Algorithms and Bureaucrats: Evidence from Tax Audit Selection in Senegal

Pierre Bachas, Anne Brockmeyer, Alipio Ferreira, Bassirou Sarr*

October 31, 2024

Abstract

Can algorithms enhance the work of bureaucrats in a developing country setting? Developing economies are often data-poor environments, where individual bureaucrats have substantial discretion to take key decisions, such as selecting taxpayers for audits. Exploiting a trove of newly digitized microdata, we conduct a field experiment across tax offices in Senegal whereby half of the annual audit program is selected by inspectors and the other half is selected by a risk-scoring algorithm. We document three sets of results. First, inspector-selected audits are 18 ppt more likely to be conducted and detect higher amounts of evasion. Second, algorithm-selected audits are less cost-effective and do not generate less corruption. Third, even an ex-post optimized algorithm trained on audit outcome data would have increased detected evasion by only 5% compared to the inspector selection. This is consistent with the inspectors' high skills, strong financial incentives and the imperfection of the available data.

*Pierre Bachas: World Bank Research and EU Tax Observatory, pbachas@worldbank.org; Anne Brockmeyer: World Bank, IFS and UCL, abrockmeyer@worldbank.org; Alipio Ferreira: Southern Methodist University, alipioferreira@smu.edu; Bassirou Sarr: Senegal Ministry of Finance. We thank Denis Cogneau, Laurent Corthay, Lucie Gadenne, Janet Jiang, Nicola Limodio, Jan Loeprick, Markus Kitzmuller, Imran Rasul, Dan Rogger, Eduardo Souza-Rodrigues, Gabriel Zucman for helpful comments and discussions, and seminar/conference audiences at Berkeley Haas, IFS, IIPF, CESifo Public Economics Week, INSPER, TARC Exeter, WB Tax Conference, CMI TaxCapDev Conference, the Paris School of Economics, Oxford Centre for Business Taxation, Oxford University Economics Department, PUC Rio, University of Muenster, University of New Mexico, Norwegian School of Economics, NTA Conference, and Universidad del Pacífico. We also thank the Senegal Tax Administration (DGID), in particular, Bassirou S. Niasse, Amadou A. Badiane, Oumar D. Diagne, Hady Dieye, Mor Fall, Serigne M. Fall, and Mathiam Thioub. We thank Samba Mbaye, Assane Sylla, and Medoune Sall from the CRDES for their collaboration, Oumy Thiandoum for excellent research assistance, the Paris School of Economics and CEPREMAP for administrative support, and the World Bank, UKAID via EDI, the Centre for Tax Analysis in Developing Countries (TaxDev), and the UKRI through Brockmeyer's Future Leaders Fellowship (grant reference MR/V025058/1) for financial support. The findings, interpretations, and conclusions do not necessarily represent the views of the World Bank, its affiliated organizations, its Executive Directors or the governments they represent, nor the Government of Senegal. All errors are our own.

1 Introduction

Audits are a key tool to ensure compliance with government policies and regulations, from water quality standards to pollution caps and tax payments. Yet resources for audits are constrained, especially in lower-income countries. The decision on which agents to audit is hence crucial for the efficiency of the enforcement system. Governments around the world use different methods for selecting cases for audits. In high-income countries, governments usually rely on data-driven and risk-based methods of audit selection (Khwaja et al., 2011; OECD, 2023). These methods leverage all available information in a systematic manner and are non-arbitrary in theory, but require high-quality data. In lower-income countries, where data is usually scarce, governments often rely on the discretion of individual bureaucrats to select cases for audits. Such discretionary audit selection can leverage the bureaucrats' private information and experience, but could open the door to biases and corruption.[1]. This paper studies whether the use of a risk-based audit selection algorithm improves audit outcomes relative to discretionary audit selection. Our setting is the tax audit program in Senegal, that aims to target firms with high amounts or high rates of tax evasion. Two features make Senegal a uniquely suitable context for our study. First, tax collections in Senegal are low and evasion is high, estimated to be at least 25% of liabilities for the VAT and over 60% for the income tax. Improving tax enforcement hence has large potential returns. Second, Senegal undertook large digital infrastructure investments, mandating electronic tax returns and digitizing third-party data sources, to improve data availability for tax enforcement. International organizations and donors are pushing for similar digitization efforts in other developing countries. Our project constitutes the first effort to systematically exploit Senegals' new data infrastructure at scale. Before the data revolution, tax inspectors in Senegal had the discretion to select cases for audits without necessarily providing a rationale for their choice.

In collaboration with the Senegalese tax authority, we designed a risk-scoring algorithm to improve audit selection, aiming to identify firms with high evasion amounts or high evasion rates. Our algorithm draws on taxpayers' self-assessment declarations for different taxes and on third-party records from customs, public procurement, and firm-transaction reports. For each taxpayer (firm), the algorithm yields a risk score, composed of multiple indicators that aim to proxy for tax evasion. One class of indicators flags inconsistencies, for instance, when a firm's self-declared sales are lower than its third-party reported sales. Another class of indicators flags anomalies, such as abnormally low profit rates compared to other firms of the same size and sector. The choice of risk indicators is based on best practices in audit selection and discussions with experts from the World Bank and the International Monetary Fund (IMF).

The audit selection algorithm was implemented at scale in the main tax audit offices in Senegal in the years 2018-2020. Each January, the tax administration sets an annual audit program that in theory

[1] International organizations hence tend to recommend reforms that limit the amount of discretion given to bureaucrats (Cordova-Novion and Sahovic, 2010; World Bank, 2005)

should be fully conducted by the audit offices during the year. We refer to the audit program for a tax office and year as an audit *list*. To allow for an experimental evaluation of the algorithm, half of the cases on each list were selected by tax inspectors (at discretion), and half were selected by the risk-scoring algorithm. Within each list, the ordering of cases was randomized to minimize the possibility that cases were treated differently due to how they were selected. By design, inspectors knew how a case was selected. Overlap between the inspector selection and algorithm selection was possible but limited in practice.

To quantify the relationship between the case selection method and audit outcomes, we run a horse race between algorithm-selected audits and inspector-selected audits within each list. Concretely, we use an OLS regression of audit outcomes on audit list (i.e. office $\times$ year) fixed effects and an indicator variable for the selection method. The outcome measures come from three different data sources. First, we digitized the universe of audit results records which contain information on the audit start and end date, inspectors working on the case, infractions uncovered, and additional tax to pay. Second, we surveyed a subset of audited firms to capture firms' perceptions of the audit process and corruption outcomes. Finally, we use information on inspector characteristics, motivation and selection methods from a survey conducted before the start of the intervention. Importantly, historical audit outcome data was not available in digital format at the audit selection stage. Indeed, in lower-income countries, audit outcome data is often kept in paper format, and usually stored outside of the general tax administration IT system. The algorithm we use is hence not a trained machine-learning algorithm, but rather one built on economic concepts and logic, exploiting the available data as well as possible.

We document three sets of results. First, we examine audit outcomes. Contrary to the protocol requiring full program implementation, we find that the overall execution rates of the annual audit program are low. Only 53% of selected firms are actually audited. Inspectors have a preference for conducting cases they (or other inspectors) selected. The execution rate is 18 percentage points lower for algorithm-selected audits. On average, algorithm-selected firms are different in terms of observable characteristics: they are smaller, less profitable and older than inspector-selected firms. Yet even after we control for firm characteristics, an audit execution gap of 14 percentage points between the two audit case types remains. Conditional on implementation, almost 90% of audits uncover tax evasion and the detection rate does not vary significantly between inspector-selected and algorithm-selected cases. However, conditional on detection, inspector-selected cases yield higher assessments of taxes evaded. Our results are robust across subsamples and time periods. A silver-lining for the use of an algorithm is the fact that algorithm-selected firms are less likely to dispute the audit result, a result that is robust to controlling for firm characteristics. The lower dispute rate reduces but does not close the gap in confirmed evasion amounts between inspector-selected and algorithm-selected firms.

The second part of our analysis puts the audit process under the microscope. Using our uniquely detailed administrative data, we find that algorithm-selected audits use less manpower, are implemented faster according to various measures, and are conducted by more junior teams. The implementation of algorithm-selected audits is evaluated less positively by firms, as captured by firms' perceptions of auditors' professionalism, competency and efficiency. We can reject that the incidence of corruption is lower in algorithm-selected cases. Combining the amount of evasion uncovered and measures of audit costs into a measure of productivity, we find that algorithm-selected audits are on average significantly less productive than inspector-selected audits. However, the productivity distributions between the two audit case types have a substantial overlap, suggesting there are many algorithm-selected audits that are more productive than the marginal inspector-selected audit.

Finally, we examine the reasons for the execution gap between inspector-selected and algorithm-selected audits and the predictors of audit implementation. We document that the risk score predicts audit implementation and detected evasion, which confirms that it contains information relevant for audit outcomes. We also show through a cross-randomized information treatment in the desk audit program that providing inspectors with additional information, i.e. risk flags and excel sheets with the microdata, does not increase their likelihood of implementing a (desk) audit. We then use the sample of inspector-selected cases to train a random forest algorithm to predict audit execution based on firm characteristics and outcomes reported in previous tax returns. Extrapolating and running the same random forest model on algorithm-selected cases does not allow us to explain the lower execution rate among algorithm-selected audits. Indeed, the model predicts an execution rate for algorithm-selected audits that is similar to the rate for inspector-selected audits, and much higher than the realized execution rate for algorithm-selected audits. However, by training a random forest model on all implemented audit cases to predict the detected evasion, and using this model to predict detectable evasion for all audit case on the annual list, we can show that inspectors are strategic in their execution decisions. The distribution of predicted and realized evasion for implemented audit case are very similar, while the distribution of predicted evasion for non-executed cases is shifted to the left. Consistent with strategic choice, the likelihood that a case is implemented increases in predicted evasion.

Overall, our results show that the type of algorithm that we were able to build with the data at hand is not able to outperform the inspectors' selection. This is despite our best efforts to draw on all available digital data and to fine-tune the algorithm in three subsequent years based on our lessons from the context, data and results; and despite the fact that inspectors chose which algorithm cases to implement, hence adding their skill to the algorithm selection. Given our close collaboration with and endorsement from senior management at the tax administration, we are confident that our results are not driven by inspectors' desire to derail the algorithm. In fact, inspectors were highly motivated to use an algorithm for audit selection before the start of the program. In our inspector survey, 80% of the respondents agreed with the statement that audit case selected should be automated and

data-driven. Inspectors have strong financial incentives to uncover evasion, as they receive financial bonuses proportional to the uncovered evasion.

Instead, we interpret our results as showing that inspectors are indeed better than the algorithm at identifying high-evasion cases. Inspectors involved in audits of firms, which we focus on, are among the top civil servants, highly trained, experienced and incentivized. The data that the algorithm runs on, however, is imperfect. For instance, match rates between different datasets are incomplete, key variables are likely measured with noise, and, most importantly, outcome data to train the algorithm was not available at the onset.

Given that a substantial fraction of algorithm-selected cases dominates inspector cases based on productivity, we conclude that, in a developing country context, a risk-scoring algorithm can be used to *supplement* audit selection. However, inspectors should retain discretion over which of the algorithm-selected cases to implement.

A few caveats to our set of results are in order. Our empirical design was intended to capture the difference in audit outcomes between algorithm-selected and inspector-selected cases. The baseline hypothesis was that the likelihood that an audit detects evasion and the detected evasion amount could be increased. As such, our design does not provide direct evidence on *optimal* audit selection, which would take into account the evasion uncovered to audits, the real and deterrence effect of audits on audited taxpayers in the medium term, and the potential real and compliance spillovers of audits on non-audited firms.

The rest of the paper is organized as follows. After a review of the relevant literature, Section 2 presents the context and data. Section 3 presents the design of our intervention. Section 4 presents results on audit outcomes, Section 5 presents results on the audit process, and Section 6 reviews potential mechanisms. Section 7 discusses policy implications and concludes.

1.1 Related Literature

Our work contributes to a recent but rapidly growing literature on how algorithms can enhance or substitute for human decision-making (see Table A.1 for an overview). Machine-learning algorithms have been shown to improve decision-making and welfare in a variety of settings, including judicial bail decisions (Kleinberg et al., 2017), the targeting of health, safety and water quality inspections (Johnson et al., 2023; Glaeser et al., 2016; Hino et al., 2018), hiring decisions for teachers and police officers (Chalfin et al., 2016), and the targeting of tax rebates or public credit guarantees to firms (Andini et al., 2018, 2022).[2] Similar to our results, (Stevenson and Doleac, 2024) find that judges retain a large degree of discretion in using algorithmic suggestions for sentencing. These examples are

[2]The evidence is not unequivocally positive, as algorithms can exhibit biases (Obermeyer et al., 2019). Some tasks, such as providing employee feedback, are more successfully done by humans than by algorithms (Margalit and Raviv, 2024).

exclusively from the US and other high-income countries. Our contribution is to test the performance of an algorithm in a low-income country context. As in previous studies, the central agents in our context (tax inspectors) are white-collar workers. However, the skill difference between these agents and the average worker in the economy is much larger than in previous studies. This can help explain why the agents in our setting perform so well compared to the algorithm.

In the context of tax audit selection specifically, machine-learning algorithms have shown promise in the US (Black et al., 2022) and Italy (Battaglini et al., 2024). In high-income countries, audit outcomes and potential predictors of these outcomes are usually available in digital format. High-income countries can also afford to have a share of their audit program randomized. In lower-income settings, in constrast, audit outcomes are rarely available in digital format, the number of audits is smaller and audits are usually not randomized.[3] In contrast to applications focused on machine learning, we hence test the performance of a risk-scoring algorithm as a potentially more suitable tool for lower-income countries. Our study connects with another experiment in Senegal by Knebelmann et al. (2024) who show, in the context of property taxation, that a simple rule to assess tax liabilities leads to more accurate assessments than bureaucrat discretion. A likely explanation for why our results differ is that the tax inspectors we study are elite government agents engaging in a complex task. In contrast, property tax assessors are relatively low-skilled bureaucrats with limited performance incentives, delivering a more mechanical task which can be automated. Our findings also differ from Haseeb and Vyborny (2022), who show that using a proxy means test instead of discretion in allocating cash transfers to poor households in Pakistan improved targeting and welfare. The decision-making agents in their setting are lower-level elected officials, subject to political pressures and with no particular training in identifying the poor.

In addition, our analysis is relevant for the environmental literature that has studied how the targeting of inspections and the allocation of inspectors to cases affect inspection outcomes. Duflo et al. (2018) show that discretionary targeting lowers pollution by more than random targeting of inspections, as random inspections found fewer extreme polluters than discretionary inspections. Building on this finding, we examine whether data-driven targeting can further improve upon discretionary targeting. Furthermore, Duflo et al. (2013) show that, if inspected firms pick their inspectors (or vice versa, as in our context), a conflict of interest emerges which can significantly bias audit outcomes. Our algorithm may hence influence audit outcomes by removing this potential conflict of interest.[4]

[3] An exception is Pakistan, where a legal vacuum lead to the randomization of audits, which Best et al. (2021) exploit to show that audits had no causal effect on medium-term compliance. This finding contrasts with Advani et al. (2023) who documents a positive compliance impact of audits in the UK, and Kotsogiannis et al. (2024) who find positive medium-term effects for full audits in Rwanda. For audits of individuals, the welfare impacts have been found to be large, especially at the top of the income distribution (Boning et al., 2023). This is consistent with the fact that tax inspectors in Senegal prefer to audit big firms and maximize detected evasion amounts rather than detected evasion rates.

[4] We also connect with another strand of the environmental literature which studies the adoption of new technologies such as satellite imagery to detect infractions, e.g. illegal mining (Saavedra, 2023) and deforestation Assunção et al. (2023).

Finally, our work contributes to two strands of the literature on state capacity and development. One strand has documented the value of information for improving tax compliance (Kleven et al., 2011; Naritomi, 2019; Pomeranz, 2015) and has studied the potential for new technologies to generate and process information for tax enforcement purposes, with sometimes mixed results (Okunogbe and Santoro, 2022; Okunogbe and Tourek, 2024; Brockmeyer and Saénz Somarriba, 2025).[5] Another strand of this literature has focused on the role of bureaucrats in building state capacity (Besley et al., 2022; Finan et al., 2017). This strand has documented the value of discretion or autonomy for bureaucrat performance (Bandiera et al., 2021; Rasul and Rogger, 2018), but also highlighted the risks associated with discretion, e.g. favoritism (Szucs, 2023). For tax collectors specifically, the literature has quantified the impact of performance incentives (Khan et al., 2015, 2019) and the value of local information on taxpayers (Balán et al., 2022; Dzansi et al., 2022). Our study pushes the frontier by comparing the systematic use of observable and digitized data (in the form of the algorithm) with the value of the tax inspectors' own knowledge and experience. Unlike the lower-level tax collectors studied in previous papers, whose performance could be enhanced through the provision of information, the inspectors in our setting outperform even the systematic use of all available information, consistent with their high qualifications and experience.

2 Context and Data

2.1 Tax Policy

Senegal has low tax revenue and experiences widespread tax evasion. According to the World Bank data, tax revenue as a share of GDP is below 20%, as in most other low and middle-income countries. Evasion is substantial, suggesting that improved enforcement could have high returns. Tax gap estimates indicate that at least 23% of potential VAT revenue (approximately 2% of GDP) and 63% of potential income tax revenue (7% of GDP) are evaded (XXX-CITES).

Senegal's tax structure is typical for a lower middle-income country (see Table B.1). The VAT is the largest source of revenue, representing slightly more than 29% of total tax revenue in 2022, followed by income taxes (29% of total tax revenue) and customs duties (15%). The Value Added Tax (VAT) applies monthly at a standard rate of 18%, with a reduced rate for selected activities (e.g. tourism). The Corporate Income Tax (CIT) is paid annually at a rate of 30% of profits or 0.5% of turnover, whichever is larger. Small firms with a yearly turnover of less than 50 million CFA Francs (about 100,000 USD) are eligible for a simplified tax (*Contribution Globale Unique*, CGU) on turnover, which replaces the previously mentioned taxes and has rates varying from 1% to 8% depending on economic sectors and turnover. As in most other countries, firms withhold the personal income tax

[5] Other studies have shown how technology can help governments manage expenditure and prevent leakages and waste (Muralidharan et al., 2016; Banerjee et al., 2020).

(PIT) at source for their employees, also referred to as Paye-As-You-Earn (PAYE).

2.2 Tax Enforcement

The tax administration (*Direction Générale des Impôts et des Domaines*, or DGID) is tasked with enforcing the tax code. Its main available enforcement tool is the annual audit program. Audits are implemented by 26 different tax offices, with taxpayers segmented by size, region and sector (Figure B.1). Audits are either full audits, carried out by a team of inspectors at the taxpayer's premises, or desk audits, conducted remotely using the firm's tax returns and third-party data. Almost all audits are targeted at firms.

At the beginning of each calendar year, audit offices are provided with a target number of audits to complete, and prepare a list of taxpayers to audit. This list is reviewed and approved by tax office management and by senior management, usually with very minor changes.

Before our intervention, the selection of cases for tax audits was discretionary in that it did not follow an explicit rule. No rationale was required for desk audit selection. For full audit selection, some units required a justification, e.g. a review of the firm's audit history or a summary of relevant indicators such as total sales and the profit margin. However, the justification could take different forms and was not systematically required for all cases. Therefore, the criteria used for case selection varied across units and inspectors.

Figure B.2 illustrates the steps in the audit process. After examining a case, inspectors list the detected irregularities and associated penalties and communicate them to the taxpayer in an initial "notification". They can also request additional information from the taxpayer. Upon receiving the notification, taxpayers have 30 days to respond.[6] The inspector then examines the response and has 60 days to prepare and send a "confirmation" with the confirmed irregularities and penalties and the final amount to pay. The inspector then generates a revenue order for the tax collection unit, which requires the taxpayer to make a payment within ten business days.[7]

2.3 Data

During the last decade, DGID invested in digitizing its tax information and required taxpayers to file electronically. As a result, data availability has expanded substantially. Our study draws on data from three administrative sources and two surveys. The administrative data include the self-assessment declarations filed by taxpayers, third-party reports, and audit reports. The tax declarations and third-party data were used in the calculation of the risk scores and had just been digitized at the onset of the

[6] Failure to respond is interpreted as the taxpayer agreeing with the inspector's findings.

[7] Taxpayers can appeal at the Ministry of Finance or a judicial court, and the appeal may suspend the payment process temporarily.

project. The audit reports were not digitized, and our research team hand-collected and digitized them as part of the project. To study how audits were implemented, we complement the administrative data with a tax inspector survey and a taxpayer survey, designed by the research team. Only aggregated results from these surveys were shared with the tax administration.

Tax Declarations. Table 1, Panel A, provides an overview of the available tax declarations. Our primary sources of information are the declarations for the Corporate Income Tax (CIT), Value Added Tax (VAT), and the Pay-As-You-Earn (PAYE) wage withholding reports for the personal income tax, covering the period of 2014-2019. The CIT data covers about 5,000-7,000 firms per year, and the VAT data about twice as many.[8] The PAYE data allows us to calculate the number of employees and the aggregate wage bill for each firm. Around 2,000 firms file tax under the simplified regime *Contribution Global Unique* (CGU), which is available only to firms with a yearly turnover below 50 Million FCFA (approximately 80,000 USD). Less than 150 financial institutions pay the *Taxe sur les Activités Financières* (TAF), a VAT-substitute for the financial sector.

Third-Party Reports. Table 1, Panel B, describes the information about taxable transactions and activities that we obtain from third parties. Imports and exports are recorded by the customs authority, procurement from state institutions is recorded by the treasury, and firm-to-firm transactions are recorded in VAT annexes that firms file since 2017.[9] While these data are provided at the transaction level, we aggregate them at the firm-year level to merge with the tax declarations.

Audit Reports. We collect audit process and results data in two ways (Table 1, Panel C). First, we digitize all audit result reports for 2017-2020. The reports cover all process steps from audit announcement to notification, confirmation, and payment request. They contain the name(s) of the inspector(s) who conducted the audit, the taxes verified in the audit, infractions detected, evaded amounts, applicable penalties, and the dates of each step in the audit process. In addition, we asked inspectors to report audit information in an excel sheet pre-filled with their list of audit cases. These excel files contain information on audit cases that is not directly observed in the administrative audit reports, such as the number of days that an inspector spent working on a case.

Tax Inspector Survey. In 2017, prior to our intervention, we conducted a detailed survey among all available tax inspectors involved in conducting audits, capturing information about their demographics, employment history, perceptions of the audit function, methods for audit selection, and use of different sources of information. The survey data covers 97 inspectors.[10] Descriptive statistics are presented in Section 2.4.

[8] The number of VAT filers is higher than the number of CIT filers because self-employed individuals and unincorporated firms may file VAT but not CIT.

[9] All VAT filers above xxx of turnover have to file annexes that detail all transactions.

[10] This includes almost all of the 120 taxpayers inspectors in the country, except those who are not working in an auditing function anymore and were thus not interviewed for the survey. XXX of these inspectors joined the audit teams after 2017 and were surveyed after our intervention started. All survey results shown in this paper are robust to excluding these XXX inspectors.

Taxpayer Survey. After the completion of the 2018 and 2019 audit programs, we surveyed 742 firms in the Dakar region, most of which had been selected for audit as part of the program. We conducted the survey in two waves, from October to December 2020 and from March to May 2021. The survey allowed us to elicit taxpayers' views on the audit process, audit risk, and their general perception of the tax administration. The first survey wave focused on the 2018 audit program. We sampled all full audit cases of our program, a matching number of desk-audit cases, randomly sampled and stratified across offices to preserve the relative distribution of planned audits across offices, and a matching number of non-audited cases, similarly stratified. This yielded a total list of 1226 targeted firms. For the second survey wave, we targeted a total of 936 firms, which are again equally divided into the full audit program cases, desk audit cases and non-audited cases, all selected by stratified random sampling. We asked surveyors to prioritize interviews of full audit cases, and the response rate for these cases was indeed higher (xxx%, vs xxx% for other cases).[11]

2.4 Tax Inspectors

The bureaucrats central to the tax enforcement process are tax inspectors and their managers, who are usually former tax inspectors themselves. Tax inspectors form an elite corps in the public administration and often proceed to hold influential political roles. At the time of writing, both Senegal's president and the prime minister are former tax inspectors. Selection into the tax administration function is highly competitive. Most inspectors have completed a Master's degree from the *Ecole Nationale d'Administration* (ENA), and over 60% hold a PhD. According to our survey, the mean and median inspector has ten years of professional experience. Over 70% of inspectors report being satisfied with their work at DGID, and a similar fraction report being motivated by their work.

Tax inspectors are very well-paid. Base salaries were typically 5-6 Million FCFA in the period we study, and stand at 7-8 Million FCFA (approximately 12,000 to 13,500 USD) in 2024. In addition, inspectors receive a bonus that depends on the amount of evasion uncovered in audits and is usually higher than the base salary. The median (mean) bonus for inspectors in our data was 7 million FCFA (10.6 million FCFA). The size of the bonus varies across inspectors (e.g. with seniority), but quantifying the exact strength of incentives is difficult, as it depends on the endogenous assignment of inspectors to cases. There is no institutional reason for incentives to be correlated with the case selection mechanism. Appendix B.3 discusses the bonuses and performance evaluation in more detail.

When selecting cases for audits, inspectors typically start from a long list of cases to consider, roughly half of which are randomly chosen cases. Inspectors then examine the long-listed cases to derive a short list. Figure 1, Panel A, shows the objectives that inspectors pursue in audit selection. The most commonly named objective was 1) a diversity of audits (i.e. a desire to create the perception

[11] The random sample was restricted to firms that filed CIT. In our current analysis, we do not use responses from firms that were not part of our audit program.

of a non-zero audit probability for most taxpayers, which we interpret as an interest in the deterrence function of audits), 2) the detected evasion amount, and 3) the detected evasion rate.[12] Panel C, shows that there is a clear ranking in the types of information used in audit selection. The most commonly used data sources (over 90% of all cases) are the self-assessment declarations, closely followed by third-party data (60-80% of cases), with soft information being the distant third data source (around 40% of cases). Over 60% of inspectors report finding quantitative information more useful than soft information, while xxx% prefer soft information, and the remaining respondents consider the two sources equally important. Most inspectors use excel or other softwares to analyze taxpayers' quantitative information. A few indicators are frequently mentioned in response to the survey's open-ended question about proxies of evasion. These include low turnover compared to other variables, frequent losses, and high VAT credits. Respondents almost unanimously agree that turnover is the single most important variable they consider.

Inspectors are critical of their status quo audit selection method and keen to improve. In our survey, 55% of respondents consider that they do not have sufficient data for audit selection, and almost all want more advice on how to select cases. Concretely, 85% of respondents agree that a more systematic analysis of data would be beneficial to audit selection and 70% think that an automated selection based on risk indicators would be a good idea.[13]

To summarize, tax inspectors are highly-trained elite civil servants, strongly incentivized via bonus payments to detect evasion through audits, and motivated to use new tools to improve audit selection and hence audit performance.

3 Experimental Design

3.1 Intervention: Risk-Based Audit Selection

Given demand for a more systematic audit selection method from tax inspectors and DGID management, we collaborated with DGID to design and introduce a risk-scoring algorithm for audit selection in Senegal. The objectives of the algorithm were a) to ensure that audit selection followed objective and quantifiable criteria, and b) to increase the detected evasion amounts (plus associated penalties and fines) and detected evasion rates.

We faced two types of constraints in designing the algorithm. First, the algorithm needed to be intuitive, transparent, and easy to communicate to policymakers and tax inspectors. Second, the

[12] Panel B illustrates the trade-off inspectors face between maximizing the detectable evasion amount and the evasion rate. Roughly 45% of inspectors prefer to audit larger firms, even if they yield lower detected evasion amounts and evasion rates, but a similar fraction of inspectors trade off these two objectives.

[13] This is consistent also with the fact that inspectors spend a non-trivial share of their time, 27%, to decide which cases to audit, while they spend 52% of their time on the actual audits, and the remainder on administrative tasks and complaints. An "effortless" selection via the algorithm would free more time to conduct audits.

algorithm could not be trained on past audit data, as outcome data was only available for a small and selected subset of audits.[14] The number of potential predictors for evasion was orders of magnitude higher than the number of cases on which to train an algorithm, making a machine-learning approach unsuitable. Given these constraints, we designed an algorithm based on intuitive indicators discussed with and validated by the tax administration. The choice of indicators drew on technical assistance work done by the World Bank in Pakistan and Turkey, best practices shared by the tax administration in Denmark, and feedback from experts at the World Bank and the IMF. Our strategy is broadly applicable in other lower-income countries where similar constraints on risk-scoring design are likely to bind.

The algorithm generated a risk score for each firm that allowed us to rank firms within a tax office. The risk score combines two classes of risk indicators at the firm level: inconsistencies and anomalies. Inconsistencies are *within-firm indicators* that flag taxpayers with inconsistent information across different datasets. For example, an inconsistency arises if the self-reported turnover is lower than the third-party reported turnover constructed as the sum of exports, procurement contracts, and purchases declared by other firms. In contrast, anomaly indicators are *across-firm indicators* that flag outlying behaviors potentially associated with tax evasion relative to the firm's peers. For example, one anomaly indicator flags firms with low profit margin relative to firms of the same economic sector and similar size. Each inconsistency and anomaly is captured in the form of a ratio. We assign "points" from 1 to 10 based on the deciles of the ratio, to take into account the severity of the irregularity. We then aggregate the points using importance weights which we assigned. We weighted inconsistencies higher than anomalies as we were more confident that inconsistencies reflect non-compliance, while anomalies could also reflect temporary economic difficulties or poor management. Appendix C describes the risk scoring algorithm in detail.

3.2 Study Design

To evaluate the performance of the risk-scoring algorithm, we introduced algorithm-selected audits in all audit lists. We asked each office to report the total number of planned audits for the year. Inspectors were then asked to select half this number of cases following their discretionary method, while the algorithm selected the other half of cases. To obtain the algorithm selection, we ranked firms by risk score within a tax office and selected the top cases on the list until the required number of cases was reached. We had access to the inspector selected cases before running the algorithm selection, allowing us to tag overlapping cases selected by both methods.

An alternative research design would have been to randomize the use of the algorithm across tax of-

[14] Accessing and systematically digitizing all paper records of audit results, which came in a variety of formats and were considered sensitive information by the administration, took several years and a substantial investment of time and resources by the research team.

fices. This was infeasible in our context, given political constraints and the small number of tax offices conducting audits (26), most of which joined our program only in the third year. The combination of both selection methods is also the more realistic policy, as fully replacing discretionary selection risks reducing revenue and can breed resistance among inspectors. Instead, with our design, both inspectors and managers were initially enthusiastic to use the algorithm, as evidenced in workshops we held with all tax offices.[15] Our design effectively allows us to study the combination of algorithm selection with inspector selection, as we discuss later in more detail.

We randomized the order of case types displayed on each list and asked inspectors to adhere to the proposed ordering. The randomized ordering was meant to ensure that both case types were treated similarly in terms of timing and effort exerted to complete the audit.[16] Inspectors were provided with a protocol emphasizing the importance of conducting all audit cases with the same rigor.

The intervention started with the four large and three medium taxpayer offices in 2018 and included four regional tax offices in 2019 and 2020. The remaining regional tax offices joined the program in 2020, but we were not able to collect outcome data for these offices. Our evaluation is hence focused on the eight offices with complete data. Each year, the algorithm's indicators and weights were slightly updated. Tables C.2 and C.3 displays the number of cases selected by year, tax office and selection method. By design, the number of inspector-selected and algorithm-selected cases was approximately the same for most years and tax units, with some deviations. We discuss these deviations in Section 4.4 and show that our results are not sensitive to them.

3.3 Additional Treatments for Desk Audits

Desk audits are more numerous than full audits, and are conducted by individual inspectors (or pairs of inspectors), so that there is no endogenous assignment of inspectors to cases. We took advantage of these features to introduce two additional treatments in desk audit lists. First, in the years 2018 and 2019 we selected some cases at random (within the pool of firms of the respective tax unit) and included them in the audit lists. These random cases serve as a benchmark to assess the quality of the algorithm because inspectors were unaware that these cases were randomly selected. Instead, both randomly and algorithm-selected cases were presented as having been selected by the "new method". This approach avoids the possibility that inspectors ignore randomly selected cases, deeming them low-return. At the same time, it may lower inspectors' enthusiasm for algorithm cases if the random cases were consistently worse in terms of audit yields.

[15] Indeed, although our project ran only from 2018 to 2020, DGID still asked us to generate the algorithm selection in 2021 and 2022.

[16] Specifically, the order of cases was randomized across all selected cases for an office/inspector in 2018. In 2019 and 2020, the first case on each list was inspector-selected, with subsequent cases alternating between the two selection methods. Cases were randomly allocated to slots on the list. For desk audits, where each inspector selects their own list, the algorithm-selected cases were randomly distributed across inspectors in a tax office.

Second, the desk audit program is accompanied by an "information treatment" cross-randomized across all case types (inspector, algorithm, and random). The treatment consisted in providing inspectors with case-specific information for two-thirds of desk audits: for one-third of the cases, inspectors received a summary report containing the three main risks flagged by the algorithm (e.g. abnormally low profit rate, turnover lower than third-party reported turnover); for another one-third of cases, they received the same information plus an excel spreadsheet with the firm's tax declarations and third-party data for the last four years (i.e. the data used by the algorithm); and for the remaining third of cases they received no additional information. With this intervention, we aimed to test whether providing high-quality, readable information — available in Senegal since the data revolution — improves audit implementation and performance.

3.4 Empirical Strategy

To assess the implementation and performance of algorithm-selected cases, we compare the outcomes of these cases to those of inspector-selected cases within the same audit list. Concretely, we estimate the following model via Ordinary Least Squares:

$$y_{i\ell} = \beta_0 + \beta_1 Algorithm_{i\ell} + \beta_2 Overlap_{i\ell} + \beta_3 Random_{i\ell} + \gamma_\ell + \varepsilon_{i\ell}, \quad (1)$$

where $y_{i\ell}$ is the outcome of an audit for case i selected in audit list ℓ, the dummies $Algorithm_{i\ell}$, $Overlap_{i\ell}$ and $Random_{i\ell}$ indicate algorithm-selected, overlap and randomly-selected cases, γ_ℓ denotes the list fixed effects and $\varepsilon_{i\ell}$ is a conditional mean zero error term. The audit lists are tax office-year specific for full audits, and can be inspector-year specific for desk audits. The dummy for random selection is applicable only for desk audit cases. Note that overlap cases appeared on the lists sent to inspectors as inspector-selected cases. The list fixed effects allow us to control for any office and year-specific factors, such as the identity of the office manager, the fact that more experienced inspectors are more likely to be assigned to the large and medium taxpayers offices, or that inspectors in regional offices spend a larger share of their time on non-audit administrative tasks, such as taxpayer service and information campaigns.

To perform inference, we cluster the error term at the level of the audit list, allowing for the correlation of unobservable characteristics within a tax office (for full audits) and within the cases selected by a single inspector (for desk audits).

The main coefficient of interest is β_1, which captures the difference in audit outcomes between algorithm-selected cases and inspector-selected cases, the omitted category. A fair comparison between the two selection methods requires that a) the number of cases selected with both methods is equal, b)inspectors' performance incentives were the same for both case types, and c) the resources available to inspectors to conduct the audits did not vary systematically by case type. We have no

reason to believe that conditions b) and c) are violated in our context. Our study design attempted to ensure that condition a) is met. While there are some deviations from this condition, we demonstrate below the robustness of our results to additional tests for audit lists that did not meet the condition.

Nonetheless, any difference in audit outcomes between algorithm-selected and inspector-selected cases we observe is due to a combination of the change in type of taxpayer selected and a possible change in inspector skill and effort. For full audit cases, office management assigns inspector teams to audit cases. Although we emphasized that all audit cases should be treated in the same way, we cannot preclude the possibility that managers take into account the case type when assigning inspectors with different skill sets to cases. We investigate this further below. Similarly, we cannot preclude that inspectors varied the degree of effort exerted by case type.[17]

3.5 Descriptive Statistics on Selected Firms

Table 3 examines the characteristics of firms selected for audit. Panel A compares the characteristics of firms selected by inspectors to the general population of firms. Inspector-selected firms are orders of magnitude larger than the average firm appearing in the administrative tax data (column 1), more likely to engage in import or export (column 4), and slightly though not significantly more profitable (column 3). Inspector-selected firms do not significantly differ from other firms in their location or age (columns 5-6).

Panel B compares firms selected by the algorithm to those that were selected by inspectors, using equation 1. Algorithm-selected firms are smaller than firms selected by inspectors, though still substantially larger than the average formal firm (column 1).[18] They also exhibit a lower profit rate than inspector-selected firms, but are more likely to engage in international trade (columns 3 and 4). The difference in profitability appears by construction, as one of the indicators for our risk score is low profitability compared to other firms in the same sector and size group. Our survey data also indicate that algorithm-selected firms have a higher share of sales in cash than inspector-selected firms (column 8). The algorithm is blind to this firm characteristic unobserved in the administrative data, but inspectors may be able to asses or proxy for it. A high perceived share of cash sales may lead them to avoid auditing a firm, as it would be difficult to detect and document tax evasion on cash sales.

4 Audit Outcomes

We now examine the association between algorithm selection and audit outcomes. First, we consider the association between the selection mechanism and audit execution, given that only a part of the

[17] Taxpayers were not systematically informed how they were selected for audit, and we have no evidence that inspectors shared this information with them. We thus do not expect a change in taxpayer behavior.

[18] Figure C.1 provides a graphical illustration of the gradient between firm size and audit selection within tax offices.

programmed audits were implemented. Second, we analyzeevasion detection, evasion amount, and taxpayer dispute. Third, we examine results for the resource cost of audits.

4.1 Audit Execution

Based on the inspectors' audit reports, we can observe which cases on the list were started and which ones were not. We consider an audit as executed if the case was started, even if it did not lead to a fine. Overall, only 53% of cases selected for full audits and 33% of cases selected for desk audits are actually audited. The implementation rates vary across years but are not lower in later years. Hence, the incomplete implementation is not due to the fact that some audits are implemented in later years than they were scheduled. We consider all audits implemented until the end of 2021, and merge this information with audit program lists for 2018, 2019 and 2020.

Table 4 summarizes the main results. Full audit cases selected by the algorithm are 18 percentage points less likely to be implemented than inspector-selected cases, corresponding to a 34% reduction compared to the mean implementation rate (column 1). "Overlap" case, which were selected by both the inspectors and the algorithm (but which appeared as "Inspector selected" on the lists) were 15 percentage points more likely to be implemented than cases selected only by inspectors.[19]

The results for the desk audit program are qualitatively similar, but the differences are much smaller (columns 2 and 3). Algorithm-selected desk audits were 4 percentage points less likely to be implemented, which corresponds to 13% of the mean. The point estimate is statistically significant only when we control for inspector fixed effects (column 3). Overlap cases are again more likely to be implemented, though the coefficient is not statistically significant. Random cases are no less likely to be implemented than algorithm-selected cases, likely because the two case types appeared in the same way to inspectors (as selected by "new methods") and the expected return of a case might not be clear to inspectors.[20]

The reduction in the probability of conducting audits is our main result. It means a case is less likely to be looked at when it is selected by the algorithm than when it is selected by the inspectors themselves. This result may be explained by the fact that inspectors inherently prefer cases chosen at discretion and devote more effort to these cases. On the other hand, inspectors might simply sort cases based on observable characteristics regardless of their selection method, and algorithm cases appear less interesting by that measure. We discuss these two hypotheses in section 6. In what follows, we

[19]On the other hand, replacement cases, which appeared at the bottom of the list sent to inspectors and which were explicitly marked as replacements for algorithm-selected cases that turned out to be void (e.g. taxpayers that had become inactive or were not reachable) were 47% less likely to be implemented than other algorithm-selected cases (coefficients not shown). These cases were labeled as replacement cases (rather than algorithm x replacement) on the list sent to inspectors, but DGID audit lists did not previously include replacement cases. So it is reasonable to assume that inspectors perceived these cases as algorithm-selected cases that were lower-ranked in the priority given their role as a replacement.

[20]Replacement cases are significantly less likely to be implemented but the magnitude of this effect is smaller than for full audits. This is because the share of replacement cases in total cases is larger for desk audits than full audits.

present evidence of how the cases selected by the algorithm differ in terms of several other outcomes, conditional on execution.

4.2 Audit Yield

Table 4, columns 4-6, show the relationship between the selection method and the probability of evasion detection, conditional on implementation.[21] Overall, 89% of full audits and 74% of desk audits detect some evasion. For full audits (column 4), algorithm-selected cases are 3.7 percentage points more likely to detect evasion, and overlap cases are 8.1 percentage points more likely. These effects are economically meaningful, but only the point estimate on overlap cases is marginally statistically significant.

For desk audits (columns 5 and 6), algorithm-selected cases are 4.7 percentage points less likely to detect evasion, and overlap cases are 7.4 percentage points less likely to do so, although only the latter point estimate is marginally statistically significant. The point estimates on random and replacement cases are relatively precisely estimated zeros, suggesting that inspector selection does not perform better than random selection. Overall, these results indicate that any conducted cases is similarly likely to detect evasion for full and desk audits. However, conditional on positive detection, we find that algorithm-selected full audits detect significantly smaller amounts of tax evasion plus fines (columns 7-9), though this result is not significant for desk audits once we control for inspector fixed effects (column 9).[22]

To summarize, we find that algorithm-selected audits are less likely to be implemented, but similarly likely to yield a detection of evasion when implemented. Conditional on detection, algorithm audits are associated with slightly smaller amounts of evasion, in the case of full audits only. In general, results are starker for full audits and much smaller for desk audits. This is consistent with the fact that full audits are more costly and involve higher stakes for both the administration and the taxpayer.

4.3 Dispute of Audit Outcomes

Once an audit concludes with a detection of tax evasion, the taxpayer receives a notification and the possibility to dispute the audit results. After dispute and negotiation with the tax administration, a confirmation of the audit results is issued. The confirmed amount of evasion is usually lower than the notified amount. This may indicate errors in how the audit was conducted, uncertainty over the exact amount evaded, or collusive behavior between taxpayers and inspectors to alleviate the penalties.

Algorithmic case selection might reduce uncertainty and collusive behavior. To test this hypothesis,

[21] Intent-to-treat analysis that do not condition on audit implementation are shown in Table D.1.

[22] Similarly, we do not detect statistically significant differences in the detected evasion rate, measured either as a share of liability or as a share of the firm's mean turnover across several years, by selection method. We show these results in Table D.3.

we express the confirmation amount as a share of the notification amount, and plot its distribution for inspector- and for algorithm-selected cases, separately for full and desk audits (Figure 3).[23] For most audits, the confirmation amount is substantially lower than the notification amount. The confirmation amount matches the notification amount in only 22% of completed full audits and 29% of desk audits. For full audits, the mean and median of the confirmation/notification share are 40% and 18% respectively. For desk audits, these figures are 62% and 58%, indicating slightly less pushback against the notification amount. For full audits, the figure suggests that inspector-selected cases are subject to more pushback against the notification amount than algorithm-selected cases.

Table 5 shows these results in regression format. Among full audits for which both a confirmation and a notification is present in our data, algorithm audits are 12 percentage points more likely to have matching confirmation and notification amounts.[24] This difference in dispute also has an impact on the final difference in the detected evasion amount between the algorithm and inspector-selected audits. For full audits, algorithm audits have a 34% lower notified amount and only a 19% lower confirmed evasion amount than inspector-selected audits (columns 3 and 5).[25] These results suggest that the lower likelihood of disputes for algorithm audits might improve audit outcomes from a revenue raising and human resource perspective.

4.4 Robustness Tests

We now examine the robustness of our results across subsamples. This allows us to show that our results are not driven by shortcomings in the implementation of our intervention, and that there is little heterogeneity in the results across tax offices and time periods.

First, we tackle the issue that the number of inspector-selected and algorithm-selected cases in a list sometimes deviated from the intended 50-50 split. The most important deviation occurred in the Large Taxpayers Office, where in 2019 and 2020 the tax administration decided that 30% of the full audit cases would be selected by the algorithm, instead of 50%. Other deviations occurred when managers adjusted the number of cases selected at discretion slightly upwards or downwards. If both inspectors and the algorithm had ranked cases by perceived risk and had selected cases following the rank order, the average riskiness of selected cases would be higher when a smaller number of cases is selected. We hence rerun our results for the following sub-samples: a) lists with an identical number of inspector-selected and algorithm-selected cases (as they appear to the researcher), b) lists with an identical number of inspector-selected and algorithm-selected cases (as they appear to the tax inspectors), and c) lists in which the number of algorithm cases is weakly smaller than the number of inspector-selected cases. The latter subsample is less restrictive than the first two, but also less clean

[23] The figure excludes cases with no confirmation, which represent about a third of these cases.

[24] The result is very similar when we control for firm size.

[25] For desk audits, there are no statistically significant differences in dispute of outcomes between inspector and algorithm-selected audits.

from an identification perspective, as it may give the algorithm cases an advantage.[26]

Second, we recognize that there is wide heterogeneity across tax units and years. The Large Taxpayer Unit has the best inspectors and the lowest number of firms per inspector. Although large firms are more complex than smaller firms, inspectors are more familiar with them. Large firms are also expecting an audit every four to five years. The association between audit selection method and outcomes may thus be different in the LTU than in other tax offices which concentrate many more firms per inspector. We hence rerun our analysis excluding the Large Taxpayers Unit.

Third, the experiment's implementation differed in each of the three years. The year 2018 was the first one, and inspectors may have felt the cost of the novelty of the algorithm. In 2020, both firms and tax inspectors were affected by the pandemic. The 2019 implementation is the cleanest from a research perspective, as we carefully prepared and tightly monitored the execution of the intervention, including by asking inspectors to report audit execution steps in pre-filled excel files. We thus estimate the results for the 2019 program only.

Figure 2 depicts the coefficients on the algorithm selection indicator from Equation 1 for the three main outcomes: the probability of starting the audit, the probability of detecting positive evasion, and the amount of detected evasion (plus fines, in log points). Panel A is for full audits, and Panel B for desk audits. The two panels show the baseline coefficient with two types of standard error computation, and the various subsamples as discussed above and in the legend.

The figure shows that the point estimates are stable across subsamples, and that the findings are qualitatively and quantitatively robust. Algorithm cases are less likely to be conducted, no more likely to detect evasion (though point estimates are always positive), and conditional on detection,uncover significantly smaller evasion amounts. For desk audits, however, the point estimates are almost always small and statistically indistinguishable from zero. The negative association between algorithm selection and audit execution is only present in the full sample, but disappears when we limit the analysis to lists with more comparable numbers of algorithm and inspector-selected observations. Our inability to detect statistically significant differences for desk audits is not because of a lack of power. In fact, the number of desk audits is higher and the variation of outcomes among these cases smaller than for full audits. Our experimental design also allows a within-inspector comparison.

[26]The difference between sub-samples a) and b) comes from a methodological change over time in how overlap cases (chosen by both the inspector and the algorithm) were dealt with. In 2018, we ran the algorithm to select a number of cases equal to the number of inspector-selected cases, allowing for overlap cases. The overlap cases appear to inspectors as inspector-selected cases. Thus, the number of "pure" algorithm cases appearing on their list is weakly lower than the number of inspector-selected cases. In 2019 and 2020, we corrected for this by adding an additional algorithm case for each overlap case. Overlap cases still appear on the inspectors' lists as inspector-selected cases.

5 The Audit Process

5.1 Time and Human Resource Costs

We next examine how the length of audits and human resource allocation varies conditional on the audit selection method. Table 6, Panel A show the results for the number of inspectors and proxies for audit length, using the sample of implemented audits. The most accurately captured measure of audit cost is the number of inspectors working on an audit, which is reported in the administrative data. Column 1 shows that algorithm-selected full audits are composed of teams that are 10% smaller than the average team of 2.9 members. In addition, teams for algorithm-selected audits have lower seniority, as measured by their age and length of service (Table E.1). These differences in team composition are consistent with the fact that, for desk audits, we find that inspectors with above-median experience have lower execution rates and are less likely to favor the use of an algorithm for audit selection (Table E.2). We do not detect any correlation between overall list-level execution rates and algorithm execution rates (Figure E.1).

Table 6, Columns 2-6, show the association between audit selection and three different measures of audit length. Overall, algorithm-selected full audits are shorter, according to three measures of audit length. In column 2, the outcome is the audit duration as reported by firms subject to a full audit.[27] The mean duration is 28 days, and algorithm-selected audits are nine days faster on average. In columns 3-4, we use as outcome the audit duration as measured by the difference between the date of notification and the audit start date. This measure captures the length of an audit but not the time inspectors actively spend working on the case, as they typically work on several cases simultaneously.[28] The average full audit takes 157 days until notification. Algorithm-selected audits are completed 25 days faster. For desk audits there is no statistically significant difference in speed between algorithm and inspector-selected audits.

Finally, in columns 5-6, the outcome is the self-reported number of days that inspectors work on a case. For full audits, we multiply the reported number of days by the number of inspectors working on a case. Full audits require 188 inspector-days while desk audits require only 9 days on average. Although the point estimates on the algorithm dummy are negative, there isnt a significant difference between algorithm and inspector cases, likely due to a reduced sample size and to noise in the outcome. For full audits, the outcome is only reported for 51 cases, and it is unclear whether inspectors were referring to their own time investment or to that of the full team.

Audit duration measures in columns 2-6 are not consistently available for all executed cases (for

[27] The sample includes firms that we can observe as having been subject to a full audit in our administrative data, although the question asks firms about the duration of the *last* full audit, so they may or may not be referring to the one that we have in our data.

[28] As discussed above, a dispute of audit results can extend the duration between notification and confirmation.

example self-reported days working on case was only filed by a subset of inspectors). However, we investigate whether the probability of reporting the data differs across selection method and found no difference, as shown in panel B of Table 6.[29]

In addition to the amount of detected evasion, audit reports indicate the specific year of the infraction and a short description.[30] To test if the scope of inspections varied by selection method, we run regressions using as outcomes the number of years covered in the audit report, the number of detected infractions, and the share of fines to evasion amount, which indicates the severity of detected infractions. The results show that in algorithm cases, inspectors systematically reported fewer years with any infraction and fewer infractions, both for full and desk audits (Table E.3, columns 1-4). All differences are highly statistically significant. There is little evidence for systematic differences in the fine-evasion ratios, as indicated both by the small coefficients and the fact that the point estimate for full audits is not statistically significantly different from zero (Table E.3, columns 5-6).[31]

5.2 Audit Productivity

Given the large differences in audit outcomes and audit process measures, the question of whether algorithm-selected audits may be more cost-effective than inspector-selected audits arises. Table 7 shows the results for our main specification in Equation 1, but uses audit productivity measures as outcomes, as per the column titles. We divide the log uncovered evasion by the different audit cost measures presented in the previous section: the number of inspectors working on a case (columns 1 and 2), the number of days from opening to closing a case from the administrative data (3 and 4), the audit length in number of days from the survey data (5 and 6) and the self-reported number of days the inspectors worked on the case (7 and 8). In all estimations except in column 8, the association between algorithm selection and audit productivity is negative. The coefficient on algorithm selection is statistically significant when using the number of inspectors as a cost measure, or the number of days from the administrative data (for full audits). This means that, although algorithm audits use fewer resources, the difference in evasion is larger than the difference in audit cost, leading to lower productivity for algorithm audits.

However, the averages captured by the regression coefficients mask substantial heterogeneity across cases. Figure E.2 shows distributions of the three main audit productivity measures, for algorithm-selected cases and for inspector-selected cases. The productivity distribution for inspector-selected cases is shifted slightly to the right for all measures. For our preferred measure of productivity — the evasion amount divided by the number of agents working on a case — the two distributions are also statistically significantly different as per the Kolmogorov-Smirnov test statistic, consistent with the

[29] The results in Table 6 are almost unchanged when we control for turnover.

[30] By Senegalese law, an audit can investigate tax declarations up to four years prior to the audit start date.

[31] Table E.4 shows that there are no statistically significant differences between algorithm and inspector-selected audits on a range of other outcomes that capture the severity of the detected infractions.

results in Table 7. Yet there is substantial overlap between the two distributions, suggesting that a large share of algorithm-selected cases have higher productivity than the marginal inspector-selected case. This rationalizes the fact that implementation of the inspector-selected set of audits is incomplete and some algorithm audits are implemented.

Audit productivity is positively correlated with firm size, consistent with the preference given to larger firms in audit selection and execution (Appendix Figure E.3).

5.3 Taxpayers' Perceptions

An alternative measure of inspectors' effort and performance comes from the taxpayer's point of view. We surveyed around 600 firms that had been selected for audits, and asked them about their perceptions of interactions with the tax authority. The objective was to observe the taxpayer's perspective of the auditing process, in terms of efficiency and incidence of corruption, as well as their beliefs concerning what behaviors might trigger an audit and the capabilities of the tax administration.

We construct two indices to capture respectively the perceived efficiency of audits and the perceived incidence of corruption of audits. Each index is built by combining three questions following the procedure in Anderson (2008). To measure efficiency, we combine the questions on taxpayers' assessment of the auditor's (i) technical knowledge, (ii) efficiency during the audit, (iii) capacity to uncover all evaded taxes. For corruption, we combine the questions on the perception of (i) the auditor's dishonesty during the audit, (ii) the frequency of bribes paid to inspectors (among firms similar to the respondent's), (iii) the existence of a preferential tax audit treatment for firms connected with the tax administration.

Obtaining answers to questions on the incidence of corruption is challenging. Yet, taxpayers report that bribes are paid in 16% of audits (64% response rate), and 27% of respondents think that connected firms receive preferential treatment (89% response rate). Inspectors obtain an average score of 6.4/10 on honesty, lower than their score of 7.3/10 on knowledge, but equal to their score on efficiency.

Table 8 shows the regression results for the taxpayers' responses, conditional on the selection method that led to their audit. Panel A uses the sample of all interviewed firms that were also in the audit selection and reported a recent audit, while Panel B conditions on firms that were audited as per our audit reports.[32] Columns (1) to (3) show results of taxpayers' evaluation concerning the efficiency of the audit. For full audits (column 1) firms selected by the algorithm report a significantly lower audit efficiency score, compared to firms selected by inspectors. However, the coefficient is less negative when we condition on firms that were actually audited as part of the program (Panel B). In either case, the point estates are not statistically significant. Firms selected by both inspectors and the algorithm

[32]Firms that were actually audited as per the administrative audit data account for about a half of firms reporting having been audited. This could mean either that firms perceive routine interactions with the tax administration as audits, or that the administrative data is incomplete, maybe because inspectors break protocol and fail to report cases with no infraction.

reported higher evaluation grades of inspectors. The lower efficiency of audits does not apply to desk audits, maybe due to the more limited interactions (column 2).

Columns (4) to (6) display the association between audit selection and the perceived incidence of corruption. We do not find evidence of differential corruption perception across selection methods, despite relatively tight standard errors for full audits (Column 4).

The taxpayer survey results are consistent with the notion that inspectors were less invested in algorithm cases. Given that the questions asked about interactions in "full audits", it is not surprising that the coefficients among firms selected for desk audits are mostly small and statistically insignificant.[33] Yet, they do not lend support to algorithmic selected cases facing a different incidence of corruption.

6 Mechanisms

The previous sections show that algorithm-selected audits were less likely to be implemented, and, when implemented, detected lower evasion. At the same time, algorithm-selected audits consumed less human resources, took less time, and their results were less often disputed by taxpayers. We now examine the reasons for why algorithm-selected cases display a lower execution rate. We hypothesize that inspectors are strategic in their choice of cases to implement, prioritizing cases that they expect, based on observable characteristics, to yield a high amount of evasion.

First, we show that the risk score is predictive of tax evasion, thus ruling out that inspectors would not want to use it due to its low performance. Second, we show that providing additional information (risk flags and micro data) for algorithm audits does not increase their execution rates, suggesting that inspectors' hesitance to implement algorithm cases is not due to a lack of information. Third, we predict execution based on observable characteristics. We find that only a third of the execution gap between algorithm and inspector-selected audits can be explained by firms' observable characteristics. Instead, we show in the final subsection that inspectors use different criteria when deciding whether to implement an inspector-selected audit versus an algorithm-selected audit, but are highly strategic: for both case types, the execution rate is strongly increasing in the predicted detectable evasion. We conclude that algorithm-selected cases are executed at lower rates because they have a lower expected return on average than inspector-selected cases.

6.1 The Risk Score Predicts Audit Performance

A necessary though not sufficient condition for inspectors to conduct algorithm-selected audits is that the risk score predicts audit outcomes, i.e. the detection of evasion, the evasion amount or evasion rate. To test whether this is the case, Table 9 reruns our main estimation (Table 4, column 1), adding

[33] For the same reason, it is difficult to explain the significant negative coefficient found for corruption among conducted desk audit cases selected by the algorithm (Panel B, Column 5).

the risk score as an additional regressor. We allow the association between the risk score and audit outcomes to differ between inspector-selected and algorithm-selected cases, as by construction the variation in risk scores is high for inspector-selected cases, but limited for algorithm-selected cases, which all have a high risk score. We find that inspector-selected cases with a higher risk score are more likely to be audited (column 1), and are slightly but not statistically significantly more likely to yield a detection of evasion (columns 3 and 4), and that they exhibit a much higher recovered evasion amount (columns 5 and 6). For both full and desk audits, a one standard deviation increase in the risk score is associated with 38% more tax evasion detected. For algorithm-selected cases, there is no detectable association between the risk score and audit outcomes. This is not surprising given that all cases have high risk scores in this sample and there is thus limited variation.

The results for inspector-selected cases, especially the association with evasion amounts, could be driven by the weighting of risk scores with turnover, which mechanically increases the risk score for larger firms. However, the results are almost unchanged when we control for baseline turnover (Table F.1) or use the unweighted risk score as the regressor (Table F.2).[34]

6.2 Providing Information does not Increase Audit Performance

Another factor that may determine whether or not inspectors implement an audit is the information they have about the case. Inspectors may have better information about the cases they selected themselves. They may have limited information about the compliance risks among algorithm-selected cases, and hence struggle to start and implement algorithm-selected audits. To test this hypothesis, we conducted a randomized information treatment among desk audit cases. For about a third of cases, inspectors were provided with the risk flags from the algorithm. For another third of cases, inspectors were provided with both the risk flags and with a user-friendly spreadsheet containing the full information available from the firms' tax declarations and third-party data. For the remaining cases, inspectors were not provided with any additional information by the research team.[35]

To test the effect of the information treatment, we add an indicator variable for the treatment to our main regression specification (Equation 1). We also add an interaction between the information treatment and the algorithm selection dummy, to consider the possibility that the additional information is more valuable for algorithm cases, which inspectors are less familiar with. Table 10 shows that the information treatment did not affect the probability of starting a case, nor subsequent audit outcomes. Columns 1-3 bundle the two treatments into one treatment indicator, while columns 4-6 show the effects of the two sub-treatments "risk indicators only" and "risk indicators and data spreadsheet". Columns 1-3 show that the negative association of algorithm selection with the main audit outcomes

[34]Table F.3 shows that the association between individual risk score components and audit outcomes largely have the expected sign, but given the small number of observations and noise in the data, we cannot detect with statistical confidence which components of the risk score are most predictive of audit outcomes.

[35]Table C.5 shows that taxpayer characteristics were balanced across treatment groups.

remains significant, whereas the interaction with the information treatment has a small and statistically insignificant coefficient. Columns 4, 5, and 6 show that neither of the treatments had any statistically or economically meaningful effect on audit outcomes. This is true for both algorithm and inspector-selected cases.[36] We conclude that a lack of information about algorithm-selected cases does not contribute to explaining low execution rates among these cases.

6.3 Observable Characteristics do not Explain the Execution Gap

We now examine whether the observable characteristics of firms selected by the algorithm help explain why these taxpayers were less likely to be audited. First, we rerun our main analysis of the association between algorithm selection and audit execution rates (Table 4, column 1) after adding firm characteristics as regressors. We show the results in Table 11, Panel A. In the first column, we reproduce our main result on audit execution as a benchmark: algorithm-selected audits are 18 percentage points less likely to be executed. In column 2, we control for firm size (log turnover), which reduces the execution gap between algorithm and inspector-selected cases, but the gap still remains highly statistically significant and large at 13 percentage points.[37] In columns 3-5, we add the profit rate, firm productivity and other firm characteristics, including the distance to the tax office in travel time (as per google maps, weekday average), and the Euclidian distance to the tax office in meters as additional regressors. Adding these regressors does not substantially reduce the remaining audit execution gap, nor does it increase the R2 of the estimation.[38]

While the results suggest that a simple OLS model cannot rationalize the execution gap based on observable characteristics, it is possible that there are non-linearities and interactions in the association between firm characteristics and audit execution. We hence turn to more flexible prediction models. As average execution rates vary substantially across tax offices and years, we first demean both the execution rate and all potential predictors by the list-level average, where a list is a tax office $\times$ year. This allows us to use data from all tax offices but focus on the role of firm characteristics in predicting within-office variation in execution rates. As predictors of execution, we consider three lags of turnover, productivity (turnover/wage bill), the profit rate, firm age and distance to the tax office. We estimate (train) a model to predict audit execution in the sample of inspector-selected cases using three different approaches: a flexible OLS, a double-lasso and a random forest. When using the OLS or double-lasso approach, we discretize all predictors into deciles to allow for non-linearity.

[36] The null result also holds in a simplified estimation where we only add the information treatment dummy to Equation 1, without interacting it with the mode of audit selection.

[37] Similarly, when we use only firms' turnover ranking to predict execution, the simulated choice of case to execute is still far from the realized execution choice, as shown in Table F.5.

[38] Panels B and C in Table 11 show similar results when the outcome variable is the detection of evasion conditional on audit execution (Panel B) or the evasion amount (Panel C). In both cases, the association between algorithm selection and the audit outcome changes little when we add firm characteristics as controls in the estimation.

Concretely, we estimate (train) the following model on inspector-selected cases:

$$\dot{y}_{i\ell} = \beta \dot{X}_i + \epsilon_{i\ell} \tag{2}$$

where $\dot{y}_{i\ell} \equiv y_{i\ell} - \bar{y}_\ell$ and ℓ indicates a tax office by year list.

The results are shown in Panel A of Figure 4, for full audits and in Panel A of Figure F.1 for desk audits. Focusing on the left part of the panel, the black solid line marks the realized execution rate for inspector-selected cases, compared to the tax-office-year mean. The black markers show the average of the prediction using different models, and the box plot illustrates the range of these predictions, when using 100 randomly selected 70% subsamples to re-estimate the model. All estimates within the p5 to p95 range suggest that the execution rate among inspector-selected cases is significantly higher than the office-year mean, and the p5-p95 range is tightly centered around the realized execution rate. This shows that the models perform well when predicting audit execution out-of-sample among the inspector-selected cases. The mean squared prediction errors printed below each box plot also suggest that, at the individual level, the random forest yields more accurate predictions than the OLS and Lasso, consistent with the fact that it allows for interactions among the predictors.

However, when we use the same models (trained on the sample of inspector-selected cases) to predict execution in the sample of algorithm-selected cases, we cannot match the realization. The right part of Panel A shows that predicted execution rates (grey markers) are far higher than realized execution rates (grey solid line). The predicted execution rates are similar to those for the inspector-selected cases, and around 18 percentage points higher than the realized execution rates for algorithm-selected cases. While turnover is an important predictor of execution, and would suggest lower execution rates for algorithm-selected firms which are smaller on average than inspector-selected firms, other firm characteristics work in the opposite direction.[39]

Hence, even a flexible modeling approach based on observable characteristics cannot help explain the audit execution gap between inspector and algorithm-selected cases. Instead, the evidence suggests that inspectors use different (mental) models when deciding whether to audit inspector-selected cases versus algorithm-selected cases. To see this more directly, consider the importance score attached to different predictors in the random forest model. When running the model on inspector-selected cases, the five most important predictors for execution are lagged log turnover, followed by log distance in meters (93% relative importance), lagged profit rate (93%), firm age (89%) and lagged productivity (86%). In contrast, when estimating the same model on algorithm-selected cases, log distance in meters ranks first among the predictors, followed by lagged log turnover (with only 86% relative

[39] We have also reversed the exercise and used the algorithm sample to train the model, applying it then to the sample of inspector-selected cases. In this case, the predicted execution rates in the sample of inspector-selected audits are similar to the realized execution rate for algorithm-selected audits and much lower than the realized execution rate for inspector-selected audits.

importance), lagged productivity (84%), duration of a trip in minutes (84%) and firm age (81%).[40]

6.4 Inspectors' Audit Execution Decisions are Strategic

Finally, we test the hypothesis that inspectors are strategic and maximize detectable evasion when deciding on which audits to implement. To do this, we train a random forest model to predict the amount of detected evasion, in the sample of all cases that were audited. In addition to the predictors discussed in the previous section, we also control for a dummy indicating whether the case was selected by the algorithm. This means we allow for observable characteristics to map into detectable evasion in different ways for algorithm-selected and inspector-selected cases.

We then use the model to predict the hypothetically detectable evasion amount for cases that inspectors chose not to audit. Figure 4, Panel B, plots the results for full audits. For cases that were audited, the distribution of the predicted and realized evasion amounts are very similar, suggesting that the model yields reasonable predictions. The distribution of the expected detectable evasion for non-audited cases, however, is shifted leftwards, indicating that these cases have much lower potential.[41] The figure also shows that the realized execution rate is clearly upward-sloping in predicted evasion amounts.

Hence, regardless of the mental model that inspectors apply to decide on audit execution (which is different for inspector and algorithm cases), inspectors are strategic in their choice of audits to implement, prioritizing cases based on expected detectable evasion. The results also imply that, although the risk score has some power to predict audit outcomes as we showed in Section 6.1, the inspectors are still more skilled than the algorithm in identifying high-evasion cases.

7 Conclusion

This paper assesses the impact of introducing automated selection in a tax enforcement agency. Collaborating with the Senegalese tax administration on an intervention at scale, we designed a selection algorithm based on risk indicators, and introduced automated selection in the audit lists during the years 2018, 2019, and 2020. We then compared the implementation and return of audit cases selected by the risk-scoring algorithm to cases selected by tax inspectors based on a traditional discretionary procedure. We find that algorithm audits were substantially less likely to be conducted relative to inspector-selected cases, especially among full audits. Among conducted audits, they were similarly

[40] Table F.4 shows the point estimates on the predictors in a linear prediction model, similarly evidencing the fact that the size and size-ranking of point estimates are different for inspector-selected cases and algorithm-selected cases. Table F.5 shows that an execution strategy that simply relies on ranking firms by turnover also yields a better match with the realized execution pattern.

[41] This pattern holds also when considering only inspector-selected audits or only algorithm-selected audits (Figure F.2). This figure also shows that, conditional on predicted evasion, the execution rates for algorithm and inspector-selected cases are not very different.

likely to detect evasion, but algorithm-selected full audit detected smaller amounts of evasion. We also study the differences in several characteristics of the audits, showing that, among the conducted cases, algorithm cases used less human resources, were shorter, had less infractions, and had infractions in fewer years.

We test whether the worse performance of the algorithm is due to underlying characteristics of the cases, low levels of information about the selected cases, or inspector characteristics. Using a rich set of observable characteristics, we still cannot rule out that inspectors were less likely to carry out algorithm cases even for similar levels of observable characteristics. We tested the information hypothesis by providing inspectors with additional information in a random set of cases, but the additional information has not impact on audit implementation or other outcomes. Finally, we find that heterogeneity in inspector experience accounts partly for the lower performance of the algorithm, with highly experienced inspectors being less likely to "penalize" the algorithm.

Our study sheds light on the important issue of introducing "good practices" into a government organization, an particularly into an enforcement agency. Despite its initial apparent advantages, the automation of audit selection using evasion indicators faced strong execution hurdles and seems to have suffered from a reduced effort by inspectors in the field. Moreover, even among desk audits, in which inspectors executed algorithm and discretionary audits at similar rates, there was no discernible advantage of using the algorithm to select audits. Our study suggests that, at least in the short run, the reduction of discretion in audit selection fails to produce benefits.

References

Advani, Arun, William Elming, and Jonathan Shaw, "The Dynamic Effects of Tax Audits," *The Review of Economics and Statistics*, May 2023, *105* (3), 545–561.

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

Andini, Monica, Emanuele Ciani, Guido de Blasio, Alessio D'Ignazio, and Viola Salvestrini, "Targeting with machine learning: An application to a tax rebate program in Italy," *Journal of Economic Behavior Organization*, 2018, *156* (C), 86–102.

_ , Michela Boldrini, Emanuele Ciani, Guido de Blasio, Alessio D'Ignazio, and Andrea Paladini, "Machine learning in the service of policy targeting: The case of public credit guarantees," *Journal of Economic Behavior Organization*, 2022, *198*, 434–475.

Assunção, Juliano, Clarissa Gandour, and Romero Rocha, "DETER-ing Deforestation in the Amazon: Environmental Monitoring and Law Enforcement," *American Economic Journal: Applied Economics*, 2023, *15* (2), 125–156.

Balán, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L. Weigel, "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*, March 2022, *112* (3), 762–97.

Bandiera, Oriana, Michael Best, Adnan Qadir Khan, and Andrea Prat, "The Allocation of Authority in Organizations: A Field Experiment with Bureaucrats*," *The Quarterly Journal of Economics*, 2021, *136* (4), 2195–2242.

Banerjee, Abhijit, Esther Duflo, Clement Imbert, Santhosh Mathew, and Rohini Pande, "E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in india," *American Economic Journal: Applied Economics*, 2020, *12* (4), 39–72.

Bank, The World, *Good practices for business inspections : guidelines for reformers*, Washington, D.C. : World Bank Group. http://documents.worldbank.org/curated/en/286811468329950178/Good-practices-for-business-inspections-guidelines-for-reformers, 2005.

Battaglini, Marco, Luigi Guiso, Chiara Lacava, Douglas L. Miller, and Eleonora Patacchini, "Refining public policies with machine learning: The case of tax auditing," *Journal of Econometrics*, 2024, p. 105847.

Besley, Timothy, Robin Burgess, Adnan Khan, and Guo Xu, "Bureaucracy and development," *Annual Review of Economics*, 2022, *14*, 397–424.

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

Black, Emily, Hadi Elzayn, Alexandra Chouldechova, Jacob Goldin, and Daniel Ho, "Algorithmic fairness and vertical equity: Income fairness with irs tax audit models," *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 2022, pp. 1479–1503.

Boning, William C, Nathaniel Hendren, Ben Sprung-Keyser, and Ellen Stuart, "A Welfare Analysis of Tax Audits Across the Income Distribution," Working Paper 31376, National Bureau of Economic Research June 2023.

Brockmeyer, Anne and Magaly Saénz Somarriba, "Electronic Payment Technology and Tax Compliance: Evidence from Uruguay's Financial Inclusion Reforms," *American Economic Journal: Economic Policy*, 2025.

Chalfin, Aaron, Oren Danieli, Andrew Hillis, Zubin Jelveh, Michael Luca, Jens Ludwig, and Sendhil Mullainathan, "Productivity and Selection of Human Capital with Machine Learning," *American Economic Review*, May 2016, *106* (5), 124–27.

Cordova-Novion, Cesar and Tari Sahovic, *Inspections reforms : do models exist*, Washington, D.C. : World Bank Group. http://documents.worldbank.org/curated/en/216181471587685722/Inspections-reforms-do-models-exist, 2010.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India," *The Quarterly Journal of Economics*, 2013, *128* (4), 1499–1545.

_ , _ , _ , and _ , "The value of regulatory discretion: Estimates from environmental inspections in India," *Econometrica*, 2018, *86* (6), 2123–2160.

Dzansi, James, Anders Jensen, David Lagakos, and Henry Telli, "Technology and Tax Capacity: Evidence from Local Governments in Ghana," Working Paper 29923, National Bureau of Economic Research April 2022.

Finan, Frederico, Benjamin A Olken, and Rohini Pande, "The personnel economics of the developing state," in "Handbook of Economic Field Experiments," Vol. 2, Elsevier, 2017, pp. 467–514.

Glaeser, Edward L, Andrew Hillis, Hyunjin Kim, Scott Duke Kominers, and Michael Luca, "Decision Authority and the Returns to Algorithms," 2021.

_ , _ , **Scott Duke Kominers, and Michael Luca**, "Crowdsourcing city government: Using tournaments to improve inspection accuracy," *American Economic Review*, 2016, *106* (5), 114–18.

Haseeb, Muhammad and Kate Vyborny, "Data, discretion and institutional capacity: Evidence from cash transfers in Pakistan," *Journal of Public Economics*, 2022, *206* (C), S0047272721001717.

Hino, M., E. Benami, and N. Brooks, "Machine Learning for Environmental Monitoring," *NAture Sustainability*, 2018, *1*, 583–588.

Johnson, Matthew S., David I. Levine, and Michael W. Toffel, "Improving Regulatory Effectiveness through Better Targeting: Evidence from OSHA," *American Economic Journal: Applied Economics*, October 2023, *15* (4), 30–67.

Khan, Adnan Q, Asim I Khwaja, and Benjamin A Olken, "Tax farming redux: Experimental evidence on performance pay for tax collectors," *The Quarterly Journal of Economics*, 2015, *131* (1), 219–271.

Khan, Adnan Q., Asim Ijaz Khwaja, and Benjamin A. Olken, "Making Moves Matter: Experimental Evidence on Incentivizing Bureaucrats through Performance-Based Postings," *American Economic Review*, January 2019, *109* (1), 237–70.

Khwaja, Muwer Sultan, Rajul Awasthi, and Jan Loeprick, *Risk-based tax audits: approaches and country experiences*, The World Bank, 2011.

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, "Human Decisions and Machine Predictions*," *The Quarterly Journal of Economics*, 08 2017, *133* (1), 237–293.

_ , **Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein**, "Discrimination in the Age of Algorithms," *Journal of Legal Analysis*, 2018, *10*, 113–174.

Kleven, Henrik Jacobsen, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark," *Econometrica*, 2011, *79* (3), 651–692.

Knebelmann, Justine, Victor Pouliquen, and Bassirou Sarr, "TDiscretion versus Algorithms: Bureaucrats and Tax Equity in Senegal," *Working Paper*, 2024.

Kotsogiannis, Christos, Luca Salvadori, John Karangwa, and Theonille Mukamana, "Do tax audits have a dynamic impact? Evidence from corporate income tax administrative data," *Journal of Development Economics*, 2024, *170*, 103292.

Margalit, Yotam and Shir Raviv, "When Your Boss is an Algorithm: The Effect of Algorithmic Management on Worker Performance," *Available at SSRN*, 2024.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar, "Building State Capacity: Evidence from Biometric Smartcards in India," *American Economic Review*, October 2016, *106* (10), 2895–2929.

Naritomi, Joana, "Consumers as tax auditors," *American Economic Review*, 2019, *109* (9), 3031–72.

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," *Science*, 2019, *366* (6464), 447–453.

OECD, *Tax Administration 2023: Comparative Information on OECD and other Advanced and Emerging Economies*, OECD Publishing, Paris, 2023.

Okunogbe, Oyebola and Fabrizio Santoro, "The Promise and Limitations of Information Technology for Tax Mobilization," *The World Bank Research Observer*, 10 2022, *38* (2), 295–324.

_ **and Gabriel Tourek**, "How Can Lower-Income Countries Collect More Taxes? The Role of Technology, Tax Agents, and Politics," *Journal of Economic Perspectives*, February 2024, *38* (1), 81–106.

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

Rasul, Imran and Daniel Rogger, "Management of bureaucrats and public service delivery: Evidence from the nigerian civil service," *The Economic Journal*, 2018, *128* (608), 413–446.

Saavedra, Santiago, "Technology and State Capacity: Experimental Evidence from Illegal Mining in Colombia," *Available at SSRN 3933128*, 2023.

Stevenson, Megan T and Jennifer L Doleac, "Algorithmic risk assessment in the hands of humans," *American Economic Journal: Economic Policy*, 2024.

Szucs, Ferenc, "Discretion and favoritism in public procurement," *Journal of the European Economic Association*, 2023, p. jvad017.

Tables

Table 1: Number of Firms by Data Source

		2014	2015	2016	2017	2018	2019	2020
A Self reported	CIT	5136	5969	6218	6594	6720	7233	0
	VAT	11181	11901	12699	13352	13969	14213	13538
	PAYE	7061	7518	7870	8513	8782	9005	8621
	CGU	1581	1827	2026	2203	2650	2671	2801
	TAF	86	105	112	122	121	111	118
B Third party	Imports	8963	12427	13068	11859	13551	13677	10591
	Exports	1398	1724	1881	1824	1697	1659	1558
	Procurement	809	735	1380	1340	1903	1897	1684
	VAT annexes	6	9	21	805	3606	3209	0
C Audits data	Digitized	0	0	1	3294	2753	2946	3714
	Self-reported (Excel)	0	0	0	102	561	664	51

Notes: This table shows the number of unique taxpayers (firms) by year of available data, in the main datasets used to construct the risk scoring algorithm and analyze its performance. The available data covers the years 2014 to 2020, and the experiment was conducted in the years 2018, 2019, and 2020. The self-reported data (Panel A) include the self-assessment declarations for the Corporate Income Tax (CIT), Value-Added Tax (VAT), Pay-As-You-Earn withheld personal income tax (PAYE), the simplified Senegalese tax for small enterprises called CGU, and the Senegalese tax for financial services enterprises called TAF. The third-party datasets (Panel B) include transactions data from customs (imports and exports), procurement records, and VAT annexes. The audit data (Panel C) come from a) an effort to digitize audit reports from 2017 onwards, covering all audits conducted under our intervention, and b) an excel spreadsheet which the research team requested tax inspectors to fill out to report audit outcomes and audit process information. This table is discussed in Section 2.3.

Table 2: Number of Cases by Selection Method (2018-2020 Programs)

	Full Audits			Desk Audits			
	Inspectors	Algorithm	Total	Inspectors	Algorithm	Random	Total
Large Taxpayer Unit	265	131	396	320	157	60	477
Medium Taxpayer Unit	155	153	308	320	290	135	610
Liberal Professions	46	50	96	235	222	85	457
SME(Regional)	71	73	144	422	401	84	823
Total	537	407	944	1297	1070	364	2367

Notes: Number of selected cases, tax office, audit type, and selection method. The sum of the rows is larger than the total because there are overlapping cases between algorithm and discretion. Random cases and replacement cases are exclusive to desk audits. This table is discussed in Section 2.3.

Table 3: Firm Characteristics of Algorithm vs Discretionary Selection

Panel A: Characteristics of Inspector-Selected Cases

	Tax Declarations				Administrative		Taxpayer Survey		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	log(Turnover)	log(Payroll)	Profit Rate	P(Trade)	Duration Trip	Firm Age	Employees	% Sales in Cash	Audit Frequency
Inspectors' Selection	3.312***	1.401***	0.0100	0.169***	-0.660	0.967*	-5.161	-7.425	-0.00808
	(0.838)	(0.271)	(0.00969)	(0.0230)	(0.435)	(0.460)	(7.050)	(5.186)	(0.151)
N	22573	7430	5922	61238	53655	51992	696	701	640
R2	0.226	0.125	0.0482	0.0227	0.119	0.101	0.0917	0.0229	0.0286
Mean outcome	10.81	14.81	-0.15	0.41	13.98	5.91	31.33	50.25	2.77

Panel B: Comparison Among Selected Cases

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Algorithm	-1.777***	-1.673***	-0.0655**	0.0549	1.216*	1.365*	15.97	12.31**	-0.0503
	(0.524)	(0.362)	(0.0312)	(0.0439)	(0.706)	(0.677)	(16.85)	(5.398)	(0.248)
Inspectors x Overlap	0.0401	1.022*	0.00313	0.0608	2.244	2.313	-98.48	9.948	0.370
	(1.211)	(0.525)	(0.0282)	(0.0801)	(2.001)	(2.113)	(64.31)	(24.13)	(0.947)
N	1186	1021	906	1323	1232	1118	268	277	268
R2	0.122	0.125	0.0872	0.137	0.0488	0.0715	0.278	0.133	0.112
Mean outcome	18.39	16.51	-0.02	0.51	14.23	13.76	44.72	45.70	2.79

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table depicts the OLS regression coefficients of characteristics on selected methods for full audits. Appendix Table C.1 shows the same exercise for desk audits. Panel A uses the sample of selected cases across the three years of the experiment, including only full audit selection. Panel B uses the sample of all firms observed in the administrative data, and the coefficient on Inspectors' Selection indicates that the firm was selected for audit at some point during 2018-2020. For Panel A, columns 1-6, the regression is a panel regression with fixed effects at the list level (year X tax office for full audits, year X inspector for desk audits), and standard errors clustered at the list level. Otherwise, the regression is cross-sectional, with tax office fixed effects and standard errors clustered at the tax office level. The characteristics of the firms stem from three sources. From the tax declarations, we use the log of the yearly declared turnover, the log of the yearly declared payroll, the profit rate, and the probability that the firm has exports or imports. We use the value for the year before the firm was selected for audit. We use data from the firm registry on the firm's age and the distance between its location and DGID, in minutes of travel time (computed using GoogleMaps for a Monday at 3 PM). Finally, we use the taxpayer survey to compute the (self-reported) number of full-time employees, the share of total sales done in cash, and the perceived yearly frequency of full audits. This table is discussed in Section 3.5.

Table 4: Algorithm Selection and Audit Outcomes

	P(Execution)			P(Detection | Execution)			log(Evasion) | Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.183***	-0.0456	-0.0436**	0.0371	-0.0465	-0.0441	-0.636***	-0.252*	-0.160
	(0.0369)	(0.0289)	(0.0208)	(0.0309)	(0.0367)	(0.0303)	(0.216)	(0.139)	(0.131)
Inspectors x Overlap	0.159**	0.0465	0.0423	0.0809*	-0.0744*	-0.0756	0.290	-0.580	-0.397
	(0.0711)	(0.0422)	(0.0365)	(0.0441)	(0.0367)	(0.0464)	(0.304)	(0.346)	(0.340)
Algorithm x Random		-0.00104	-0.00139		0.00284	0.00534		-0.216	-0.211
		(0.0395)	(0.0311)		(0.0417)	(0.0331)		(0.133)	(0.167)
Tax center x Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector x Year	No	No	Yes	No	No	Yes	No	No	Yes
N	944	2731	2731	507	1016	997	453	751	732
R2	0.263	0.227	0.317	0.151	0.324	0.430	0.311	0.292	0.427
Mean outcome	0.53	0.37	0.37	0.89	0.73	0.73	19.29	17.74	17.71

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of the main audit outcomes on the selection method of the case for full and desk audits separately. We show specifications for full audits controlling for tax office x year fixed effects, which corresponds to the selection list for full audits. For desk audits, we show specifications with tax office x year fixed effects and inspector x year fixed effects, the latter corresponding to the list for desk audits. The first outcome (Columns 1 to 3) is the probability that the case on the list was conducted by the inspectors. The second outcome (Columns 4 to 6) is the probability that an audit results in a penalty (initial or final notice), conditional on the case being conducted. The third outcome (Columns 7 to 9) is the log amount of the evasion plus penalties, conditional on the case ending in positive penalty. The evasion amount is obtained from the final notice and complemented with the initial notice when the final notice is not available. The sample includes all cases selected in the 2018, 2019, and 2020 audit programs. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. Standard errors are clustered at the tax office x year levels, and inspector x year levels when inspector fixed effects are included. This table is discussed in Section 4.1. Table D.1 shows intent-to-treat effects, Table D.2 shows robustness to controlling for audit slot fixed effects, and Table D.3 shows the results when using the evasion rate as the outcome.

Table 5: Dispute of Audit Outcomes

	P(Confirmation= Notification)		log(Notification)| Both>0		log(Confirmation)| Both Not. and Conf. >0	
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	0.130***	0.0162	-0.346	-0.161	-0.194	-0.193
	(0.0407)	(0.0515)	(0.204)	(0.173)	(0.258)	(0.236)
Inspectors x Overlap	-0.0321	0.0320	0.665	-0.569	0.542	-0.0602
	(0.0651)	(0.100)	(0.401)	(0.447)	(0.385)	(0.479)
Algorithm x Random		-0.0113		-0.113		0.0541
		(0.0954)		(0.262)		(0.248)
N	264	372	260	346	260	346
R2	0.163	0.238	0.323	0.547	0.189	0.509
Mean outcome	0.16	0.21	19.46	17.96	18.68	17.33

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case, controlling for list fixed effects (year x tax office for full audits, year x inspector for desk audits). The sample includes all cases with **non-missing confirmation and notification**, and for the second and third outcome, the sample is conditional on **both notification and confirmation being non-zero**. Standard errors are clustered at the list level. This table is discussed in Section 4.3.

Table 6: Resources Allocated to Audit Execution

Panel A: Resource Outcomes

	Number of Agents	Duration in Days (Firm Survey)	Days From Start to Conf. (Admin Data)		Days Working on Case (Self-reported)	
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Full audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	-0.222***	-8.659**	-24.92**	12.95	-7.226	-3.500
	(0.0800)	(3.956)	(11.13)	(19.53)	(18.76)	(2.401)
Inspectors x Overlap	0.0207	-26.52	-41.44	32.37		-22.76
	(0.133)	(26.90)	(27.16)	(38.79)		(18.24)
Algorithm x Random				0.296		-2.038
				(18.87)		(1.711)
N	507	214	281	236	51	108
R2	0.233	0.255	0.291	0.390	0.263	0.568
Mean outcome	2.87	28.16	159.09	112.81	188.03	9.06

Panel B: Data Availability

	(1)	(2)	(3)	(4)	(5)	(6)
Algorithm		-0.00802	0.00891	0.0000339	-0.0370	-0.0301
		(0.0193)	(0.0102)	(0.00660)	(0.0372)	(0.0252)
Inspectors x Overlap		-0.0265	-0.0154	-0.000154	-0.00539	0.0121
		(0.0506)	(0.0350)	(0.00420)	(0.0459)	(0.0368)
Algorithm x Random				0.00166		0.0477*
				(0.00249)		(0.0250)
N		507	507	997	507	997
R2		0.214	0.956	0.974	0.491	0.443
Mean outcome	1.00	0.21	0.55	0.26	0.10	0.11

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of resources used in an audit on the selection method of the case. The measures of resources are the number of agents working on the case (for full audits only, column 1) and the audit duration measured in three different ways. The sample includes all cases selected and conducted from the 2018, 2019, and 2020 audit programs, for which we had data on the outcome variables. Panel A shows the results for the outcomes, and Panel B shows the results of regressions for dummies of data availability. Column 1 shows the results for the number of agents involved in the audit, which is obtained from the audit records and is only relevant for full audits because desk audits are typically conducted individually. Column 2 shows the results for audit duration as reported by taxpayers in the taxpayer survey. The survey asked them about their last full audit experience. Columns 3 and 4 use the difference between the date of initial notice and the date of the start of the audit (either verification announcement or information request or start date of audit) as a proxy for the duration of the audit. This measure is available only for cases that had an initial notice. We drop negative values, winsorize the data at the 99th percentile, and control for dummies indicating availability of date of initial notice and final notice. These controls are helpful because the definition of "start date" may vary depending on the availability of information. Columns 5 and 6 measure the duration of the audit using the self-reported number of days worked on a case, reported on a spreadsheet that inspectors needed to fill out during the 2019 audit program. We multiplied the answer by the number of inspectors that worked in the case. Standard errors are clustered at the tax office x year (list) levels, and inspector x year levels when inspector-fixed effects are included. This table is discussed in Section 5.1.

Table 7: Algorithm Selection and Audit Productivity

	Evasion/Number of Agents		Evasion/Duration in Days (Survey)		Evasion/Days Start to Conf. (Admin.)		Evasion/Days Working (Self reported)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	-0.537**	-0.254*	-0.132	-0.379	-0.545*	-0.518	-0.397	0.357
	(0.200)	(0.130)	(0.319)	(0.512)	(0.273)	(0.408)	(0.562)	(0.294)
Inspectors x Overlap	0.310	-0.365	0.997	0.534	0.165	-2.677***		0.482
	(0.292)	(0.339)	(2.623)	(0.852)	(0.459)	(0.617)		(0.306)
Algorithm x Random		-0.108		-0.446		-0.0963		0.177
		(0.169)		(0.675)		(0.338)		(0.381)
N	453	732	95	93	251	154	39	60
R2	0.307	0.442	0.214	0.445	0.221	0.592	0.262	0.411
Mean outcome	18.22	17.47	15.95	14.82	15.04	13.58	13.61	15.91

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of four measures of audit productivity on the selection method of the case for full and desk audits separately. We show specifications for full audits controlling for tax office x year fixed effects, which corresponds to the selection list for full audits. For desk audits, we show specifications with inspector x year fixed effects. The first outcome (Columns 1 and 2) measures productivity as the log of evasion (plus fines) divided by the number of agents involved in the case; the second outcome (Columns 3 and 4) uses log of evasion divided by the number of days the last full audit lasted as declared by the taxpayer in the taxpayer survey (notice that the result for desk audits may not be relevant here); the third outcome (Columns 5 and 6) uses log of evasion divided by the number of days the audit lasted as defined by the date of start and date of notification; and the fourth outcome (Columns 7 and 8) uses log of evasion divided by the number of days the audit lasted as self-reported by the inspector (only available for 2019). The sample includes all cases selected in the 2018, 2019, and 2020 audit programs. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. Standard errors are clustered at the tax office x year levels, and inspector x year levels when inspector fixed effects are included. This table is discussed in Section 5.2.

Table 8: Selection Method and Survey Results

Panel A: Surveyed Firms (Self-Reporting a Recent Audit)

	Efficiency Index			Corruption Index		
	Full Audits	Desk Audits	All Audits	Full Audits	Desk Audits	All Audits
	(1)	(2)	(3)	(4)	(5)	(6)
Algorithm	-0.377	-0.0292	-0.171	0.0389	-0.109	-0.0375
	(0.182)	(0.0326)	(0.105)	(0.0233)	(0.130)	(0.0374)
Inspectors x Overlap	0.467**	0.0783	0.202	-0.548*	-0.257	-0.266
	(0.0743)	(0.0775)	(0.112)	(0.181)	(0.232)	(0.148)
Algorithm x Random		-0.00729	-0.000554		-0.110	-0.0137
		(0.0943)	(0.0737)		(0.155)	(0.121)
N	216	330	546	218	329	547
R2	0.0867	0.0227	0.0220	0.128	0.0658	0.0457
Mean outcome	0.01	0.00	0.00	-0.08	0.05	0.00

Panel B: Only Audited Firms (as per Administrative Audit Data)

	(1)	(2)	(3)	(4)	(5)	(6)
Algorithm	-0.270	0.253	-0.0235	0.190	-0.259*	-0.0681
	(0.343)	(0.120)	(0.134)	(0.126)	(0.0833)	(0.116)
Inspectors x Overlap	0.243*	0.0321	-0.0599	-0.337	0.0182	0.0460
	(0.0652)	(0.175)	(0.121)	(0.319)	(0.186)	(0.0880)
Algorithm x Random	0	0.0908	0.118*	0	-0.182	-0.119
	(1.65e-14)	(0.0864)	(0.0346)	(1.77e-14)	(0.170)	(0.0590)
N	104	150	256	104	152	258
R2	0.0799	0.0566	0.0265	0.171	0.0952	0.0620
Mean outcome	0.02	-0.05	-0.01	-0.15	-0.04	-0.08

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results with fixed effects at the year and tax office level for outcomes from the taxpayer survey. Given the lower number of observations in the survey, we do not run the desk audit regression with year x inspector fixed effects, as is done with outcomes from the administrative data. Robust standard errors are in parenthesis, also clustered at the year and tax office level. The table shows the association between audit selection method and taxpayers' perception of the efficiency of audits (columns 1-3), and of the incidence of corruption during audits (columns 4-6). The sample of surveyed taxpayers is a subsample of the audit cases selected by the yearly audit programs between 2018 and 2020. Panel A shows the results for all firms interviewed in the survey, conditional on them self-reporting an audit in recent years. Panel B is restricted to firms that were actually audited as part of the yearly audit program according to the administrative audit reports. The two outcome are (1) an index of perceived efficiency of audits, and (2) an index of perceived corruption of audits. Each measure is constructed by combining three questions into a normalized index following Anderson (2008). The coefficients are measured in terms of standard deviation of the respective indexes. This table and the survey questions the indices are based on are discussed in Section 5.3. Evidence that respondents' likelihood of answering the questions is not associated with the audit selection method is presented in Table E.5.

Table 9: Association of Risk Score with Audit Outcomes

	P(Execution)		P(Detection|Execution)		log(Evasion)|Detection	
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	-0.295***	-0.0735**	-0.00242	-0.0762*	-1.076***	-0.318
	(0.0656)	(0.0305)	(0.0258)	(0.0399)	(0.307)	(0.213)
Risk score	0.0482*	-0.00130	0.00338	0.00380	0.376**	0.379***
	(0.0278)	(0.0191)	(0.0210)	(0.0219)	(0.163)	(0.141)
Alg. x Risk score	-0.0147	0.0214	0.0118	0.0144	-0.316*	-0.415***
	(0.0296)	(0.0200)	(0.0271)	(0.0206)	(0.170)	(0.136)
N	944	2731	507	997	453	732
R2	0.269	0.318	0.153	0.431	0.323	0.436
Mean outcome	0.53	0.37	0.89	0.73	19.29	17.71

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table shows the correlation of the algorithm risk score with the three main audit outcomes: the probability of starting the selected audit, the probability of detecting evasion among started cases, and the log of detected evasion amount plus fines. The risk score is a continuous variable computed at within "clusters" of firms of similar economic activity and size. We standardize the risk score values at the cluster-year level, i.e., we subtract the mean and divide by the cluster's standard deviation. We do that because the risk score scale changed in the three iterations of the experiment. The coefficients show how the risk score correlates with the outcomes for discretionary cases (non-interacted term) and for algorithm cases (interacted term). The results stem from OLS regressions with fixed effects at the list level (year X tax office for full audits, and year x inspector for desk audits). Standard errors are clustered at the list level. This table is discussed in Section 6.1. Tables F.1 to F.3 show robustness tests for these results.

Table 10: Treatment Effect of Information Intervention

	(1) P(Execution)	(2) P(Detection| Execution)	(3) log(Evasion)| Detection	(4) P(Start)	(5) P(Detection| Execution)	(6) log(Evasion)| Detection
Algorithm	-0.0516	-0.0816	-0.224	-0.0546	-0.0818	-0.222
	(0.0428)	(0.0506)	(0.207)	(0.0430)	(0.0507)	(0.207)
Algorithm x Random	-0.00326	0.00862	-0.244	-0.00317	0.00879	-0.246
	(0.0318)	(0.0342)	(0.171)	(0.0319)	(0.0340)	(0.173)
Information	0.0205	-0.0505	-0.223			
	(0.0376)	(0.0416)	(0.209)			
Algorithm x Information	0.0107	0.0638	0.267			
	(0.0425)	(0.0485)	(0.274)			
Info. (Indicators)				0.0305	-0.0516	-0.0730
				(0.0472)	(0.0506)	(0.241)
Info. (Indicators+Data)				-0.0198	0.00208	-0.295
				(0.0474)	(0.0565)	(0.261)
Alg. x Info. (Indicators)				0.0211	0.0684	0.0745
				(0.0582)	(0.0583)	(0.301)
Alg. x Info. (Indicators+Data)				-0.0114	-0.00891	0.374
				(0.0630)	(0.0611)	(0.360)
N	2136	896	631	2136	896	631
R2	0.317	0.407	0.425	0.318	0.407	0.427
Mean outcome	0.42	0.70	17.66	0.42	0.70	17.66

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table shows the estimation of the information treatment intervention's effect on the probability audit execution, and the correlation of the information treatment with subsequent audit outcomes. Only desk audit cases were used in the intervention, so the sample does not include full audit cases. The treatment was cross-randomized across algorithm and inspector cases in 2018 and 2019, but only used for algorithm cases in 2020. Therefore, we excluded the inspector-selected cases in 2020. Columns 1, 2, and 3 show the results for a specification containing a dummy indicating whether the case was treated. Columns 4, 5, and 6 distinguish between two modalities of the treatment: providing only indicators of risk about the case to the inspectors and providing risk indicators plus a spreadsheet with data on the taxpayers' tax declarations and third-party data. OLS results with fixed effects at the year X inspector level. Standard errors are clustered at the inspector level. This table is discussed in Section 6.2.

Table 11: Explaining the Execution Gap for Full Audits: The Role of Observable Characteristics

Panel A: Outcome P(Execution)					
	(1)	(2)	(3)	(4)	(5)
Algorithm	-0.183***	-0.131***	-0.130***	-0.126***	-0.127***
	(0.0369)	(0.0347)	(0.0341)	(0.0338)	(0.0344)
N	944	944	944	944	944
R2	0.263	0.336	0.337	0.341	0.343
Mean outcome	0.53	0.53	0.53	0.53	0.53
Panel B: Outcome P(Detection|Execution)					
	(1)	(2)	(3)	(4)	(5)
Algorithm	0.0371	0.0410	0.0441	0.0440	0.0432
	(0.0309)	(0.0293)	(0.0293)	(0.0289)	(0.0284)
N	507	507	507	507	507
R2	0.151	0.159	0.165	0.165	0.169
Mean outcome	0.89	0.89	0.89	0.89	0.89
Panel C: Outcome log(evasion)|Detection					
	(1)	(2)	(3)	(4)	(5)
Algorithm	-0.636***	-0.489**	-0.498*	-0.478*	-0.487*
	(0.216)	(0.230)	(0.247)	(0.244)	(0.245)
Turnover	No	Yes	Yes	Yes	Yes
Profit rate	No	No	Yes	Yes	Yes
Productivity	No	No	No	Yes	Yes
Firm char.	No	No	No	No	Yes
N	453	453	453	453	453
R2	0.311	0.367	0.385	0.404	0.407
Mean outcome	19.29	19.29	19.29	19.29	19.29

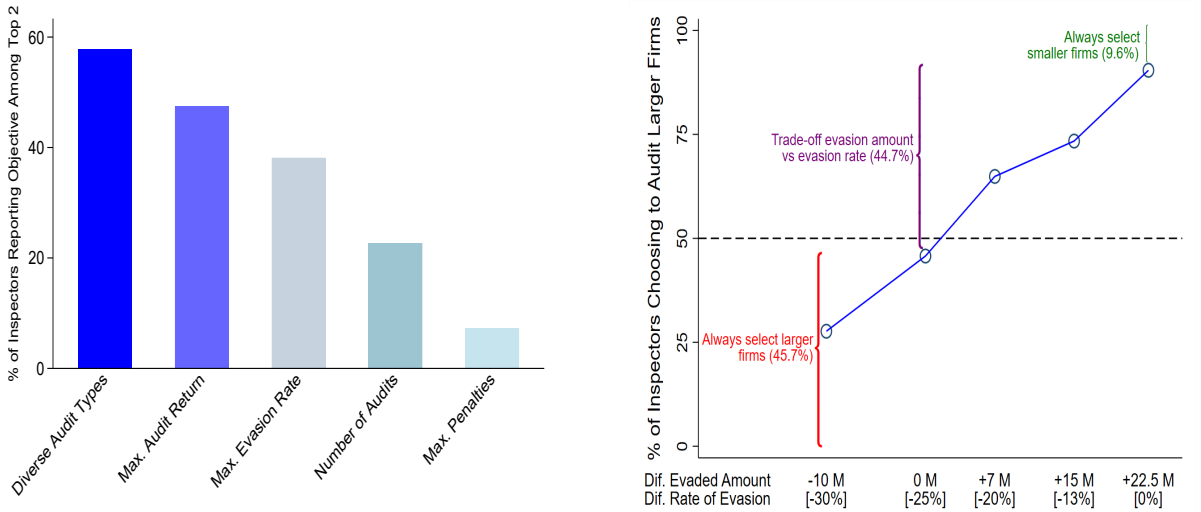
Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of the main audit outcomes on the selection method of the case for full and desk audits separately. We show specifications for full audits controlling for tax office x year fixed effects, which corresponds to the selection list for full audits. For desk audits, we show specifications with tax office x year fixed effects and inspector x year fixed effects, the latter corresponding to the list for desk audits. The first outcome (Columns 1 to 3) is the probability that the case on the list was conducted by the inspectors. The second outcome (Columns 4 to 6) is the probability that an audit results in a penalty (initial or final notice), conditional on the case being conducted. The third outcome (Columns 7 to 9) is the log amount of the evasion plus penalties, conditional on the case ending in positive penalty. The evasion amount is obtained from the final notice and complemented with the initial notice when the final notice is not available. The sample includes all cases selected in the 2018, 2019, and 2020 audit programs. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. Standard errors are clustered at the tax office x year levels, and inspector x year levels when inspector fixed effects are included. This table is discussed in Section 6.3.

Figures

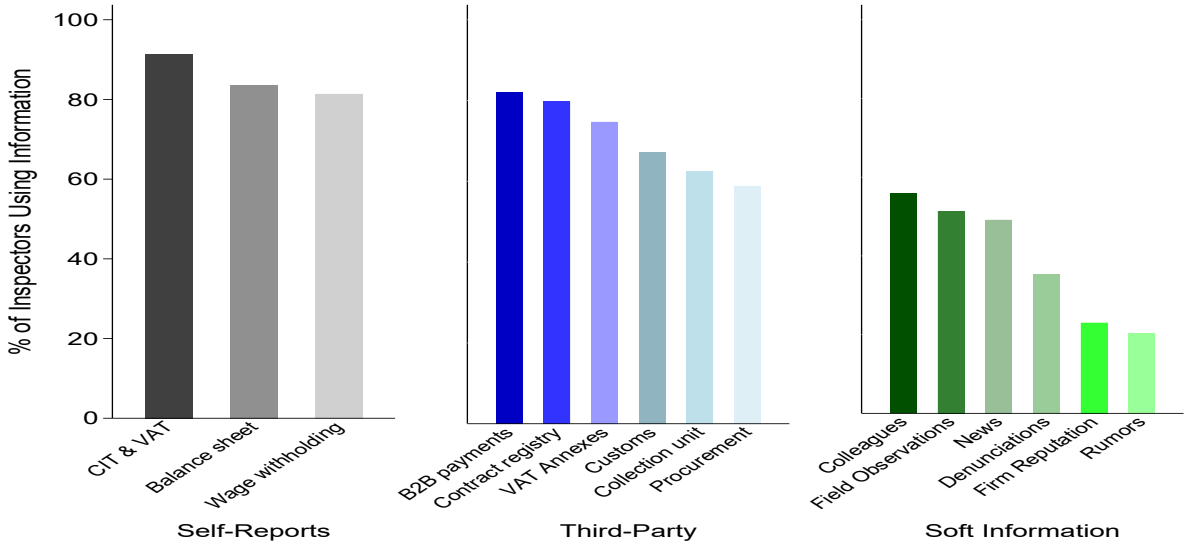
Figure 1: Tax Inspector Objectives and Practices in Audit Selection

(a) Objectives in Audit Selection

(b) Trade-Off: Evasion Amount vs Rate

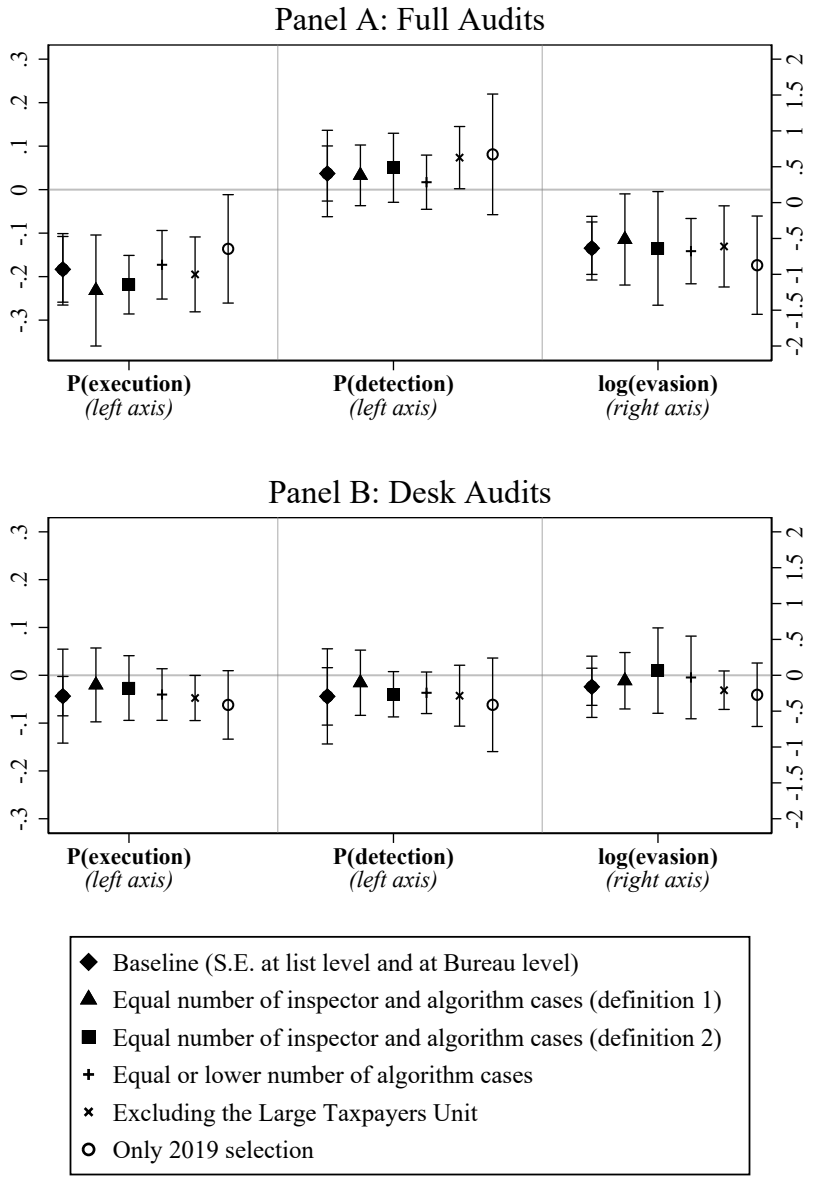


(c) Data Used in Audit Selection



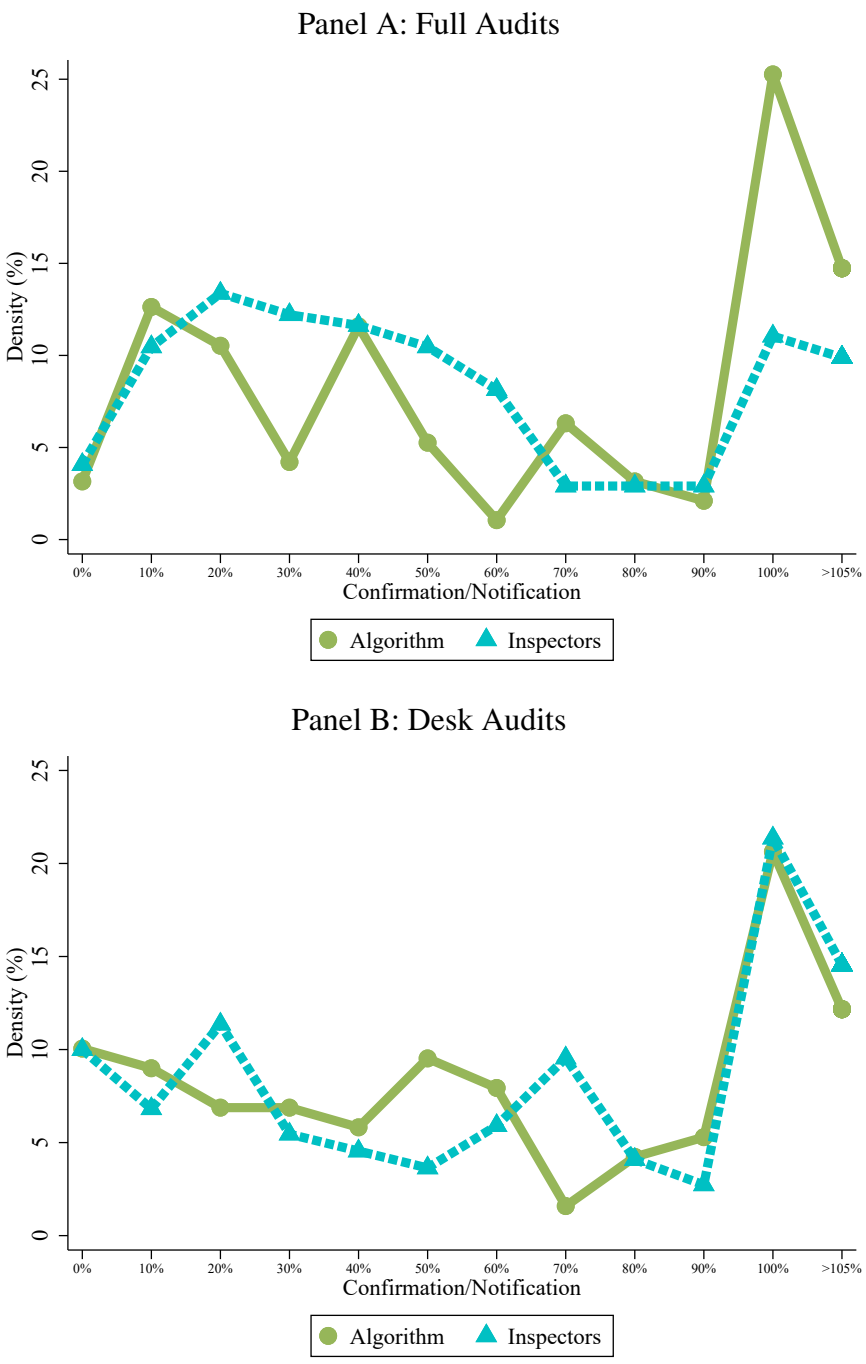
Notes: This figure shows inspectors' objectives, trade-offs and the information they use for audit selection. The data is collected at baseline from our survey of 97 tax inspectors actively involved in conducting audits. Panel (a) shows the top two reported personal objectives of audits, from the following choice set: diversity of audit types, maximize audit returns, maximize the evasion rate, conduct the largest number of audits and maximize the assessed penalties. Results are similar when reporting the top objective. Panel (b) shows the trade-off inspectors make between aiming for larger audit returns in absolute amounts versus higher detected evasion rates. This is based on a survey experiment in which inspectors were provided with the choice of auditing one of two firms, one large and one small, with different evasion amounts and hence different evasion rates. The figure plots the share of inspectors who prefer to audit the larger of the two firms (y-axis), as a function of the degree of tax evasion differential between the firms (x-axis). Each dot corresponds to a different scenario. In the first scenario (bottom left dot), auditing the large firm is worse both for the expected evasion amount recovered (-10M FCFA) and for the evasion rate differential (-30%); by the last question (top right dot) the expected evasion amount is much larger when auditing the large firm (+22.5 M FCFA) and the evasion rate is equal across the large and small firm. Panel (c) shows the frequency of using different data types for audit selection, classifying data into three categories: firms' self-reports, third-party reported data, and soft information which is hard to codify. This figure is discussed in Section 2.4.

Figure 2: Algorithm Selection and Audit Outcomes: Robustness of Estimates Across Subsamples



Notes: This figure shows the robustness of our results in different subsamples, for full audits (panel A) and desk audits (panel B), and for the three main outcomes as indicated in the x-axis titles. In each sub-panel, the first coefficient reproduces our main result from Table 4, which we call the baseline. We show standard errors when clustering at the list level, as in our main analysis, and at the bureau level. The remaining coefficients using the same empirical specifications as Table 4, employing inspector fixed-effects for the desk audit estimations, and limit the sample in the following way: 2) lists that feature and equal number of inspector-selected and algorithm-selected cases (as observed by the researchers), 3) lists that feature and equal number of inspector-selected and algorithm-selected cases (as observed by the inspectors), 4) office-year observations where the number of algorithm-selected cases is equal to or lower than the number of inspector-selected cases, 5) small and medium taxpayer offices only (i.e. excluding the large taxpayer office), 6) audit lists for 2019. For the distinction between methods 1 and 2 for equalizing the numbers of algorithm and inspector-selected cases, see footnote 26. Tables D.4 to D.9 show the details of all these estimations with numbers of observations in table format. This figure is discussed in Section 4.4.

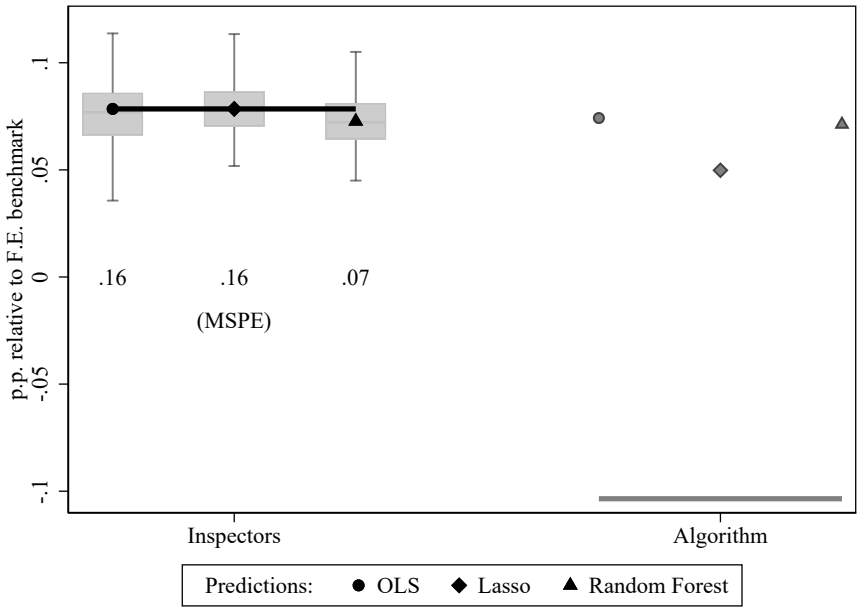
Figure 3: Dispute of Audit Results
Distribution of Confirmed Amount/Notified Amount of Evasion



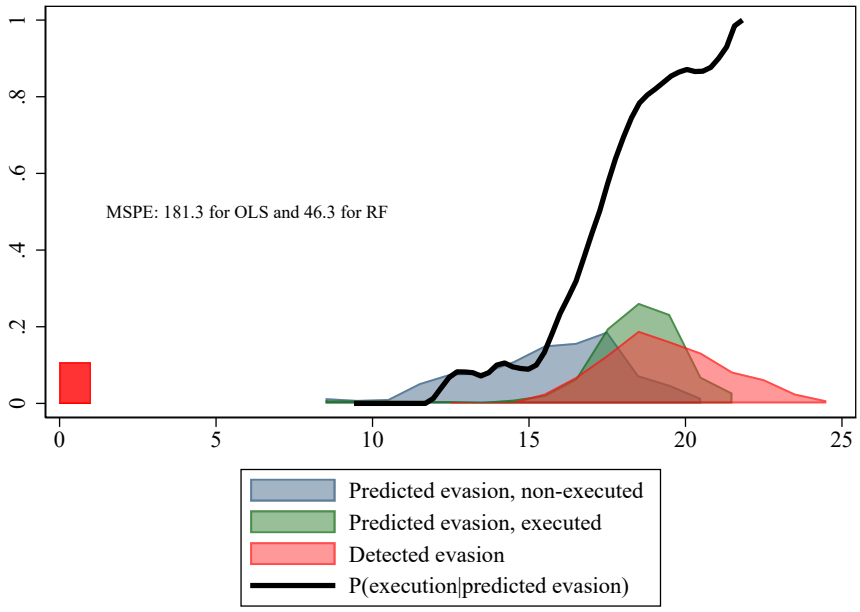
Notes: This figure shows the distribution of the confimed amount over the notifified amount of evasion, fore ach selection method and audit type. The share is calculated for firms that have a non-zero notified amount and that have a non-missing confirmed amount. The share of confirmation to notification measures the extent to which taxpayers bargain to lower the detected evasion. This figure is discussed in Section 4.3.

Figure 4: Predicting Inspector Choices and Audit Outcomes Using Observable Characteristics

Panel A: Predicted and Realized Execution of Casess



Panel B: Predicted and Realized Detected Evasion Amount



Notes: Panel A of this figure shows the predicted mean execution of audits based on observable characteristics of the firms. Several models are estimated using only inspector-selected firms, and the model is used to predict the execution rate of algorithm cases based on the firms' characteristics. The difference between the realized execution rate and the predicted execution rate represents the part of the algorithm's effect that cannot be explained by firm characteristics. This figure is discussed in Section 6.3. Panel B plots the distribution of detected evasion for all full audits (log of FCFA amounts), as well as the results from a random forest model predicting detected evasion based on firms' observable characteristics. The black solid line plots the distribution of the realized audit execution rate by level of predicted evasion. This figure is discussed in Section 6.4. This figure is for full audits only. Figure F.1 shows that same analysis for desk audits.

Online Appendix: Not for Publication

This appendix contains additional information and analyses. Appendix A provides an overview of the relevant machine learning literature. Appendix B provides further information on the context of our study. Appendix C presents the design of the risk-scoring algorithm. Appendix E presents additional results on audit outcomes. Appendix F presents additional results on the audit process. Appendix G presents additional results on mechanisms.

A Related Literature

Table A.1: Machine Learning Applications to Enhance the Work of Government Officials

Authors, Year	Paper	Country	Journal	Policy Challenge	Tool	Result
Andini et al. (2018)	Targeting with Machine Learning: Tax Rebates	Italy	JEBO	Tax rebates for individuals	Decision Tree, k-NN, Random Forest	The study showed that 'consumption constrained' households, as predicted by the ML algorithm, spent 36.9% of the bonus on food, while non-constrained households showed an insignificant effect. Not targeting non-constrained households could save 29.5% of the funds, which could be better allocated to those in need, improving the program's impact. .
Andini et al. (2022)	Machine Learning for Public Credit Guarantees	Italy	JEBO	Credit guarantees for firms	Decision Tree, Random Forest, Lasso	The study uses ML to improve the targeting of public credit guarantees in Italy, focusing on firms that are both creditworthy and credit-constrained and shows that ML selected firms had a greater benefits in terms of access to credit.
Battaglini et al. (2024)	Refining Public Policies with Machine Learning: The Case of Tax Auditing	Italy	JoEctrics	Tax audit selection	Random Forest	The focus is on using machine learning to improve the efficiency of tax audits in Italy. The paper demonstrates that ML algorithms can rank audits based on expected tax evasion and recovered evasion. Specifically, it shows that replacing the 10% worst-performing audits with taxpayers selected by the algorithm can increase detected evasion by 38% and recovered evasion by 29%.
Black et al. (2022)	Algorithmic Fairness and Vertical Equity: IRS Tax Audit Models	US	ACM	Tax audit selection	LDA, Random Forest, Grad Boosted	This study explores trade-offs in algorithmic design for tax audit fairness. Comparing different machine learning models, it shows that flexible models shift the audit burden from high- to middle-income taxpayers. Fairness constraints can balance audits but reduce revenue. Switching from classification models (predicting tax misreporting) to regression models (estimating underreporting) shifts focus back to high-income individuals, increasing revenue. However, prioritizing cost efficiency risks undermining vertical equity.
Chalfin et al. (2016)	Productivity and Selection of Human Capital	US	AER	Police hiring, teacher promotions	Stochastic Gradient Boosting, Lasso	This paper shows that ML-based algorithms can improve predictions of worker productivity. Specifically, it demonstrates that using ML to hire police officers could reduce incidents like shootings by 4.8%. In the context of granting tenure to teachers, it presents that ML-driven selections could improve test results by 0.0167σ for Math and 0.0111σ for ELA.
Glaeser et al. (2021)	Crowdsourcing City Government	US	AER P&P	Restaurant inspections	Random Forest, Grad Boosted Trees	The study uses a prediction tournament to explore how tournaments can substitute traditional consultancy. The authors show that implementing the top-performing algorithm from the tournament could increase restaurant inspection productivity by 30 to 50 percent, demonstrating the potential benefits of (i) crowdsourced solutions and (ii) using machine learning to better target inspections.
Hino et al. (2018)	Machine Learning for Environmental Monitoring	US	Nature	Water pollution inspections	Regression Forest	In the context of water-pollution inspections, the study shows that machine learning algorithms could help detect over seven times the number of violations compared to current methods. Even when accounting for local budgetary constraints and maintaining a minimum probability of inspection for all facilities, the algorithm could still double the number of violations detected
Johnson et al. (2023)	Improving Regulatory Effectiveness through Better Targeting	US	AEJ App	Workplace safety inspections	Super Learner	The study shows that targeting inspections based on ML-predicted outcomes could have averted up to 2.2 times more injuries than the current random assignment. This approach could generate up to 876 million in social value by optimizing inspection allocations, highlighting how better targeting can significantly improve safety outcomes.
Kleinberg et al. (2018)	Human Decisions and Machine Predictions	US	QJE	Grant bail to defendants	Gradient Boosted Trees	Using bail decisions in the U.S., the authors show that machine learning algorithms can generate significant welfare gains by improving the accuracy of crime risk predictions. Their policy simulation suggests that the algorithm can reduce crime by up to 24.7% without increasing jail rates, or alternatively, decrease the number of people jailed by up to 41.9% without raising crime rates

Notes: This table summarizes the literature on how machine learning algorithms can enhance or substitute for the work of government officials. This table is mentioned in Section 1.1.

B Context Appendix

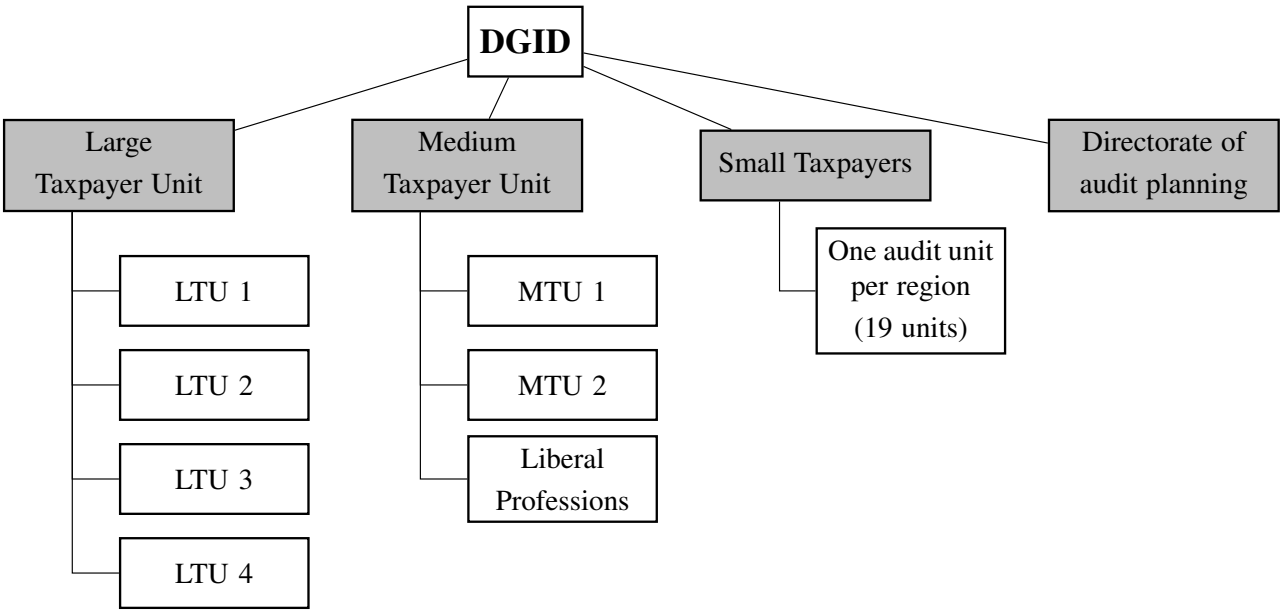
B.1 Tax Policy and Enforcement in Senegal

Table B.1: Senegal's Tax Revenue Composition in Comparison

	(1) Senegal	(2) Sub-Saharan Africa	(3) Lower-Middle Income Countries	(4) High Income Countries
Total tax revenues (% of GDP):	19.63	15.32	18.32	32.31
Sources of revenue (% of total):				
VAT	29.36	26.59	28.49	19.98
Excise taxes	9.27	10.12	10.76	5.34
Personal income tax	15.80	18.76	14.75	21.80
Corporate income tax	13.28	19.09	21.57	12.81
Taxes on international trade	14.78	15.73	12.07	0.64
Other taxes	12.80	8.23	7.13	13.49
Social security contributions	4.70	1.48	5.21	25.94

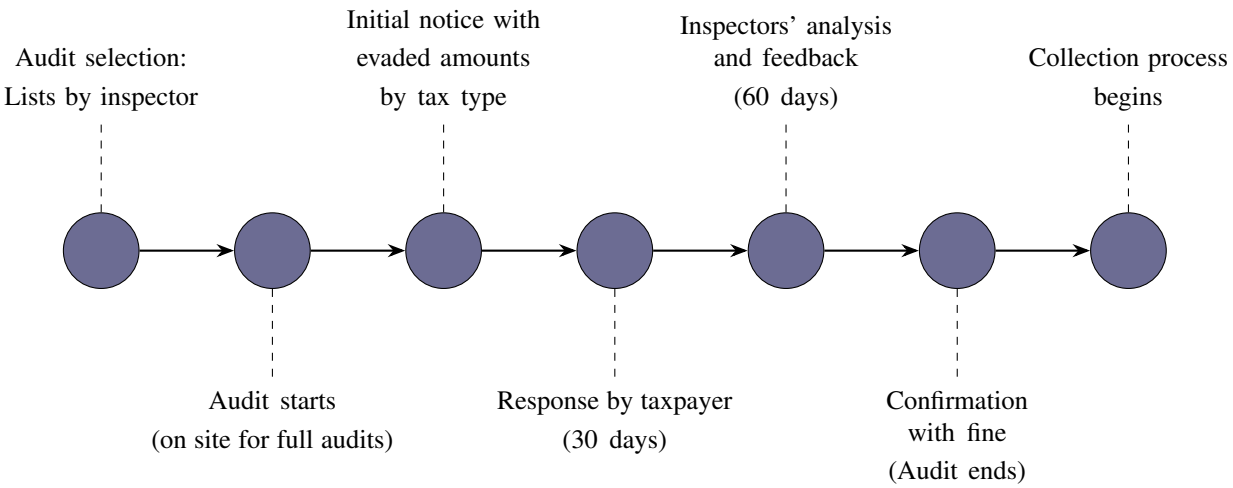
Notes: This table shows total tax revenues, including social contributions, as a percentage of GDP, along with the breakdown of total revenue by category (as a percentage of total revenue) for Senegal, Sub-Saharan Africa, lower-middle-income countries, and high-income countries. The data includes observations from 40 out of 47 Sub-Saharan African countries, 46 out of 53 lower-middle-income countries, and 57 out of 83 high-income countries. We use data from the latest available year (2021 or 2022) from the Government Finance Statistics. The category "Other taxes" includes taxes on payroll and workforce, taxes on property, and any other taxes. This table is discussed in Section 2.1.

Figure B.1: Tax Offices Involved in Conducting Audits



Notes: This figure shows the units within the tax administration (DGID) in charge of conducting audits. The Large Taxpayer Unit (LTU) oversees firms with a turnover greater than 3 billion CFA francs (approximately 5.3 million USD). The LTU is divided into four units specialized by economic activity: LTU1 for mining and energy sectors, LTU2 for financial services and the telecommunications industry, LTU3 for real estate, and LTU4 Unit 4 for all other sectors. The medium taxpayer unit (MTU) oversees firms with a turnover between XXX million CFA francs and 3 billion CFA francs (MTU1) and those with a turnover between 100 million and xxx million CFA francs (MTU2), as well as regulated liberal professions, such as lawyers, notaries and medical practitioners. Other taxpayers are overseen by one of the 19 small taxpayer centers, organized by region. This figure is discussed in Section 2.2.

Figure B.2: Audit Process at the Senegalese Tax Authority



Notes: This figure shows the steps taken to inspect a taxpayer. At the beginning of the year, inspectors are given a list of audits agreed with their superior in the hierarchy. Upon conducting the audit, they draft an initial notice containing apparent infractions. The initial notice contains the value of presumed evasion and the corresponding fine. The taxpayer can respond to the notice providing evidence that they have complied, which the inspector analyses before sending a final notice. Shortly after the final notice, the taxpayer receives a request to pay the evaded amount plus fines. This figure is discussed in Section 2.2.

B.2 Audit Selection Around the World

Table B.2: Tax Audit Selection Methods in Selected Countries

Country	Discretionary selection	Risk analysis	Random selection
Kenya	Yes ; For all except large taxpayers	Yes ; Only for large taxpayers	No
Senegal	Yes	Yes, Introduced in FY 2018	Introduced in FY 2018
Zimbabwe	Yes; Inspectors rated on selection.	Yes; based on turnover variances	No
Lesotho	No	No	Yes ; Randomly by managers
Tanzania	Abandonned in 2007	Yes	
United Kingdom	Yes; For 55% of audit cases	Yes; Risk scoring	Yes ; Simple random sample
Switzerland	Yes for all cases	No	Yes, periodically for some taxes
United States	No	Yes	
France	Yes; For intelligence gathering	Yes; statistical techniques, data-mining	No
Bulgaria	Yes ; According to set criteria	Yes; Central risk analysis	No
Turkey	No	Yes; Analysis by tax type	Yes ; to collect unbiased data

Notes: This table is based on Khwaja et al. (2011) and our survey of select country tax officials.

B.3 Tax Inspector Performance Incentives in Senegal

Tax inspectors have a strong financial performance incentives. **Their remuneration is based on three elements: a base salary, a share from a "common fund", and an individual bonus. The common fund and the individual bonuses are paid based on the fines** (including also penalties, confiscations etc) that non-compliant taxpayers pay. The fines are typically 50% of the recovered amount of evasion. Less commonly, they can be 25% or 100% of the recovered amount.

For each audit, 10% of the fines paid are set aside for the **individual performance bonus** for auditors who worked on the specific case. Of these 10%, 2% goes the unit manager, 6% to the inspectors, and 2% to the recovery agents. When inspectors work as a team, the division of the sum in the team is determined in discussion with the manager. The remaining 90% of the fines enters the common fund, which is distributed among all agents in the relevant tax administration offices. This means all agents receive a share, from secretaries to managers.

An agent's **share of the common fund** depends on their function and level of seniority in the administration (rewarded based on "statutory points") and on their performance (rewarded based on "performance points"), as evaluated quarterly by their manager against pre-defined performance objectives. Managers and high-level technical advisors automatically receive the maximum number of 150 performance points. Exceptional performers (with performance of 80-100 out of 100) receive an increased shared of the common fund, while low performers (performance below 80) see their share of the common fund reduced or even eliminated (when performance is evaluated to be below 10). Detailed rules determine how seniority levels are determined and how performance evaluations are conducted. The rules also govern a set of special circumstances, e.g. moves between tax offices. In general, more senior inspectors have stronger performance incentives as they receive a larger share of the common fund.

This discussion is linked to the discussion on bureaucrat characteristics and incentives in Section 2.4.

C The Risk Scoring Algorithm

C.1 Motivation

A key feature of this project is to assist the Senegalese tax administration (DGID) to design a tool which assesses firms' tax evasion risk. Starting in 2017, the team held consultations with DGID leadership and former tax inspectors to map the compliance risks of Senegalese firms and to exploit all available data sources to assess this risk. Moreover, we discussed with experts in the field of taxation and risk management, who worked on tax evasion risk assessment in middle-income countries. With these inputs, we designed a risk-scoring tool, following best international practice, as implemented by the World Bank and its partner institutions.

Although the use of advanced machine-learning tools for prediction has exploded in economic analysis, it was decided together with DGID that the risk-score would be guided by simple variables which logically should predict evasion risk. The simplicity of the design is motivated by several factors, ranked by order of importance. First, the tool needed to be transparent, such that underlying compliance risks could be understood by tax inspectors, and explained to taxpayers when required. Second, the available data on historical audit results was sparse and not digitized, which limited the scope of our model calibration and model selection exercises (further details below). Finally, all cases concluded by 2017 were selected in a discretionary manner.

Thus, one should consider the risk-scoring tool as a transparent best-practice risk assessment, given the administrative capacity, rather than a fined-tool fully optimized algorithm. We note that the constraints faced by DGID are likely to bind in many low income countries, and especially in other West African countries, which often look at Senegal for administrative innovations.

Table C.1 below summarizes the steps we took in deriving the risk score.

Table C.1: Steps of Risk-Score Calculation

Step	Description
(1) Prepare database	The tax declarations of each taxpayer are merged across taxes (VAT, CIT, Payroll) and across years. Data from third parties is then merged in (customs, procurement, transaction network).
(2) Calculate inconsistency ratios	Inconsistencies are situations in which a self-reported tax liability can be considered as misreported or incomplete, by comparing different data sources. An example of an inconsistency ratio is third-party reported sales over self-reported sales.
(3) Calculate anomaly ratios	Anomalies correspond to abnormal reporting behavior, compared to peers. Anomalies may be associated with tax evasion, but do not indicate tax evasion behavior with certainty. An example of an anomaly ratio is the inverse of the profit rate.
(4) Define comparison clusters	Clusters regroup firms in the same economic sector and of comparable size. Peer comparisons of anomaly ratios are done within clusters.
(5) Transform ratios into risk indicators	For inconsistencies, the magnitude of the ratio is used to assign a value, ranging from one to ten (using deciles). For anomalies, firms within the top decile of a particular ratio within their cluster are assigned a value of one.
(6) Assign weights to indicators	Weights are assigned to each indicator reflecting our beliefs about their relative importance.
(7) Aggregate indicators and years	The weighted risk indicators are first aggregated for each year. Then the yearly scores are summed up to form a total risk score covering the past four years. More recent years are weighted higher than more distant years.
(8) Weigh risk score by declared turnover	The aggregated risk score is weighted by the log of turnover to give more importance to larger firms.

Notes: This table describes the steps taken in calculating the risk score based on which the algorithm selects firms for tax audits. This is discussed in Section 3.1.

C.2 Choosing Indicators and Weights

As explained above, the algorithm computes some ratios from the data of firms (declarations and third party data) and then calculates the value of the indicator based on the distribution of this ratio within a cluster of comparable firms. We tried several combinations of indicators before stabilizing the algorithm in a reduced set of them. The goal was to have a set of indicators that was sensible and correlated with evasion, but at the same time simple and understandable for the tax inspectors.

Table C.1 summarizes the steps that we took to conceptualize the algorithm. We tried out several possible indicators that could suggest under-declaration of tax liability. We discarded most based on some analysis of data availability or statistical relevance. In the end, we discarded indicators that required information that was available for a reduced set of firms and indicators that did not seem to have any correlation with evasion, as per past evasion data. We tested these indicators on data from historical audits data. We performed out of sample regressions with LASSO and OLS and computed the out of sample mean squared prediction errors to compare different models. This allowed us to assert that the ranking normalization performed well with respect to alternatives (meaning that it

presented a lower prediction error).

We decided to restrict the algorithm to a small list of indicators. Three of them are inconsistencies, plus a flag for inconsistent filing of taxes. On top of that, we have seven anomalies, of which two refer to value added tax, two refer to corporate income tax, one refers to third party data comparisons, one to share of imports from low tax countries and one refers to the financial services tax (only applicable to a reduces set of firms). The final list of indicators that is used in the algorithm, and the respective weights is summarized in the following table.

Some details for the calculation of the indicators are worth mentioning. In some cases of anomalies, the top decile within a cluster comprises more than 10% of cases. As long as the value is not zero, we include all these firms. Whenever there is not enough non-zero values that can fill un 10% of the firms, we only flag the non-zero values. We also top code (999 999 999) all values for which the denominator of te underlying ratio of the indicator is zero or missing. Therefore they belong by definition to the top decile. We also top code all values of negative tax liability, to make sure they also get flagged. The idea of the indicators is always that the larger the ratio, the less taxes the firm is paying.

We designed the risk-scoring scheme using best practices, drawing on policy documents from the World Bank (tax administration projects in Pakistan and Turkey), SKAT in Denmark, and the IMF's recommendations to DGID. We provide a high-level description of this process to preserve confidentiality around audit selection processes. We compute risk scores using information sets/tax returns submitted to DGID on corporate income taxes, VAT, personal income tax withholding remittance, as well external data from customs (imports/exports) and public procurement contracts, for the period 2013-2016.[42] The score relies on two types of risk indicators: discrepancies and anomalies. Discrepancy indicators flag taxpayers whose self-reported information according to their tax returns differs from information in datasets obtained from customs or the government budget department in charge of paying state procurement. For instance, a discrepancy indicator is logged when taxpayers' reported turnover over multiple years is lower than its aggregate costs, that its imports plus its wage bill over the same period. Anomaly indicators use industry/sector benchmarking to flag firms with unusual behavior relative to their peers. An example would be a firm in petroleum retail with low profit rate compared to its peers, which might be associated with evasion. Discrepancies and anomalies are aggregated to produce a risk-score for each taxpayer.

[42] We also attempted to apply predictive analytics from the machine learning literature on these datasets and on previous audit results was conducted to check whether risk indicators could predict DGID audit returns. This exercise was inconclusive because of the selected nature of the sample for whom audit returns are available, the small number of observations and noise in the data.

C.3 Statistics on Audit Program

Table C.2: Count of Selected Full Audits by Year, Tax Office, and Selection Method

		Algorithm	Discretion	Overlap	Total
DGE	2018	81	94	13	182
	2019	25	96	11	121
	2020	25	75	5	100
CME 1	2018	31	33	2	67
	2019	27	27	1	54
	2020	25	25	0	50
CME 2	2018	25	25	0	53
	2019	20	20	0	40
	2020	25	25	0	50
CPR	2018	15	15	0	32
	2019	15	15	1	30
	2020	20	16	2	36
Dakar P.	2018	0	0	0	0
	2019	14	15	2	29
	2020	7	7	0	14
Ngor A.	2018	0	0	0	0
	2019	11	10	0	21
	2020	8	8	0	16
Pikine G.	2018	0	0	0	0
	2019	8	7	0	15
	2020	8	8	0	16
G. Dakar	2018	0	0	0	0
	2019	9	8	0	17
	2020	8	8	1	16
All	2018	152	167	15	334
	2019	129	198	15	327
	2020	126	172	8	298
Total		407	537	38	959

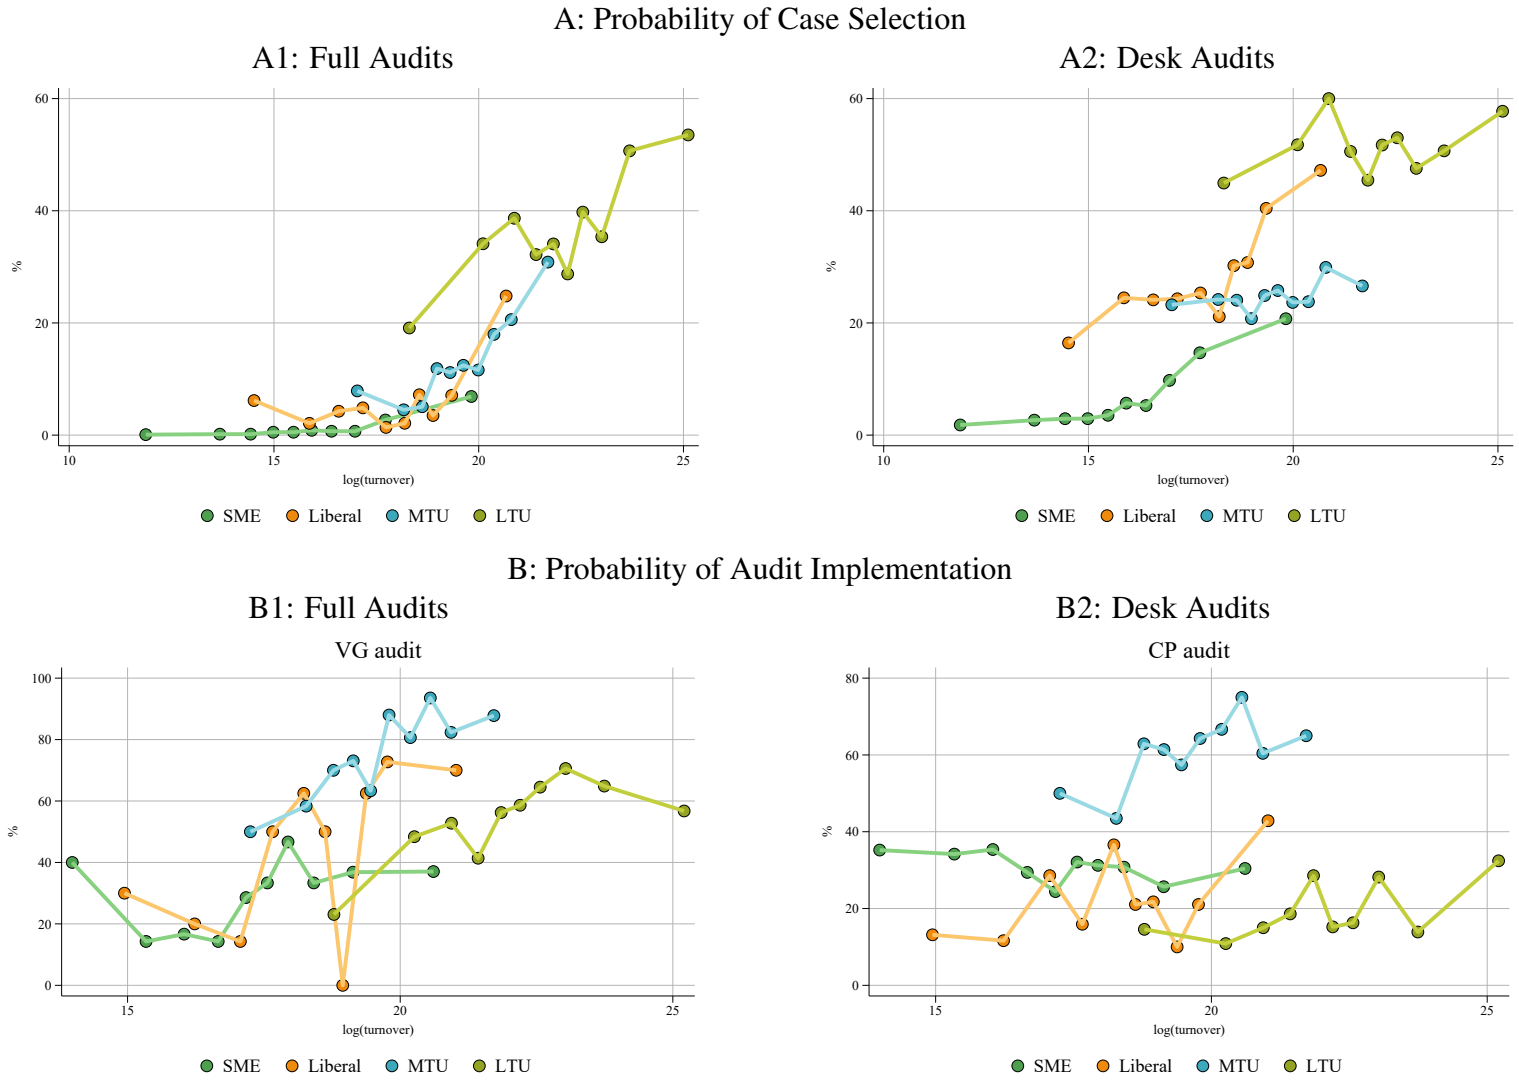
Notes: Number of selected full audits by year and tax office. The sum of the rows is larger than the total because there are overlapping cases between algorithm and discretion. This table is discussed in Section 3.2.

Table C.3: Count of Selected Desk Audits by Year, Tax Office, and Selection Method

		Random	Algorithm	Discretion	Overlap	Replacement	Total
DGE	2018	60	72	81	12	0	213
	2019	0	0	0	0	0	0
	2020	0	85	239	76	85	409
CME 1	2018	34	49	52	5	0	135
	2019	14	42	52	10	7	115
	2020	0	0	0	0	0	0
CME 2	2018	49	78	83	6	0	210
	2019	38	83	95	12	16	232
	2020	0	38	38	1	38	114
CPR	2018	59	86	95	10	0	240
	2019	26	66	70	4	12	174
	2020	0	70	70	7	70	210
Dakar P.	2018	0	0	0	0	0	0
	2019	37	71	78	7	15	201
	2020	0	72	72	1	72	216
Ngor A.	2018	0	0	0	0	0	0
	2019	19	57	59	2	10	145
	2020	0	49	59	2	49	157
Pikine G.	2018	0	0	0	0	0	0
	2019	14	26	26	0	8	74
	2020	0	63	63	2	63	189
G. Dakar	2018	0	0	0	0	0	0
	2019	14	10	12	2	8	44
	2020	0	53	53	2	53	159
All	2018	202	285	311	33	0	798
	2019	162	355	392	37	76	985
	2020	0	430	594	91	430	1454
Total		364	1070	1297	161	506	3237

Notes: Number of selected desk audits by year and tax office. The sum of the rows is larger than the total because there are overlapping cases between algorithm and discretion. This table is discussed in Section 3.2.

Figure C.1: Selection and Implementation of Audits Across the Firm-Size Distribution



Notes: This figure plots the share of selected and audited firms within the period 2018-2020 conditional on their decile of turnover and tax office. The values on the x-axis correspond to the mean of the decile within the four tax offices. The deciles were computed based on the mean declared turnover of the firms in the period 2017-2020, excluding firms that have zero turnover. For Panel A, we compute the deciles within tax office using the population of firms, and plot the share of firms within that decile-tax office. For Panel B, we condition the computation on the set of selected firms. This figure is mentioned in Section 3.5.

C.4 Balancing Tests

Table C.4: Balancing Test for Randomization of Ordering: Probability of Being on Top of the List

	(1) P(top)	(2) P(middle)	(3) P(top)	(4) P(middle)	(5) P(top)	(6) P(middle)	(7) P(top)	(8) P(middle)
Algorithm case	0.00264	0.00390						
	(0.0204)	(0.0139)						
log(Mean Turnover)			0.000889	0.000527				
			(0.00150)	(0.00148)				
log(Mean Tax Liability)					0.00115	0.00100		
					(0.00159)	(0.00156)		
Profit rate							0.0490	-0.104*
							(0.0592)	(0.0594)
N	3675	3675	3675	3675	3675	3675	3675	3675
R2	0.000523	0.000424	0.000609	0.000440	0.000664	0.000521	0.000720	0.00133
Mean outcome	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table shows the coefficients of a regression to predict the position of a case on the inspectors' list, conditional on characteristics. It predicts the probability that the case is located at the top third of the list, or in the middle third of the list. The table shows OLS results with fixed effects at the year X inspector level (tax office level for full audits). Standard errors are clustered at the inspector level (tax office level for full audits). This table is discussed in Section 4.

Table C.5: Balancing Tests of Information Intervention

Panel A: Firm Characteristics								
	Profit Rate	log(Turnover)	log(Payroll)	log(N. Employees)	log(Exports)	log(Tax Liability)	log(Firm's Age)	log(Distance to Firm)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Indicators	-0.0223	-0.0183	-0.0778	-0.123	-0.164	-0.238	0.0492	0.0903
	(0.0223)	(0.438)	(0.524)	(0.232)	(0.267)	(0.668)	(0.0350)	(0.0769)
+ Data Spreadsheets	0.0139	0.140	0.793	-0.00177	0.0463	-0.215	-0.0678**	0.0564
	(0.0223)	(0.465)	(0.603)	(0.289)	(0.318)	(0.814)	(0.0336)	(0.0808)
N	1134	1862	1452	1410	234	1420	1840	1561
R2	0.171	0.209	0.196	0.227	0.376	0.381	0.199	0.213
Mean outcome	-0.06	16.50	13.81	2.84	27.42	4.37	2.44	8.49
Panel B: Availability of Observations								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Indicators	-0.00376	0.000956	0.0106	0.00999	-0.0232	-0.0134	-0.0311*	0.00298
	(0.0218)	(0.0198)	(0.0273)	(0.0266)	(0.0179)	(0.0255)	(0.0173)	(0.0268)
+ Data Spreadsheets	0.00277	0.00528	-0.0195	-0.0150	0.0203	0.0185	0.00510	0.00837
	(0.0278)	(0.0210)	(0.0270)	(0.0268)	(0.0177)	(0.0251)	(0.0205)	(0.0269)
N	2136	2136	2136	2136	2136	2136	2136	2136
R2	0.260	0.115	0.179	0.169	0.273	0.0886	0.0962	0.0903
Mean outcome	0.53	0.87	0.68	0.66	0.12	0.66	0.86	0.73

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table shows results of OLS regressions of firm characteristics on the information treatments. The sample only includes desk audit cases for the years 2018, 2019, and 2020, which were the ones used in the intervention. The treatment was cross-randomized across algorithm and inspector cases in 2018 and 2019, but only used for algorithm cases in 2020. Therefore, we excluded the inspector-selected cases in 2020, which stemmed from a different selection method that is not directly comparable with the algorithm cases. The outcomes are based on tax declarations or firm registry data. For the tax data, we used information from the year before the firm was selected for audit. The outcomes are defined as follows: the profit rate is defined as total profits divided by total sales, obtained from CIT declarations (Column 1), log(turnover) is the natural logarithm of total sales (plus 1 to avoid dropping with zero turnover) obtained from CIT and VIT declarations (Column 2), log(payroll) is the log of total payroll (plus 1) obtained from the Pay-As-You-Earn declarations (Column 3), log(N. employees) is the log of the number of employees obtained from the Pay-As-You-Earn declarations (Column 4), log(exports) is the log of total value of exports (plus 1) obtained from customs data (Column 5), log(tax liability) is the log of the sum of VAT, CIT and PAYE liability as computed from the tax declarations (Column 6), log(Firm's age) is the log of the age in years of the firm obtained from the firm's date of creation (Column 7), and log(distance to firm) is the log of the distance from the tax office's location to the firm's premises in minutes computed using Google Maps on a Monday afternoon (Column 8). The regressions include fixed effects at the list level (inspector x year). Standard errors are clustered at the list level. This table is discussed in Section 6.2.

D Characteristics of selection for Desk Audits

Table C.1: Firm Characteristics of Algorithm vs Discretionary Selection

Panel A: Characteristics of Inspector-Selected Cases									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Inspectors' Selection	4.709***	0.989***	0.0414*	-0.282***	0.0994	2.367***	-0.402	0.783	0.147
	(0.826)	(0.197)	(0.0204)	(0.0739)	(0.349)	(0.496)	(8.550)	(2.449)	(0.171)
N	22573	7430	5922	61238	53655	51992	696	701	640
R2	0.265	0.126	0.0496	0.0692	0.119	0.111	0.0912	0.0162	0.0297
Mean outcome	10.81	14.81	-0.15	0.41	13.98	5.91	31.33	50.25	2.77

Panel B: Comparison Among Selected Cases									
	Tax Declarations				Administrative		Taxpayer Survey		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	log(Turnover)	log(Payroll)	Profit Rate	P(Trade)	Duration Trip	Firm Age	Employees	% Sales in Cash	Audit Frequency
Algorithm	0.956***	-0.729**	-0.0137	0.0840***	0.996**	2.026***	1.532	3.492	-0.221
	(0.329)	(0.293)	(0.0196)	(0.0182)	(0.437)	(0.348)	(11.76)	(4.165)	(0.215)
Inspectors x Overlap	1.528***	-1.115**	-0.0123	0.132***	-1.065	2.792***	5.537	1.605	-0.518
	(0.553)	(0.551)	(0.0305)	(0.0420)	(0.975)	(0.927)	(8.406)	(17.10)	(0.737)
Algorithm x Random	-0.795*	-0.120	-0.0137	-0.105***	-1.704***	-1.113*	-5.268	4.952	0.0421
	(0.424)	(0.454)	(0.0266)	(0.0283)	(0.579)	(0.613)	(8.254)	(5.902)	(0.306)
N	6369	4015	3304	13642	8568	7669	649	651	574
R2	0.295	0.173	0.155	0.263	0.190	0.159	0.294	0.229	0.224
Mean outcome	14.65	14.59	-0.06	0.17	14.80	11.63	22.61	51.43	2.71

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table depicts the OLS regression coefficients of characteristics on selected methods for desk audits. It reflects Table 3 on the main text, which shows the same results for full audits. Panel A uses the sample of selected cases across the three years of the experiment, including only desk audit selection. Panel B uses the sample of all firms observed in the administrative data, and the coefficient on Inspectors' Selection indicates that the firm was selected for audit at some point during 2018-2020. For Panel A, columns 1-6, the regression is a panel regression with fixed effects at the list level (year X tax office for full audits, year X inspector for desk audits), and standard errors clustered at the list level. Otherwise, the regression is cross-sectional, with tax office fixed effects and standard errors clustered at the tax office level. The characteristics of the firms stem from three sources. From the tax declarations, we use the log of the yearly declared turnover, the log of the yearly declared payroll, the profit rate, and the probability that the firm has exports or imports. We use the value for the year before the firm was selected for audit. We use data from the firm registry on the firm's age and the distance between its location and DGID, in minutes of travel time (computed using GoogleMaps for a Monday at 3 PM). Finally, we use the taxpayer survey to compute the (self-reported) number of full-time employees, the share of total sales done in cash, and the perceived yearly frequency of full audits. This table is discussed in Section 3.5.

D.1 Deviations from Pre-Analysis Plan

Our analyses in Sections 4, 5, and 6 of the paper follow the pre-analysis plan, with small exceptions:

- We pre-specified the use of audit return outcomes (e.g. evasion) and the use of audit cost outcomes, but not the cost-effectiveness or productivity of audits, i.e. the ratio of the two pre-specified classes of outcomes. We nonetheless consider audit productivity in Section 5.2, as it is an important outcome from a policy perspective, especially given our results on audit returns and audit costs.
- We pre-specified but do not report self-reported difficulty of the audit and self-reported challenges encountered during the audit as outcomes, as inspectors' compliance with the reporting

requirement for these variables (submitted in pre-filed excel files) was limited. The data points are hence too few and selected to draw meaningful conclusions. We also did not use some audit process variables, such as whether inspectors requested additional information or a revised tax return, as the reporting of these outcomes is incomplete.

- We do not report heterogeneity, spillovers and learning effects of the information treatment in the desk audit sample, as the treatment had no statistically detectable effect on the main outcomes.
- The analysis in Sections 6.3 and 6.4 was not pre-specified. We think this extension of the analysis is justified as the sections are not about analyzing additional outcomes or additional heterogeneity dimensions, beyond those specified in the pre-analysis plan. Rather, the sections shed more light on the underlying reasons for the main results presented in Sections 4 and 5.
- We leave the analysis of medium-term outcomes, such as taxpayers' future compliance behavior, for a separate paper, given space constraints.

E Additional Results on Audit Outcomes

Table D.1: Algorithm Selection and Audit Outcomes, Intent-to-Treat Analysis

Panel A (Linear Regression)						
	P(Execution)		P(Detection | Execution)		log(Evasion) | Detection	
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	-0.183***	-0.0436**	-0.141***	-0.0469**	-3.052***	-0.899**
	(0.0369)	(0.0208)	(0.0366)	(0.0214)	(0.726)	(0.376)
Inspectors x Overlap	0.159**	0.0423	0.196**	0.0285	4.062**	0.462
	(0.0711)	(0.0365)	(0.0794)	(0.0368)	(1.638)	(0.671)
Algorithm x Random		-0.00139		-0.0123		-0.277
		(0.0311)		(0.0302)		(0.512)
N	944	2731	944	2731	944	2731
R2	0.263	0.317	0.244	0.187	0.247	0.197
Mean outcome	0.53	0.37	0.47	0.27	9.26	4.87
Panel B (Poisson Regression)						
Algorithm	-0.346***	-0.120**	-0.299***	-0.174**	-0.339***	-0.189**
	(0.0820)	(0.0583)	(0.0891)	(0.0805)	(0.0902)	(0.0795)
Inspectors x Overlap	0.267**	0.124	0.349**	0.106	0.352**	0.0931
	(0.116)	(0.103)	(0.141)	(0.130)	(0.146)	(0.132)
Algorithm x Random		0.0192		-0.00860		-0.0169
		(0.0676)		(0.0952)		(0.0924)
N	925	2504	925	2493	925	2493
R2						
Mean outcome	0.54	0.40	0.48	0.30	9.45	5.34

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table is identical to Table 4 discussed in Section 4.1, except it shows intent-to-treat effects, where outcomes are set to zero for audits that are not implemented. Panel A shows ordinary least square results, and Panel B shows Poisson Pseudo-Maximum Likelihood estimates. Standard errors are clustered at the level of the fixed effects or all specifications.

Table D.2: Algorithm Selection and Audit Outcomes, Controlling for List Slot Fixed Effect

	P(Execution)			P(Detection | Execution)			log(Evasion) | Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.179***	-0.0468	-0.0447**	0.0411	-0.0465	-0.0443	-0.617***	-0.258*	-0.163
	(0.0355)	(0.0289)	(0.0208)	(0.0308)	(0.0373)	(0.0305)	(0.213)	(0.137)	(0.131)
Inspectors x Overlap	0.155**	0.0473	0.0431	0.0800*	-0.0751*	-0.0766	0.273	-0.595*	-0.414
	(0.0715)	(0.0419)	(0.0365)	(0.0445)	(0.0367)	(0.0476)	(0.299)	(0.347)	(0.336)
Algorithm x Random		-0.00155	-0.00182		0.00262	0.00517		-0.230*	-0.221
		(0.0394)	(0.0310)		(0.0408)	(0.0329)		(0.133)	(0.169)
Tax center x Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector x Year	No	No	Yes	No	No	Yes	No	No	Yes
N	944	2731	2731	507	1016	997	453	751	732
R2	0.265	0.227	0.317	0.156	0.325	0.430	0.315	0.294	0.428
Mean outcome	0.53	0.37	0.37	0.89	0.73	0.73	19.29	17.74	17.71

This table is identical to Table 4 discussed in Section 4.1, except that we now control for the placement of a case on the audit list, by adding list order quartile fixed effects. We use quartiles rather than list slot fixed effects, because the lists vary in length. The fixed effects help control for differential effort or attention over time within an audit program. The intent-to-treat effects are in Table D.1.

Table D.3: Algorithm Selection and the Evasion Rate

	% of Liability			% of Pre-Audit Turnover		
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	0.00658	-0.0266	-0.0364	0.0217	0.0156	0.0126
	(0.0416)	(0.0300)	(0.0308)	(0.0400)	(0.0253)	(0.0191)
Inspectors x Overlap	0.192***	-0.109**	-0.158***	0.140*	-0.0512	-0.0640*
	(0.0676)	(0.0467)	(0.0561)	(0.0765)	(0.0316)	(0.0325)
Algorithm x Random		0.0120	0.00960		0.00217	0.000819
		(0.0479)	(0.0369)		(0.0390)	(0.0310)
Tax center x Year	Yes	Yes	Yes	Yes	Yes	Yes
Inspector x Year	No	No	Yes	No	No	Yes
N	497	980	960	494	978	957
R2	0.0996	0.191	0.337	0.110	0.142	0.307
Mean outcome	0.65	0.50	0.50	0.25	0.21	0.21

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case. The sample includes all cases selected in the audit programs of 2018, 2019, and 2020. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. Standard errors are clustered at the tax office x year levels, and inspector x year levels when inspector fixed effects are included. These results are mentioned in Section 4.2.

E.1 Estimates Underlying the Robustness Figure

Table D.4: Robustness of Main Results - Cluster S.E. at Tax Office level

	P(Execution)			P(Detect)|Execution			log(Evasion)|Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.183***	-0.0456	-0.0436	0.0371	-0.0465	-0.0441	-0.581**	-0.178	-0.0815
	(0.0347)	(0.0392)	(0.0415)	(0.0420)	(0.0435)	(0.0421)	(0.226)	(0.195)	(0.199)
Inspectors x Overlap	0.159***	0.0465	0.0423	0.0809**	-0.0744*	-0.0756**	0.0681	-0.213	-0.0262
	(0.0330)	(0.0471)	(0.0475)	(0.0311)	(0.0366)	(0.0237)	(0.353)	(0.363)	(0.390)
Algorithm x Random		-0.00104	-0.00139		0.00284	0.00534		-0.0927	-0.0838
		(0.0254)	(0.0281)		(0.0472)	(0.0388)		(0.0820)	(0.108)
Tax center X Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector X Year	No	No	Yes	No	No	Yes	No	No	Yes
N	944	2731	2731	507	1016	997	453	751	732
R2	0.263	0.227	0.317	0.151	0.324	0.430	0.214	0.271	0.402
Mean outcome	0.53	0.37	0.37	0.89	0.73	0.73	18.84	17.42	17.39

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case. This table show the results for the main outcomes restricting to the lists in which the number of algorithm and discretionary cases were exactly the same. The sample includes all cases selected in the audit programs of 2018, 2019, and 2020. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. Standard errors are clustered at the tax office x year levels, and inspector x year levels when inspector fixed effects are included. This table provides the estimates shown in Figure 2.

Table D.5: Robustness of Main Results - Equal Number of Cases (Definition 1)

	P(Execution)			P(Detect)|Execution			log(Evasion)|Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.232***	-0.0201	-0.0202	0.0329	-0.0248	-0.0156	-0.414	-0.170	-0.181
	(0.0573)	(0.0390)	(0.0386)	(0.0313)	(0.0322)	(0.0338)	(0.346)	(0.214)	(0.244)
Inspectors x Overlap		0.0671	0.0582		0.0388	-0.00719		-1.209***	-0.0753
		(0.226)	(0.260)		(0.0378)	(0.0156)		(0.234)	(0.239)
Algorithm x Random		0.479***	0.464***		0.0637	0.00898		-1.521**	-0.725*
		(0.0597)	(0.109)		(0.0495)	(0.0195)		(0.640)	(0.416)
Tax center X Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector X Year	No	No	Yes	No	No	Yes	No	No	Yes
N	414	798	798	264	209	199	243	201	191
R2	0.194	0.173	0.235	0.0925	0.0662	0.177	0.177	0.378	0.466
Mean outcome	0.63	0.26	0.26	0.92	0.96	0.95	18.52	17.59	17.58

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case. This table show the results for the main outcomes excluding the Large Taxpayer Unit. The sample includes all cases selected in the audit programs of 2018, 2019, and 2020. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. Standard errors are clustered at the tax office x year levels, and inspector x year levels when inspector fixed effects are included. This table provides the estimates shown in Figure 2.

Table D.6: Robustness of Main Results - Equal Number of Cases (Definition 2)

	P(Execution)			P(Detect)|Execution			log(Evasion)|Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.218***	-0.0272	-0.0267	0.0503	-0.0647**	-0.0397*	-0.525	-0.297	-0.169
	(0.0302)	(0.0339)	(0.0338)	(0.0356)	(0.0280)	(0.0235)	(0.461)	(0.279)	(0.311)
Inspectors x Overlap	0.00628	0.103	0.0896	0.208*	-0.0815	-0.116	1.266**	-1.813***	-1.409***
	(0.137)	(0.0957)	(0.110)	(0.0995)	(0.0828)	(0.1000)	(0.554)	(0.462)	(0.505)
Algorithm x Random		0.0266	-0.0150		0.0852*	0.0410		-0.128	-0.224
		(0.107)	(0.125)		(0.0429)	(0.0310)		(0.867)	(0.710)
Tax center X Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector X Year	No	No	Yes	No	No	Yes	No	No	Yes
N	366	838	838	224	180	173	200	175	168
R2	0.210	0.131	0.192	0.118	0.0659	0.252	0.113	0.260	0.380
Mean outcome	0.61	0.21	0.21	0.89	0.97	0.97	18.25	17.58	17.61

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case. This table show the results for the main outcomes only for the year 2019, which we consider the implementation of the experiment to be best. The sample includes all cases selected in the audit programs of 2018, 2019, and 2020. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. Standard errors are clustered at the tax office x year levels, and inspector x year levels when inspector fixed effects are included. This table provides the estimates shown in Figure 2.

Table D.7: Robustness of Main Results - More Algorithm Cases

	P(Execution)			P(Detect)|Execution			log(Evasion)|Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.173***	-0.0448	-0.0403	0.0171	-0.0537**	-0.0366*	-0.617**	-0.306	-0.245
	(0.0382)	(0.0276)	(0.0271)	(0.0301)	(0.0241)	(0.0217)	(0.292)	(0.267)	(0.298)
Inspectors x Overlap	0.131*	0.0752	0.0716	0.0798	-0.0790	-0.0904	0.0687	-1.184***	-1.096*
	(0.0694)	(0.0454)	(0.0465)	(0.0471)	(0.0565)	(0.0768)	(0.337)	(0.416)	(0.553)
Algorithm x Random		0.0333	-0.00994		0.0791*	0.0471		-0.0110	-0.0785
		(0.106)	(0.119)		(0.0445)	(0.0360)		(0.878)	(0.706)
Tax center X Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector X Year	No	No	Yes	No	No	Yes	No	No	Yes
N	855	1180	1180	474	226	208	428	220	203
R2	0.261	0.123	0.202	0.116	0.0674	0.252	0.208	0.299	0.440
Mean outcome	0.55	0.19	0.19	0.90	0.97	0.97	18.86	17.87	17.81

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case. This table show the results for the main outcomes computing the standard errors at the bureau level, as opposed to bureau-year level. The sample includes all cases selected in the audit programs of 2018, 2019, and 2020. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. This table provides the estimates shown in Figure 2.

Table D.8: Robustness of Main Results - Excluding Large Taxpayer Unit

	P(Execution)			P(Detect)|Execution			log(Evasion)|Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.195***	-0.0488**	-0.0475**	0.0736**	-0.0486	-0.0427	-0.528	-0.224	-0.108
	(0.0406)	(0.0238)	(0.0238)	(0.0337)	(0.0303)	(0.0321)	(0.333)	(0.153)	(0.159)
Inspectors x Overlap	0.220*	0.00784	-0.00135	0.147*	-0.0803	-0.0817	1.115**	0.181	0.310
	(0.123)	(0.0546)	(0.0554)	(0.0823)	(0.0589)	(0.0536)	(0.470)	(0.339)	(0.352)
Algorithm x Random		-0.00860	-0.00866		0.0131	0.00958		-0.0482	-0.00588
		(0.0332)	(0.0335)		(0.0369)	(0.0352)		(0.185)	(0.177)
Tax center X Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector X Year	No	No	Yes	No	No	Yes	No	No	Yes
N	548	2194	2194	319	906	900	285	645	639
R2	0.277	0.222	0.304	0.205	0.307	0.414	0.136	0.152	0.289
Mean outcome	0.58	0.41	0.41	0.89	0.71	0.71	18.32	17.13	17.13

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case. This table show the results for the main outcomes computing the standard errors at the bureau level, as opposed to bureau-year level. The sample includes all cases selected in the audit programs of 2018, 2019, and 2020. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. This table provides the estimates shown in Figure 2.

Table D.9: Robustness of Main Results - Only 2019 Program

	P(Execution)			P(Detect)|Execution			log(Evasion)|Detection		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits	Full audits	Desk audits	Desk audits
Algorithm	-0.136**	-0.0620*	-0.0620*	0.0811	-0.0693	-0.0618	-0.911*	-0.349	-0.196
	(0.0559)	(0.0356)	(0.0355)	(0.0622)	(0.0481)	(0.0485)	(0.457)	(0.239)	(0.230)
Inspectors x Overlap	0.234*	-0.0215	-0.00910	0.143**	-0.146	-0.154*	0.197	-0.267	0.134
	(0.111)	(0.0872)	(0.0885)	(0.0487)	(0.0913)	(0.0832)	(0.459)	(0.402)	(0.440)
Algorithm x Random		0.0554	0.0623*		0.00187	0.00197		-0.0552	-0.00826
		(0.0378)	(0.0364)		(0.0544)	(0.0514)		(0.275)	(0.269)
Tax center X Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inspector X Year	No	No	Yes	No	No	Yes	No	No	Yes
N	327	909	909	195	547	547	163	299	299
R2	0.165	0.0978	0.242	0.183	0.142	0.281	0.220	0.153	0.317
Mean outcome	0.59	0.60	0.60	0.83	0.54	0.54	18.80	17.14	17.14

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case. This table show the results for the main outcomes computing the standard errors at the bureau level, as opposed to bureau-year level. The sample includes all cases selected in the audit programs of 2018, 2019, and 2020. The data includes the selection and audits of the Large Taxpayer Unit, Medium Taxpayer Units 1 and 2, Liberal Professionals, and the regional SME units of Dakar Plateau, Grand Dakar, Pikine Guediawaye, and Ngor Almadies. This table provides the estimates shown in Figure 2.

F Additional Results on the Audit Process

Table E.1: Inspector Characteristics by Case Selection Method, Full Audits

Panel A: Full-Audit Cases in Which Some Inspector Reported Information							
	(1) Mean Age	(2) Max Age	(3) Mean Years of Experience	(4) Max Years of Experience	(5) Share with Masters/PhD	(6) Max Education	(7) Share in Favor of Algorithm
Algorithm	-0.275	-0.864*	-0.0299	-0.383*	0.0361	-0.0771	-0.0241
	(0.304)	(0.490)	(0.194)	(0.214)	(0.0210)	(0.0645)	(0.0282)
N	459	459	422	422	422	422	422
R2	0.358	0.399	0.365	0.419	0.327	0.336	0.324
Mean outcome	37.10	40.52	8.29	9.54	0.81	3.23	0.82
Panel B: Only Cases in Which All Inspectors Reported Information							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Algorithm	-0.174	-0.649	0.155	-0.120	0.0130	-0.0425	-0.0182
	(0.403)	(0.665)	(0.341)	(0.312)	(0.0495)	(0.0613)	(0.0483)
N	230	230	117	117	117	117	117
R2	0.423	0.358	0.295	0.212	0.600	0.854	0.178
Mean outcome	37.39	41.85	9.35	11.49	0.76	3.35	0.75
Panel C: Probability that All Inspectors Reported Information							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Algorithm	-0.0143	-0.0143	0.0153	0.0153	0.0153	0.0153	0.0153
	(0.0438)	(0.0438)	(0.0300)	(0.0300)	(0.0300)	(0.0300)	(0.0300)
N	507	507	507	507	507	507	507
R2	0.275	0.275	0.407	0.407	0.407	0.407	0.407
Mean outcome	0.45	0.45	0.23	0.23	0.23	0.23	0.23

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table shows the OLS estimation results for the mean characteristics of the team members working on cases. The regression is done at the case level only for started full audits. Full audits are assigned at the tax office level, and tax office leaders compose teams to work on the cases. Since there is no assignment of cases to inspectors at the selection phase, we can only run this regression for the cases that were effectively conducted. The outcomes are obtained from the inspector survey, and are averaged across the inspectors that worked in the case. Columns 1 and 2 show the results using as outcomes the mean and max age of the team members. Columns 3 and 4 use the share of team members with a Masters/PhD and the max of a categorical education variable that goes from 1 (only high school) to 4 (Masters/PhD). Column 5 uses the mean response of team members to whether they would favor the use of an algorithm to automate selection of cases. Columns 7, 8, and 9 use the share of highly experience (more than median) team members in the team, the mean experience, and the max experience among team members. The regressions control for fixed effects at the tax office-year level, and we cluster standard errors at the same level. Coefficients on the inspectorxoverlap dummy are included in the regression but omitted from the display. This table is discussed in Section 5.1.

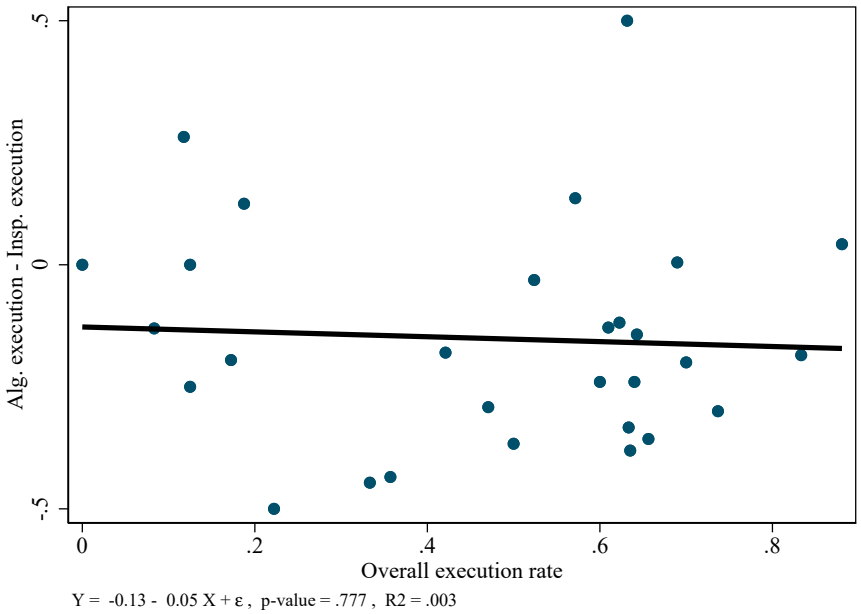
Table E.2: Inspector Characteristics and Performance, Desk Audits

	(1) # cases Executed	(2) P(execution)	(3) P(detection|execution)	(4) log(evasion) (Mean)	(5) Executed Alg. cases	(6) P(execution| Alg.)	(7) Executed Alg. > Executed Insp.	(8) P(In favor of Alg.)
Masters/PhD	-0.722	-0.00107	0.0592	0.839	-0.0717	0.0262	0.174	0.135
	(0.831)	(0.0311)	(0.0615)	(1.071)	(0.493)	(0.0522)	(0.102)	(0.0795)
Above median age	0.543	0.0478	0.0191	0.262	0.117	0.0442	-0.0570	-0.0660
	(0.490)	(0.0267)	(0.0478)	(1.079)	(0.298)	(0.0272)	(0.0835)	(0.122)
Above median experience	-0.628	-0.0534*	-0.00590	-0.589	0.108	-0.00762	0.128	-0.102**
	(0.565)	(0.0274)	(0.0629)	(1.132)	(0.356)	(0.0372)	(0.0809)	(0.0365)
N	86	86	76	76	86	86	86	86
R2	0.816	0.542	0.561	0.561	0.746	0.542	0.470	0.0651
Mean outcome	4.77	0.32	0.90	16.92	2.53	0.29	0.34	0.77

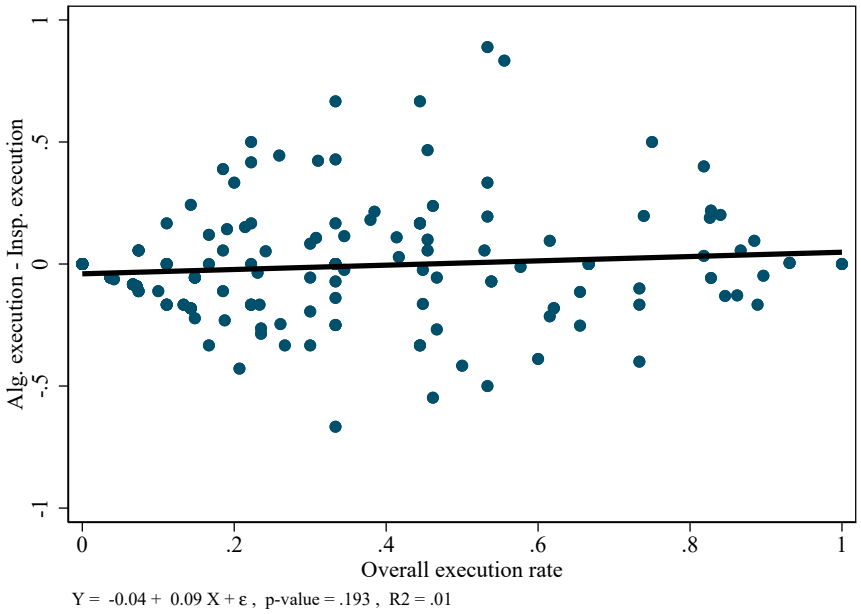
Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table shows the estimation of the OLS regressions of outcomes of inspectors on inspector characteristics. The analysis is done at the list (inspector x year) level and only includes the lists of desk audits. These lists were assigned individually to inspectors. The regressors are three individual characteristics of the inspectors collected using the inspector survey (whether the inspector had a Masters/PhD and whether they had higher than median experience working at the tax authority) and administrative data (age). Column 1's outcome is the total started cases for each list, Columns 2-4 are averages within list of the main outcomes (see Table 4), Column 5 uses the number of started algorithm cases, Column 6 the share of started algorithm cases among algorithm-selected cases, Column 7 uses an indicator of whether the number of started algorithm cases was larger than the number of started inspector cases in the list, and Column 8 uses the response of the inspector to whether they would favor the automation of case selection using an algorithm. The regressions control for fixed effects at the tax office-year level, and we cluster standard errors at the same level. This table is discussed in Section 5.1.

Figure E.1: Absence of Correlation Between Audit Execution Rate and Execution Rate for Algorithm-Selected Cases

(a) Full Audits



(b) Desk Audits



Notes: This figure plots outcomes at the audit-list level (i.e. the tax office × year level for full audits and the tax inspector × year level for desk audits). The X-axis shows the mean execution rate for all cases on the annual list. The Y-axis shows the difference in the execution rate between algorithm-selected cases and inspector-selected cases. This figure is mentioned in Section 5.1.

Table E.3: Algorithm Selection and Scope of Audits (1/2)

	N. infractions		N. years		Fine/Evasion	
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	-1.194***	-0.450***	-0.342**	-0.259*	0.00976	0.0341***
	(0.339)	(0.145)	(0.166)	(0.134)	(0.0132)	(0.0119)
Inspectors x Overlap	0.0953	-0.166	0.247	-0.240	0.00299	-0.00296
	(0.649)	(0.523)	(0.256)	(0.271)	(0.0290)	(0.0214)
Algorithm x Random		0.380*		0.385**		-0.00337
		(0.211)		(0.178)		(0.0130)
N	456	788	456	788	454	773
R2	0.318	0.338	0.185	0.293	0.210	0.269
Mean outcome	4.93	3.05	3.48	2.52	0.41	0.44

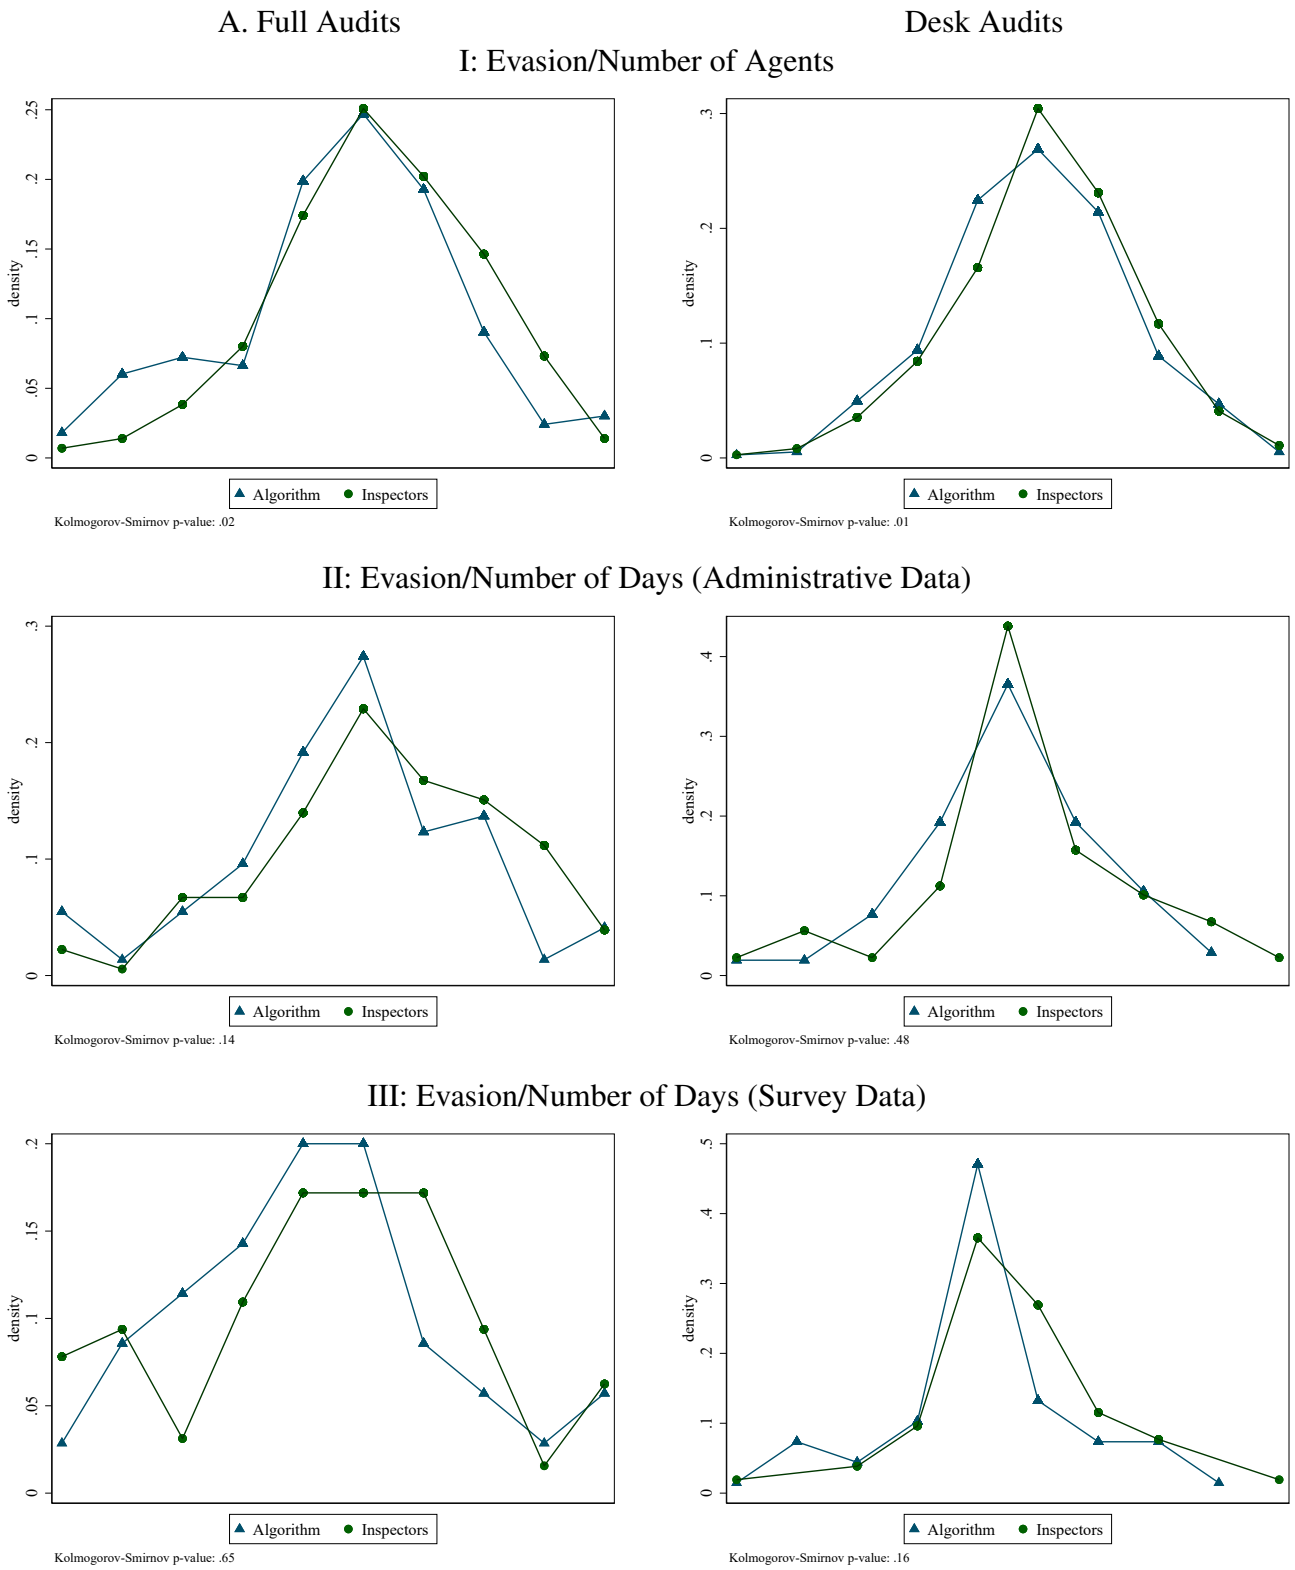
Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case, controlling for list fixed effects (year x tax office for full audits, year x inspector for desk audits). The outcomes are extracted from the audit reports. The three outcomes are i) the number of years from the taxpayers' declarations in which the inspector found an infraction (notice that the inspection can investigate tax declarations up to four years before the audit date according to Senegalese law), ii) the number of different infractions found by the inspector, and iii) the severity of the infraction as indicated by the ratio of fine to the evaded amount. Standard errors are clustered at the list level. Sample is conditioned on cases that started an audit. This table is discussed in Section 5.1.

Table E.4: Algorithm Selection and Scope of Audits (2/2)

	Share severe infractions (main taxes)		Share severe infractions (all taxes)		P(infraction in main taxes)		P(severe infraction in main taxes)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	1.381	3.372	1.724	5.199*	0.0172	-0.0352	-0.0185	-0.000122
	(4.346)	(3.806)	(4.026)	(3.033)	(0.0257)	(0.0310)	(0.0492)	(0.0385)
Inspectors x Overlap	-2.963	5.597	-13.69***	4.632	0.0548	0.0454	-0.0253	0.0554
	(6.347)	(8.750)	(3.626)	(6.653)	(0.0337)	(0.0678)	(0.0775)	(0.0930)
Algorithm x Random		3.998		3.068		-0.0122		0.0281
		(4.913)		(4.173)		(0.0440)		(0.0530)
N	457	789	457	789	457	789	457	789
R2	0.175	0.318	0.288	0.371	0.182	0.312	0.109	0.273
Mean outcome	60.89	64.62	60.69	72.55	0.94	0.85	0.78	0.71

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results of a regression of audit outcome on the selection method of the case, controlling for list fixed effects (year x tax office for full audits, year x inspector for desk audits). The outcomes are extracted from the audit reports. The three outcomes are i) the number of years from the taxpayers' declarations in which the inspector found an infraction (notice that the inspection can investigate tax declarations up to four years before the audit date according to Senegalese law), ii) the number of different infractions found by the inspector, and iii) the severity of the infraction as indicated by the ratio of fine to the evaded amount. Standard errors are clustered at the list level. Sample is conditioned on cases that started an audit. This table provides additional outcomes for the analysis discussed in Section 5.1.

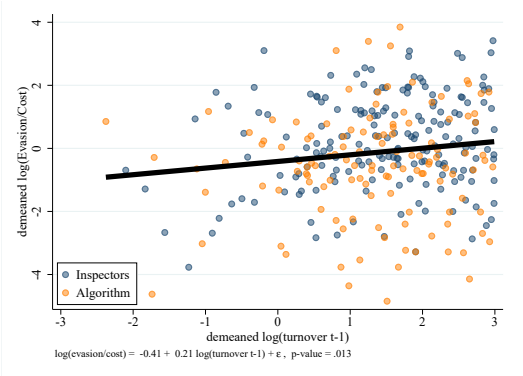
Figure E.2: Distribution of Audit Productivity



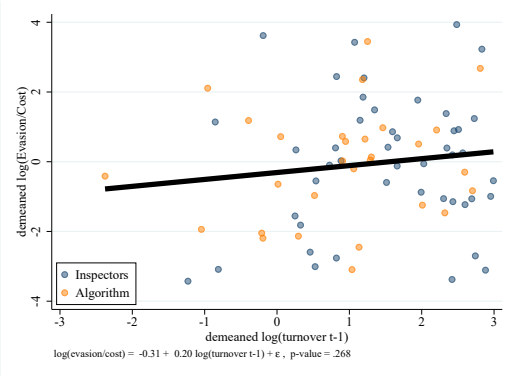
Notes: This figure plots density distributions of audit productivity outcomes, as per the panel titles. The outcomes are demeaned at the list level. The data is split into ten equally-spaced bins for each audit outcome and type. This figure is discussed in Section 5.1.

Figure E.3: Productivity of Uncovering Evasion Using Different Definitions of Cost

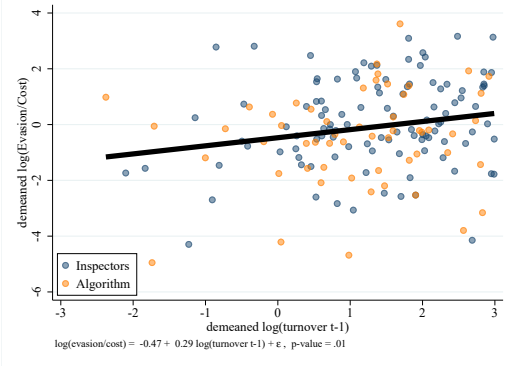
(a) Number of Agents



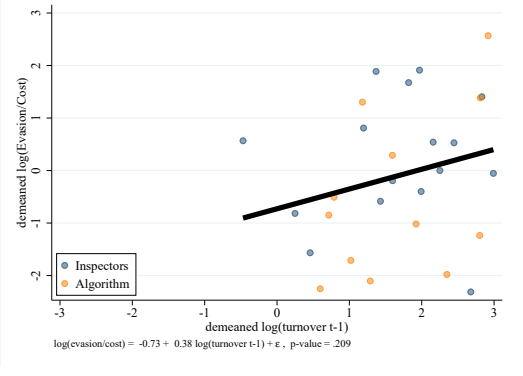
(b) Duration of Audit (Taxpayer Survey)



(c) Days from Start to Notification



(d) Duration of Audit (Self-Reported)



Notes: This figure plots the data points of audit cases. The Y-axis represents the productivity of uncovering evasion dividing evasion by three different cost measures. The X-axis represents the log of the firm's turnover as declared one year before the audit. All the variables were demeaned at the list level. The lines show a local non-parametric regression around the points, run separately for the algorithm-selected and inspector-selected cases. This figure is discussed in Section 5.1.

Table E.5: Selection Method and Survey Results for Corruption - Probability of Answering the Question

Panel A: All Interviewed Firms

	Corruption in General			Corruption Experience			Grade on Honesty			Friend at Tax Auth.		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Full audits	Desk audits	All	Full audits	Desk audits	All	Full audits	Desk audits	All	Full audits	Desk audits	All
Algorithm	-0.0218	0.00369	-0.00758	-0.0134	-0.00530	-0.00904	0.00375	0.0196	0.0122	0.0332	0.0469	0.0409
	(0.0743)	(0.0700)	(0.0508)	(0.0161)	(0.0233)	(0.0145)	(0.0518)	(0.0574)	(0.0386)	(0.0353)	(0.0361)	(0.0252)
Inspectors x Overlap	-0.138	-0.263	-0.209	0.0213	0.0140	0.0168	-0.0237	-0.0114	-0.0176	0.142*	0.0609	0.0959
	(0.302)	(0.189)	(0.168)	(0.0176)	(0.0169)	(0.0120)	(0.147)	(0.139)	(0.101)	(0.0690)	(0.124)	(0.0754)
Algorithm x Random		-0.0346	-0.0259		0.0123	0.0145		-0.0870	-0.0832		0.00195	0.00686
		(0.0831)	(0.0789)		(0.0227)	(0.0205)		(0.0574)	(0.0549)		(0.0626)	(0.0606)
N	254	369	623	254	369	623	254	369	623	254	369	623
R2	0.0847	0.231	0.174	0.133	0.156	0.146	0.146	0.252	0.226	0.0866	0.137	0.117
Mean outcome	0.64	0.60	0.61	0.96	0.97	0.97	0.81	0.71	0.75	0.87	0.86	0.86

Panel B: Only Audited Firms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Algorithm	-0.0336	-0.137	-0.0820	-0.0314	-0.0173	-0.0240	-0.0128	-0.0345	-0.0250	0.0251	0.0962*	0.0671*
	(0.111)	(0.0826)	(0.0664)	(0.0371)	(0.0354)	(0.0252)	(0.0700)	(0.0802)	(0.0535)	(0.0487)	(0.0567)	(0.0389)
Inspectors x Overlap	0.529***	-0.186	0.105	-0.00192	0.0328	0.0173	-0.000782	0.00671	0.00701	0.226*	0.0489	0.107
	(0.164)	(0.310)	(0.229)	(0.00264)	(0.0278)	(0.0182)	(0.00432)	(0.178)	(0.105)	(0.115)	(0.241)	(0.158)
Algorithm x Random		0.115	0.0974		0.0417	0.0447**		-0.00276	-0.00787		0.0572	0.0753
		(0.110)	(0.105)		(0.0265)	(0.0214)		(0.0821)	(0.0756)		(0.0892)	(0.0858)
N	116	179	295	116	179	295	116	179	295	116	179	295
R2	0.0932	0.275	0.207	0.233	0.172	0.203	0.0720	0.260	0.236	0.102	0.198	0.168
Mean outcome	0.70	0.60	0.64	0.95	0.97	0.96	0.87	0.73	0.78	0.87	0.82	0.84

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. OLS results with fixed effects at the year x tax office level for outcomes from the taxpayer survey. Given the low number of observations in the taxpayer survey, it is not possible to run the regression with year x inspector fixed effects for desk audits, as is done with outcomes extracted from administrative data. Panel A is for all selected firms interviewed in the survey, whereas Panel B is restricted to firms that were audited according to the administrative audit reports. Standard errors are clustered at the year x tax office level. The outcomes are as follows: "Corruption in general" (Columns 1 to 3) is the declared belief of the percentage of audits in Senegal that are affected by corruption; "Corruption experience" (Columns 4 to 6) means that the respondent declared to have experienced an instance of bribery to obtain tax favors; "Grade on honesty" (Columns 7 to 9) is a grade from 0 to 10 to the honesty of the inspector the respondent last interacted with during a full audit; and "Friend at tax authority" captures the degree to which the respondent agrees with the statement that "if the firm's boss has a friend at the tax authority, the firm will rarely be audited". This table is connected to Table 8 discussed in Section 5.3.

G Additional Results on Mechanisms

Table F.1: Association of Risk Score with Audit Outcomes, Controlling for Firm Size

	P(Execution)		P(Detect)|Execution		log(Evasion)|Detection	
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	-0.249***	-0.0699**	0.0121	-0.0729*	-0.987***	-0.319
	(0.0616)	(0.0303)	(0.0250)	(0.0395)	(0.296)	(0.213)
Risk score	0.0200	-0.0149	-0.00407	-0.00259	0.337**	0.380***
	(0.0263)	(0.0187)	(0.0184)	(0.0212)	(0.160)	(0.141)
Alg. x Risk score	0.0105	0.0364*	0.0174	0.0222	-0.288*	-0.416***
	(0.0273)	(0.0196)	(0.0246)	(0.0197)	(0.165)	(0.136)
Lagged log(turnover)	0.0163***	0.00535***	0.00582*	0.00594***	0.0390**	-0.00140
	(0.00263)	(0.00121)	(0.00321)	(0.00216)	(0.0174)	(0.00885)
N	944	2731	507	997	453	732
R2	0.328	0.324	0.164	0.439	0.333	0.436
Mean outcome	0.53	0.37	0.89	0.73	19.29	17.71

This table is identical to Table 9, except that we had lagged log turnover as an additional regressor. This table is discussed in Section 6.1.

Table F.2: Association of Risk Score with Audit Outcomes, Using the Unweighted Risk Score

	P(Execution)		P(Detect)|Execution		log(Evasion)|Detection	
	(1)	(2)	(3)	(4)	(5)	(6)
	Full audits	Desk audits	Full audits	Desk audits	Full audits	Desk audits
Algorithm	-0.250***	-0.0871***	0.0842***	-0.0673	-0.856***	-0.318
	(0.0603)	(0.0280)	(0.0291)	(0.0423)	(0.285)	(0.228)
Unw. Risk score	0.0600*	0.00682	0.0164	0.0151	0.313	0.0909
	(0.0301)	(0.0184)	(0.0204)	(0.0211)	(0.185)	(0.167)
Alg. x Unw. Risk score	-0.0417	0.0346	-0.0423	0.00258	-0.279	0.0218
	(0.0377)	(0.0236)	(0.0344)	(0.0298)	(0.214)	(0.177)
N	944	2731	507	997	453	732
R2	0.266	0.318	0.154	0.431	0.315	0.428
Mean outcome	0.53	0.37	0.89	0.73	19.29	17.71

This table is identical to Table 9, except that we use the raw risk score, prior to reweighting in by turnover, as a regressor. This table is discussed in Section 6.1.

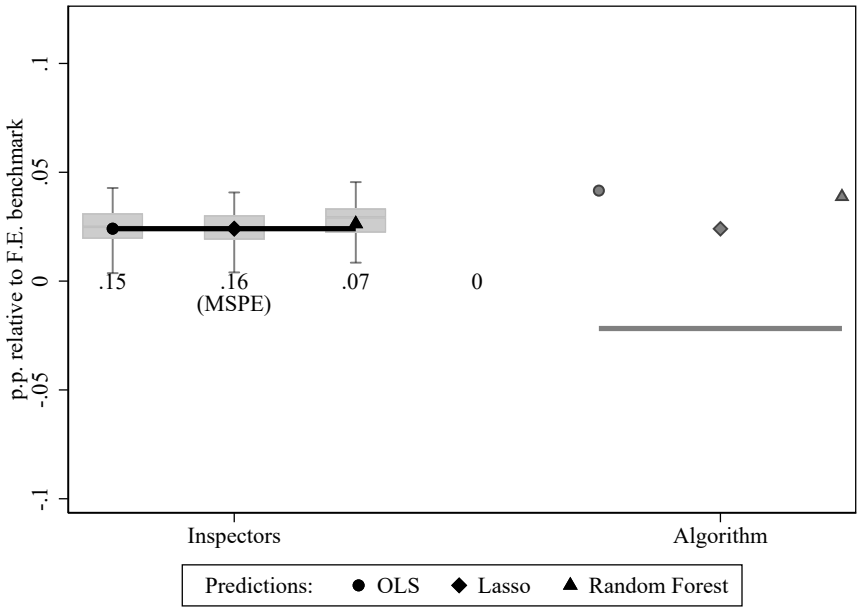
Table F.3: Association of Risk Score Components with Audit Outcomes

	P(Execution)		P(Detect)|Execution		log(Evasion)|Detection							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Full audits	Full audits	Desk audits	Desk audits	Full audits	Full audits	Desk audits	Desk audits	Full audits	Full audits	Desk audits	Desk audits
Algorithm	-0.276***	-0.252***	-0.0697**	-0.0735**	-0.0131	-0.00704	-0.0274	-0.0302	-1.091***	-1.011**	0.0323	0.182
	(0.0553)	(0.0493)	(0.0281)	(0.0310)	(0.0376)	(0.0451)	(0.0419)	(0.0420)	(0.359)	(0.405)	(0.208)	(0.227)
VAT anomalies risk score		0.0203		-0.00346		0.0256		0.00295		0.0896		0.0135
		(0.0257)		(0.0168)		(0.0174)		(0.0174)		(0.149)		(0.144)
CIT anomalies risk score		-0.0362		0.0162		-0.00784		0.00170		0.109		-0.128
		(0.0363)		(0.0186)		(0.0222)		(0.0231)		(0.217)		(0.170)
Inconsistencies risk score	0.0351	0.0380	0.0134	0.0151	0.0184	0.0189	-0.000136	-0.00350	0.264	0.251	-0.134	-0.121
	(0.0245)	(0.0254)	(0.0145)	(0.0144)	(0.0213)	(0.0221)	(0.0244)	(0.0249)	(0.193)	(0.209)	(0.138)	(0.145)
Anomalies risk score	0.0339	0.0490	0.0128	0.00700	0.0269	0.0164	-0.0148	-0.0216	0.180*	0.0826	-0.0264	0.0406
	(0.0273)	(0.0427)	(0.0141)	(0.0198)	(0.0170)	(0.0215)	(0.0214)	(0.0253)	(0.104)	(0.197)	(0.125)	(0.214)
Algorithm x Inconsistencies	-0.0163	-0.0252	0.0117	0.0103	-0.0154	-0.0176	-0.000898	0.00429	-0.375*	-0.385*	0.179	0.190
	(0.0337)	(0.0337)	(0.0196)	(0.0192)	(0.0375)	(0.0382)	(0.0297)	(0.0302)	(0.191)	(0.199)	(0.177)	(0.185)
Algorithm x Anomalies	-0.0445	-0.0534	-0.0200	-0.0152	-0.0414*	-0.0301	0.0127	0.0161	-0.0136	0.0631	0.0558	-0.157
	(0.0313)	(0.0524)	(0.0183)	(0.0255)	(0.0222)	(0.0251)	(0.0257)	(0.0310)	(0.148)	(0.227)	(0.157)	(0.252)
Algorithm x VAT anomalies		-0.0416*		-0.00268		-0.0304		0.0247		-0.162		0.00913
		(0.0244)		(0.0202)		(0.0205)		(0.0255)		(0.169)		(0.157)
Algorithm x CIT anomalies		0.0606		-0.00544		0.0148		-0.0341		0.0483		0.400**
		(0.0456)		(0.0219)		(0.0294)		(0.0278)		(0.237)		(0.196)
N	944	944	2731	2731	507	507	997	997	453	453	732	732
R2	0.265	0.271	0.318	0.319	0.153	0.157	0.430	0.435	0.325	0.330	0.425	0.432
Mean outcome	0.53	0.53	0.37	0.37	0.89	0.89	0.73	0.73	19.29	19.29	17.71	17.71

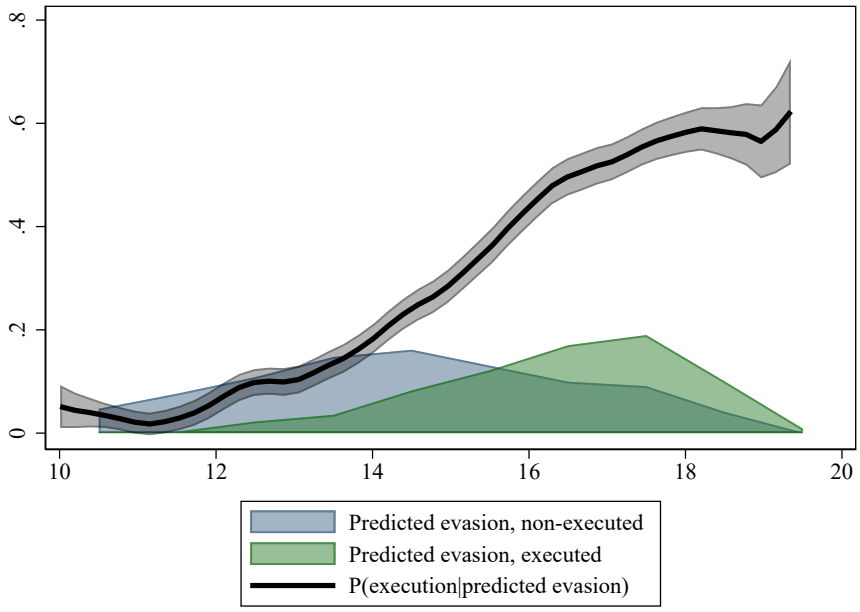
This table is similar to Table 9, but shows the association between audit outcomes and components of the risk score. This table is discussed in Section 6.1.

Figure F.1: Predicting Inspector Choices and Audit Outcomes Using Observable Characteristics

Panel A: Predicted and Realized Execution of Cases



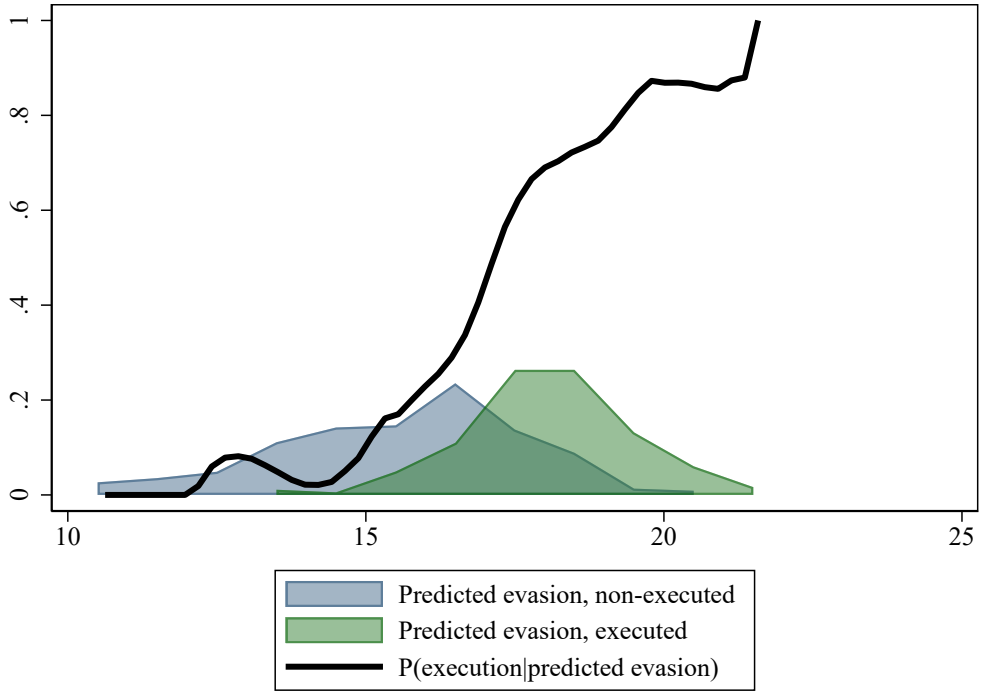
Panel B: Predicted and Realized Detected Evasion Amount



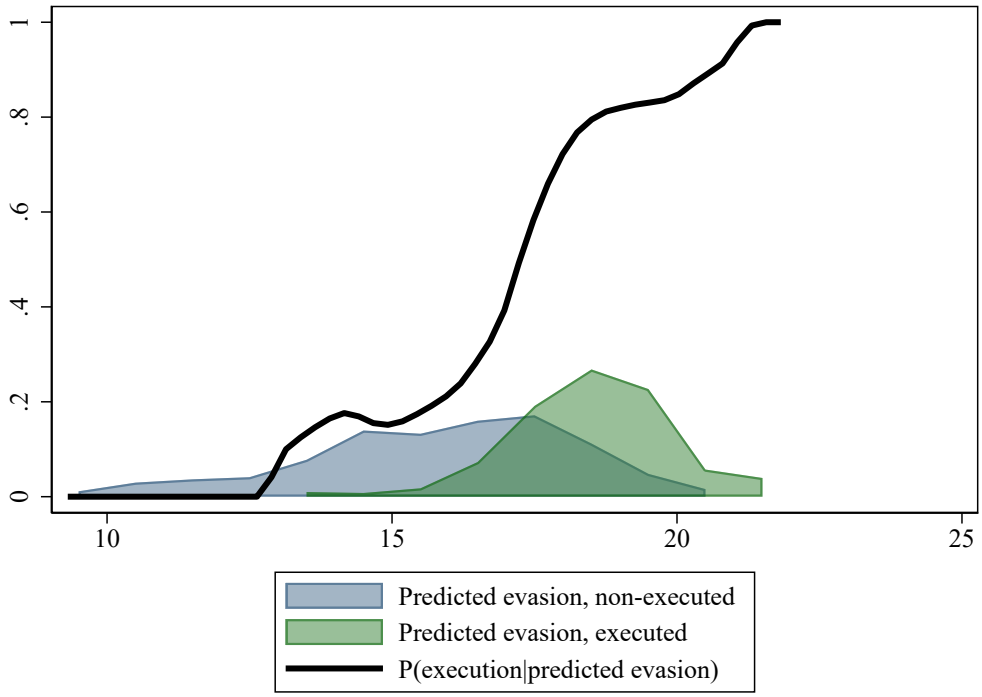
Notes: This figure is identical to Figure 4 but focuses on desk audits rather than full audits.

Figure F.2: Predicted and Realized Detected Evasion Amount

(a) Only Algorithm Cases

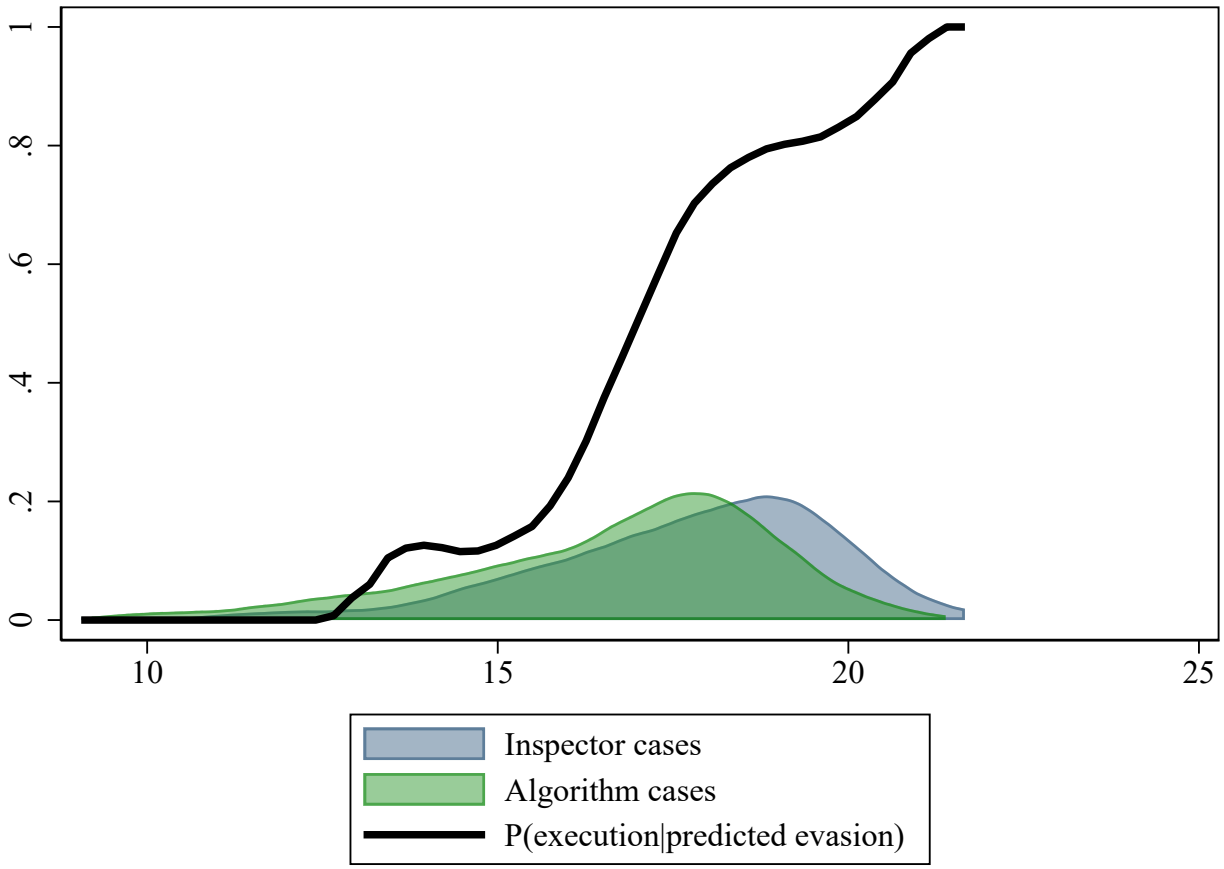


(b) Only Inspector Cases



Notes: This figure is the same as Figure 4, Panel B, but distinguishes between audits selected by the algorithm (panel a) and audits selected by inspectors (panel b). This figure is discussed in Section 6.4.

Figure F.3: Predicted Detectable Evasion Amount for All Cases in the Lists



Notes: This figure plots the distribution of detectable evasion for all full audits (log of FCFA amounts) estimated with a random forest model trained on the executed cases. The graph plots the distributions of the predicted amounts for the algorithm and the inspectors' lists. The black solid line plots the distribution of the realized audit execution rate by level of predicted evasion. This figure is discussed in Section 6.4.

Table F.4: Predictive Coefficients of Execution for Algorithm and Inspector Cases

	Algorithm	Inspector
	(1)	(2)
L1 turnover	-0.184**	-0.208***
	(0.0768)	(0.0692)
L1 turnover sq.	0.0171**	0.0197***
	(0.00784)	(0.00633)
L1 turnover cu.	-0.000368*	-0.000428***
	(0.000196)	(0.000143)
L1 profit rate	-0.0711	0.0790
	(0.224)	(0.156)
L1 profit rate sq	-0.0782	-0.00493
	(0.126)	(0.163)
L1 profit rate cu	-0.0321	-0.111
	(0.277)	(0.238)
L1 productivity	0.0473	0.0382
	(0.0523)	(0.0370)
L1 productivity sq	-0.00501	-0.00735
	(0.0107)	(0.00806)
L1 productivity cu	0.000121	0.000340
	(0.000537)	(0.000421)
Distance km	0.0474	0.102
	(0.0918)	(0.0645)
Distance min	-0.0287	-0.0512**
	(0.0387)	(0.0245)
Firm's age	-0.000860	0.0212
	(0.0157)	(0.0201)
N	407	537
R2	0.347	0.373

Notes: * 0.10 ** 0.05 *** 0.01 significance levels. This table shows the correlation of firm characteristics and the execution rate of the full audit in two different OLS regressions. Column 1 shows the relationship between firm characteristics and the execution of audit cases selected by the algorithm. Column 2 shows the same for cases selected by the inspectors. The regressions include year and tax office fixed effects. Standard errors are clustered at the year and tax office level. This table is discussed in Section 6.3.

Table F.5: Prediction Model Using Ranking of Turnover Only

	Full Audits			Desk Audits		
	All executed	Algorithm	Inspectors	All executed	Algorithm	Inspectors
DGE	69	58	73	43	62	48
CME1	86	82	88	65	66	70
CME2	82	79	81	73	71	80
CPR	73	57	88	44	43	53
DSF	54	53	63	62	68	61

Notes: This table shows the percentage of audited cases that are correctly predicted by a simple turnover ranking within tax office and year, using declared turnover of the year prior to selection. The columns show the percentages of the N executed cases within each office that are also among the N largest cases in the selection list. For the algorithm and inspectors' columns, N is defined as the number of executed cases among algorithm-selected and inspector-selected cases. This table is discussed in Section 6.3.