaucratic fragmentation and systematic gaps in enforcement, we obtain access to business income tax audit data (URA) and trade license inspection data (KCCA).

For URA, we analyze data on the universe of business income tax audits conducted between 2015 and 2024. The data contain information on all assigned audits, whether and when they were conducted and completed, and the audit type.[19] We link this data to the TIN registry to provide a window into the URA's enforcement efforts. Specifically, it allows us to determine the share of taxpayers who were ever assigned an audit[20] and the share of assigned audits that were conducted, as well as differences in these figures by firm size.

For KCCA, we gain access to data on all trade license inspections conducted in 2025. KCCA conducts geographically targeted inspections by assigning inspectors to parishes

[18]Our main specification compares tracking rates between firms with a TL renewal since FY23 (i.e., since January 2022) and firms with a less recent TL renewal (January 2020 – December 2021). We also compare tracking rates between firms with a TL renewal since missing the FY23 filing deadline (i.e., since January 2024) and firms who last renewed between January 2020 and December 2023). This setup will likely understate the extent to which cross-checking (recent) KCCA records improves tracking rates due to the nature of the sample we can conduct the tracking exercise on. To maintain a minimum of logistical efficiency for the survey fieldwork and because we only had access to reasonably granular location data via KCCA's records, we are limited to tracking a sample of firms that held an active trade license at some point between FY 2021 and FY 2023 (see Appendix B). Therefore, what we estimate is the effect of a firm having *recently* renewed its trade license compared to firms that also had a trade license in the past but less recently renewed it. To the extent that even somewhat dated KCCA information improves tracking rates, our estimates will provide a lower bound of the value of KCCA records for tracking rates.

[19]The URA distinguishes between two types of audits: issue audits that are triggered when a specific inconsistency is found and flagged for deeper investigation, and comprehensive audits that are more extensive and typically encompass all returns filed by a firm across different tax types. In addition, the URA conducts 'desk audits', which are considered compliance advisory rather than audits and are not part of the audit data accessed by us. For a more complete discussion of the URA audit system, see Henning and Okello Ayo (2025).

[20]We only consider audits of business income taxes, that is, of firms' corporate or presumptive income tax returns. Income taxes are the second most commonly audited tax type by the URA (after VAT) and make up 20% of all assigned audits.

within Kampala for a set time (e.g., several months). Inspectors then roam their assigned parish, covering a different part every day and check all businesses they encounter. Every interaction—irrespective of whether a business was found to have a valid trade license or not—is logged in a mobile application and assigned a transaction type which described the action taken as a result of the inspection. Where a business is found to have no valid trade license, it is asked to pay for a renewed license on the spot. Businesses who are found without a trade license and who do not pay for renewal on the spot are sealed (closed) by the inspector. Table A5 provides an overview of all inspections conducted by the KCCA in 2025. We limit our main analysis to firms included in our cross-linked dataset of tax records and focus on two outcomes: whether a firm received any inspection and whether a firm received enforcement in the form of (temporary) business closure. As with URA audits, we compare KCCA inspection and enforcement rates by firm size.

4 Tax nudge experiment

The URA and KCCA communicate with taxpayers through text messages, but neither authority operated a centralized or standardized reminder system. URA messaging is delegated to local tax offices, which have discretion over content and timing. Similarly, the KCCA can send trade license renewal reminders to those with expiring licenses, but this is not automated and depends on local office capacity. Crucially, firms that have not previously filed with an authority—and are therefore unknown to it—would not receive any routine communications.

To align with these communication streams, we co-designed a centralized tax message intervention with both tax authorities, but based on the cross-linked tax records. Our intervention differs from baseline communication in three ways: (i) it was centrally administered to a large experimental sample, which also included firms not previously targeted by the authority in question; (ii) it was coordinated between both tax authorities, including a message jointly signed by both; (iii) it used explicit deterrence language threatening enforcement consequences. Following consultations with both tax authorities, each message was sent twice, on November 30 and December 14, 2023, about one month before the URA filing deadline. By this time, about 75% of firms had not yet filed CIT or PT, and approximately two-thirds of trade licenses were due for renewal between January and March.

The experiment is designed to test the effects of different messages on URA and KCCA tax filing, as well as potential spillover effects between tax layers. Firms were randomly assigned to one of four treatment arms, in addition to a pure control group. The content of the messages followed the most effective message in Cohen (2024), adapted to be shorter and consistent with each authority's mandate. Following KCCA practice, their messages were sent in both English and Luganda, whereas URA messages were sent in English only.

Table 1 summarizes the treatment combinations. Treatments that do not signal any coordination are those in which firms receive a message from only one tax tier. In the 'Dual message' treatment, firms receive separate messages from both authorities, thereby cueing them that they are visible to both authorities without explicitly signaling coordination. Firms in the 'Coordination' treatment arm receive a jointly signed message on behalf of both tax authorities, thus making coordination explicit.

Table 1: Treatment Arms and Text Messages by Sender

Treatment	**URA**	**KCCA**
Control	No message	No message
Deterrence URA	Dear esteemed client, file your income tax return and pay the tax due to avoid payment of interest, penalties, and possible closure of business. URA	No message
Deterrence KCCA	No message	Dear esteemed client, renew your trade license to avoid enforcement action and/or prosecution. KCCA
Dual message (separate)	Dear esteemed client, file your income tax return and pay the tax due to avoid payment of interest, penalties, and possible closure of business. URA	Dear esteemed client, renew your trade license to avoid enforcement action and/or prosecution. KCCA
Coordination (joint)	Dear esteemed client, file your income tax return and renew your trade licenses to avoid penalties and enforcement that can lead to business closure. URA & KCCA	

Notes: Messages were sent on November 30th and again on December 14th 2023. Sample size of the Control group is 12,079, Deterrence URA is 12,044, Deterrence KCCA is 12,082, Dual message (KCCA+URA) is 12,054, and Coordination is 12,058. Half of the Coordination messages were sent by the KCCA and half by the URA, reversing the order of their respective signatures.

Randomization was conducted at the firm level, ensuring that a firm with multiple trade licenses or branches received only one message to avoid treatment contamination.

The intervention sample includes all businesses verifiably active between 2020 and 2022, comprising 60,317 active firms in the fiscal year 2022. We stratify randomization by the tax scheme to which businesses pay (or would pay) at the national level, i.e., CIT or PT. This implicitly stratifies randomization by turnover size. Each treatment arm included approximately 12,000 firms. Half of the coordination messages were sent via KCCA and half via URA.

The KCCA bulk message report indicates that only 120 messages failed to deliver. While we cannot verify whether firms read the messages, from an earlier survey we learned that about 78% of respondents read URA messages and 63% KCCA messages.[21] Our analysis is therefore conducted on an intent-to-treat basis.

We also match our cross-linked tax records with a similar deterrence nudge intervention that URA sent to 7,502 presumptive tax payers in 2019 (Cohen, 2024) who we can identify as active firms in our data.[22] This allows us to investigate effects over time, sending messages at a different time of the year (June vs. November/December), as well as the effect of repeated nudge exposure (possible for 2,358 firms included in both interventions).

4.1 Empirical strategy

As in Holz et al. (2023), we analyze the impacts of our deterrence message experiment using a difference-in-differences approach. In the main analysis, we estimate the average treatment effect of our intervention on tax filing behavior. Specifically, we estimate the following model:

$$y_{i,t} = \alpha + \sum_{j=1}^{5} \beta_j \, T_{ij} \times \mathrm{FY23}_t + \sum_{j=1}^{5} \gamma_j \, T_{ij} + \delta \, \mathrm{FY23}_t + \varepsilon_{i,t} \tag{1}$$

where $y_{i,t}$ refers to whether firm i in year t filed national business income tax or obtained a municipal trade license, respectively. T_{ij} refers to the treatment j assigned to firm i, distinguishing between the control and the treatment conditions. The binary indicator $FY23_t$ takes the value zero when the outcome is measured in the fiscal year 2022, the year before the intervention, and one for the fiscal year 2023. The coefficient β_j measures the interaction between treatment status and the post-intervention indicator, capturing

[21] This survey was conducted in 2023 with similar sample selection criteria to the main survey used in this study. The lower share reported for KCCA messages could be explained by the lower frequency of messages by the KCCA, and, more generally, the fact that only firms that already have a license receive messages when it is close to expiry.

[22] The data come from a subset of firms that are observably similar to the sample in our experiment, operating before 2019, and that were part of the experiment described in Cohen (2024). Randomization in the 2023 experiment was orthogonal to treatment in 2019.

the causal impact of the treatment on changes in tax compliance relative to the control group.

We first consider the effects on URA tax declarations and whether the firm filed a tax return (CIT or PT). As an extension, we analyze the effects on the firm's calculated tax liability, as an attempt to capture compliance at the intensive margin. Subsequently, we employ the same estimation strategy but consider the effects on having an active KCCA trade license to investigate the effects at a different tax tier. To estimate whether a message sent regarding the business tax of one authority affects compliance with the business tax of the other authority (i.e., spillovers across tiers), we estimate the effect of URA compliance messages on KCCA compliance and vice versa.

We then test the effects of two messages that raise the information signal regarding cross-authority collaboration. First, we test the effect of receiving separate messages from both authorities. This indicates to firms that they are visible to both authorities without explicitly signaling coordination. Second, we test the effect of the jointly signed message by URA and KCCA that makes coordination salient. Lastly, we estimate the effects of tax nudges over time using the URA intervention of 2019. We estimate equation (1) but using the fiscal year 2018 as a baseline and estimate the treatment effects separately for each year from 2019 until 2023.

We use OLS models for all outcome variables by default and test the robustness of binary indicators using logit models, with marginal effects evaluated at the means. We cluster standard errors at the parent-firm level, the level of randomization.

5 Results

5.1 Fragmentation and detecting tax evasion

5.1.1 Non-compliance and tax record complementarity

We start by presenting novel estimates of the extent of tax non-compliance among businesses in the urban center of a low-income economy. As described in section 3.1, our dataset links firms across the records of the municipal (KCCA) and national (URA) tax authority—and across different taxes filed to URA—allowing us to identify businesses who filed any tax[23] to any of the two authorities for any year between 2018 and 2023 (Table A2). Among this set of verifiably active businesses, we identify those that failed to file national business tax returns to URA (CIT or PT) or that did not hold an active

[23] Presumptive tax (PT), corporate income tax (CIT), Value-Added Tax (VAT), Pay-As-You-Earn Income Tax (PAYE), or trade license (TL) fees

Table 2: Tax non-compliance among verifiably active firms

	URA (CIT/PT)	KCCA (TL)	Either (CIT/PT or TL)	N
Panel A: Non-compliance rates				
Average (2018–2023)	0.431	0.488	0.856	432,951
	(0.001)	(0.001)	(0.001)	
Ever (2018–2023)	0.584	0.563	0.979	155,260
	(0.001)	(0.001)	(0.000)	
	URA	KCCA	N_{URA}	N_{KCCA}
Panel B: Complementarity of records				
Never filed/paid	0.319	0.397	155,260	155,260
	(0.001)	(0.001)		
Never filed/paid (non-compliers)	0.534	0.771	186,727	211,114
	(0.001)	(0.001)		
Never filed/paid previously (non-compliers)	0.657	0.891	186,727	211,114
	(0.001)	(0.001)		

Note: Panel A shows the extensive margin tax non-compliance rate for the sample of verifiably active businesses between 2018 and 2023, separately for URA national business tax (CIT or PT), KCCA municipal tax (TL), and for non-compliance with either of the two taxes. Panel B shows the share of businesses who never filed to one of the authorities. The first row of panel B shows the complementarity between national and municipal records; that is, the share of total businesses that never occur in URA's filing records or KCCA's payment records, respectively. These firms are known taxpayers of one but not the other authority. The second row of panel B moves to the firm-year level. It shows what share of businesses who do not comply with a given authority's taxes never filed or paid to this authority. They are therefore only identifiable through the respective other authority's records. The third row of panel B shows the same share but only considers whether a firm had occurred in URA's (KCCA's) records up until the year in question. Standard errors in parentheses.

municipal trade license with KCCA.

Even among verifiably active firms, tax non-compliance is the norm rather than the exception (Table 2, Panel A). On average, 86% of firms do not pay national business tax to URA or do not have a valid trade license in a given year. Non-compliance is pervasive for both national and municipal taxes. Every other active firm (49%) does not have a valid KCCA trade license, whereas a similar share of active firms (43%) do not file national business income tax to URA. Between 2018 and 2023, 58% of firms ever failed to file business tax to the URA, and 56% of firms ever operated without a valid trade license. Among the 155,260 firms in our dataset, nearly every firm (98%) failed to comply at least once with its tax obligations.[24]

A firm that does not file or pay taxes to one authority has a high chance of being a serial offender. Almost three in four firms (72%) never occur in one of the authorities' records. Any given firm, therefore, tends to file with one authority and one authority only

[24]Even when focusing on firms that are already in both authorities' records and excluding the first year a firm starts filing any tax (allowing for some discretion while a business is starting to operate), we still estimate that 90% of firms ever miss a tax declaration or payment.

over the six-year time horizon. In a random effects logistic regression model, between-firm variation explains 87% of URA non-compliance over time and 93% of KCCA non-compliance, capturing the extent to which firms are persistent in either filing or not filing to a specific authority (see Table A3).

This makes URA's and KCCA's records highly complementary for catching tax non-compliers (Table 2, Panel B). About a third of firms (32%) never file taxes to the URA and 40% never obtain a trade license with KCCA. Cross-record linkage, therefore, dramatically expands the number of firms that any tax authority could detect if it were to solely consult its own records. Among all instances of business tax non-compliance to the URA that we detect, 53% involve firms that never occur in URA's filing records. Non-compliance detection increases even more under cross-record linkage for the KCCA, for whom three in four non-compliers (77%) never occur in its own records.

Put differently, cross-record linkage doubles the number of non-compliers the URA can detect and quadruples it for the KCCA. Taken together, we identify over 260,000 *additional* cases of tax non-compliance that could not be detected when relying on a single authority's records. This represents almost a tripling in the total number of detectable tax non-compliance cases in Kampala between 2018 and 2023.[25] Appendix C.4 conducts a range of sensitivity checks for these estimates, confirming their robustness.[26]

The value of cross-linked records is not limited to detecting additional cases of tax evasion but also has implications for a tax authority's ability to enforce on them. A single authority's records provide only very limited information on whether a firm was active or not.[27] Yet, the ability to discern active but non-compliant firms from inactive firms is essential in a context with limited enforcement resources: only if following up with non-*filers* credibly leads to the detection of non-*compliers* may it be worthwhile for a tax authority to allocate resources to the leads generated by cross-record linkage.

[25]From 135,276 cases of tax non-compliance (firms in records of the authority whose tax they fail to comply with) to 397,814 (total number of detected cases of non-compliance between 2018 and 2023). Records overlap even less from a real-time enforcement perspective in which an authority only consults its historical records—as opposed to reverse engineers the non-compliance history of a firm upon appearing in its records for the first time (Table 2, final row).

[26]Specifically, we probe the robustness of our results to incorrect links between two records (Section C.4.1) and overlooked links (Section C.4.2). Neither test meaningfully changes our conclusions. In a sensitivity exercise (Section C.4.3), we estimate that cross-record linkage still facilitates a 2.8-fold increase in the number of detectable non-compliance cases after accounting for possible linking errors. Only those cases where the firm in question *had also filed or paid tax* result in a waste of audit resources. We estimate that closing gaps in KCCA's TIN recording would result in 5.6% fewer non-compliance leads for KCCA's enforcement— cases that would currently result in a waste of resources. URA, on the other hand, would stand to gain at least 19% *more* leads if KCCA closed its gaps in TIN-documentation.

[27]The exception being cross-linkage with other taxes the URA collects. This captures a meaningful but comparatively small share of active but non-compliant businesses (Table A2).

Table 3: Success rate for tracking down non-filing businesses, by trade license activity

	Firm successfully tracked down	
	(1)	(2)
Renewed TL since 01/2022	0.124**	
	(0.049)	
Renewed TL since 01/2024		0.113***
		(0.026)
Constant	0.580***	0.637***
	(0.047)	(0.019)
Observations	1270	1270
Change [%]	21.4	17.7

Notes: The table reports the probability to successfully track down a firm that failed to file national business tax in FY 2023, depending on when a firm last renewed its trade license. The tracking exercise was implemented in November/December 2024, about a year after a firm missed the 2023 filing deadline for the national business tax. Column 1 reports the change in probability if the firm held an active trade license since FY 2023 (January 2022 or later). Column 2 reports the change in probability if the firm had renewed their trade license since missing the FY23 filing deadline (January 2024 or later). In both cases, the comparison group are firms whose last trade license renewal was less recent (January 2020 – December 2021 in column 1; January 2020 – December 2023 for column 2). All estimates from OLS regressions. Robust standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

To illustrate this, we conduct a physical tracking exercise in Kampala in which we attempt to locate a sample of 1,270 businesses that did not file national business tax to URA in FY 2023. Without the benefit of cross-checking KCCA records, these could be tax non-compliers that an audit would catch, *or* they could be businesses that were inactive, closed permanently, or moved locations—resulting in a waste of audit resources. When we try to physically track these firms down one year after they missed the filing deadline, tracking rates are 11.3–12.4 percentage points (18–21%) higher for firms with a recent trade license renewal (Table 3). In other words, we estimate that cross-checking activity against recent KCCA records allows to focus on a set of non-filers for which an audit is more likely to be successfully deployed. The value of reduced fragmentation between authorities' records, therefore, operates through two complementary channels. First, it allows each authority to cast a wider net, detecting previously unknown non-filers. Second, it provides the means to ascertain non-compliers' activity status in independent records, thereby rendering enforcement more efficient.

Estimating non-compliance rates by verifying activity across authorities has clear advantages in a context with high rates of firm turnover. However, it has one obvious drawback: by construction, we miss firms that did not file or pay *any* tax in a given year since we treat them as if they were inactive. In reality, some of these firms may have been operating but simply did not file any taxes to any authority. This would render our estimates of non-compliance lower bounds.

To shed light on the prevalence of this group of firms, we conduct another in-person tracking exercise in Kampala (Table A4). We sample 149 firms from our dataset that meet the criteria for inclusion in our survey (see Section B) and had no sign of activity in 2023. These are firms we would classify as inactive in 2023 and would therefore exclude from our estimates of non-compliance for this year. We then try to track down these firms, find, and speak to 93 of them (62%). Among those we could not find, we were able to confirm, via phone calls and neighboring firms, that 32% of firms closed and 11% had relocated. Over two-thirds of firms we find (70%) agree to an interview in the process of which we ask the firm whether it operated in 2023.[28] In this way, 62 out of 65 interviewed firms (95%), all of which had not filed any tax in 2023, confirm to us that they were operating in 2023. A conservative estimate, treating all firms we did not manage to locate as inactive and excluding firms that refused to be interviewed, therefore suggests that among a sample of firms without any sign of activity in 2023, half (51%) were operating and not complying with any of their tax obligations. This highlights that focusing only on verifiably active firms is a relevant yet conservative restriction for reliably detecting tax evasion among small- to medium-sized firms in a low-income context.

5.1.2 Non-compliance by business size

Size matters for understanding non-compliance patterns. In Table 4, we distinguish between smaller firms and larger firms. We find non-compliance patterns that are consistent with the enforcement threat that firms of different sizes face (see Section 2): smaller firms are much more likely than larger firms to not have filed or paid tax at all in a given year (55% vs. 23%). When a smaller firm files taxes, it tends to pay for a municipal trade license to KCCA (66% non-compliance vs. 84% for national business tax to URA). This pattern is reversed for larger firms, which are more likely to file national business income tax (40% non-compliance) than pay for a municipal trade license (80% non-compliance).[29] Limiting the sample to verifiably active firms—and therefore excluding

[28] Firms were interviewed by enumerators not affiliated with the tax authorities and clearly identifiable (via badges and shirts) as survey staff of the local university. Firms were not informed about our ability, or desire, to link their activity to filing behavior in administrative tax records, and we asked about activity as part of a longer interview on business practices and the economic climate in Kampala. The question was therefore likely perceived as innocuous by most respondents.

[29] Keep in mind that we do not observe fully informal firms not appearing in any of the two authorities' records. Anecdotally, these firms tend to be overwhelmingly of a smaller size. Comparing non-compliance

Table 4: Tax non-compliance by firm size

	All firms			Active firms	
	(1) Filed/paid nothing	(2) URA not filed	(3) KCCA not paid	(4) URA not filed	(5) KCCA not paid
Smaller firms	0.548	0.836	0.656	0.637	0.239
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Larger firms	0.229	0.396	0.804	0.217	0.746
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Observations	763,197	763,197	763,197	432,951	432,951

Note: Non-compliance rates among firms of different size. The table compares larger firms (filing corporate income tax) to smaller firms (filing presumptive tax or never having filed business income tax to URA). Columns 1–3 compare non-compliance rates among all firms, that is, including firms that did not file to URA or pay to KCCA at all in a given year. These firms were either inactive or non-compliant with both taxes. Columns 4–5 limit the sample to firms who were verifiably active, that is, firms who filed/paid at least one of the taxes in a given year. Clustered standard errors at the firm level in parentheses.

all firms that did not file or pay tax at all in a year—confirms these patterns.[30]

5.1.3 Revenue shortfall

We estimate revenue shortfalls for each authority that approximate the magnitude of fiscal costs arising from selective non-compliance. For KCCA, where trade license fees are levied at predetermined flat rates, we estimate counterfactual trade license payments using the firm's trade license costs in years in which it does pay. For firms that never held a trade license, we map them to the trade license schedule using information on their business type from the TIN registry. For URA, on the other hand, liabilities are a more complex function of turnover, deductions, and tax scheme (CIT or PT). We therefore train a machine-learning imputation model that predicts counterfactual liabilities from observable firm characteristics and tax filings in other years.[31] In each case (KCCA and URA shortfall), our interest is in approximating *total aggregate* shortfall well, rather than

estimates between different authorities—but within a firm size stratum (i.e., within rows of Table 4)—is therefore more informative than comparing larger to smaller firms for a given tax.

[30]Non-compliance estimates shrink more strongly for smaller firms compared to larger firms because the sample of active, smaller firms is more strongly selected, excluding a larger share of firms that do not file or pay any tax at all in a year.

[31]For URA, our estimates exclude the largest 2.5% of firms. Even though their tax payments are significant, their liabilities are noisy and difficult to predict from observable features. As these firms are systematically more likely to be compliant on the extensive margin, their exclusion mainly affects what denominator to compare our shortfall estimates to rather than the shortfall estimates themselves (see Appendix D for more detail).

determining any single firm's counterfactual liability with high accuracy. Appendix D describes our approach in detail and confirms its validity in a number of checks.

We estimate that KCCA misses out on approximately as much trade license revenue every year as it already raises in total revenue from trade license payments. Across the 2018 to 2023 period, shortfalls sum to 102% of realized revenues on average. We also estimate large shortfalls of on average 42% of total observed tax liabilities for URA. Shortfalls are lower for URA for two main reasons. First, non-compliance with URA's business taxes is lower, on average, than with KCCA trade licenses in our sample (see Table 2). Second, there is a much greater difference in URA business tax liabilities between small and large firms than there is for KCCA trade license rates, and non-compliers with URA's taxes tend to be smaller firms (see Table 4).

These estimates demonstrate that bureaucratic fragmentation, and the associated high tax non-compliance rates, leave a substantial, unexploited fiscal dividend for both authorities. Cross-authority data linkage makes this dividend visible, quantifiable, and subject to enforcement.

5.2 Fragmentation and Enforcement

5.2.1 Targeting of enforcement efforts

How is it possible for such a large share of businesses to evade taxes? We obtain access to URA's audit records for the universe of assigned income tax audits between 2015 and 2024 and find evidence of large and systematic gaps in tax enforcement (Table 5). First, over the ten-year period, only 25 out of every 1,000 taxpayers were selected for an audit. Conditional on being assigned, only a fraction of audits are eventually conducted and completed (6.1%). This means that for every audit that is conducted, 15 other audits assigned to suspicious cases by the URA are not followed through with.[32] Taken together, about 0.2% of taxpayers registered with the URA ever receive an income tax audit over the ten-year period.[33]

These averages hide that audit assignment and completion are highly skewed toward

[32]The high share of assigned but not completed audits is not driven by ongoing cases. Most audit processes take only a few months. Focusing on audits assigned before 2023 (at least 2 years before the data cutoff) yields a share of 5.3% completed audits among all audits assigned. Rather, these numbers seem to reflect low enforcement capacity. The URA's Central Operations Office only carries out the most intense forms of audits assigned to the highest risk cases. The remaining cases are assigned to local URA offices in line with their capacity or not conducted once capacity is exhausted (Henning and Okello Ayo, 2025)

[33]Once an audit is completed, it is likely to result in a claim against the audited taxpayer (73% of completed audits).

Table 5: URA Audits

	(1) Assigned audits	(2) Completed audits
Panel A: All taxpayers		
Share of taxpayers	0.025	0.086
Share of audits		0.061
Observations (Firms)	4,759,674	119,100
Observations (Audits)		418,279
Panel B: Firm size		
Smaller firms	0.050	0.011
Larger firms	0.306	0.241
Observations (Smaller firms)	177,610	8,897
Observations (Larger firms)	123,598	37,799
Panel C: Taxpayer type		
Unincorporated businesses	0.016	0.016
Corporations	0.185	0.197
Observations (Unincorporated)	4,510,682	73,124
Observations (Corporations)	248,992	45,976

Notes: The table summarizes audit rates across firms based on the universe of conducted audits in Uganda between 2015 and 2024. There are a total of 4,759,674 registered taxpayers in Uganda. Between 2015 and 2024, there were 418,279 completed income tax audits which are spread across 119,100 unique taxpayers (one taxpayer can be assigned multiple audits). Column 2 is conditional on having been assigned an audit.

larger, corporate taxpayers. About three in ten larger firms filing corporate income tax ever are selected for an audit in contrast to just 5% of smaller firms filing presumptive income tax. Similarly, 19% of TINs registered as corporations ever are assigned an audit, whereas this is true for only 1.6% of unincorporated firms registered by individual entrepreneurs. The same skew is also visible in audit completion rates. One in four audits assigned to larger firms is completed (24%) compared with just 1.1% for smaller businesses; and one in five audits assigned to corporations is completed (20%) but only 1.6% of audits assigned to unincorporated firms.[34] Effectively, this means that less than one in a thousand smaller businesses ever receive an audit over a ten-year period.[35]

We next compare URA's auditing strategy to that of the KCCA. Concentrating its enforcement resources entirely on Kampala, and aided by the simplicity of confirming trade license validity, the KCCA is able to cast a wider net with its tax enforcement. In 2025, about one in six firms in KCCA's records received any kind of inspection or enforcement activity. In the same year, the KCCA temporarily closed ("sealed") a total of 3,870 businesses (2.2% of all taxpayers in its records) due to non-compliance with trade license regulations (Table A5). Among firms in our cross-linked dataset, 8.6% were subject to any inspection or enforcement activity by KCCA (Table 6, Panel A). About one in a hundred firms (11% of firms with any inspection or enforcement) was temporarily closed by KCCA in 2025.

KCCA's enforcement focused precisely on those firms that URA audits systematically neglect. Smaller and unincorporated businesses are 78% (4.5pp) and 72% (4.1pp) more likely to be subject to any inspection or enforcement activity than larger businesses and corporations, respectively (Table 6, Panels B and C). As smaller businesses are also more likely to be closed conditional on being subject to any inspection or enforcement activity (column 3), their total probability to receive enforcement action in the form of business closure is more than twice that of larger firms (column 2).[36]

These patterns make one authority's enforcement efforts in principle highly complementary to those of the other: where URA systematically misses small- to medium-sized firms with its audits, KCCA sets the focus of its own enforcement efforts;

[34] Differences in audit assignment and completion between larger and smaller firms also hold conditional on sector, registration year, and location (Kampala vs. rest of Uganda). Based on another proxy of size, mean turnover, we also confirm this relationship holds within firm size strata.

[35] This excludes low-intensity 'desk audits' which the URA considers 'compliance advisory' rather than an audit and which are therefore not part of the data we access (Henning and Okello Ayo, 2025).

[36] This could be because smaller firms are more likely to be unable to pay for their trade license conditional on receiving an inspection or because they are disproportionately targeted by KCCA as they often lack legal representation to contest the KCCA's actions in court (see Section 2). Either way, the data show a large concentration of KCCA's inspection and enforcement actions on exactly the small-to-medium-sized group of firms that URA audits neglect.

Table 6: KCCA Inspections

	(1) Any activity	(2) Business sealed	(3) Business sealed (any inspection)
Panel A: All firms			
Share of firms	0.086	0.009	0.109
Observations	162,852	162,852	13,973
Panel B: Firm size			
Smaller firms	0.102	0.012	0.114
Larger firms	0.057	0.005	0.096
Δ Smaller vs. Larger	.045***	.006***	.018***
Observations (Smaller firms)	104,085	104,085	10,607
Observations (Larger firms)	58,767	58,767	3,366
Panel C: Taxpayer type			
Unincorporated businesses	0.099	0.011	0.114
Corporations	0.058	0.005	0.092
Δ Unincorporated vs. Corporations	.041***	.006***	.022***
Observations (Unincorporated)	110,111	110,111	10,924
Observations (Corporations)	52,741	52,741	3,049

Note: The table summarizes inspection rates across firms based on the universe of conducted inspection activities in Kampala in 2025. The analysis sample is limited to firms in our harmonized dataset. The first two columns are unconditional whereas the third column conditions on any inspection activity having been conducted.

where KCCA covers relatively few larger firms, URA targets the vast majority of audits. Yet, in practice, a lack of coordination and fragmented record-keeping leaves this complementarity entirely unexploited.

5.2.2 Firm responses to enforcement gaps

How do firms exploit the enforcement gaps left by bureaucratic fragmentation and administrative loopholes? Our data allow us to document several strategies. First, firms appear to 'ghost' the URA by choosing partial informality: they stop to file tax to URA but maintain an active trade license with KCCA for years to come. We identify 17,040 firms that previously filed tax returns with URA but have not filed again. Among these firms, all of which we are able to link to KCCA records, a third (32%) still maintained an active trade license three years after vanishing from URA filing records, and a quarter (26%) still did so four years onward (Figure 2). One plausible explanation for why these firms stopped filing to URA relates to their accumulated tax debt. We link all firms to their tax ledgers as of June 2023 and find that over two-thirds of 'ghost firms' (69%) were owing money to the URA. This compares to 46% of other firms.

Figure 2: Ghost firms

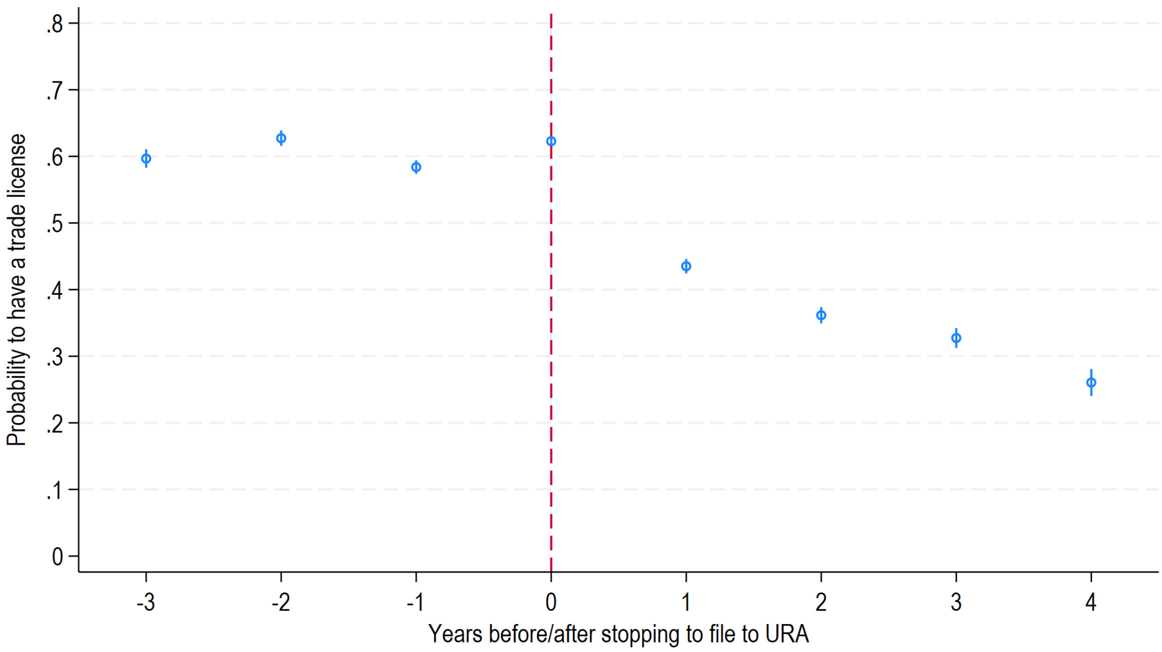

Notes: Probability that a firm has an active KCCA trade license before and after stopping to file tax to URA. The sample comprises of firms that used to file CIT or PT to URA but permanently stopped filing and who at any point had a KCCA trade license. Regression includes financial year fixed effects. Whiskers demarcate 95%-confidence intervals.

Second, we find that firms exploit the lack of a universal documentation requirement upon registration. We identify 749 unique cases in which a business owner stopped filing national business tax for their firm, only to register an observably identical business under a new TIN soon after (Figure A2). In nearly half of these cases (49%), firms obtained a new TIN within just one to three years of their last filing. The speed of re-entry is consistent with owners who abandon a TIN not because their business has genuinely ceased, but because accumulated tax obligations have made continued operation under the existing TIN more costly than starting anew. Because we are limited to identifying only the most glaring cases, this number likely represents a substantial underestimate of the true prevalence of the issue.[37] In fact, in our survey of small- to medium-sized firms

[37]To identify these cases, our team and the URA rely on only very basic means. We flag cases in which a taxpayer with the same phone number and/or email address registered a business in the same industry and with a near identical name after stopping to file taxes for their previous business. This will miss cases, for example, in which a business owner used new contact details or a different firm name when re-registering, or in which the new business was instead registered by a partner or family member. Anecdotally, we have been told that such cases are likely much more common than what either party can capture and that the URA is similarly limited in its ability to identify and prevent them. Our data also

that we conduct in Kampala, business owners estimate that about 30% of peer firms would find it preferable to reopen as a new business rather than paying old tax debt when it is enforced.

Finally, we find that even those firms that are compliant at the extensive margin exploit administrative loopholes to lower their effective tax debt. Trade license renewal is required once every twelve months, but trade license validity is not backdated when renewed late. This incentivizes late renewals by effectively granting firms a 'tax holiday' for each day they renew their license late (Figure A1). Businesses systematically exploit this loophole by renewing their license, on average, half a month (15 days) late. Put differently, for every trade license that is renewed *at all* within a year of expiry, firms obtain on average 15 days of additional trade license validity by gaming the system. This costs KCCA an estimated 3.2% of its total trade license revenue annually.

5.3 Responding to fragmentation: the tax nudge experiment

We now investigate the effectiveness of a light-touch approach to increasing tax compliance and reaping a fiscal coordination dividend in the context of Kampala's fragmented tax administration: merely signaling inter-authority coordination.

We begin with deterrence messages sent by each authority in isolation (Figure 3, Table E2). Across both taxes and authorities, we find robust and precisely estimated null effects. For the message sent through the URA, the point estimate of the effect on URA tax filling is 0.7 percentage points (pp) with a 95% confidence interval (CI) of [–0.4 pp, 1.8 pp]. For KCCA, the effect of having an active trade license is -0.2 pp (95% CI: [–1.4 pp, 0.9 pp]).

Next, we test for spillover effects of the message sent by one authority on compliance with the other authority's tax. Consistent with the descriptive evidence of firms treating both authorities as operating in silos, we find little evidence of cross-tax spillovers. The exception is the effect (1.2 pp) of the URA messages on KCCA trade licenses, which is statistically significant with a 95% CI of [0.0 pp, 2.3 pp]. As this effect is small and not robust, we do not interpret it as meaningful evidence in favor of spillover effects.[38]

We then test two treatment arms, raising the information signal about inter-authority coordination for treated firms. First, a group of firms received separate deterrence

indicate that, in many cases, the newly registered TIN operates in the very same line of business and the same division (location) as the previous one.

[38]The effect loses significance when estimated via a logit model, does not hold when limiting the sample to firms that had a trade license due in a three months period after the intervention, and is zero for chronically trade license non-compliant firms (Table E4). It is also not robust to controlling for the family-wise error rate using the Westfall and Young resampling method (Jones et al., 2019) with an adjusted p-value of p=0.24.

messages from both the URA and KCCA. This cues firms that they are visible to both authorities, but does not explicitly signal coordination between them. Second, we sent one joint message signed by both authorities, which makes coordination explicit. Neither approach increases the likelihood of tax compliance, revealing that merely *signaling* coordination is insufficient to effect behavioral changes among firms and reap the revenue potential of cross-authority coordination.

We confirm the robustness of these results across several dimensions. Intensive-margin tax liabilities show precisely estimated null effects throughout (Appendix E.3). Firms that were chronically non-compliant with URA or KCCA prior to the intervention are similarly unaffected by the nudges (Appendix E.4), consistent with the interpretation that coordination signals are insufficient to shift the behavior of firms that were persistently non-compliant. Heterogeneity by firm size broadly confirms the main results: null effects hold for larger firms across all treatment arms (Appendix E.5). Among smaller firms, we again observe a small, statistically significant effect of URA messages on trade license payments, mirroring the pattern already noted in the main results. Given its limited magnitude, its absence among chronically non-compliant firms, and its lack of robustness in the main specification, we do not interpret this as evidence of a meaningful spillover effect. The results are also robust to alternative specifications, including binary choice models, restricting the sample to firms with licenses issued shortly after the intervention, excluding firms that had already begun filing taxes prior to our intervention, and applying more restrictive TIN-matching conditions (Appendix E.7).

Figure 3: Treatment effects on extensive margin tax compliance

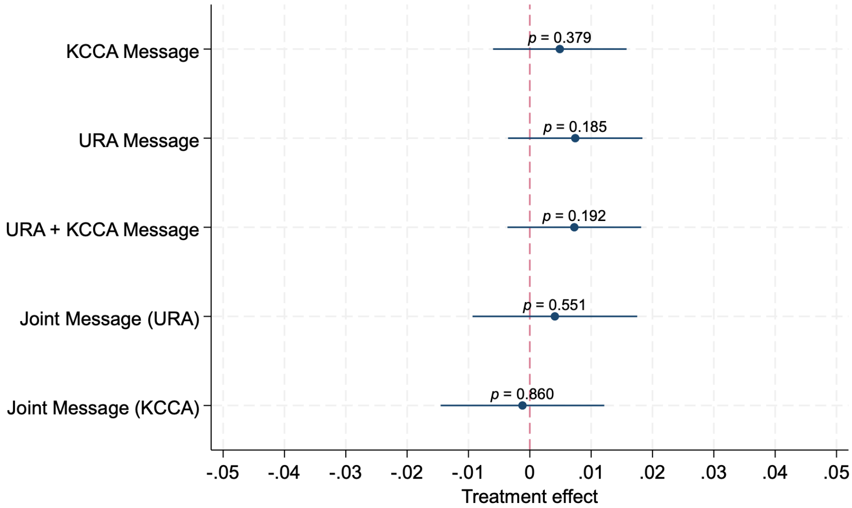

(a) Treatment Effects on URA Tax Filing

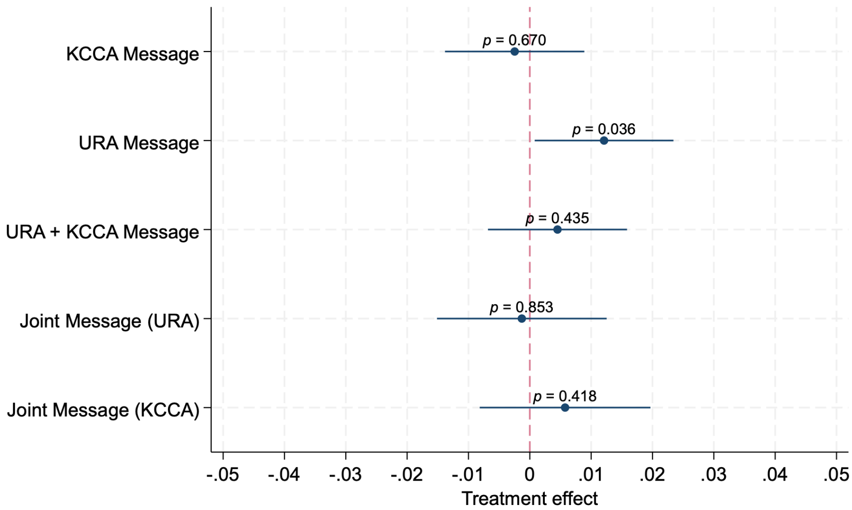

(b) Treatment Effects on Having an Active Trade License

Notes: Figure reports treatment effects from OLS regressions on URA tax filing (a), and of an indicator for having at least one active trade license (b) in 2023 on treatment assignment, following the specification in Equation 1. Estimates include 60,317 firms active in 2022. Standard errors are clustered at the parent firm level, and bars show 95% confidence intervals. For full regression results, see Table E2.

In the final step, we confirm that the null effects we find for tax nudges are not driven by implementation issues specific to our study. Applying the same criteria to identify active firms in 2019, and tracing effects in subsequent years, we investigate the effects of the URA's 2019 tax nudge intervention.[39] Figure 4 shows the effect estimates for every year from 2019 until 2023 on URA business income tax filing. Consistent with the 2023 findings, the 2019 intervention had no direct or indirect impact on tax compliance among Kampala firms in any year (Figure 4 and Figure E8).[40] We find no evidence of repeated exposure effects using a sample of 2,358 firms exposed to the 2019 and 2023 intervention (Figure E7).

Figure 4: Treatment effects of 2019 tax messages over time

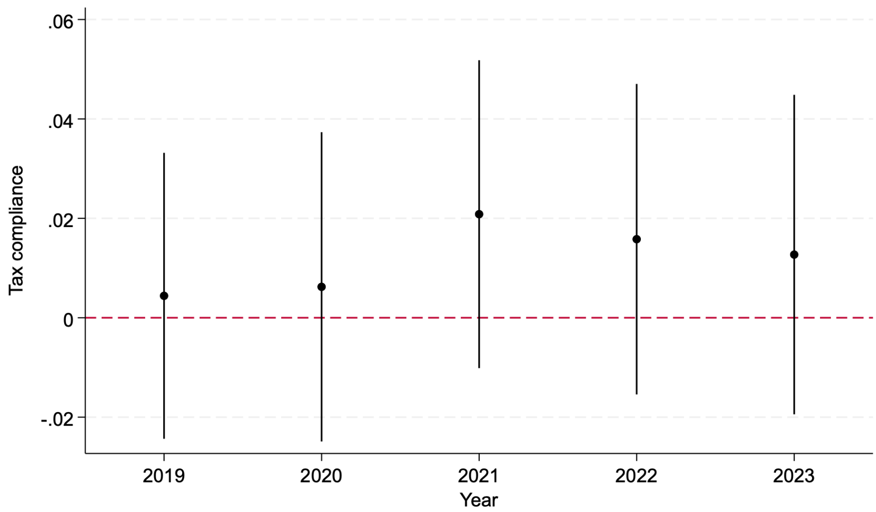

Notes: Figure reports treatment effects from OLS regressions of an indicator for filling URA business income tax for year 2019 through 2023 on treatment assignment, following the specification in Equation 1. The base year for the difference-in-difference estimator is 2018. Estimates are based on treatment assignment in the 2019 intervention on 7,502 active firms in 2019. Standard errors are clustered at the parent-firm level, and bars show 95% confidence intervals. For full regression results, see Table E3.

These results demonstrate the robust null effects of a 'quick fix' approach to addressing bureaucratic fragmentation and raising revenue in an environment with weak tax administration. This includes messages that signal inter-authority collaboration, and the findings hold regardless of which authority sends or signs the message, or which type of business or tax is targeted.

[39]In the estimations, we pool treatment arms into a single treatment variable comprising deterrence, encouragement, and informative messages. All results hold when focusing on the deterrence message only. As the 2019 intervention did not include a treatment arm on inter-authority coordination, we are limited to assessing the effectiveness of the standard variant of deterrence nudges.

[40]Effects are somewhat less precisely estimated which is unsurprising given the reduced sample size (N=7,502) due to the smaller number of Kampala firms in the 2019 intervention.

6 Discussion and conclusion

Low-income countries face immense pressure to improve their fiscal sustainability, yet often lack the administrative capacity to implement policies boosting domestic revenue mobilization. Further, what exists in capacity is often fragmented across multiple government tiers and authorities. We find such bureaucratic fragmentation to be both, an important enabler of low revenue mobilization and a source of significant revenue potential.

By partnering with two distinct government tiers, we offer a unique window into the tax enforcement and compliance environment in fragmented tax administrations. We integrate the previously siloed tax records of the national and municipal tax authorities in Kampala, Uganda, and conduct a cross-government-tier experiment signaling increased coordination via tax nudges. We find firms to systematically exploit the gaps that limited administrative capacity leaves in any single authority's tax net. This results in pervasive and systematic tax evasion that any authority alone struggles to detect or enforce but that cross-authority collaboration reveals.

Our findings identify a significant source of passive waste in fiscal policy, albeit in the way revenue is *collected*, not in the way it is spent (Bandiera et al., 2009). Our experimental results stress that, in an environment in which taxpayers have learned to expect and exploit structural enforcement gaps, tax authorities are unlikely to simply nudge such waste away. The bottom line of our paper is therefore that there is substantial revenue potential in making more efficient use of the capacity tax administrations already possess—but that eschewing administrative reforms that increase collaboration in favor of mere signals to taxpayers offers no shortcut.

How can the requisite reforms succeed when previous efforts to increase collaboration have faltered? Prior to our study, data sharing and harmonized record-keeping entailed bureaucratic costs for both authorities, with unclear economic benefits. As in Kampala, government authorities often differ in size, mandate, and enforcement focus and maintain parallel bureaucratic structures. This increases the fixed costs of coordination and reduces the perceived scope for mutual gains.

Our intervention lowers these barriers in three concrete ways. First, we incurred the non-trivial fixed costs of producing a cross-linked registry of 155,000 firms that is sufficiently large to be systemically relevant. Second, we document meaningful economic gains for both authorities. Based solely on data cross-checks, either authority can identify and confirm tens of thousands of cases of verifiable non-compliance every year, often involving firms previously absent from its own records. Our estimates place the associated revenue shortfall at over 40% of URA's current business tax revenues. For KCCA, we estimate it misses out on about the same amount of trade license revenue it currently collects from compliant firms every year. Third, we identify pathways to convert this potential into realized gains, even under constrained capacity. Cross-checking activity records allows enforcement resources to be targeted at recently active non-compliers rather than non-filers who may have closed. This substantially improves the expected return on each audit or inspection deployed. Incorporating even basic compliance checks against the other authority's records into existing enforcement activities would help close each other's gaps with minimal additional resources. Administrative fixes of loopholes we identify, for example, by recording the owner's national ID during registration or by backdating trade license validity, may recover revenue without expanding capacity at all.

Making use of this potential requires administrative reforms and sustained commitment to collaboration practices that involve real political and bureaucratic challenges. These should not be trivialized. However, alternative policies to raise revenue are at least as challenging to implement, and increased domestic revenue mobilization in LICs has no substitute in the long term. It is therefore economically sound and mutually viable to give improved coordination serious consideration in Kampala and similar contexts. Strengthening tax administration in this way may also create the enabling environment for other, ambitious policy interventions to gain traction toward fiscal sustainability in low-income countries. Whether coordination reforms deliver on this potential in practice remains an open question that warrants rigorous evaluation.

Bibliography

Almunia, M., Hjort, J., Knebelmann, J., and Tian, L. (2024). Strategic or confused firms? evidence from "missing" transactions in Uganda. Review of Economics and Statistics, 106(1):256–265. https://doi.org/10.1162/rest_a_01180.

Antinyan, A. and Asatryan, Z. (2025). Nudging for Tax Compliance: A Meta-Analysis. The Economic Journal, 135(668):1033–1068. https://doi.org/10.1093/ej/ueae088.

Balán, P., Bergeron, A., Tourek, G., and Weigel, J. L. (2022). Local Elites as State Capacity: How City Chiefs Use Local Information to Increase Tax Compliance in the Democratic Republic of the Congo. American Economic Review, 112(3):762–797. https://doi.org/10.1257/aer.20201159.

Bandiera, O., Prat, A., and Valletti, T. (2009). Active and Passive Waste in Government Spending: Evidence from a Policy Experiment. American Economic Review, 99(4):1278–1308. https://doi.org/10.1257/aer.99.4.1278.

Basri, M. C., Felix, M., Hanna, R., and Olken, B. A. (2021). Tax Administration versus Tax Rates: Evidence from Corporate Taxation in Indonesia. American Economic Review, 111(12):3827–3871. https://doi.org/10.1257/aer.20201237.

Bergeron, A., Tourek, G., and Weigel, J. L. (2024). The State Capacity Ceiling on Tax Rates: Evidence From Randomized Tax Abatements in the DRC. Econometrica, 92(4):1163–1193. https://doi.org/10.3982/ECTA19959.

Besley, T. and Persson, T. (2014). Why Do Developing Countries Tax So Little? Journal of Economic Perspectives, 28(4):99–120. https://doi.org/10.1257/jep.28.4.99.

Brockmeyer, A., Mascagni, G., Nair, V., Waseem, M., and Almunia, M. (2024). Does the Value-Added Tax Add Value? Lessons Using Administrative Data from a Diverse Set of Countries. Journal of Economic Perspectives, 38(1):107–132. https://doi.org/10.1257/jep.38.1.107.

Brockmeyer, A., Smith, S., Hernandez, M., and Kettle, S. (2019). Casting a Wider Tax Net: Experimental Evidence from Costa Rica. American Economic Journal: Economic Policy, 11(3):55–87. https:doi.org//10.1257/pol.20160589.

Carrillo, P., Donaldson, D., Pomeranz, D., and Singhal, M. (2023). Ghosting the Tax

Authority: Fake Firms and Tax Fraud in Ecuador. American Economic Review: Insights, 5(4):427–444. https://doi.org/10.1257/aeri.20220321.

Carrillo, P., Pomeranz, D., and Singhal, M. (2017). Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement. American Economic Journal: Applied Economics, 9(2):144–164. https://doi.org/10.1257/app.20140495.

Cohen, I. (2024). Technology and the state: Building capacity to tax via text. Journal of Public Economics, 236:105154. https://doi.org/10.1016/j.jpubeco.2024.105154.

Collin, M., Di Maro, V., Evans, D. K., and Manang, F. (2025). Property tax compliance in Tanzania: Can nudges help? Economic Development and Cultural Change, 73(4). https://doi.org/10.1086/734186.

de Neve, J.-E., Imbert, C., Spinnewijn, J., Tsankova, T., and Luts, M. (2021). How to improve tax compliance? Evidence from population-wide experiments in Belgium. Journal of Political Economy, 129(5):1425–1463. https://doi.org/10.1086/713096.

Gil, P., Holz, J., List, J. A., Simon, A., and Zentner, A. (2024). Toward an understanding of tax amnesty take-up: Evidence from a natural field experiment. Journal of Public Economics, 239:105245. https://doi.org/10.1016/j.jpubeco.2024.105245.

Hallsworth, M., List, J. A., Metcalfe, R. D., and Vlaev, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. Journal of Public Economics, 148:14–31. https://doi.org/10.1016/j.jpubeco.2017.02.003.

Henning, D. and Okello Ayo, J. (2025). Tax Audits and Their Distortionary Effects. https://djhenning.github.io/assets/pdf/Henning_JMP.pdf (last accessed on 08 April 2026).

Holz, J. E., List, J. A., Zentner, A., Cardoza, M., and Zentner, J. E. (2023). The 100 million nudge: Increasing tax compliance of firms using a natural field experiment. Journal of Public Economics, 218:104779. https://doi.org/10.1016/j.jpubeco.2022.104779.

Jensen, A., Brockmeyer, A., and Gadenne, L. (2024). Taxation and development. VoxDevLit, 12(1):3. https://voxdev.org/voxdevlit/taxation.

Jensen, A. D. and Weigel, J. L. (2025). No Taxation without Administration:

Bringing the State Back into the Public Finance of Developing Countries. Journal of Economic Literature. https://doi.org/10.1257/jel.20251760.

Jones, D., Molitor, D., and Reif, J. (2019). What do workplace wellness programs do? evidence from the illinois workplace wellness study. The Quarterly Journal of Economics, 134(4):1747–1791. https://doi.org/10.1093/qje/qjz023.

Jouste, M., Agaba, G., and Ssebuyira, S. (2024). The Uganda Revenue Authority presumptive tax data. Technical report, UNU-WIDER. https://doi.org/10.35188/UNU-WIDER/WTN/2024-1.

Jouste, M., Nalukwago, M. I., and Waiswa, R. (2021). Do tax administrative interventions targeted at small businesses improve tax compliance and revenue collection? Evidence from Ugandan administrative tax data. https://doi.org/10.35188/UNU-WIDER/2021/951-8.

KCCA (2026a). Revised Trading License Rates. https://www.kcca.go.ug/trading-licence-rates (last accessed on 8 April 2026).

KCCA (2026b). Trading License: Assessment, Payment, Issuance. www.kcca.go.ug/media/docs/Trading_Licence_FAQs2.pdf (last accessed on 8 April 2026).

KCCA (2026c). Trading License Description. https://ecitie2.kcca.go.ug/portal/revenue-sources-business-licenses/ (last accessed on 8 April 2026).

Khan, A. Q., Khwaja, A. I., and Olken, B. A. (2016). Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors. The Quarterly Journal of Economics, 131(1):219–271. https://doi.org/10.1093/qje/qjv042.

Mascagni, G. (2018). From the Lab to the Field: A Review of Tax Experiments. Journal of Economic Surveys, 32(2):273–301. https://doi.org/10.1111/joes.12201.

Mascagni, G., Dom, R., Santoro, F., and Mukama, D. (2023). The VAT in practice: equity, enforcement, and complexity. International Tax and Public Finance, 30(2):525–563. https://doi.org/10.1007/s10797-022-09743-z.

Mascagni, G. and Nell, C. (2022). Tax compliance in Rwanda: Evidence from a message field experiment. Economic Development and Cultural Change, 70(2):587–623. https://doi.org/10.1086/713929.

Mascagni, G., Santoro, F., Mukama, D., Karangwa, J., and Hakizimana, N. (2022).

Active Ghosts: Nil-filing in Rwanda. World Development, 152:105806. https://doi.org/10.1016/j.worlddev.2021.105806.

McKenzie, D. and Paffhausen, A. L. (2019). Small firm death in developing countries. Review of Economics and Statistics, 101(4):645–657. https://doi.org/10.1162/rest_a_00798.

Moore, M., Prichard, W., and Fjeldstad, O.-H. (2018). Taxing Africa: Coercion, Reform and Development. African Arguments. Zed Books, London, 1st edition.

Naritomi, J. (2019). Consumers as Tax Auditors. American Economic Review, 109(9):3031–3072. https://doi.org/10.1257/aer.20160658.

Okunogbe, O. and Tourek, G. (2024). How can lower-income countries collect more taxes? the role of technology, tax agents, and politics. Journal of Economic Perspectives, 38(1):81–106. https://doi.org/10.1257/jep.38.1.81.

Parliament of Uganda (1997). Income Tax Act. https://ulii.org/en/akn/ug/act/1997/11/eng@2024-12-23.

Pomeranz, D. (2015). No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax. American Economic Review, 105(8):2539–2569. https://doi.org/10.1257/aer.20130393.

Pomeranz, D. and Vila-Belda, J. (2019). Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities. Annual Review of Economics, 11(Volume 11, 2019):755–781. https://doi.org/10.1146/annurev-economics-080218-030312.

Santoro, F. (2024). Income Tax Payers Are Not All the Same: A Behavioral Letter Experiment in Eswatini. Economic Development and Cultural Change, 72(2):771–799. https://doi.org/10.1086/722332.

Slemrod, J. (2019). Tax Compliance and Enforcement. Journal of Economic Literature, 57(4):904–954. https://doi.org/10.1257/jel.20181437.

Vincent, R. C. (2023). Vertical taxing rights and tax compliance norms. Journal of Economic Behavior & Organization, 205:443–467. https://doi.org/10.1016/j.jebo.2022.11.003.

Waiswa, R., Lastunen, J., Wright, G. C., Noble, M., Ayo, J. O., Nalukwago, M. I., Barugahara, T. K., Kavuma, S. N., Arinaitwe, I., Mwesigye, M., et al. (2021). An assessment of presumptive tax in uganda: Evaluating the 2020 reform and four alternative reform scenarios using ugamod, a tax-benefit microsimulation model for uganda. Technical report, WIDER Working Paper 2021/163. https://doi.org/10.35188/UNU-WIDER/2021/103-7.

Weigel, J. (2024). Supermodular Bureaucrats: Evidence from Randomly Assigned Tax Collectors in the DRC. Working Paper. https://static1.squarespace.com/static/59832fbcf9a61eb15deffb3c/t/6750bf3ec970411ba53f89a4/1733345089635/Collectors_Paper_20241022.pdf (last accessed on 10 April 2026).

A Additional results

Table A1: Differences in trade license grade and cost by strata

	(1) Grade 2-4 (dummy)	(2) log TL cost	(3) log TL cost
Presumptive	0.004	-0.331***	-0.326***
	(0.003)	(0.007)	(0.006)
KCCA only	0.140***	-0.472***	-0.354***
	(0.003)	(0.006)	(0.005)
Constant	0.125***	12.598***	12.696***
	(0.002)	(0.005)	(0.005)
p-val, Presumptive vs. KCCA only	0.000	0.000	0.000
Observations	106,919	92,825	92,825
TL grade controls	-	No	Yes

Notes: OLS regressions of a dummy for a firm's trade license being of Grade 2-4 (as opposed to Grade 1; col. 1) and of trade license costs (cols. 2–3) on firm size strata. The strata compared are firms registered for corporate income tax with URA (larger firms, base category), firms registered for presumptive income tax with URA (smaller firms), and those not registered with URA at all (KCCA only). Robust standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

Table A2: Verification of activity status by tax record

	CIT/PT only	TL only	Both	VAT/PAYE only	N
2018	0.370	0.400	0.140	0.089	65,753
	(0.002)	(0.002)	(0.001)	(0.001)	
2019	0.397	0.375	0.153	0.075	69,570
	(0.002)	(0.002)	(0.001)	(0.001)	
2020	0.351	0.441	0.139	0.069	73,412
	(0.002)	(0.002)	(0.001)	(0.001)	
2021	0.438	0.352	0.136	0.075	67,204
	(0.002)	(0.002)	(0.001)	(0.001)	
2022	0.441	0.349	0.137	0.072	77,968
	(0.002)	(0.002)	(0.001)	(0.001)	
2023	0.533	0.300	0.159	0.007	79,044
	(0.002)	(0.002)	(0.001)	(0.000)	
Average (2018–2023)	0.424	0.368	0.144	0.063	432,951
	(0.001)	(0.001)	(0.001)	(0.000)	

Notes: For each year the table shows the share of businesses whose activity we verified through different registries or tax types. It discerns between businesses only filing national business tax (corporate or presumptive income tax), those only paying for an active municipal trade license, those filing CIT/PT *and* paying for an active trade license, and those who did neither but whose activity we verified through VAT and PAYE returns. At the time of access, VAT and PAYE records for 2023 were only available for the first months of the fiscal year which explains the lower share of businesses verified in this way in 2023. We also present averages pooling over the time period. Standard errors in parentheses.

Table A3: Persistence in choice of tax filed

	URA	KCCA
Share between-firm variation	0.867	0.930
	(0.001)	(0.001)
N	405,553	405,553
N (firms)	152,117	152,117

Note: Intra-cluster correlations from a logistic regression model with random effects at the firm level. To aid interpretability, we exclude firms only filing VAT or PAYE from the analysis. Results therefore capture to what extent a firm is either a chronic non-complier or reliable complier to URA (KCCA) in the years it files to any authority. Standard errors in parentheses.

Table A4: Tracking seemingly inactive firms

Firm tracking	
# Firms w/o recorded activity	149
# Firms found	93
# Firms interviewed	65
# Firms confirming activity in FY23	62
Reasons firm not spoken to (N=56) [%]	
Failed to locate firm	57.2
Firm Closed	32.1
Firm Relocated	10.7
Share of firms found	62.4% (93/149)
Share of firms active in 2023 (excl. refusals)	**51.2%** (62/121)

Notes: The table reports on the results of the tracking exercise, which verified the activity status of firms without any activity in the URA's or KCCA's records in 2023. We tried locating a sample of 149 firms that did not pay national business tax in 2023, nor had an active trade license, and that we would therefore exclude from our compliance estimates as 'inactive'. Twenty-eight firms we were able to find and speak to refused to be interviewed.

Table A5: KCCA Inspection and Enforcement Activities (All firms)

	All activities	Firm level (all taxpayers)	Firm level (any inspection)
Reminder	75.87	12.45	80.41
Message	10.31	2.09	13.57
Seal Business	7.60	2.15	11.63
Meeting	4.10	1.05	6.25
Unseal Business	1.32	0.30	1.88
Sensitize	0.62	0.17	0.99
Demand notice	0.16	0.04	0.27
Email	0.01	0.00	0.02
Issue Agency Notice	0.01	0.00	0.02
Any inspection activity		0.16	
Observations	61,985	180,151	37,147

Notes: The table reports all inspection and enforcement activities taken by the KCCA in 2025. One firm can receive multiple actions. The first column shows the frequency of all inspection activities at the activity level. The second column shows the frequency at the firm level for all taxpayers in KCCA's records (up until 2024, firms established later are excluded). The third column shows the frequency at the firm level for all taxpayers receiving any inspection activity. All values in percent.

Table A6: Tax non-compliance by firm size (three strata)

	All firms			Active firms	
	(1) Filed/paid nothing	(2) URA not filed	(3) KCCA not paid	(4) URA not filed	(5) KCCA not paid
Corporate	0.229	0.396	0.804	0.217	0.746
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Presumptive	0.366	0.544	0.667	0.281	0.475
	(0.002)	(0.001)	(0.002)	(0.002)	(0.003)
KCCA only	0.650		0.650		
	(0.001)		(0.001)		
Observations	763,197	763,197	763,197	432,951	432,951

Note: Non-compliance rates among firms of different size. The table compares larger firms (registered for corporate income tax) to smaller firms (registered for presumptive tax or not registered with URA at all) and firms who are not registered with URA at all. Columns 1–3 compare non-compliance rates among all firms, that is, including firms that did not file to URA or pay to KCCA at all. These firms were either inactive or non-compliant with both taxes. Columns 4–5 limit the sample to firms who were verifiably active, that is, firms who filed/paid at least one of the taxes in a given year. Firms who only ever paid to KCCA are, trivially, non-compliant with URA filing requirements in all cases (columns 2 and 4) and held an active trade license when conditioning on verified activity (column 5). Clustered standard errors at the firm level in parentheses.

Figure A1: Delayed trade license renewal

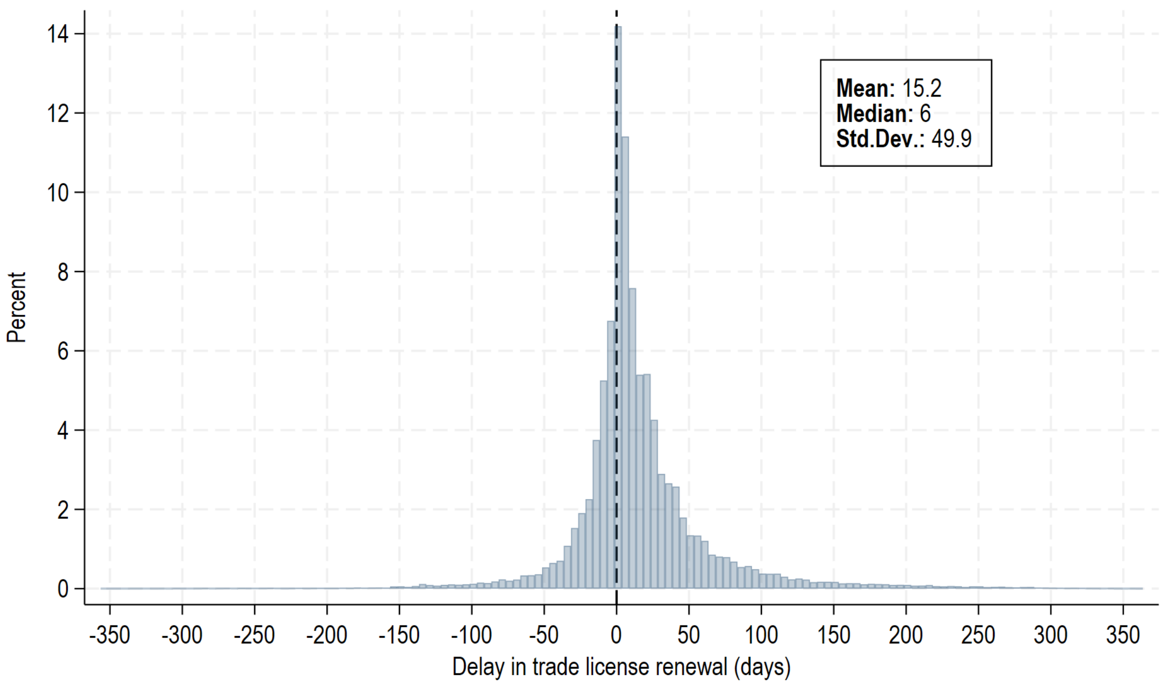

Notes: Days after trade license expiry that trade license was renewed.

Figure A2: Fake Death and New TIN Registration: Descriptive Overview

(a) Years until re-registration with same details

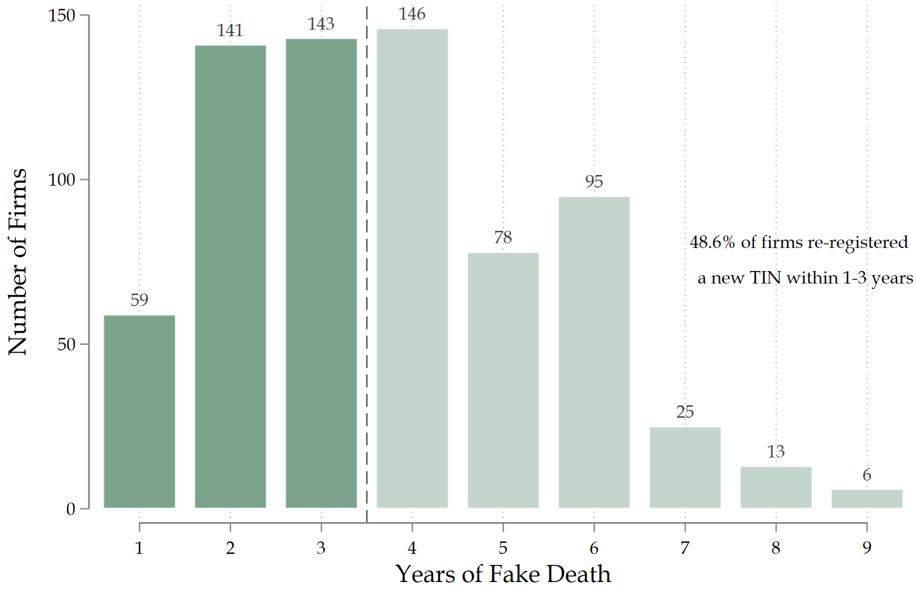

(b) Similarity between Old and New TINs

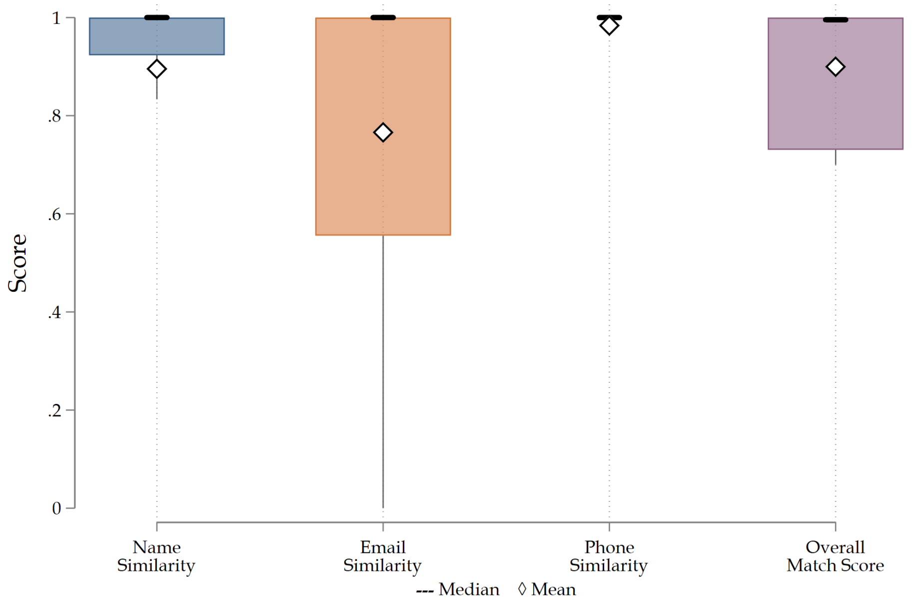

Notes: This figure describes firms identified as having "faked their death"—firms whose owners stopped filing taxes under one TIN and subsequently registered an observably identical business under a new TIN. Firms are matched based on shared phone numbers, email addresses, and near-identical firm names, with at least one exact match required for phone or email. Panel (a) shows the distribution of years between de-registration and re-registration; the dashed line separates firms that re-registered within 1–3 years (dark bars) from those with longer gaps (light bars). 48.6% of these firms registered a new TIN within 1-3 years after they stopped filing taxes under the previous TIN. Panel (b) reports the distribution of similarity scores across matching variables, showing high agreement on names and phone numbers in addition to the exact phone or email match required for inclusion. Among matched firms, 53% re-register in the same sector of activity and 65.2% in the same division (location) within Kampala.

B Survey

We complement our analysis of administrative data with survey data from businesses in Kampala collected in November and December 2024. The survey involved 1,923 firms contained in our administrative dataset with a turnover between UGX 20 million (approximately 5,600 USD) and UGX 300 million (approximately 84,000 USD) and collected information on business characteristics, practices, attitudes, and knowledge of the tax system.

Firms were sampled at random within Kampala from the list of businesses that had paid for a KCCA trade license since 2020 and met the turnover requirement. We obtained turnover information from either recent tax filings of the firm or used information on the grade of the firm and trade license amount as proxies where turnover information was not available. Businesses selected for our study were recruited with the help of local leaders from the village office and asked to schedule a date and time for the interview. In case of non-response or refusal, we recruit a replacement firm from a replacement list. The target respondent of the interview is the business owner or another person affiliated with the business who is knowledgeable about the business's financial and tax affairs.

Table B1: Survey Summary Statistics

	Summary
N	1,922
Larger firm (predominantly CIT)	0.188 (0.391)
Number of employees	4.5 (16.9)
Self-reported sales estimate (thousand UGX)	82.175 (124.146)
Number of business branches	1.3 (3.0)
KCCA business grade (self-report)	
1	1,134 (88.0%)
2	111 (8.6%)
3	31 (2.4%)
4	12 (0.9%)
Trade license amount 2022 (thousand UGX)	299.028 (206.187)
Division	
Central	1,267 (65.9%)
Kawempe	121 (6.3%)
Makindye	184 (9.6%)
Nakawa	224 (11.7%)
Rubaga	126 (6.6%)
Years in operation	9.2 (6.5)
Owner female	0.4 (0.5)
Person that handles tax: owner	0.9 (0.3)
Business industry	
Manufacturing	104 (5.4%)
Construction	53 (2.8%)
Retail	1,014 (52.8%)
Wholesale trade	168 (8.7%)
Transport and storage	47 (2.4%)
Accommodation and food	144 (7.5%)
Information and Communication	33 (1.7%)
Finance and insurance	37 (1.9%)
Real estate	6 (0.3%)
Professional, scientific and technical	260 (13.5%)
Education	29 (1.5%)
Other	27 (1.4%)

Notes: Sales estimates are winsorized at 5%. Standard deviations in parentheses for continuous variables; percentages for categorical variables.

C Cross-Linking of Administrative Tax Data

We collaborate with KCCA and URA to harmonize and integrate their tax records for the entire city of Kampala. Our approach to linking business taxpayers is via their unique taxpayer identification numbers (TINs), the statutory unique firm-level identifier across tax authorities in Kampala. The TIN allows us to consistently track parent firms across both authorities' records and all fiscal years. However, the availability of TINs differs systematically across the tax records of URA and KCCA. All firms in URA's records possess a TIN as registration with URA requires a TIN. On the other hand, TIN documentation in KCCA's records is incomplete. In parts, this may reflect incomplete record-keeping by KCCA rather than a firm not being registered for a TIN. Restricting our dataset to firms with TINs readily documented in KCCA's data may thus exclude active and, in principle, linkable firms.

To address this, we attempt to retrieve the TINs of firms lacking this information in KCCA records by matching their contact details to the universal TIN registry. The following subsections describe how we recover missing TINs, document the complete linkage pipeline, and assess the sensitivity of our results to potential errors.

C.1 Recovering TINs in KCCA records

We begin with 210,652 unique trade licenses from KCCA's records covering the period of 2016 to 2025. Of these, 85,359 licenses (40.5%) have a TIN directly recorded in KCCA's database. The remaining 125,293 licenses (59.5%) have no documented TIN in KCCA's records. These likely include a large share of firms not registered for a TIN, but they may also reflect incomplete records for firms that do have a TIN. The latter group is, in principle, linkable and the focus of our effort to recover missing TINs.

To recover TINs for licenses missing this information, we access the universal TIN registry and apply a similarity-based matching procedure using firm names, phone numbers, and email addresses from KCCA records. We impose conservative matching requirements. Candidate matches must exceed a name similarity threshold of 80% (see C.2) and match exactly on at least one contact identifier (phone number or email). This helps ensure that recovered TINs reflect genuine matches rather than coincidental similarity.

We successfully recovered TINs for 41,026 licenses (32.7% of licenses without recorded TINs), corresponding to 39,040 unique firms. The remaining 84,267 licenses (65.8%) did not match any TIN in the registry—either because the firm never registered for a TIN, or because it registered under different contact details than those recorded by KCCA. These licenses are excluded from our analysis as we cannot link them across the records of KCCA and URA. Our matching criteria are intentionally conservative to prioritize data quality over coverage (see Appendix C.4 for robustness checks on non-compliance estimates).

After restricting to firms verifiably active in fiscal years 2018–2023, the final analysis sample from KCCA records comprises 90,840 unique firms. We restrict to 2018–2023 in line with the period for which there exist complete business income tax records from URA. Of these, 59,639 firms have a TIN recorded directly by KCCA, and 33,972 have TINs we recovered. Table C1 summarizes the complete pipeline.

Table C1: Record Linkage Pipeline: From KCCA Trade Licenses to Identified TINs

Category	# Trade Licenses	# unique TINs	Final Sample[d]
Unique Trade Licenses	210,652	–	–
A. TIN in KCCA records[a]	85,359 [40.5%]	69,548	**59,639**
B. TIN not in KCCA records	125,293 [59.5%]	–	–
B1. TIN recovered[b]	41,026	39,040	**33,972**
B2. No TIN match found[c]	84,267	–	–

Notes: This table documents the record linkage pipeline for KCCA trade licenses in Kampala. Columns (1) and (2) cover the full universe of trade license records available to us (2016–2025). Column (3) restricts to firms verifiably active in fiscal years 2018–2023 to align with our analytical sample (see Table C3). Trade license counts and unique parent firms (identified by TIN) are reported separately. Percentages are calculated relative to the total number of licenses. Licenses with no recorded payment amount are excluded.

[a] A: licenses with a TIN directly recorded by KCCA. Column (2) reports unique parent firms, as one firm may hold multiple licenses.

[b] B1: licenses without a TIN in KCCA for which a TIN was recovered via similarity-based matching on firm name, phone number, and email address against the universal TIN registry.

[c] B2: licenses for which no TIN match was found — these firms cannot be linked to URA records and are excluded from the analysis.

[d] The final sample of KCCA records (93,611) comprises 59,639 firms with a TIN directly recorded by KCCA and 33,972 firms whose TIN was recovered via similarity-based matching. Column (3) restricts to firms verifiably active in fiscal years 2018–2023.

C.2 TIN recovery methodology

Identifier standardization and similarity metrics. All identifiers are first standardized to minimize spurious mismatches. Firm names are converted to uppercase, punctuation is removed, and common legal suffixes (e.g., Ltd, Limited) are harmonized. Phone numbers are normalized by stripping whitespace and non-numeric characters, enforcing a consistent numeric format. Email addresses are converted to lowercase and stripped of leading or trailing spaces. This pre-processing ensures that similarity scores reflect substantive differences rather than formatting inconsistencies, and that the exact match filters applied downstream operate on harmonized strings.

Similarity scores. We then compute similarity scores for firm names and emails using string-distance metrics and generate candidate matches between KCCA records and the TIN registry. To ensure high-confidence linkage, we impose two requirements. First, candidate pairs must share an exact match on at least one criterion (the blocking stage). Given that registration names may be duplicates (e.g., two individuals named John Doe), one contact identifier—i.e., phone or email—must be an exact match beyond the anchor provided by the blocking stage. Second, the weighted composite similarity score across name and email must exceed 80 percent. These requirements substantially reduce the risk of false positives: phone numbers and email addresses are unique identifiers and therefore at lower risk of duplication than names, and no link can be established based on fuzzy similarity alone.

Match classification. Based on similarity scores and contact identifier alignment, we classify candidate matches into four categories:

- *Perfect matches*: Records aligned across all three identifiers (name, phone, email).
- *High-confidence matches*: Firms aligning very closely in names, combined with an exact match on email or phone number.
- *Low-confidence matches*: Firms with similar names but mismatched contact details. Manual inspection revealed a high risk of false matches in this group due to the non-uniqueness of names. This also goes for firms that match on phone or email, but that carried very different names. Such cases may relate to separate businesses of the same owner and would therefore not constitute a valid match.
- *No match*: Firms for which no candidate match satisfying the above criteria could be identified.

We only recover the TINs of perfect and high-confidence matches with the TIN registry.

Low-confidence matches for whom we cannot establish the correct TIN with high certainty, and unmatched firms from KCCA's records without a TIN are not included in the harmonized dataset. Table C2 provides a number of examples for illustration.

Manual verification and validation. In addition to the algorithmic thresholds described above, candidate matches are subject to manual verification to ensure that observed similarities reflect the same underlying economic entity rather than coincidental identifier alignment. Manual checks focus in particular on cases with common firm names, partial identifier overlaps, multiple candidate matches, and instances in which a single TIN is associated with multiple firm names or in which identical contact details appear across multiple firm records. These cases are screened for potential duplicates or registry inconsistencies and resolved conservatively.

A further concern sometimes raised with fuzzy string matching on short numeric identifiers, such as phone numbers, is that bigram-based similarity compares strings based on overlapping character pairs rather than on numeric distance. As a result, unrelated numbers can yield high similarity scores through incidental overlaps (such as repeated digits). As described above, our procedure addresses this concern by ensuring that fuzzy similarity is never sufficient to establish a link on its own: every retained pair is anchored by an exact match on either the firm's non-missing phone number or email address, with fuzzy scores contributing only as additional evidence on top of that anchor.

Table C2: Illustrative Examples of Similarity-Based Record Linkage

Category	Included	Details	Tax Records	TIN Registry	Rationale
Perfect match (1.000)	Yes	Name	Wobbly Penguin Co.	Wobbly Penguin Co.	Exact agreement on firm name and contact identifiers.
		Email	wobbly.penguin@example.com	wobbly.penguin@example.com	
		Phone	901234567	901234567	
Close-to-perfect match (≈ 0.96)	Yes	Name	Galactic Pickles	The Galactic Pickles	Firm names differ only by additional tokens; both phone and email match exactly.
		Email	galactic.pickles@example.org	galactic.pickles@example.org	
		Phone	912345678	912345678	
High-confidence match (≈ 0.89)	Yes	Name	Velvet Octopus Studios	Velvet Octopus Studios	Exact agreement on firm name and phone number; email addresses differ across registries, but the anchor condition is satisfied.
		Email	contact@velvetoctopus.example	velvetoctopus.office@example.org	
		Phone	934567890	934567890	
Low-confidence match (≈ 0.80)	No	Name	Sleepy Tortoise Bakery	Sleepy Tortoise Bakery	Firm name matches exactly and composite similarity clears the 0.80 threshold, but neither phone nor email matches exactly.
		Email	sleepy.tortoise@bakery.com	info@sleepybakery.com	
		Phone	956789012	956789013	
Low-confidence match (≈ 0.75)	No	Name	Dancing Llama Ltd	Spinning Llama Books	Phone matches exactly, but firm names and emails do not match; composite similarity also falls below the 0.80 threshold.
		Email	dancing.llama@example.com	books.sllama@example.com	
		Phone	945678901	945678901	
No match (—)	No	Name	Quantum Pineapple Inc.	*no close match*	No candidate match retrieved from the TIN registry.
		Email	hello@quantumpineapple.example	*no match*	
		Phone	978901234	*no match*	

Notes: All firm names, email addresses, and phone numbers shown in this table are entirely synthetic and used solely for illustrative purposes; they do not correspond to real firms, individuals, or contact information. The composite similarity scores reported correspond to the weighted bigram-based similarity produced by `reclink2` across firm name, phone, and email. Phone numbers carry more weight in the composite than firm names and email addresses, reflecting their role as unique identifiers less prone to duplication across firms. A TIN is recovered only when the composite similarity score clears the 0.80 threshold and at least one of phone or email matches exactly across the two registries. The "Rationale" column explains, for each case, why the record is retained or excluded under these rules.

C.3 Study samples

Table C3 describes the composition of our analysis sample (Panel A) and the subset selected for the tax nudge experiment (Panel B). Both samples consist of firms from two sources: KCCA's trade license payment records for firms with TINs (either included in KCCA's data or recovered by us) and URA's corporate and presumptive income tax filing records.

Panel A: Full analysis sample. The full analysis sample comprises 155,260 unique firms observed over fiscal years 2018–2023, restricted to firms verifiably active in the administrative data during this window. Firms are classified into three mutually exclusive categories based on the record in which they appear. The first group (31.9%) consists of firms recorded in KCCA's trade license payment data but absent from URA's filing records. The second (39.7%) consists of firms recorded in URA but not in KCCA based on their TIN. The third group (28.4%) appears in both records, though not necessarily in the same fiscal year.

Panel B: experimental sample. The experimental sample was constructed in three steps. First, we selected firms present in both KCCA and URA records in 2022 with annual turnover between UGX 20 million and UGX 300 million, targeting small and medium-sized enterprises subject to either corporate or presumptive income tax obligations. Second, we drew firms present in either authority's records between 2019 and 2021 but not in 2022, applying the same turnover criterion averaged over the years observed; this step captures firms appearing only in URA. Third, we drew KCCA-only firms—never occurring in URA's records—restricted to grade-1 firms whose trade license rate fell between the 25th and 90th percentiles within their economic category. We exclude firms not required to hold a trade license, such as foreign exchange bureaus, legal firms, clinics, pharmacies, and private educational institutions (see for e.g. KCCA, 2026c). Sampling was stratified by tax regime (CIT, PT, neither), and randomization into treatment and control was performed within strata.

This procedure yielded 62,029 firms. Sampling preceded full access to the universal TIN registry, so not all firms could be verified as credibly linked to a TIN at selection. Applying our sample selection criteria yields 60,317 validated firms (97.2% of the initial sample): 30% appear only in KCCA records, 34% only in URA records, and 36% in both.

Table C3: Study Samples

Category	Unique TINs	%
Panel A: Full Analysis Sample		
Total	155,260	100
A. KCCA only (no URA filing records)	49,501	31.9
B. URA only (no KCCA payment records)	61,649	39.7
C. In both (KCCA & URA records)	44,110	28.4
Panel B: Experimental Sample		
Initially drawn	62,029	
Retained sample	60,317	100
By record source:		
A. KCCA only (no URA filing records)	17,951	29.8
B. URA only (no KCCA payment records)	20,578	34.1
C. In both (KCCA & URA records)	21,788	36.1

Notes: This table describes the composition of the analysis sample (Panel A) and the subset selected for the 2023 field experiment (Panel B). Panel A covers fiscal years 2018–2023, restricted to firms verifiably active in the administrative data. Firms are classified into three mutually exclusive categories based on their presence across authorities: KCCA only (never appear in URA filing records), URA only (never appear in KCCA payment records), and in both (appear historically in both KCCA and URA records, not necessarily in the same fiscal year). Panel B reports firms selected for the 2023 field experiment; of the 62,029 firms initially drawn, 1,713 (2.8%) could not be linked to the TIN registry and are excluded, leaving the retained sample of 60,316 firms (97.2%) for the analysis. The retained firms are classified using the same three categories.

C.4 Robustness to TIN-retrieval errors

Here, we probe the robustness of our main results to the possibility of errors in TIN retrieval. For firms whose TINs we had to recover from the TIN registry, there are two possible errors: First, there could be cases in which we retrieve an incorrect TIN from the TIN registry (false positives in TIN retrieval). Second, there could be cases in which we fail to recover the TIN of a firm even though the firm in reality has a TIN (false negatives). Such errors would reflect the same real-world constraints that fragmented tax authorities face due to incomplete record keeping and a lack of harmonization in business registries, but they may bias our estimates. We therefore investigate them in a number of robustness checks.

C.4.1 False positives: incorrectly retrieved TINs

We replicate our main results using the sample of firms whose TINs were readily provided in the records of KCCA and URA and who we therefore did not have to recover from the TIN registry ourselves (Table C4). This rules out the possibility of incorrectly recovered TINs.

This restriction drops firms who did not have a TIN provided in KCCA's records from the analysis sample. In our main sample, these firms either only occur in KCCA's records or are found in both authorities' records after cross-linkage. As a logical consequence of dropping these firms, we estimate a higher non-compliance rate for KCCA's tax (fewer firms that are verifiably active and pay for a trade license) and lower URA non-compliance rate (fewer firms who pay for a trade license but are not filing national business tax). Our main conclusions about high non-compliance rates and substantial complementarity between records are unchanged.

Table C4: Tax non-compliance among firms whose TINs did not have to be retrieved

	URA (CIT/PT)	KCCA (TL)	Either (CIT/PT or TL)	N
Panel A: Non-compliance rates (without recovered TINs)				
Average (2018–2023)	0.380	0.550	0.857	360,503
	(0.001)	(0.001)	(0.001)	
Ever (2018–2023)	0.502	0.663	0.982	121,288
	(0.001)	(0.001)	(0.000)	
	URA	KCCA	N_{URA}	N_{KCCA}
Panel B: Complementarity of records (without recovered TINs)				
Never filed/paid	0.225	0.508	121,288	121,288
	(0.001)	(0.001)		
Never filed/paid (non-compliers)	0.463	0.821	136,889	198,337
	(0.001)	(0.001)		
Never filed/paid previously (non-compliers)	0.590	0.913	136,889	198,337
	(0.001)	(0.001)		

Note: Panel A shows the extensive margin tax non-compliance rate for the sample of verifiably active businesses between 2018 and 2023, separately for URA national business tax (CIT or PT), KCCA municipal tax (TL), and for non-compliance with either of the two taxes. Panel B shows the share of businesses who never filed to one of the authorities. The first row of panel B shows the complementarity between national and municipal records; that is, the share of total businesses that never occur in URA's filing records or KCCA's payment records, respectively. These firms are known taxpayers of one but not the other authority. The second row of panel B shows what share of businesses who do not comply with URA national business tax (KCCA municipal trade licenses) never occur in URA's filing records (KCCA's payment records) when pooling across years. The third row of panel B shows the same share but only considers whether a firm had occurred in URA's (KCCA's) records up until the year under question. The sample drops firms whose TINs had to be retrieved from the TIN registry via similarity matching. Standard errors in parentheses.

C.4.2 False negatives: missed TINs

We replicate our analysis using only firms found in both authorities' records (Table C5). This eliminates the potential for false negatives in cross-record linkage.

This restriction drops firms from the sample who only ever occur in KCCA's records and firms who only ever occur in URA's records. It therefore results in a sample that is more heavily selected on compliant firms across both authorities' taxes. As a consequence, we estimate lower non-compliance rates for KCCA and URA. KCCA non-compliance rates change more strongly due to the greater number of dropped firms who are only in URA's records.[1] Even among this group of firms that had at least once complied with URA's

[1]These are firms that, in a year in which they file URA, always result in detected non-compliance with

and KCCA's taxes, tax non-compliance is pervasive in any given year and across the entire study horizon. In Panel B, we additionally remove any firms whose TINs we had to recover (as in Section C.4.1). Results are similar to Panel A.

Table C5: Tax non-compliance among firms in both authorities' records

	URA (CIT/PT)	KCCA (TL)	Either (CIT/PT or TL)	N
Panel A: Firms in both records				
Average (2018–2023)	0.381	0.284	0.633	170,386
	(0.001)	(0.001)	(0.001)	
Ever (2018–2023)	0.678	0.585	0.926	44,110
	(0.002)	(0.002)	(0.001)	
Panel B: Firms in both records & without recovered TINs				
Average (2018–2023)	0.384	0.265	0.616	134,415
	(0.001)	(0.001)	(0.001)	
Ever (2018–2023)	0.690	0.578	0.933	32,362
	(0.003)	(0.003)	(0.001)	

Note: The table shows the extensive margin tax non-compliance rate for the sample of verifiably active businesses between 2018 and 2023, separately for URA national business tax (CIT or PT), KCCA municipal tax (TL), and for non-compliance with either of the two taxes. The sample in Panel A is limited to firms who occur in both authorities' records, i.e., excluding firms that never file tax to URA or that never obtain a trade license from KCCA. The sample in Panel B additionally drops firms whose TINs had to be retrieved from the TIN registry via similarity matching. Standard errors in parentheses.

C.4.3 Sensitivity of Complementarity Results to TIN-retrieval Errors

How might possible linkage errors affect measured complementarity between the records of KCCA and URA? We conduct a sensitivity exercise focused on bounding the effects of firms who genuinely have a TIN but where this information is missing in KCCA records and not retrieved by us.

Specifically, we conduct the TIN retrieval exercise described in Section C.2 but using 59,639 firms that had their TIN already documented in KCCA records. Using the same TIN-retrieval algorithm we employ to recover the TINs of firms that lack this information

KCCA's taxes as they never obtain a trade license.

in KCCA records, this allows us to form an estimate of TIN-retrieval errors. We then bound the impact of firms we fail to link to a TIN (false negatives) on (i) the firm-level and firm-year-level complementarity estimates in Table 2 Panel B, and (ii) the rate at which this would lead to missed firms that should have been flagged as non-compliant or firms falsely flagged as non-compliant.

Panel A: Estimating the prevalence of TIN-retrieval errors In Panel A of Table C6, we estimate that for 38% of firms that genuinely have a TIN, our TIN recovery method fails to retrieve it. Such entries are dropped from KCCA records for our analysis (since we do not have a TIN to link across records by) but may still be in our analysis sample via their URA record. To estimate how such entries bias the main estimates in Table 2, we therefore need to discern two cases.

Case 1 (dropped firm only in KCCA's records). First, a dropped KCCA record may relate to a firm that is just in KCCA-records. Such firms drop entirely from the analysis sample. They are therefore not counted at all in our estimates of complementarity and bias the estimate of firms never filing to KCCA *upward* and of firms never filing to URA *downward.*[2]

Case 2 (dropped firm in both records). Second, a dropped KCCA record may relate to a firm that is genuinely in both authorities' records. Such firms appear in our main analysis sample as only occurring in URA's records. Thereby, they bias our estimate of firms never paying to KCCA *upward.*

To estimate the effect of dropped firms, we need to recover two unobservables:

(a) the share of KCCA records we drop from the sample even though they genuinely have a TIN

(b) the share among this group that is also in URA's records.

Let $p_{\text{match}} \equiv P(\text{success} \mid \text{has TIN}, \text{rec})$ denote the validation match rate (Table C6, Panel A), $p_{\text{success}} \equiv P(\text{success} \mid \text{rec})$ the observed recovery success rate in the main analysis, $p_{\text{TIN}} \equiv P(\text{has TIN} \mid \text{rec})$ the share of firms without a TIN in KCCA's records but that genuinely have a TIN, and $\varepsilon \equiv P(\text{success} \mid \text{no TIN}, \text{rec})$ the false-positive rate among firms without a TIN.

We can recover (a) from

[2] Firms that never registered for a TIN are by construction KCCA-only. Their exclusion reflects our sample definition (firms identifiable across registries) rather than a linkage error. Including them would push the never-URA share further upward but is outside the scope of this sensitivity exercise.

$$p_{\text{success}} = p_{\text{match}} \cdot p_{\text{TIN}} + \varepsilon \cdot (1 - p_{\text{TIN}}). \tag{C1}$$

under two assumptions. First, we assume that $\varepsilon = 0$, that is, we focus on firms that were not re-linked at all and do not separately model the effect of firms that were re-linked but to the wrong TIN. Such cases make up a minority of linking errors in Table C6, Panel A. Modeling their impact would come with additional assumptions not well anchored in the data and it is ambiguous whether such cases would bias our aggregate estimates downward, upward, or not at all. Second, we assume that the validation estimates from Table C6, Panel A carry over to the group of firms whose TINs we genuinely had to recover.[3]

For (b), we use the share of firms that are in both records among firms whose TINs we successfully recovered. This will provide an accurate proxy for (b) if firms whose KCCA records were falsely dropped are as likely to be in both records as firms whose TINs were successfully recovered.[4]

Putting the pieces together, we can estimate the share of recovery-needed firms that are dropped but that genuinely have a TIN as

$$P(\text{dropped}, \text{has TIN} \mid \text{rec}) = p_{\text{TIN}} \cdot (1 - p_{\text{match}}). \tag{C2}$$

To bound the impact this group has on our main estimates, we need to decompose it into firms that are only in KCCA's records (Case 1, D_K) and firms that are in both authorities' records (Case 2, D_B):

$$D_K = N_{\text{rec}} \cdot p_{\text{TIN}} \cdot (1 - p_{\text{match}}) \cdot (1 - p_{\text{inboth}}), \tag{C3}$$

$$D_B = N_{\text{rec}} \cdot p_{\text{TIN}} \cdot (1 - p_{\text{match}}) \cdot p_{\text{inboth}}, \tag{C4}$$

with D_K the number of dropped KCCA-only firms with a TIN, D_B the number of dropped

[3] The validation sample consists of firms whose TINs were documented in KCCA records; the recovery-needed sample consists of firms whose TINs were not. The two populations may differ in contact-information quality or other characteristics affecting our ability to successfully recover the TINs of firms who possess one. The TIN-retrieval rate for firms whose TINs we had to recover therefore may be lower than what the validation sample suggests. In a robustness check, we find cross-record complementarity to remain substantive for both authorities even at the theoretical lower bound of p_{match} at which every recovery-needed firm has a TIN.

[4] An alternative proxy would be the in-both share among firms with documented TINs in KCCA (the validation sample). As this involves a potentially stronger selection issue, we do not use it as our primary proxy for (b).

firms truly in both records, N_{rec} the number of firms without a TIN in KCCA's records, and p_{inboth} the in-both share among successfully recovered firms. In this way, we estimate that there are 13,648 firms who are excluded from our analysis sample even though they have a TIN (firms in KCCA's records only, D_K) and 7,214 cases in which we drop the KCCA record of a firm that is in URA's records but remained unlinked (D_B).

Panel B: Cross-record complementarity adjusting for TIN-retrieval errors Expressing the firm-count bias as fractions of the observed sample: $\alpha = D_K/N_{\text{obs}}$ and $\beta = D_B/N_{\text{obs}}$ allows to form a bias-corrected estimate of firms never filing to URA and firms never filing to KCCA (first row of Table C6, Panel B).

$$\widetilde{s}^{\text{never URA}} = \frac{s^{\text{never URA}} + \alpha}{1+\alpha}, \tag{C5}$$

$$\widetilde{s}^{\text{never KCCA}} = \frac{s^{\text{never KCCA}} - \beta}{1+\alpha}. \tag{C6}$$

where $\widetilde{s}^{\text{never URA}}$ and $\widetilde{s}^{\text{never KCCA}}$ are the new, bias-corrected estimates and $s^{\text{never URA}}$ and $s^{\text{never KCCA}}$ are the main estimates previously reported in Table 2.

Table C6, Panel B reports our bias-corrected estimates of tax record complementarity. Relative to our previous main estimates (reproduced in brackets), we estimate a slightly higher share of firms never occuring in URA's records (37% vs. 32%) and a lower share of firms that never occur in KCCA's records (32% vs. 40%).

In row 2 of Panel B, we scale this to the firm-year level by assuming that falsely dropped firms who only occur in KCCA's records (Case 1) have the same compliance patterns as observed firms who are only in KCCA records; and equivalently for falsely dropped firms occurring in both records (Case 2). We estimate that after correcting for unrecovered TINs, 57% of URA non-compliers involve firms never occurring in URA's records (previously: 53%) and that KCCA would still miss 72% of its non-compliers if it were to only consult its on records (previously: 77%). After accounting for linking errors, our main conclusion is qualitatively and quantitatively nearly unchanged: cross-record linkage increases detectable non-compliance in Kampala by a factor of 2.78 (previously: 2.94).

Panel C: Missed and wrong leads from TIN-retrieval errors. Finally, we investigate how TIN-recovery errors could impact the efficiency of audits that were deployed based on flags raised via cross-record linking. Without a fully harmonized registration system that documents TINs across all records of URA and KCCA, missing TINs—and the associated errors in cross-record linkage—could reduce the reliability of integrated records for enforcement. The key policy question for KCCA and URA is what share of non-compliers flagged for enforcement via cross-record linkage would result in a waste of audit resources and what share of possible flags for enforcement go undetected. The former is the case when a non-complier is flagged but had in reality complied with taxes, thus providing an incorrect lead for enforcement. The latter occurs when a firm is not flagged even though it did not comply with taxes.

Incomplete documentation of TINs is an issue concentrated on KCCA's records and therefore asymmetrically affects the enforcement efforts of URA and KCCA: incorrectly dropped records could, if anything, represent *missed* leads for URA but could provide *incorrect* leads for KCCA's enforcement. Using our estimates from Panel B, we find that URA would stand to gain at least 19% *more* leads for enforcement if KCCA were to close its gaps in TIN recording. This means that under improved TIN documentation in KCCA's records, the value of cross-record comparison would further increase for URA's non-compliance detection and enforcement efforts. For KCCA's enforcement, we estimate that closing TIN recording gaps would result in about 5.6% fewer flagged non-compliance cases (5.6%), cases that currently provide a wrong lead.

Table C6: Complementarity Between Tax Records
Accounting for TIN-retrieval Errors

Panel A: Estimates for TIN-retrieval errors

	N	%
TINs tested	59,639	100
Re-linked via similarity matching	30,259	50.7
Re-linked (different TIN)	6,690	11.2
Not re-linked (false negative) [a]	29,380	38.0

Panel B: Complementarity adjusting for TIN-retrieval errors

	URA	KCCA	N_{URA}	N_{KCCA}
Never filed/paid	0.374	0.322	168,908	168,908
	[0.319]	[0.397]	[155,260]	[155,260]
Never filed/paid (non-compliers)	0.573	0.718	222,337	199,250
	[0.534]	[0.771]	[186,727]	[211,114]

Panel C: Missed and wrong leads from TIN-retrieval errors

	URA	KCCA	N_{URA}	N_{KCCA}
Missed leads [% of leads in main estimates]	19.1%	0.0%	186,727	211,114
Wrong leads [% of leads in main estimates]	0.0%	5.6%	186,727	211,114

Notes: Adjusted estimates appear on the top line of each row in Panel B; the corresponding main estimates from Table 2 are shown in square brackets. Panel B sample sizes reflect adjusted denominators (with main-estimate denominators in brackets). Panel C reports the count of missed leads (cases of non-compliance not currently flagged but identified after correcting for linkage errors) and wrong leads (cases currently flagged as non-compliant but truly compliant), with shares relative to total non-compliance leads generated in the main estimates shown in parentheses.
[a] The category includes firms for which the procedure returned no candidate TIN, as well as the cases that did not meet our matching criteria.

C.5 Linking Tax Records by Year

For our analysis, we need to link records at the firm-year level. This requires some harmonization between URA's and KCCA's recording convention, as each authority's taxes follow a different filing period. We harmonize all tax records to align with URA's standard fiscal year (July 1–June 30) using a majority-period assignment rule. Below are some examples to illustrate how this harmonization works in practice.

URA CIT/PT. While 94% of firms use the standard URA financial year (July 1–June 30) as their accounting period, approximately 6% use alternative accounting periods, including the calendar year or firm-specific operational timelines. We harmonize these heterogeneous reporting windows using a majority-months rule.

- *Standard financial year*: Firm A files CIT for July 1, 2020–June 30, 2021. This return is assigned to FY 2020-21 in accordance with the URA financial year
- *Calendar year*: Firm B follows an accounting period from January 1, 2020–December 31, 2020. Six months fall into FY 2019-20 (January–June 2020) and six into FY 2020-21 (July–December 2020). Following URA practice, we assign this return to FY 2020-21 (the later fiscal year when periods are split evenly).
- *British fiscal year*: Firm C follows an accounting period from April 1, 2021–March 31, 2022. Nine months fall into FY 2021-22 (July 2021–March 2022) and three months fall into FY 2020-21 (April–June 2021). We assign this return to FY 2021-22 because the majority (nine of twelve months) fall within that fiscal year.

KCCA Trade license. Trade licenses are valid for twelve months from the payment date. We apply the same majority-period rule to assign each individual trade license to a fiscal year. Since a single parent firm may hold multiple trade licenses simultaneously (for different branches or business activities), we then aggregate to the parent-firm level: a firm is considered compliant in a given fiscal year if it holds at least one valid trade license assigned to that year.

- *On-time renewal*: Firm D purchases a trade license on May 1, 2021, valid until April 30, 2022. Two months fall into FY 2020-21 (May–June 2021) and ten months fall into FY 2021-22 (July 2021–April 2022). This license is assigned to FY 2021-22.
- *Late renewal, gap in coverage*: Firm E pays for a trade license in January 2020 (valid until December 2020, assigned to FY 2020-21). The firm renewed in April 2021, three months late. The new license is valid until March 2022: three months fall into FY 2020-21 and nine into FY 2021-22. This trade license is assigned to FY 2021-22. The

firm's next renewal is in January 2023, nine months after expiry. This license covers equal parts of FY 2022-23 and FY 2023-24 and is assigned to FY 2023-24. It is inferred that the firm lacks a valid license for FY 2022-23, as it operated without coverage for most of that year. It would be labeled as non-compliant if it is active in other records, such as filing CIT/PT with the URA.

- *Late renewal, continuous fiscal year compliance*: Firm F pays for a trade license in January 2020 (assigned to FY 2020-21). The firm renewed in October 2021, ten months late. This license is valid until September 2022: nine months fall into FY 2021-22 and three into FY 2022-23, assigned to FY 2021-22. The firm renews on time in October 2022 (assigned to FY 2022-23). It is noted that, despite operating without a license from January to September 2021, the firm is considered compliant for all fiscal years in our data, as each fiscal year has at least one license assigned to it under the majority-months assignment rule.
- *Multiple trade licenses*: Firm G operates two branches and holds two trade licenses. License 1 was purchased in November 2020 (valid until October 2021, assigned to FY 2021-22) but was not renewed. License 2 was purchased in March 2022 (valid until February 2023, assigned to FY 2022-23). At the firm level, Firm G is considered compliant for both FY 2021-22 and FY 2022-23 because it held at least one valid license in each year, even though the licenses cover different branches and time periods. Approximately 14% of firms in our sample hold multiple trade licenses. The combination of the majority-period assignment rule and firm-level aggregation means firms can be labeled compliant for a fiscal year provided at least one of their licenses covers the majority of that fiscal year, even if individual licenses lapse or are renewed late at different times.

D Revenue Shortfall from Non-Compliance

To quantify the fiscal implications of selective non-compliance, we estimate revenue shortfalls as the counterfactual tax liabilities of non-compliant firms in a given year. We then express this relative to the observed tax revenue from compliant firms in that year.

KCCA Trade License Shortfalls. Trade license rates are determined based on the nature of business (i.e., its main activity) and location within Kampala (divided into four grading areas). For a given business, both characteristics are typically stable over time. For firms who have ever held a trade license, we can therefore calculate the average amount the firm pays in trade license costs per year whenever it complies and use this to estimate the firm's counterfactual trade license payment in years it was active but did not pay. Since trade license fees are levied at predetermined flat rates which have not changed since 2017, this corresponds very closely to firm's true counterfactual payment.[1]

For firms who never held a trade license with KCCA, we cannot calculate counterfactuals based on the firm's payment history. We therefore need to determine the business' nature and grade via other firm characteristics we have from the TIN registry. First, we use business type information recorded in the TIN registry. Since the business type classifications in the TIN registry follow a different scheme than what the KCCA's trade license schedule uses, we map firms to the KCCA business nature that their business type typically corresponds to (based on firms for whom we have both, a business type from the TIN registry and their KCCA business nature). Second, as we lack fine grained location information from the TIN registry, we take firms as either being of Grade I or Grade II. Together, these grades make up over 99% of all trade licenses (80% Grade I and 19% Grade II). We regard the specification assuming all firms without KCCA information are of Grade I as our main specification—with the Grade II version serving as a lower bound robustness check—for two reasons. First, 80% of all trade licenses are of Grade I, making this the most likely option. Second, a validation check suggests this leads to much more accurate estimates of counterfactual trade license payments.[2]

To estimate the total revenue shortfall from non-compliance with trade licenses, we sum the total estimated counterfactual payments of firms that were verifiably active but did not hold an active trade license, separately for each year. We then express this relative to the KCCA's total trade license revenues in the same year. Table D1, summarizes the result. In a typical year, we estimate the KCCA to miss out on about the same amount of revenue due to non-compliance as it already raises in total trade license revenue from compliant firms (102% averaged between 2018 and 2023). We estimate shortfalls to range from 80% of total revenue in 2020 to 123% in 2023.[3]

[1] In a validation exercise, we compare total trade license payments *as observed* to total trade license payments *as estimated* from each firm's leave-one-out average of trade license payments. Estimated payments based on past and future payment amounts of the firm sum to between 100% and 102% of total observed payments.

[2] For the group of firms for whom we observe actual trade license payments, we compare the total *observed* payment amount to the total *estimated* payment amount in each year. Independently determining business nature as described and assuming all firms are of Grade I leads to broadly accurate, yet conservative estimates ranging between 88% and 94% of the benchmark observed total (91% on average). On the other hand, assuming these firms are of Grade II results in a vast underestimate, with total payments summing to only 42%–45% of the benchmark total. Both approaches underestimate benchmark totals because they inherently have to assume a firm would only have to pay for a single license whereas some firms need to pay for multiple trade licenses (see Section 2.1).

[3] If we assume firms who have never held a trade license to all be of Grade II, shortfall estimates still average 64% of annual trade license revenues.

Table D1: Annual observed trade license revenue and imputed shortfall share

Fiscal year	Observed revenue	Mean share	95% CI
2018	9,941.49	0.924	[0.901, 0.947]
2019	10,447.25	0.959	[0.933, 0.985]
2020	12,107.76	0.798	[0.776, 0.818]
2021	9,979.79	1.083	[1.050, 1.113]
2022	10,743.63	1.152	[1.117, 1.182]
2023	10,221.15	1.233	[1.200, 1.267]

Notes: Observed revenue is the total trade license revenue collected from compliant firms in the given fiscal year. Mean share is the imputed shortfall expressed as a share of the same year's observed revenue (in UGX millions). 95% confidence intervals are constructed by resampling firms with replacement 500 times and taking the 2.5th and 97.5th percentiles of each year's imputed-shortfall ratio across the bootstrap draws.

URA Business Tax Shortfalls. URA business tax liabilities depend on turnover and vary substantially across firms, making imputation more complex than for KCCA, whose flat trade-license structure permits direct application of observed amounts. We therefore train a machine-learning imputation model that predicts missing URA liabilities from features available in the merged registry. Because our object of interest is the aggregate revenue shortfall rather than firm-level liability, the model is optimized for aggregate calibration, i.e. the sum of predicted liabilities divided by the sum of observed liabilities, with values close to one indicating good aggregate calibration.

We train the model on the subset of all firm-year observations with observed URA liabilities. We exclude the top 2.5% of firms by liability for three reasons: (i) their liabilities are noisy and difficult to predict from the available features; (ii) they have systematically high compliance rates at the extensive margin, so the imputation problem for them is small; and (iii) their inclusion dominates aggregate performance metrics, masking model behavior on the majority of firms. The same exclusion applies at the imputation stage, so our reported shortfall is conditional on the bottom 97.5% of firms by liability. As predictors, we include URA information on past and future liabilities, turnover, and nil-filing behavior (lags and leads of up to five years, together with the nearest non-missing observation before and after the target year), the firm's registration history, past filing dates, and sector. From the KCCA registry, we use the trade-license amount, the number of trade licenses, the license category, payment dates, sector, and geographic information.

Given the high prevalence of missing values in the predictors, we train gradient-boosting models using the LightGBM framework, which handles missing values natively. Since the distribution of URA liabilities has a substantial cluster at zero and a long right tail, we use a Tweedie deviance loss function a criteria to build trees. We split the data 60/40 at the firm level. Hyper-parameters are selected via 6-fold cross-validation within the 60% training partition; the final model is then applied to the 40% held-out test partition for performance evaluation. The results for the best-performing configuration are reported in Table D2.

On the held-out test data, the predicted and observed aggregate liabilities align closely with an overall calibration ratio of 0.99, indicating an underestimation of roughly 1 pp. While this suggests that the model predicts well for firms similar to those in the test data, the firms whose liabilities we ultimately impute are non-compliant and may differ from the test firms. We therefore examine conditional calibration across four mutually exclusive blocks defined by each firm's past compliance status: firms that always comply with the URA, firms that were partially compliant over the years, firms that were never compliant except for the observed year, and firms entering the registry for the first time, for which no historical features are available.

The conditional calibration results in Table D2 are closely aligned across groups, with values close to one. The main exception is a calibration ratio of 1.21 in 2022 for first-time entrants, which is likely associated with the substantial increase in TIN registrations that year following the introduction of the instant TIN registration process. Overall, the model is well-calibrated, with a slight tendency to underestimate liabilities.

A fundamental limitation, which no in-sample diagnostic can fully address, is that we cannot directly validate predictions for firms that never filed with URA during the panel period, since by construction we have no ground truth for them. The closest comparison in our data are firms that filed in exactly one year, providing a single observation for evaluation per such firm. The calibration ratio of 0.95 on this proxy subgroup provides some reassurance about model performance on the never-compliant population, under the untestable assumption that single-filers are sufficiently similar in their liability-feature relationship to firms that never filed. More broadly, our imputation is conditional on the observable features available in the merged registry: if non-compliant firms differ from compliant firms in ways not captured by these features the imputed values are biased in the corresponding direction. The estimates that follow should be read with this caveat in mind.

Table D2: Conditional calibration by compliance group and fiscal year

Past-compliance bucket	Year	N	Ratio
No past observations	2018	12,947	1.03
	2019	4,760	1.07
	2020	2,991	0.98
	2021	2,858	0.99
	2022	4,028	1.21
	2023	5,375	0.98
Never compliant	2019	2,034	0.88
	2020	872	0.94
	2021	689	0.93
	2022	854	1.01
	2023	970	0.98
Partially compliant	2020	1,652	0.99
	2021	2,654	1.02
	2022	3,402	0.94
	2023	4,617	1.00
Always compliant	2019	7,832	0.96
	2020	8,236	1.05
	2021	8,539	0.95
	2022	8,860	1.00
	2023	9,679	0.97

Notes: Calibration performance on held-out validation cells, stratified by the firm's past-compliance bucket and fiscal year. Ratio is the sum of predicted liabilities divided by the sum of observed liabilities within each cell. No past observations refers to firms first observed in the indicated year (lag features unavailable); Never compliant to firms with zero compliant years prior to the indicated year; Partially compliant to firms with some compliant years; Always compliant to firms compliant in every prior year.

Based on the trained model, we impute URA tax liabilities for firms that hold an active trade license but do not file URA business income tax that year. KCCA non-compliant firms that are only identified as active through PAYE or VAT payments are not included in the shortfall estimates because the activity signal does not originate from the cross-registry linking. The counts by year and firm category are reported in Table D3. We aggregate the imputed liability shortfalls by year and relate them to the observed liabilities (excluding the top 2.5%) in the corresponding year. The estimated annual shortfall ranges from 16% of observed liability in 2023 to 57% in 2020, with a declining trend over the panel; the 2020 peak coincides with the onset of the COVID-19 pandemic.

Table D3: Imputation row counts by past-compliance bucket and fiscal year

Past-compliance bucket	2018	2019	2020	2021	2022	2023	Total
No past observations	32,792	9,771	10,868	6,656	7,721	4,752	72,560
Never compliant (0%)	0	16,802	17,330	13,507	15,079	13,170	75,888
Partially compliant	0	0	4,576	6,265	7,520	5,523	23,884
Always compliant (100%)	0	5,421	5,413	3,105	3,528	1,911	19,378
Total	32,792	31,994	38,187	29,533	33,848	25,356	191,710

Notes: Number of imputed firm-year observations by the firm's past-compliance bucket and fiscal year. "No past observations" refers to firms first observed in the indicated year (lag features unavailable); "Never compliant" to firms with zero compliant years prior to the indicated year; "Partially compliant" to firms with some compliant years; "Always compliant" to firms compliant in every prior year.

Table D4: Annual observed liability and imputed shortfall share

Fiscal year	Observed liability	Mean share	95% CI
2018	29,197.08	0.546	[0.537, 0.555]
2019	36,457.81	0.452	[0.444, 0.460]
2020	34,052.48	0.567	[0.557, 0.577]
2021	37,866.16	0.442	[0.433, 0.451]
2022	45,831.75	0.357	[0.350, 0.364]
2023	60,663.38	0.160	[0.157, 0.165]

Notes: Observed liability is the total sum of URA's tax liability across compliant firms in the given fiscal year (in UGX millions). Mean share is the imputed shortfall expressed as a share of the same year's observed liability. 95% confidence intervals are constructed by resampling firms with replacement 500 times and taking the 2.5th and 97.5th percentiles of each year's imputed-shortfall across the bootstrap draws.

E Tax Nudge Experiment

E.1 Pre-trends

Figure E1: Pre-trends by treatment arm (URA)

Notes: OLS coefficients from interactions of treatment arm indicators with year dummies, relative to fiscal year 2022 on URA tax filing. Vertical bars show 95% confidence intervals based on standard errors clustered at the parent-firm level.

Figure E2: Pre-trends by treatment arm (KCCA)

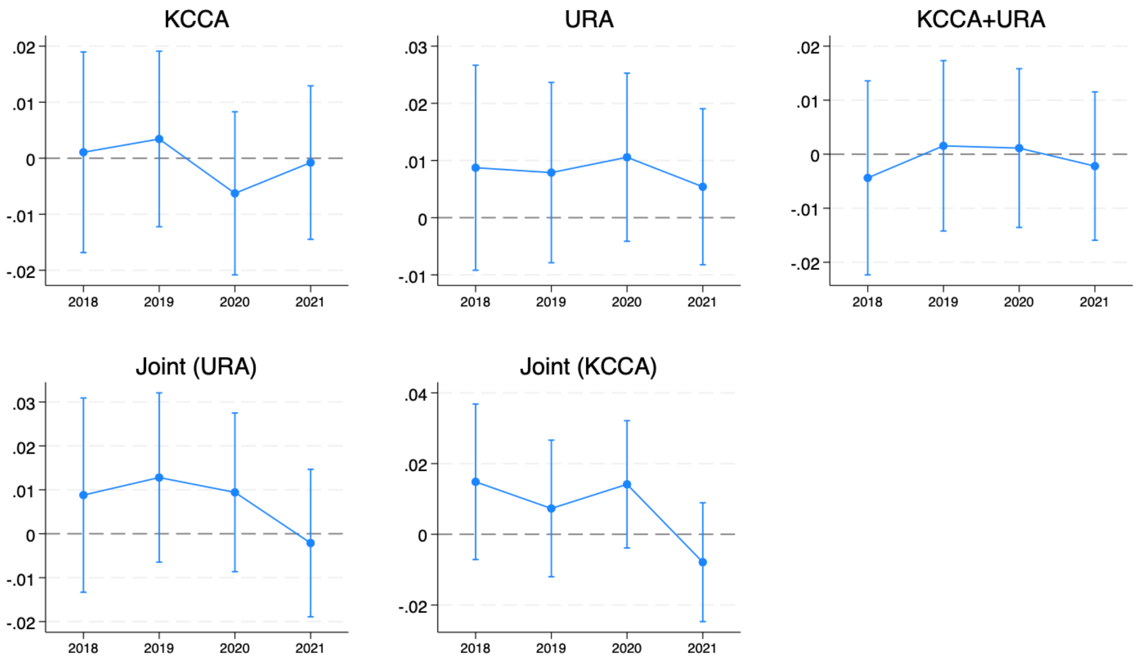

Notes: OLS coefficients from interactions of treatment arm indicators with year dummies, relative to fiscal year 2022 on having a KCCA trade license. Vertical bars show 95% confidence intervals based on standard errors clustered at the parent-firm level.

Table E1: Balance Table

	Control	KCCA	URA	KCCA+URA	Joint (URA sends)	Joint (KCCA sends)	p-value
	12,078 (20.0%)	12,082 (20.0%)	12,044 (20.0%)	12,054 (20.0%)	6,049 (10.0%)	6,009 (10.0%)	
Presumptive Tax 2022							
No	11,339 (93.9%)	11,336 (93.8%)	11,295 (93.8%)	11,308 (93.8%)	5,664 (93.6%)	5,615 (93.4%)	0.896
Yes	739 (6.1%)	746 (6.2%)	749 (6.2%)	746 (6.2%)	385 (6.4%)	394 (6.6%)	
Corporate Income Tax 2022							
No	7,751 (64.2%)	7,681 (63.6%)	7,721 (64.1%)	7,708 (63.9%)	3,913 (64.7%)	3,869 (64.4%)	0.746
Yes	4,327 (35.8%)	4,401 (36.4%)	4,323 (35.9%)	4,346 (36.1%)	2,136 (35.3%)	2,140 (35.6%)	
URA Business Tax Liability (std)	-0.012 (0.079)	-0.011 (0.142)	-0.012 (0.097)	-0.011 (0.081)	-0.011 (0.087)	-0.012 (0.072)	0.989
Paid KCCA (1+trade licenses)	0.335 (0.472)	0.334 (0.472)	0.330 (0.470)	0.336 (0.472)	0.333 (0.471)	0.336 (0.472)	0.933
KCCA payment amount (std)	0.036 (0.975)	0.033 (1.239)	0.043 (1.317)	0.044 (2.030)	0.007 (0.406)	0.012 (0.376)	0.873
Part of 2019 intervention							
No	156 (30.2%)	144 (25.8%)	123 (23.2%)	126 (24.4%)	70 (25.8%)	68 (25.4%)	0.168
Yes	360 (69.8%)	414 (74.2%)	408 (76.8%)	391 (75.6%)	201 (74.2%)	200 (74.6%)	
Kampala Divisions							
Central	5,531 (45.8%)	5,404 (44.7%)	5,253 (43.6%)	5,369 (44.5%)	2,683 (44.4%)	2,742 (45.6%)	0.414
Kawempe	1,148 (9.5%)	1,170 (9.7%)	1,177 (9.8%)	1,223 (10.1%)	583 (9.6%)	582 (9.7%)	
Makindye	1,448 (12.0%)	1,503 (12.4%)	1,514 (12.6%)	1,455 (12.1%)	740 (12.2%)	705 (11.7%)	
Nakawa	2,625 (21.7%)	2,701 (22.4%)	2,748 (22.8%)	2,672 (22.2%)	1,385 (22.9%)	1,347 (22.4%)	
Rubaga	1,326 (11.0%)	1,304 (10.8%)	1,352 (11.2%)	1,335 (11.1%)	658 (10.9%)	633 (10.5%)	

Note: Tests: For continuous variables, p-values are from F-tests of equal means across groups; for categorical variables, p-values are from Pearson's chi-squared tests of independence.

E.2 Estimation tables

Table E2: Treatment Effects on Tax Filing and Trade License Status (complete estimates)

	(1) URA	(2) KCCA
KCCA Message	0.007	-0.001
	(0.006)	(0.006)
URA Message	0.002	-0.004
	(0.006)	(0.006)
URA+KCCA Message	0.003	0.002
	(0.006)	(0.006)
Joint Message (URA)	-0.003	-0.000
	(0.008)	(0.007)
Joint Message (KCCA)	0.002	0.002
	(0.008)	(0.007)
Fiscal Year 23	-0.023***	-0.064***
	(0.004)	(0.004)
KCCA Message × Fiscal Year 23	0.005	-0.002
	(0.006)	(0.006)
URA Message × Fiscal Year 23	0.007	0.012*
	(0.006)	(0.006)
URA+KCCA Message × Fiscal Year 23	0.007	0.005
	(0.006)	(0.006)
Joint Message (URA) × Fiscal Year 23	0.004	-0.001
	(0.007)	(0.007)
Joint Message (KCCA) × Fiscal Year 23	-0.001	0.006
	(0.007)	(0.007)
N	120,634	120,634

Notes: Table reports treatment effects from OLS regressions of an indicator for filing URA business income tax (column 1) and for having at least one active KCCA trade license (column 2) in 2023 on treatment assignment, following the specification in Equation 1. Estimates include 60,317 firms active in 2022. Standard errors in parentheses are clustered at the parent firm level. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table E3: Treatment Effects on Tax Filing over time (complete estimates)

	(1) URA	(2) URA	(3) URA	(4) URA	(5) URA
treatment 2019	-0.02				
	(0.01)				
2018	0.01	-0.26***	-0.31***	-0.34***	-0.34***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
treatment 2019 × 2018	0.00				
	(0.01)				
treatment 2020		-0.02			
		(0.01)			
treatment 2020 × 2018		0.01			
		(0.02)			
treatment 2021			-0.02		
			(0.01)		
treatment 2021 × 2018			0.02		
			(0.02)		
treatment 2022				-0.02	
				(0.01)	
treatment 2022 × 2018				0.02	
				(0.02)	
treatment 2023					-0.02
					(0.01)
treatment 2023 × 2018					0.01
					(0.02)
Observations	14,575	14,575	14,575	14,575	14,575

Notes: Table reports treatment effects from OLS regressions of an indicator for filing URA business income tax for years 2019 through 2023 on treatment assignment, following the specification in Equation 1. The base year for the difference-in-difference estimator is 2018. Estimates are based on treatment assignment in the 2019 intervention on 7,501 active firms in 2019. Treatment includes all treatment arms comprising deterrence, encouragement, and informative messages. Standard errors in parentheses are clustered at the parent firm level. $^{*}\ p < 0.05$, $^{**}\ p < 0.01$, $^{***}\ p < 0.001$.

E.3 Intensive margin effect estimates

Figure E3: Treatment effects on tax liability

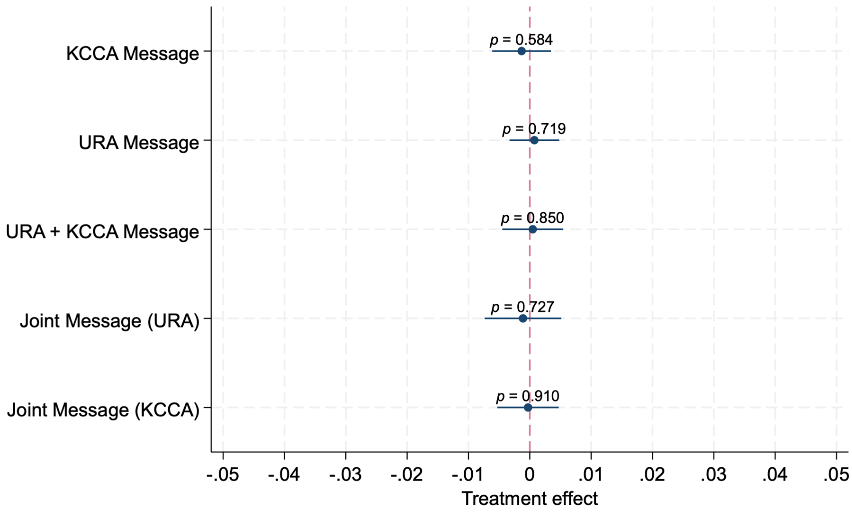

(a) URA tax liability

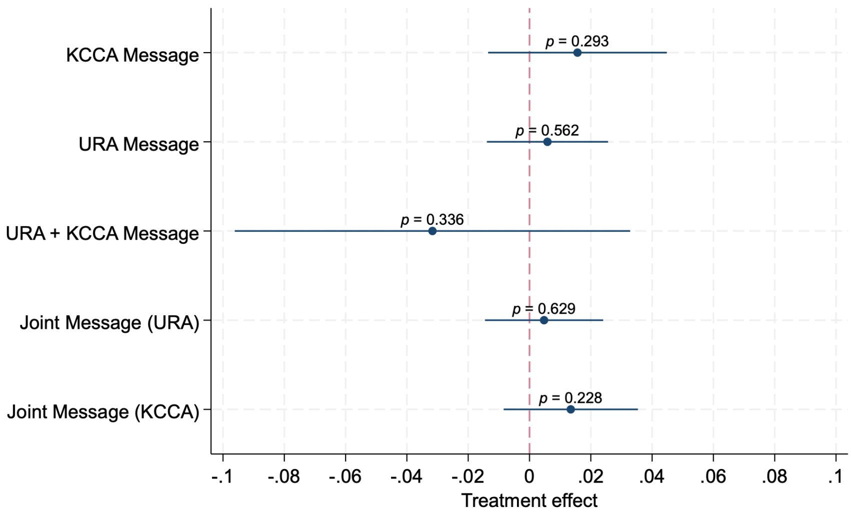

(b) KCCA trade licence payment

Notes: Figure reports treatment effects from OLS regressions of an indicator a) for the amount of URA business income tax declared and b) the trade licence amount paid to the KCCA in 2023 (standardized using the 2019 mean) on treatment assignment, following the specification in Equation 1. Estimates include 29,811 (a) and 24,577 (b) firms active in 2023. Standard errors are clustered at the parent firm level, and bars show 95% confidence intervals.

E.4 Effects on previously non-compliant businesses

Figure E4: Treatment Effects on previously non-compliant firms

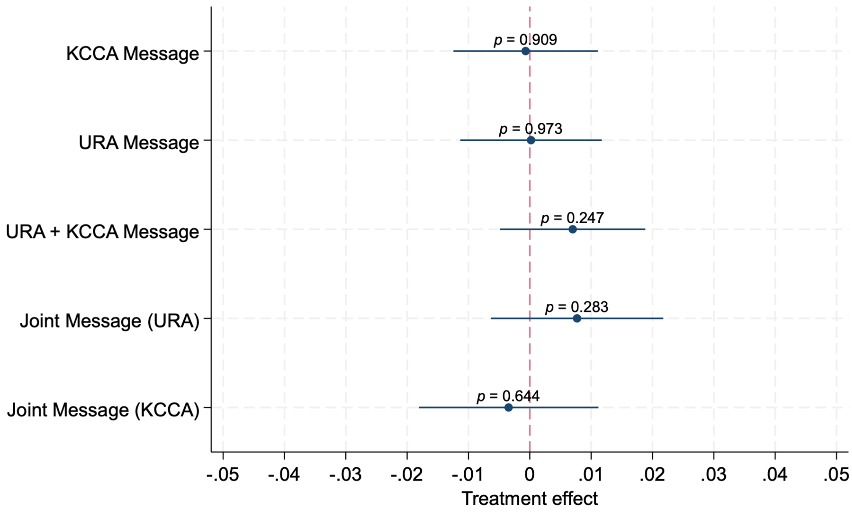

(a) URA

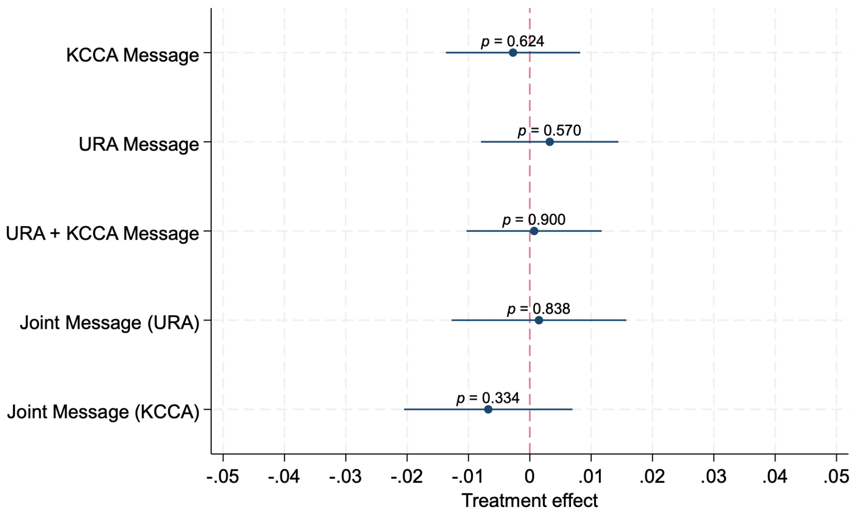

(b) KCCA

Notes: Figure reports treatment effects from OLS regressions on a) URA compliance of previously URA non-compliant firms (18,094) and b) KCCA compliance of previously KCCA non-compliant firms (20,100), following the specification in Equation 1. Standard errors are clustered at the parent firm level, and bars show 95% confidence intervals.

E.5 Effects by business size

Figure E5: Treatment effects on URA tax compliance by business size

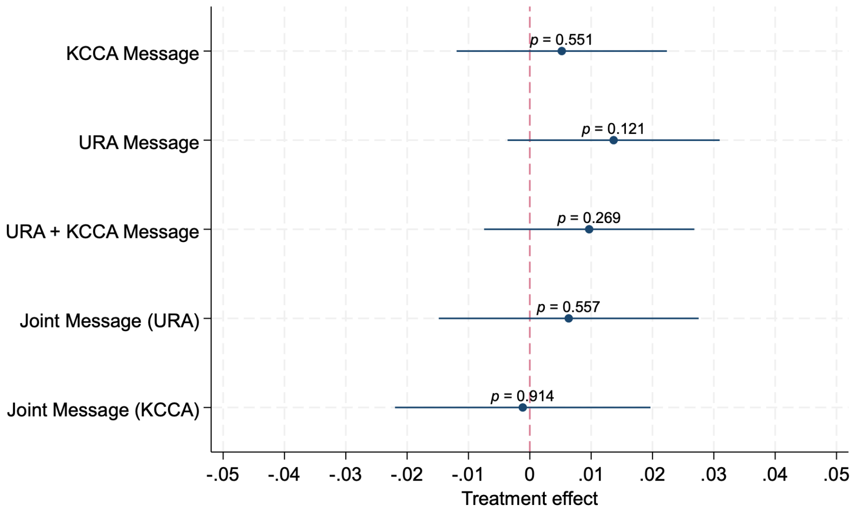

(a) larger businesses

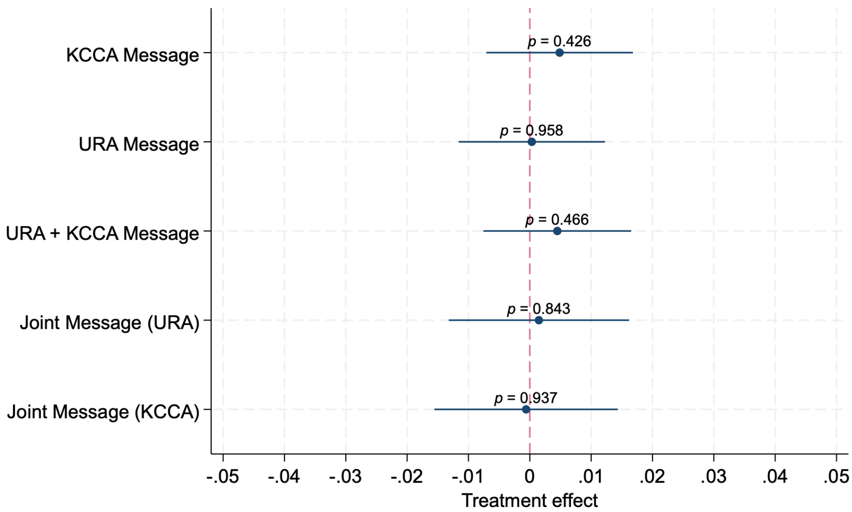

(b) smaller businesses

Notes: Figure reports treatment effects from OLS regressions on URA tax filing for larger (33,279) (a) and smaller (27,038) businesses (b) in 2023 on treatment assignment, following the specification in Equation 1. Business size is defined as described in subsubsection 3.1.4. Standard errors are clustered at the parent firm level, and bars show 95% confidence intervals.

Figure E6: Treatment Effects on KCCA trade licenses by business size

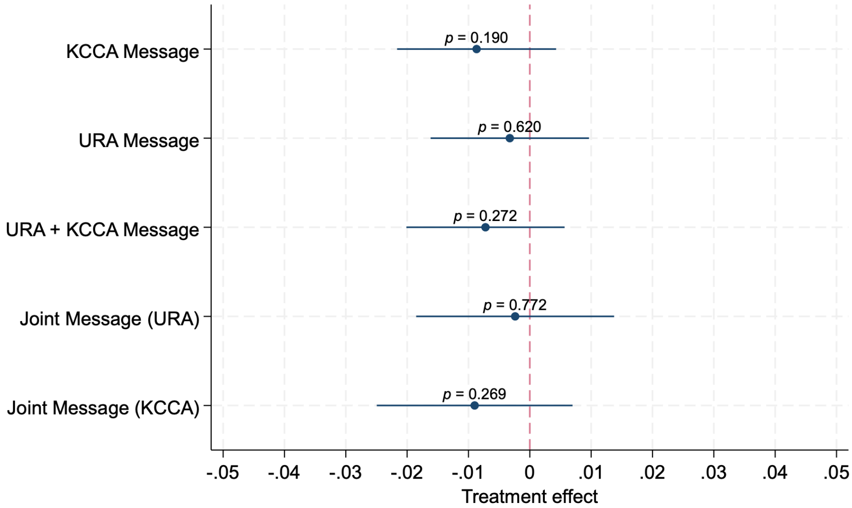

(a) larger businesses

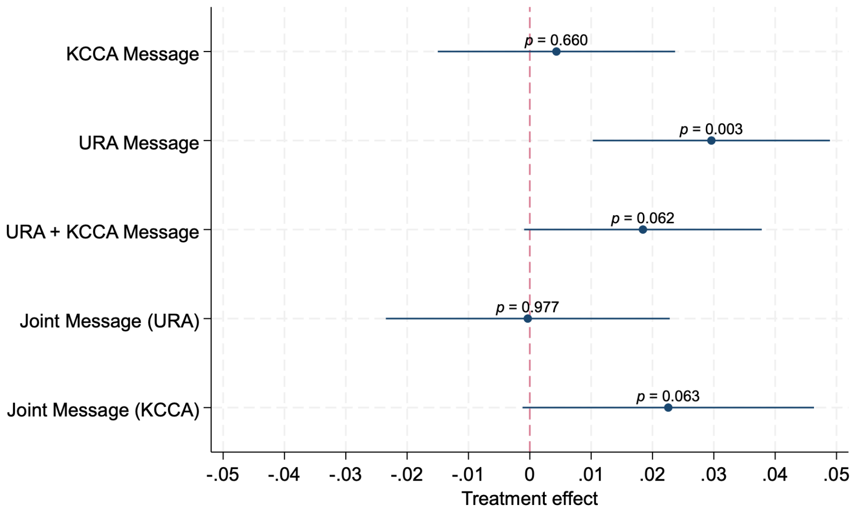

(b) smaller businesses

Notes: Figure reports treatment effects from OLS regressions on KCCA trade lices compliance for larger (33,279) (a) and smaller (27,038) businesses (b) in 2023 on treatment assignment, following the specification in Equation 1. Business size is defined as described in subsubsection 3.1.4.. Standard errors are clustered at the parent firm level, and bars show 95% confidence intervals.

E.6 Effects over time: interaction of interventions 2019 and 2022 and indirect effects

Figure E7: Interaction effects of treatment in 2019 and 2023

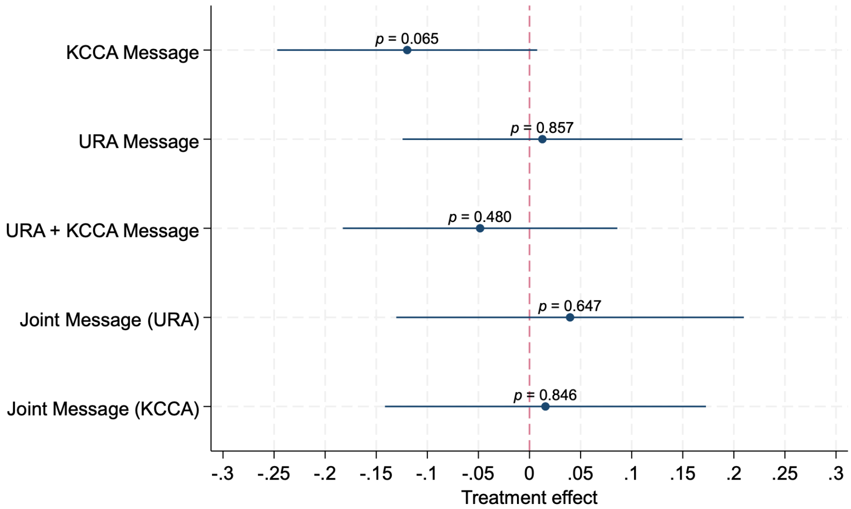

Notes: Figure reports differential treatment effects from OLS regressions of an indicator for treatment assignment in 2023 and treatment in 2019 on filling URA business income tax in 2023. Estimates include 2,358 firms that were part of the 2019 and 2023 intervention. The 2019 treatments are pooled into a single treatment category. Standard errors are clustered at the parent firm level, and bars show 95% confidence intervals.

Figure E8: Treatment effects of 2019 tax messages over time on trade licenses

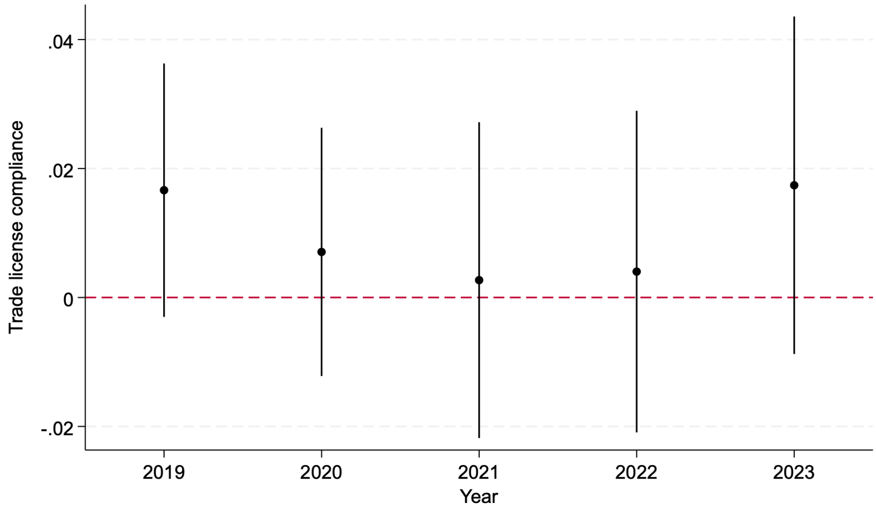

Notes: Figure reports treatment effects from OLS regressions of an indicator for having an active trade license for year 2019 through 2023 on treatment assignment in the 2019 intervention. The base year for the difference-in-difference estimator is 2018. Estimates are based on on 7,502 active firms in 2019 that were part of the 2019 intervention. Standard errors are clustered at the parent firm level, and bars show 95% confidence intervals.

E.7 Robustness tests

Table E4: Robustness to model specification and sample restrictions

	Logit (AME)		Early filers excl.		Firms in both records		In both records & without recovered TIN	
	KCCA	URA	KCCA	URA	KCCA	URA	KCCA	URA
KCCA Message	-0.003	0.011	-0.007	-0.003	0.000	0.006	-0.010	0.010
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.013)	(0.012)
URA Message	0.008	0.009	0.006	0.008	0.015*	0.007	0.010	0.004
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.012)	(0.012)
URA + KCCA Message	0.006	0.010	-0.002	0.005	0.009	0.008	0.006	0.007
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.012)	(0.012)
Joint Message (URA)	-0.002	0.001	-0.002	0.002	0.002	0.005	-0.006	0.002
	(0.007)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)	(0.015)	(0.015)
Joint Message (KCCA)	0.007	0.001	0.005	-0.003	0.013	-0.002	0.013	-0.001
	(0.007)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)	(0.015)	(0.014)
Observations	120,634	120,634	102,668	113,546	115,168	115,168	38,118	38,118

Notes: Each column reports the treatment effect from a difference-in-differences specification on KCCA and URA compliance. Columns (1)–(2) report average marginal effects (AME) from a logit model. Columns (3)–(4) exclude firms that filed before the intervention date or had no due trade license before March 2023. Columns (5)–(6) restrict to firms with a confirmed KCCA-URA TIN match. Columns (7)–(8) further restrict to firms observed in both KCCA and URA registers. Standard errors clustered at the firm level in parentheses. * $p < 0.05$