CORRUPTION IN CUSTOMS*

Cyril Chalendard
Ana M. Fernandes
Gael Raballand
Bob Rijkers

This article presents a new methodology to detect corruption in customs and applies it to Madagascar's main port. Manipulation of assignment of import declarations to inspectors is identified by measuring deviations from random

*We thank the editor and five referees for excellent comments that greatly improved the article. We thank Madagascar Customs and GasyNet for sharing their data, feedback, and willingness to collaborate. We are especially grateful to former director general Éric Rabenja, who courageously implemented a host of anticorruption reforms, at high personal cost. The support of Dannick Fiononana, Antsa Rakotoarisoa, Cédric Catheline, Stéphane Manouvrier, Prisca Michea, Tolotra Hasinantenaina Ramarosandratana, and Tahina Rabesalama was critical to the success of this project. We thank Paula Irene Londono Suarez, Nicolas Santos, and Tomas Rossetti for superb research assistance. Aaditya Mattoo provided extremely useful comments and guidance, especially in the early stages of the project. We thank Thomas Cantens, Vianney Dequiedt, Simeon Djankov, Henrik Kleven, Beata Javorcik, Joana Naritomi, Tristan Reed, Sandra Sequeira, Joana Silva, Joel Slemrod, René Tapsoba, Roy van der Weide, Shang-Jin Wei, and seminar/conference participants at the 2018 ABCDE conference, the 2021 Canadian Economic Association Annual Meeting, CESIfo, FREIT and Georgia Institute of Technology Empirical Trade Online Seminar 2020, NUS, the Tinbergen Institute, UNU-WIDER, the 2018 Midwest Economic Theory and International Trade Conference, the 2020 NBER Development Economics Program Meeting, the World Bank, and the 2017 World Customs Organization PICARD Conference, for useful comments and discussions. This article has been partly supported by the Umbrella Facility for Trade, a World Bank–administered trust fund (financed by the governments of the Netherlands, Norway, Sweden, Switzerland, and the United Kingdom) and the World Bank Strategic Research Partnership on Economic Development. We also acknowledge financial support from the World Bank research support budget and the Knowledge for Change Program (KCP), a trust-funded partnership in support of research and data collection on poverty reduction and sustainable development housed in the office of the Chief Economist of the World Bank. The findings, interpretations, and conclusions expressed herein are entirely those of the authors. They do not necessarily represent the views of the International Bank of Reconstruction and Development/World Bank and its affiliated organizations, or those of the executive directors of the World Bank or the countries they represent. The authors were employed by the World Bank at the time of writing this paper and collaborated with Madagascar customs in the context of a World Bank-financed project (Project ID: P150116). The views expressed herein are those of the authors and do not necessarily reflect the views of the International Trade Centre and the United Nations. All errors are our responsibility.

The Quarterly Journal of Economics (2023), 575–636. https://doi.org/10.1093/qje/qjac032.
Advance Access publication on September 9, 2022.

assignment prescribed by official rules. Deviant declarations are more at risk of tax evasion, yet less likely to be deemed fraudulent by inspectors, who also clear them faster. An intervention in which inspector assignment was delegated to a third party validates our approach, but also triggered a novel manifestation of manipulation that rejuvenated systemic corruption. Tax revenue losses associated with the corruption scheme are approximately 3% of total taxes collected and are highly concentrated among a select few inspectors and brokers. *JEL Codes:* F13, D73, H26.

I. Introduction

State capacity to raise tax revenue is an important enabler of development (Besley and Persson 2009). Poorer countries mobilize less tax revenue as a share of GDP (Gordon and Li 2009) and suffer higher levels of corruption. While tax evasion and weak bureaucratic performance are salient drivers of the differences in revenue mobilization across the development spectrum (Finan, Olken, and Pande 2017, Khan, Khwaja, and Olken 2016, 2019), less is known about who evades (how much) and to what extent evasion is facilitated by (which) bureaucrats. Evidence on the effectiveness of reforms to remedy systemic corruption is also scant.

This article presents a new methodology to detect and quantify the prevalence and costs of a type of corruption scheme in customs and assesses the effectiveness of an intervention intended to eliminate such corruption. Around the world, customs information technology (IT) systems usually prescribe random assignment of incoming declarations to inspectors, conditional on their productivity (in the task of clearing declarations), as a way to deter corruption. Our approach identifies potential manipulation of inspector assignment by evaluating whether certain inspectors are paired excessively frequently with certain customs brokers, deviating from what conditional random assignment would predict. To assess whether these deviations reflect corruption, we subsequently examine whether excess interaction between inspectors and brokers is associated with an increased risk of tax evasion and whether deviant declarations are treated preferentially by inspectors. We quantify the resulting tax revenue losses and their distribution across inspectors and brokers. The methodology is validated by studying the impact of an intervention that delegates the (randomization of) inspector assignment to a third-party organization.

We apply our approach to Madagascar's main port, Toamasina, which provides a suitable setting for studying corruption in customs. First, like many other developing countries, Madagascar is heavily reliant on tax revenues collected at the border (Baunsgaard and Keen 2010), which account for 48% of total tax revenues. Toamasina collects more than three-quarters (78%) of nonoil tax revenues and employs a limited number of inspectors. Each inspector oversees the collection of 1.3% of total yearly taxes in Madagascar. Second, corruption appears rife in customs. A survey of inspectors reveals that only 6% believe that nonethical conduct is sanctioned, and only 23% believe that their colleagues act with integrity. Third, inspectors repeatedly interact with a limited number of brokers, with whom they also share social ties. The combination of high stakes, a small number of players, limited sanctions for improper conduct, and extensive repeated interactions is conducive to corruption. Last but not least, Madagascar's senior customs management were willing to undertake reforms to curb corruption and provided us with unprecedented data access. They shared data for the period 2015–2018, covering rich details on each import declaration, including declared value and weight, weight measured upon arrival at the port, taxes paid, the identity of the broker who registered it and the inspector assigned to it, whether fraud was recorded, all revisions to inspector assignment, value, weight, and tax liability made during the clearance process, as well as risk management information (inspection channel, risk scores, and valuation advice).

Our methodology comprises three steps. First, we detect potential manipulation of inspector assignment by identifying pairings of inspectors with brokers that occur much more frequently than would be expected on the basis of conditional random assignment. In Toamasina, 10% of all declarations are handled by inspectors whose assignment contravened the random inspector assignment prescribed by official rules. Second, these deviant declarations are shown to have characteristics commonly associated with an elevated risk of tariff evasion and to embody sizable potential tax revenue losses. Third, we demonstrate that inspectors treat preferentially the declarations registered by brokers with whom they interact excessively frequently, ceteris paribus. They clear them faster, are less likely to deem them fraudulent, and impose lower weight, value, and tax adjustments, thus exacerbating disparities in tax revenue losses between deviant and nondeviant declarations. These findings are robust to a variety of checks, including the use of inspector-specific binomial logit models to detect deviations from random assignment while accounting for fluctuations in inspectors' schedules, using a propensity score matching approach to account for selection, using different samples and controlling for various sets of fixed effects. According to back-of-the-envelope calculations, average tax revenue per nonrandomly assigned declaration would have been 26% higher in the absence of excess interaction. Total tax revenues collected in Toamasina would have been 3% higher.[1]

We argue that these patterns are consistent with a corruption scheme in which brokers bribe staff in the customs IT department and/or the customs port manager to be paired with their preferred inspector, who agrees to clear the declarations that are the object of corruption faster, not to impose tax adjustments and penalties, and not to insist on upward adjustment (or to request just a marginal one) of the customs declared value. The resulting tax savings are presumably shared with inspectors. Although we do not directly observe bribe payments, our findings are consistent with extensive circumstantial evidence collected during repeated field visits, IT audits, and a survey of customs inspectors. Based on our findings, Madagascar's customs management sanctioned inspectors for corruption, suspended the head of the IT department, and reformed inspector assignment by divesting it to a third party outside customs. This delegated randomization made the third party responsible for inspector assignment. Using its own software, the third party randomly assigned declarations. It was so successful in eliminating deviations from random inspector assignment that delegated randomization became standard practice.

Explanations other than corruption are difficult to reconcile with the totality of the observed patterns. They also fail to explain why the delegated randomization intervention virtually eliminated the prevalence of deviations from random inspector assignment. IT manipulation resurfaced after a few months, albeit in a different guise. Customs IT staff figured out a new way to manipulate inspector assignment and bypass the delegated randomization. This bypassing was identified by assessing whether the entire set of declarations registered by brokers was shared with the third party for inspector random assignment. We show that 7.2% of all import declarations were withheld from delegated randomization.[2] The circumvention of the delegated

1. As discussed in Section IV, our approach does not capture all forms of corruption, hence these estimated losses reflect only the tax revenue losses associated with the specific corruption scheme we document.

2. In practice, such bypassing appears to have been operationalized through the temporary disabling of a randomization trigger, such that all declarations

randomization not only attests to the difficulties inherent in dislodging systemic corruption but also provides variation in exposure to the delegated randomization intervention.

The bypassing resulted in the resurgence of excess interaction between inspectors and brokers, driven exclusively by withheld declarations.[3] Interestingly, withheld declarations were disproportionately assigned to inspectors with whom brokers had interacted excessively frequently in the period before the delegated randomization intervention, suggesting persistence in the corruption scheme we unveil. These withheld declarations were on average more risky, subject to higher taxes, more undervalued, and embodied larger tax revenue losses, especially when their eventual (nonrandom) assignment resulted in excess interaction between inspectors and brokers. Inspectors only provided preferential treatment to withheld declarations if registered by brokers with whom they interact excessively frequently. These findings validate our methodology and are consistent with our interpretation that the documented patterns reflect corruption.

This article builds on and aims to contribute to several strands of literature. First, by presenting a methodology that can help detect tampering with random assignment, we aim to contribute to the literature on the detection and measurement of corruption and its development consequences (Shleifer and Vishny 1993, 2002; Bardhan 1997; Olken and Pande 2012; Zitzewitz 2012). Random assignment of declarations to inspectors is not only the norm in customs agencies across the globe but is also used to prevent corruption in many other settings, including the assignment of cases to judges and prosecutors.[4] We believe our approach can fruitfully be adapted to other contexts.

registered during specific time intervals when this trigger was deactivated were withheld from being sent to the third party to be randomized (including those that were the subject of a corruption agreement).

3. Excess interaction was not observed for declarations handled by inspectors whose assignment was randomized.

4. More than 100 customs agencies have adopted the customs clearance IT system used by Madagascar (Automated System for Customs Data, ASYCUDA) for which the default option for assignment of declarations to inspectors is random assignment (according to workload). Customs agencies that do not use ASYCUDA also typically use random inspector assignment. Random assignment of cases to judges to deter judicial corruption has been adopted by 162 countries (World Bank, 2020).

Second, we contribute to the nascent literature on the performance of bureaucrats as a determinant of state effectiveness and tax collection (Olken and Pande 2012; Dincecco and Ravanilla 2017; Pepinsky, Pierskalla, and Sacks 2017; Xu 2018; Xu, Bertrand, and Burgess 2018), by highlighting the granularity of tax evasion and showing how the behavior of a select few actors has macro-fiscal ramifications. While corruption was systemic and enabled by most inspectors (10 out of 16 in a typical semester), IT staff, and the port manager, the tax revenue losses were very concentrated among a select few inspectors and brokers. In any given semester, the top two most corrupt inspectors accounted for 55% of the tax revenue losses associated with the corruption scheme we document.

Third and related, we contribute to the literature on the determinants of tax enforcement (Kleven et al. 2011; Pomeranz and Vila Belda 2019; Slemrod 2019), and specifically the literature on tariff evasion (Bhagwati 1964; Fisman and Wei 2004; Yang 2008a, 2008b; Dutt and Traca 2010; Sequeira and Djankov 2014; Sequeira 2016; Rijkers, Baghdadi, and Raballand 2017; Wier 2020) by pinpointing which brokers and inspectors cheat and which import declarations are most likely to be undervalued. The propensity to participate in the corruption scheme is higher for brokers based in Toamasina and rises with inspectors' tenure in the port, suggesting that private information and personal relationships are important enablers of evasion. Despite accounting for larger tax revenue losses, the most corrupt inspectors paradoxically collected more tax per declaration than less corrupt ones because manipulation of assignment enabled them to control the assessment of the most lucrative declarations with the highest potential tax yield. Corruption is thus positively correlated with (naively measured) tax yield.

Fourth, our results also dovetail with the literature on the effectiveness of anticorruption interventions (e.g., Ferraz and Finan 2008; Niehaus and Sukhtankar 2013) by demonstrating that IT solutions can help curb corruption (see also Laajaj, Eslava, and Kinda 2019), but are not a panacea (see also Casaburi, Kremer, and Ramrattan 2019) because they can also serve as a conduit to it. Our evidence of a new form of IT manipulation after the reform to inspector assignment is consistent with Shleifer and Vishny's (1993) observation that corruption is difficult to dislodge when both parties benefit (as in corruption with theft). It also complements Yang (2008a), who shows how a customs reform that increased enforcement against a specific type of tariff duty evasion resulted in the use of an alternative duty-avoidance method (shipping via duty-exempt export processing zones).[5]

Finally, our findings are relevant for the understanding of trade costs, market distortions, and competition in developing countries (Atkin and Khandelwal 2020) and the debate as to whether corruption greases the wheels of the economy (e.g., Leff 1964; Kaufmann and Wei 1999; Banerjee, Hanna, and Mullainathan 2012; Freund, Hallward-Driemeier, and Rijkers 2016). We complement Sequeira (2016)'s findings that in the presence of corruption, tariffs and other import taxes may not be as burdensome as they appear on paper by showing that corruption is also associated with faster clearance. Expedited clearance limits the risk of detection and is another margin by which corruption affects trade costs and competition.[6]

The remainder of this article is organized as follows. Section II describes the context and the customs clearance process, while Section III presents our data. Section IV describes our methodology to detect deviations from official rules in inspector assignment to declarations. Section V examines whether deviant declarations are at a higher risk of tax evasion. Section VI assesses whether there is differential treatment of deviant declarations by inspectors. Section VII provides estimates of the costs of corruption in terms of tax losses. Section VIII characterizes the inspectors and brokers who are corrupt and the distribution of tax revenue losses. Section IX validates our approach by analyzing the impact of the delegated randomization intervention. Section X concludes.

II. Context: Customs Clearance Process in Madagascar

This section describes the customs clearance process and argues that the conditions in Toamasina are conducive to systemic

5. Our study's displacement is mediated by interactions between private-sector parties (brokers) and bureaucrats (inspectors) and is thus close to Lichand and Fernandes (2019) who document selection in the pairing of vendors and bureaucrats in response to changes in (perceptions of) enforcement of anticorruption measures.

6. Note, however, that our results are not informative about the overall impact of corruption on clearance times. In theory, inspectors who participate in the scheme could attempt to extort nonparticipating firms by protracting the clearance process. In practice, however, clearance times are fairly short even for declarations that are not covered by the scheme—about 20 hours on average—so prima facie the data are not suggestive of substantial extortion.

corruption: there are few players who interact repeatedly, the stakes are high, and there is almost no punishment for improper conduct.

Taxes and duties collected by customs accounted for 48% of overall tax revenue in Madagascar in 2019, despite substantial tariff evasion (Chalendard, Raballand, and Rakotoarisoa 2019). Most of this revenue was collected in Toamasina, which accounted for 78% of nonoil tax revenue and 52% of nonoil imports and employed on average 16 inspectors per year during our sample period. Each inspector oversees the collection of US$17 million worth of tax revenue per year on average, representing 1.3% of total taxes collected.

Jobs in the customs administration—especially inspector jobs in Toamasina—are among the most sought-after jobs in Madagascar. They are secure, well paid, and offer several benefits. Inspectors earn a salary of roughly US$11,000 per year (21 times annual GDP per capita of US$527) and receive as bonus 5% to 20% of the tax adjustments they impose when they detect noncompliance.[7] They can also earn performance bonuses of up to $1,000 per quarter if they are among the top inspectors in terms of clearance speed, fraud detection, and tax revenue mobilization. Inspectors thus get paid efficiency wages and have strong personal financial incentives to detect noncompliance, which should help deter corruption.

However, these performance rewards may not sufficiently incentivize inspectors to act with integrity. Corruption appears pervasive, possibly due to the virtual absence of sanctions for improper conduct, threats from economic operators, and because compensation is low relative to opportunities for graft (Chalendard et al. 2020). According to a nationwide survey of inspectors that we conducted in 2017, only 23% believe that their colleagues act with integrity, only 6% claim that nonethical behavior is sanctioned, and only 12% believe that promotions are merit-based. Close to one-third of inspectors claim being subjected to threats from economic operators on a regular basis. Undervaluation of imports was widely agreed to be the main type of customs fraud in Madagascar.[8]

7. Note that inspectors' pay does not vary mechanically with the total taxes they collect.

8. Administrative data on fraud records classify 67.2% of all fraud in Madagascar customs as underreporting of value, 27.4% as underreporting of quantities,

The inspectors in Toamasina interact with a limited number of customs brokers (*commissionnaires agrées en douane*). In a typical semester, there are on average 45 brokers who each handle 173 declarations from 33 different importers.[9] The overwhelming majority of the 3,660 importers in our sample work exclusively with one broker, as is shown in Online Appendix Table A2. Brokers must have a license, which is issued by the customs administration, and they administer the customs clearance process on behalf of the importer by fulfilling customs formalities and submitting documentation.[10] Brokers are accountable for paying taxes, duties, and potential tax adjustments and are penalized (with a fine) in case of noncompliance. In principle, repeated noncompliance can result in the revocation of the broker's license. In practice, suspension of brokers due to misconduct is rare. Customs officials and brokers frequently socialize and are part of the narrow elite in the small town of Toamasina. Many brokers either have served as customs officials themselves or deliberately recruit former customs officials because of their expertise and networks. Thus, there is extensive repeated interaction between inspectors and brokers, both inside and outside of the customs premises.

There is significant information asymmetry between importers and brokers given that the latter are much better informed about customs procedures and are the first point of contact for customs in case disputes arise. Some brokers have transparent pricing schemes that typically depend on the size and contents of the cargo, but others charge a fixed amount (inclusive of potential tax liabilities) per container cleared, irrespective of its content, which implies that their profits directly depend on the amount of tax they remit on behalf of the importer.

To understand how corruption may happen, it is instructive to consider the customs clearance process, a stylized version of which is depicted in Figure I.

---

and the remainder as product misclassification (4.9%) or misreporting the country of origin (0.5%).

9. These averages do not consider small brokers (i.e., those handling fewer than 50 declarations per semester) since they will not be part of our estimating sample (described in Section III).

10. A very small number of importers—six in our sample—obtained their own broker's license and simultaneously act as importer and broker. They handle 3.2% of the declarations.

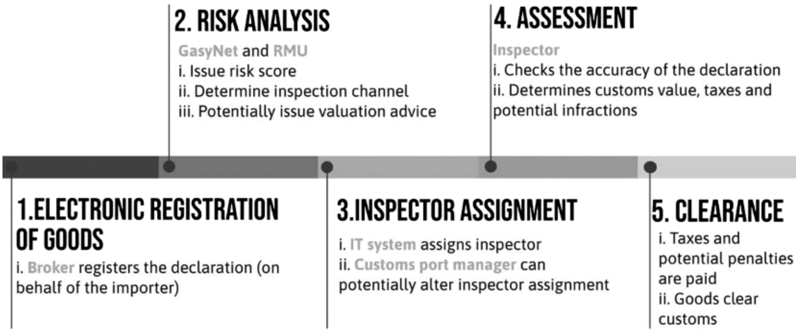

FIGURE I

Stylistic Representation of the Clearance Process

The figure depicts a stylized representation of the customs clearance process. RMU is the risk management unit of customs. GasyNet is a third party that assists customs with risk analysis and logistics.

i. **Registration.** The first step in the process is the electronic registration of an import declaration by the broker on behalf of the importer via the Automated System for Customs Data (ASYCUDA)++ customs clearance IT system.[11]
ii. **Risk analysis.** The second step consists in risk analysis conducted by both GasyNet, a third-party service provider that assists Madagascar customs with risk analysis and logistics, and the customs risk management unit.[12] For each declaration, (a) a risk score is issued based on GasyNet's proprietary risk model, (b) a clearance channel is recommended along with a qualitative justification. If the yellow channel is selected, the inspector only needs to check the documentation. If the red channel is selected, the inspector is expected to physically inspect the cargo. However, the inspector is at liberty to change the clearance channel based on their own judgment. In addition, (c) for a very small subset of high-risk declarations for which the accuracy of the declared import value is questionable,

11. ASYCUDA is an integrated customs management system developed by United Nations Conference on Trade and Development.

12. In reality the second step (risk analysis) and the third step (inspector assignment) happen simultaneously.

GasyNet issues a valuation advice: a detailed report on what the value of the specific declaration is likely to be.

iii. **Inspector assignment.** The third, and for our purposes crucial, step is the assignment of the declaration to a particular inspector by the ASYCUDA IT system. Official rules prescribe that a newly registered declaration should be assigned to whichever inspector has the lowest workload (i.e., has the fewest pending declarations on their desk) and is active (i.e., is connected to the IT system and can therefore receive new declarations). Official rules do allow for productivity differences across inspectors: a highly productive inspector will get, on average, more declarations than a poorly productive inspector. Yet the assignment of declarations to inspectors is supposed to be random conditional on their productivity. We exploit this feature of the official rules for identifying deviant declarations in Section IV. However, the customs port manager, the Chef des Opérations Commerciales (COPCO), has the authority to override the IT system's initial assignment and reassign a declaration to a different active inspector. Such reassignments are warranted in case of unanticipated absenteeism (due to illness for example) and should, a priori, happen only randomly.[13]

iv. **Assessment.** The fourth step is the assessment of the declaration by the assigned inspector based on the documentation submitted by the broker on behalf of the importer, the risk analysis diagnostics provided by the risk management unit and GasyNet, and the results of a potential physical inspection. She has to decide which (if any) adjustments to the import value, quantity, product classification, and/or origin are to be made and report whether fraud was perpetrated. She then assesses what duties, taxes, and potential penalties are to be paid based on the (potentially revised) final value and product classification of the import declaration.

v. **Clearance.** In the final step in which goods are cleared, the importer (or the broker on behalf of the importer) pays the taxes, duties, and potential penalties, and goods are released from customs.

13. Such reassignments occur for 6% of the import declarations.

Our analysis of corruption will focus on manipulation of the assignment of declarations to inspectors (by IT department staff and/or the customs port manager) done in step iii, and on differential treatment of manipulated declarations by inspectors during assessment in step iv.

III. Data

Our study combines the following databases.

- **Customs transactions data.** From Madagascar customs we obtained administrative data tracking imports at the transaction level for the period January 2015–November 2018. For each import declaration, the data cover the HS eight-digit products included (designated as items), their source country, the dates/times of registration, assessment, clearance, the broker, the importing firm, and the customs inspector assigned to handle the declaration. For each item, the data contain information on the initially declared and the finally registered import value, weight, and taxes paid (tariff and value-added tax as well as exemptions). These variables enable us to evaluate inspector modifications of value, weight, and tax liabilities. In addition, for each declaration we can track any modifications made to the IT system's initial inspector assignment by the customs port manager.[14] This will allow us to disentangle the role of IT department staff from that of the customs port manager in generating deviations from official rules in inspector assignment.
- **Fraud records.** Fraud records were provided by the Legal Department (*Service des Affaires Juridiques et du Contentieux*). For each declaration, we know both whether and what type(s) of fraud was detected and the amount of taxes recovered (if any). Information on whether and how much inspectors modified tax yield is important for assessing the role of inspectors.
- **Risk management data.** From the customs risk management unit, we received for each import declaration information on the initial and finally used clearance channel

14. These data were obtained from the customs administration's internal control systems and merged with the transaction data.

(documentary control/yellow channel, physical inspection/red channel, or no inspection/blue channel). From GasyNet we received the risk score assigned to each import declaration (related to the risk of noncompliance with customs regulations ranging from 1 to 9) and valuation advice in case it was issued.

- **Container weight measurement data.** We obtained from the company in charge of managing Toamasina's container terminal—Madagascar International Container Terminal Services Limited (MICTSL)—data on the weight of containers that arrive in Toamasina as measured by weighing at a scale upon arrival for the period 2015–2017. This port authority weight data is merged with the customs data at the declaration level, for declarations whose goods fill completely one or more containers. For declarations that share containers with other declarations this information is missing. These port authority weight data provide a useful benchmark for verifying whether the weight registered by the broker is correct.
- **UN COMTRADE data.** We rely on an international trade data source UN COMTRADE to obtain export flows—values and quantities (weight)—at the country-HS six-digit-year level for all of Madagascar's trading partners in 2015–2018. We use these mirror data for flows imported by Madagascar to construct exogenous benchmark/reference prices, to which we will compare the unit prices of the items included in the import declarations in the Madagascar customs data (as will be described later).
- **Delegated randomization of inspector assignment and IT manipulation.** On November 18, 2017 the assignment of inspectors to declarations was delegated to GasyNet. By comparing daily their list of declarations (that their system randomly assigned to some inspector) to the list of declarations that cleared customs from the customs administration, GasyNet was able to identify declarations that were withheld from the delegated randomization—as discussed in Section IX. They provided us with the list of withheld declarations.
- **Human resources data.** Information on inspectors' gender, education, age, and date of entry into work for the customs administration were provided by the Human

Resources Department (*Direction des ressources et de la formation*).

- **Inspectors' survey.** In 2017 we conducted a nationwide survey of inspectors that contained questions on job and pay satisfaction, corruption, ethics, fairness, and interactions with brokers and importers.

Madagascar's raw customs data covers all formal import transactions made under several regimes: final imports for consumption (imports for home use), reimports, temporary admissions, inward processing, warehouse, and other. Our analysis focuses on import declarations subject to taxation and to a physical or a documentary control by customs inspectors in Toamasina.[15] This implies focusing only on imports for home use and reimports and excluding declarations from importers that are members of the *Procédure Accélérée de Dédouanement* (PAD), a trusted trader program that allows member firms to benefit from expedited clearance procedures with minimal controls at the border. To minimize the risk of identifying as likely to be suspect of corruption declarations that are not, we remove from the sample (i) declarations registered by brokers that do not interact frequently with customs (i.e., brokers that register fewer than 50 declarations per semester); (ii) declarations assigned to inspectors that relocate to or move away from Toamasina during a given semester but are active for fewer than two consecutive months in that semester.[16] Our final sample accounts for an average of 76.9% of declarations, 78.9% of collected taxes, and 76.5% of total import value for import declarations subject to taxation and to a physical or a documentary control cleared in Toamasina across the period ranging from January 1 2015 to November 17 2018.[17]

To analyze which declarations are most likely to be subject to corruption agreements, we will use measures of excess interaction

15. Imports subject to specific clearance procedures (oil and vehicles) are excluded.

16. Our sample also excludes observations in the top and bottom 2.5% of the yearly distribution of initial average internal reference price, defined as the weighted average of internal reference prices for all items included in the import declaration with weights being the initially submitted weights. The internal reference price for each item is the median unit price (ratio of value to weight) reported across Malagasy importers for a given HS eight-digit–origin country–year.

17. Our sample ends one year after the start of the delegated randomization of inspector assignment to the third party and a few days before the unveiling of the IT manipulation taking place during this delegated randomization.

between inspectors and brokers as proxies for IT manipulation described in Sections IV and IX. The definition of all variables is provided in Online Appendix Table A1. Here we briefly describe the declaration-level customs outcomes on which we will estimate the effects of corruption. These are clearance time (measured as the log number of hours from the time the declaration was (last) assigned to an inspector to her assessment of the declaration), a dummy for whether or not fraud was recorded, the change in log value (finally registered − initially declared), tax adjustment, and hypothetical tax revenue losses described below. As additional declaration-level customs outcomes used in robustness exercises we consider the change in log weight (finally registered − initially declared) and the gap between the port authority weight and the initially declared weight (for simplicity called weight gap).

Hypothetical tax revenue losses for a declaration are computed based on the difference between hypothetical tax yield and actual tax yield. Measuring hypothetical tax yield is notoriously challenging given that it is unobserved. Our baseline measure of a declaration's hypothetical tax yield considers as a reference price for each of its items the median unit price (ratio of value to weight) reported across Malagasy importers for the same origin country and year. For each item included in the declaration, the relevant reference price is multiplied by the item's weight and the item's tax rate. Summing the resulting hypothetical item-level tax yield across all items included in the declaration yields the declaration-level hypothetical tax yield. This is a conservative measure, for it assumes that the median unit price is not itself underreported. Our alternative measure of a declaration's hypothetical tax yield considers as a reference price for each of its HS six-digit products the unit price reported by the exporting country in that year in UN COMTRADE multiplied by the products' weights and by the products' actual tax rates and sums these across all products in the declaration.[18] This measure has the advantage of using prices that are more likely to be exogenous to tax evasion in Madagascar.[19] Two additional measures of hypothetical tax revenue losses

18. An HS six-digit product's weight is obtained by summing across the weights of all corresponding items. An HS six-digit product's tax rate is obtained as the ratio between the sum of actual taxes and the sum of finally declared import value across all corresponding items.

19. Firms behind a given export flow might conspire with importers in Madagascar issuing fake invoices for them to minimize their tax liabilities. In addition,

are constructed for two subsets of declarations. For declarations for which port authority weight data are available, hypothetical tax yield is constructed also correcting for underreporting of quantities assuming that the measured port authority weight is correct.[20] For declarations for which GasyNet's valuation advice was issued, hypothetical tax yield is constructed as the declaration's advised value multiplied by the average tax rate. As determinants of corruption (and subsequently as controls for evasion risk) we rely on the following ex ante risk characteristics of the import declaration: the tax rate (tariffs and other taxes), the risk score, a dummy for the red channel, a dummy for being a mixed shipment (i.e., one that includes different items), the share of differentiated products as per Rauch (1999)'s classification, and a dummy for receiving GasyNet's valuation advice.[21] In robustness exercises, we consider other declaration characteristics: the log of the initially declared value, the log of the initially declared weight, the initial unit price relative to the median import unit price, and the initial hypothetical tax revenue loss (using as reference price the median import price). Summary statistics on all customs outcomes and declaration characteristics are shown in Online Appendix Tables A3 and A4.

IV. Identifying Deviant Declarations

Our identification of declarations suspected of corruption relies on detecting deviations from official rules in the assignment of incoming declarations to customs inspectors. Recall from Section II that according to official rules, incoming declarations should be randomly assigned to inspectors conditional on their productivity. For each inspector, the likelihood of being assigned any given declaration is proportional to her productivity. These

export unit prices may be downward biased since they are typically recorded as free on board (FOB) whereas import prices are recorded cost insurance freight (CIF) and hence include transportation and insurance costs.

20. We cannot correct quantities declared at the item level because port authority weight is available only for the declaration as a whole. By implication we are assuming that the weight of all items in a declaration is underreported to the same extent.

21. These variables are supposed to be predetermined from the point of view of the inspectors handling the declaration since they are not the ones lodging the declaration on behalf of the importer, nor are they in charge of issuing a risk score or making the first inspection channel recommendation.

rules imply that the process of assigning declarations to inspectors follows a multinomial distribution. Each declaration is assigned to one of K_t possible inspectors, where K_t is the total number of inspectors active in semester t, with corresponding probabilities p_{1t}, p_{2t},., p_{Kt}.[22] These probabilities sum up to 1 because the K_t outcomes are mutually exclusive and can be thought of as reflecting inspectors' relative productivity in semester t. An inspector that is more productive will, on average, handle more declarations than a less productive inspector. Because the marginal distribution of a multinomial distribution is binomial, for each inspector i, the probability of receiving x_{ibt} import declarations from the total number of declarations n_{bt} (where $n_{bt} = \sum_{i=1}^{K_t} x_{ibt}$) registered by broker b in semester t is given by the binomial probability mass function: $P(x_{ibt}|p_{it}, n_{bt}) = \binom{n_{bt}}{x_{ibt}} p_{it}^{x}(1 - p_{it})^{n_{bt}-x_{ibt}}$.

Based on these rules, the share of all declarations that a given inspector handles in a given semester, which we refer to as her inspection share (analogous to the concept of market share in industrial organization), is expected to vary across inspectors, as it depends on their productivity. However, for a given inspector, it should not vary systematically across brokers, unless inspector assignment performed by the IT system did not follow official rules—i.e., was manipulated. All inspectors should have, for a given broker, an inspection share close to their average inspection share.

To assess whether this is indeed the case we consider the import declarations registered by a specific broker during a semester—corresponding to an inspection "market"—and we define for that broker the inspection share of an inspector as the proportion of its declarations handled by that inspector:

$$S_{ibt} = \frac{x_{ibt}}{\sum_{i=1}^{K_t} x_{ibt}}. \tag{1}$$

22. The probability of observing a particular distribution $(x_{1bt}, x_{2bt}, ..., x_{kbt})$ of declarations of a given broker b across inspectors 1, 2, ..., k in semester t (where $\sum_{i=1}^{K_t} x_{ibt} = n_{bt}$, the total number of declarations in semester t registered by broker b) given their productivities $(p_{1t}, p_{2t}, ..., p_{kt})$ is:

$$(x_{1bt}, x_{2bt}, ..., x_{kbt}|p_{1t}, p_{2t}, ..., p_k) = \frac{n_{bt}!}{\prod_{i=1}^{K_t} x_{ibt}!} \prod_{i=1}^{K_t} p_{it}^{x_{ibt}}.$$

Our measure of potential manipulation of inspector assignment is the deviation between actual assignment and random assignment of declarations to inspectors. Specifically, we define the *excess interaction share* ES_{ibt} as the difference between the actual share of broker b's declarations handled by inspector i in semester t (S_{ibt}) and the predicted share ($\overline{S_{ibt}}$) she would be expected to handle if declaration assignment to inspectors followed official rules:

$$ES_{ibt} = S_{ibt} - \overline{S_{ibt}}. \tag{2}$$

To calculate measures of excess interaction between inspectors and brokers we adopt two procedures described in what follows.[23]

IV.A. Calibrating Excess Interaction

A simple procedure to calculate measures of excess interaction is to calibrate predicted inspection shares using the share of all declarations cleared in semester t that were handled by inspector i, and evaluating whether observed inspection shares deviate from these predicted shares. Formally, we set

$$\overline{S_{ibt}} = \overline{p_{it}} = \frac{\sum_{b=1}^{B_t} x_{ibt}}{\sum_{j=1}^{K_t}\sum_{b=1}^{B_t} x_{jbt}}, \tag{3}$$

where $\overline{p_{it}}$ is the predicted probability that a declaration registered in semester t will be handled by inspector i and B_t is the number of brokers having registered at least one declaration in semester t. Figure II illustrates this procedure: it shows overlaid histograms of the observed distribution of the share of declarations of a given broker cleared by a specific inspector in a given semester (the lighter bars) and the calibrated predicted inspection shares just described (the darker bars). A Kolmogorov-Smirnov test rejects

23. An alternative strategy to identify deviations from official rules would have been to rely on the workload of each inspector at any point and to evaluate whether an incoming declaration was indeed assigned to the active inspector with the lowest workload when that declaration was registered. This would in principle enable us to identify which specific declarations were nonrandomly assigned. Unfortunately, the IT system in Madagascar customs does not keep a log of which inspectors were connected at what time nor of the exact time of assignment of a declaration to an inspector, which makes implementing this strategy infeasible.

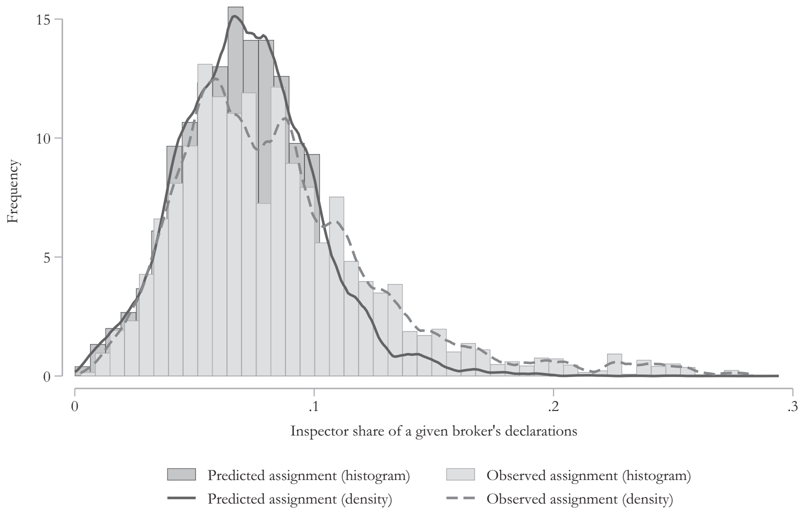

FIGURE II

Deviations from Official Rules in Assignment of Declarations to Inspectors

The figure shows the distribution of the share of declarations of a given broker handled by a given inspector in the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention). The darker-colored bars show the histogram of predicted inspection shares calibrated by setting the productivity of each inspector equal to the share of all declarations she handled in a given semester (see Section IV for details), and the solid line shows the overlaid kernel density plot of such predicted inspection shares. The light-colored bars indicate the distribution of observed inspection shares, with the long-dashed line showing the overlaid kernel density plot.

the equality of these two distributions at the 1% significance level. Clearly, the observed density distribution of inspector shares by broker is characterized by higher dispersion and more mass in the upper tail than the predicted distribution. This implies that relative to the distribution of expected inspection shares, the observed assignment of declarations is characterized by excess interaction between some inspectors and some brokers.

To assess whether, for a given broker, the observed inspection share of a given inspector is significantly different from the expected inspection share based on random assignment we must take into consideration that these expected inspection shares are not population parameters but estimates thereof. We therefore obtain standard errors for those shares using simulation methods that take five steps.[24] First, we obtain 99% confidence intervals for inspectors' productivities using Sison and Glaz's (1995) method of constructing confidence intervals for multinomial proportions. Second, we simulate the productivity distribution across inspectors 1,000 times by drawing from the 99% confidence interval of observed productivities (obtained in step 1), conservatively assuming that these productivities are uniformly distributed. Third, for each productivity simulation (obtained in step 2), we take the total number of declarations as given and simulate which inspectors are assigned to her declarations 10,000 times assuming multinomial assignment. Fourth, we test whether the observed number of declarations of a given broker handled by a given inspector is larger than the 99th percentile of the respective simulated multinomial assignment. Finally, we classify an inspector-broker pair in a given semester as being in significant excess interaction if for at least 99% of the productivity simulations we reject the null hypothesis of random assignment.[25]

Table I documents the prevalence of nonrandom assignment in Toamasina over the period ranging from January 1, 2015 to November 18, 2017, the day before the start of the delegated randomization to the third party. Panel A shows that for 10.3% of declarations the (final) inspector handling them interacts significantly more frequently with a broker than would be predicted based on conditional random assignment according to the definition. Prima facie, this is evidence of deviations from official rules in the assignment of import declarations of a given broker across inspectors. This nonrandom assignment is pervasive, with 10 out of 16 inspectors in a typical semester handling at least one nonrandomly assigned declaration. This nonrandom assignment is also persistent over time. For a given pairing of a broker with a particular inspector the excess share of declarations she handles in a given semester is correlated with her excess share in the previous semester, as is shown in Online Appendix Table A5, suggesting that the excess interactions are not accidental.[26]

24. Because we estimate excess interaction between inspectors and brokers by semester, the five steps to construct standard errors are repeated as many times as there as semesters in the sample.

25. We consider as a robustness check definitions of significant excess interaction based on at least 95% or at least 99.9% of the productivity simulations rejecting random assignment.

26. An inspector-broker pair that is deviant (e.g., characterized by excess interaction relative to nonrandom assignment) has a 50.6% probability of also being

TABLE I

Prevalence of Excess Interaction (Nonrandom Assignment) Before Delegated Randomization of Inspector Assignment

Panel A: Prevalence of excess interaction (nonrandom assignment)—calibrated			
	Number Nonrandomly assigned	Number Total	%
Declarations—after initial assignment	4,459	45,058	9.9
Declarations—after final assignment	4,661	45,058	10.3
Average per semester	At least one nonrandomly assigned declaration	Total	
Inspectors	10	16	60.2
Brokers	14	45	32.0
Inspector-broker pairs	23	690	3.3
Panel B: Prevalence of excess interaction (nonrandom assignment)—inspector logits			
	Number Nonrandomly assigned	Number Total	%
Declarations—after initial assignment	4,800	45,058	10.7
Declarations—after final assignment	4,545	45,058	10.1
Average per semester	At least one nonrandomly assigned declaration	Total	
Inspectors	10	16	61.2
Brokers	15	45	34.2
Inspector-broker pairs	25	690	3.6
Average per semester	Broker fixed effects jointly significant	Total	
Inspectors—initial assignment	7	16	44.1
Inspectors—final assignment	7	16	41.8

Notes. Declarations are characterized by significant excess interaction if they are handled by an inspector whose excess interaction share (the difference between the share of a given broker's declarations handled by the inspector in question and the hypothetical share she would be expected to handle if the allocation of declarations to inspectors were random conditional on their productivity, as prescribed by official rules, is positive and statistically significant (see Section IV). In Panel A, excess interaction measures are constructed using calibration methods. In Panel B, excess interaction measures are based on estimates from inspector-semester logit models. Initial assignment refers to the assignment originally made by the customs IT system. Final assignment takes into account subsequent potential reassignment(s) made and therefore corresponds to the last assignment that selected the inspector that cleared the declaration. The sample covers the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention).

IV.B. Estimating Excess Interaction Using Logit Models

The procedure outlined in the previous subsection has the advantage of being transparent and easy to implement, but it does not account for potentially innocent explanations for excess interaction such as fluctuations in inspectors' work schedules or differences in inspectors' productivity across days of the week. Descriptive evidence in Online Appendix Figure A1 and Online Appendix Table A6 suggests that inspectors' and brokers' work schedules are fairly stable and that fluctuations in inspectors' productivity over the workweek are limited. Nonetheless, we address this potential limitation by estimating inspector-semester specific binomial logit assignment models with several control variables, including broker fixed effects, which should have no explanatory power if inspector assignment conformed to official rules.[27]

These models estimate the probability P^{logit}_{dbit} that a declaration d registered by broker b during semester t will be assigned to inspector i instead of all other inspectors active in that semester:

$$P^{logit}_{dibt} = \frac{\exp\left(\beta'_{Xit}X_d + \pi_{pt} + \mu_{it} + \nu_{ibt}\right)}{1 + \exp\left(\beta'_{Xit}X_d + \pi_{pt} + \mu_{it} + \nu_{ibt}\right)}, \tag{4}$$

where the vector of declaration characteristics X_d contains fixed effects for the day of the week on which the declaration was registered, the tax rate, the risk score, whether the declaration was initially assigned to the physical inspection (red) channel, whether the shipment was mixed, and whether valuation advice was issued; π_{pt} contains product type fixed effects, μ_{it} contains inspector fixed effects, and ν_{ibt} contains broker fixed effects.[28] In this setup, μ_{it} represents a measure of the productivity of inspector i during

deviant in one (or both) of the subsequent semesters. This conditional probability is approximately 15 times larger than the unconditional probability of a pair being deviant (3.3% as reported in Table I).

27. In principle, a multinomial logit model of inspector assignment would be theoretically superior to the inspector-specific logit models we estimate because it would model assignment to all inspectors at once, accounting for the fact that the assignment outcomes of different inspectors are fundamentally interdependent. We tried to estimate such a model but failed to achieve convergence, presumably because including inspector-broker fixed effects in multinomial logit models leads to parameter proliferation.

28. Product type consist of 16 groups of HS two-digit codes following a classification of the World Trade Organization.

semester t (a higher value for this parameter means that, all else equal, i will be assigned more declarations). Under the null hypothesis of random assignment, only day of week fixed effects might potentially have explanatory power. The coefficients on all other explanatory variables should be zero. Most notably, broker fixed effects should be insignificant, since the fact that declaration d was registered by broker b should not significantly change the assignment probabilities to a particular inspector.[29]

We estimate binomial logit models using Bayesian Markov chain Monte Carlo (MCMC) methods to predict the inspection share S_{ibt}^{logit} of inspector i for broker j in semester t, as well as the inspection share inspector i should have for broker j if random assignment was abided by $\overline{S_{ibt}^{logit}}$, and construct estimates of excess interaction ES_{ibt}^{logit}, as the difference between the two.[30] To estimate S_{ibt}^{logit} we use a binomial logit model with all controls shown in equation (4). To estimate $\overline{S_{ibt}^{logit}}$ we use a binomial logit model that mimics random assignment given by equation (4) including only day of the week fixed effects in the vector of declaration characteristics. Confidence intervals for S_{ibt}^{logit} and $\overline{S_{ibt}^{logit}}$ are constructed using the Krinsky and Robb (1986) simulation method.[31]

29. A very minor practical drawback of estimating separate logit models for each inspector is that the sum of the predicted probabilities from equation (4) across inspectors—designated here as P_{dibt}^{logit}—is usually very close to, but not strictly equal to, 1. To address this issue, we normalize the predicted probabilities by reweighting their sum to be equal to 1:

$$P_{dibt}^{logit\ normalized} = \frac{P_{dibt}^{logit}}{\sum_{j=1}^{K_t} P_{djbt}^{logit}}.$$

In practice, this does not considerably change assignment probabilities.

30. One difference relative to the calibration procedure is that here we use estimates of S_{ibt}^{logit} instead of observed S_{ibt}, which facilitates hypothesis testing. Namely, it allows us to test the significance of broker fixed effects using likelihood ratio tests.

31. This simulation method draws a set of parameters from a normal distribution centered around the point estimates with a covariance matrix equal to the estimated covariance matrix. These parameters are used to calculate assignment probabilities to each inspector for each declaration. Taking the average of these probabilities across all declarations registered by broker b yields the share that is expected to be assigned to inspector i: S_{ibt}^{logit} in the model with all controls and $\overline{S_{ibt}^{logit}}$ in the model that mimics random assignment. The difference between S_{ibt}^{logit} and $\overline{S_{ibt}^{logit}}$ gives ES_{ibt}^{logit} for each iteration. This simulation method is repeated

Estimates of ES_{ibt}^{logit} (the difference between S_{ibt}^{logit} and $\overline{S_{ibt}^{logit}}$) are therefore driven by broker fixed effects and declaration characteristics other than the day of the week when it was registered.

This procedure has three advantages relative to our first procedure. To start with, it accommodates alternative explanations—such as differences in work schedule and specialization in specific products—for frequent pairings of the same broker with the same inspector. Second and related, it allows us to obtain consistent estimators for broker fixed effects and to test whether they are significant and have explanatory power. Third, by using Bayesian estimation methods and applying shrinkage, we account for the fact that in a small sample, one might expect some inspector-broker pairs to exhibit apparent excessive interaction even under the null hypothesis.[32] Note that we estimate one logit model per inspector per semester, resulting in a total of 116 different models. Table I, Panel B shows the prevalence of nonrandom assignment in Toamasina based on the measures of excess interaction from inspector-specific binomial logits. For 10.1% of declarations the (final) inspector handling them interacts significantly more frequently with a broker than would be predicted based on conditional random assignment. For these declarations the fixed effect of the broker registering them is individually significant in the inspector-specific binomial logit model. We also use a likelihood ratio test to assess whether broker fixed effects are jointly significant in each inspector-semester-specific binomial logit model.[33] The results of these tests are reported at the bottom of Table I,

1,000 times. For such a large number of iterations, the simulated distribution of assignment probabilities approximates the real distribution of assignment probabilities. This simulation method provides (i) the point estimate for ES_{ibt}^{logit} obtained as the average of ES_{ibt}^{logit} across the 1,000 iterations and (ii) an indicator for whether ES_{ibt}^{logit} is significantly larger than zero (i.e., whether S_{ibt}^{logit} is significantly larger than $\overline{S_{ibt}^{logit}}$).

32. In principle, random inspector-specific logit models would be preferable to the standard logit models we estimate as they would allow random broker effects and productivity parameters. We tried to estimate such models but failed to obtain convergence, as in the case of the multinomial logit models, possibly due to parameter proliferation.

33. To conduct these tests, we estimate a third type of logit model with all controls but excluding only broker fixed effects. This ensures that the significance of these fixed effects can be assessed by comparing estimates from this third type of logit model to those from the model with all controls (including broker fixed effects).

Panel B. On average across semesters, broker fixed effects are jointly significant for close to half the inspectors.

Ultimately, the two different procedures yield very similar measures of excess interaction. The correlation between excess-interaction measures based on inspector-specific logit models (ES^{logit}) and those calibrated from observed inspection shares (ES) is 0.97 (as is shown in Online Appendix Figure A2). Given this high correlation between measures derived from the two procedures, we will predominantly focus on the measures of excess interaction based on calibration. However, we replicate our main results using estimates from the binomial logit models and show that our main findings are robust to the use of that procedure (see Online Appendix Tables A9 and A10).

We end this section by emphasizing that detecting deviations from random assignment is neither necessary nor sufficient to establish potential corruption, and by highlighting some properties of our excess-interaction measures that are relevant for their appropriate interpretation. First, our measures vary across pairs of inspectors and brokers within a semester, but all declarations of a given broker handled by a particular inspector in a given semester are characterized by the same excess interaction share. Inevitably some of those declarations may not have been manipulated but will be characterized by excess interaction, which implies that we may be overestimating the prevalence of manipulation of inspector assignment but underestimating differences between manipulated and nonmanipulated declarations.

Second, our excess interaction measures have potential for false positives. This concern is partially mitigated by our use of simulation methods to identify excess interaction that is statistically significant and the use of binomial logit estimation procedures. Nonetheless, detecting potential deviations is merely the first step in the process of uncovering potential corruption. The findings in Sections V and VI showing that the declarations of brokers interacting excessively with some inspectors are at a significantly higher risk of tax evasion and that inspectors preferentially treat the declarations of brokers with whom they interact excessively frequently reduce the concern about false positives and suggest that deviations are not accidental.

Third, our excess-interaction measures also have potential for false negatives because they identify only one particular form of corruption. Our measures do not capture the (rather plausible) possibility that corrupt dealings are made between randomly assigned inspectors, brokers, and/or importers. In an extreme scenario a cartel led by the manager (that heads the inspectors) could set the terms of the bribes and make the identity of the inspector assigned to each declaration irrelevant. Our excess-interaction measures would not reject the null of excess interaction while in fact corruption would be ubiquitous. This extreme scenario is not supported by our evidence in Table I, which suggests that not all inspectors participated. But this caveat points to a requirement for our measures of excess interaction to be able to identify corruption: the need for differences across inspectors in their propensity to enter into corrupt deals with brokers. By agreeing deals ex ante, brokers can be sure that their declaration ends up with the "right" inspector. Moreover, it minimizes the risk of detection as inspectors clear undervalued shipments very quickly (with negotiations about bribes and the division of surplus taking place before the goods arrive). The presence of false negatives implies that our approach may provide an underestimate of the prevalence and consequences of corruption in Madagascar customs.

Finally, we conduct a placebo test by constructing calibrated excess-interaction measures for two sets of import declarations excluded from our main sample: imports that entered under nontaxable customs regimes and imports made by importers as part of a trade facilitation program for accelerated clearance that are not inspected (although an inspector is assigned). For both these types of declarations, incentives to manipulate inspector assignment are limited because they are either exempt from taxes or from inspector assessment. The share of declarations subject to excess interaction for those two types of declarations is extremely low: 0% for nontaxable customs regimes and 1% for the trade facilitation program.[34]

IV.C. Who Manipulates?

Who is responsible for this nonrandom assignment: the IT team that manipulates the IT system's initial assignment or the customs port manager who manually and voluntarily erases the initial assignment and reassigns declarations? To answer this question, Figure III, Panel A plots the density distributions of the initial inspector assignment made by the IT system (the long-dashed line) as well as the final assignment (the short-dashed line), which reflects both the initial assignment and potential

34. The number of declarations under nontaxable customs regimes is 4,693 and that under the trade facilitation program is 12,181.

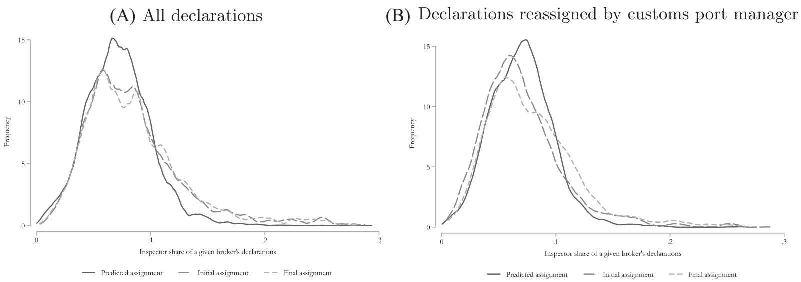

FIGURE III

Initial versus Final Inspector Assignment

The figure shows the kernel density distributions of the share of declarations of a given broker handled by a given inspector in the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention). The solid density plot shows the distribution of predicted inspection shares calibrated by setting the productivity of each inspector equal to the share of all declarations she handled in a given semester (see Section IV for details). The long-dashed line shows the distribution of the observed initial assignment of a declaration to a given inspector by the IT system (before the customs port manager potentially intervenes). The short-dashed line shows the distribution of the observed final assignment of a declaration to an inspector after potential reassignments made by the customs port manager. In Panel A the sample includes all declarations (both those that were reassigned by the customs port manager and those that were not), while in Panel B the sample includes only declarations that were reassigned by the customs port manager.

reassignments of declarations to inspectors made by the customs port manager. The distributions of initial and final inspector assignment are very similar, and both deviate markedly from that of predicted assignment if official rules were adhered to (the solid line). This is perhaps not surprising since reassignments are rare, only happening in 6% of cases. Thus, manipulation of the IT system appears to be the predominant driver of nonrandom assignment.

The port manager nonetheless appears to be complicit in the corruption scheme. Focusing only on declarations reassigned by the port manager, Figure III, Panel B reveals that these reassignments exacerbate, rather than reduce, nonrandom assignment. Instead of offsetting excess interaction, the port manager appears to be reinforcing it. If he were to choose inspectors randomly when reassigning declarations, one might have expected the final distribution to be less skewed.[35]

35. If the port manager instructed the IT department before the initial inspector assignment is made, he would be playing a more meaningful role in the

The fact that certain brokers' declarations are not randomly assigned to inspectors was confirmed in inspector interviews in Toamasina. One inspector mentioned "I have been here seven months, but there are certain brokers whose declarations I have never handled." Another complained, "I never get the good declarations." Our interpretation that such nonrandom assignment results from IT manipulation is consistent with the remarks by an external auditor of Madagascar's customs IT system of an "over-reliance on IT administrator account, which is typically used at most a few times a year to make major systemic changes, but was used multiple times a day in Madagascar. The IT administrator account allows you to override basic settings" and of "surprising and suspiciously long queues outside the office of the head of the IT department, which normally is not a client-facing function." When we confronted the port manager with our initial analysis, he acknowledged that manipulation of inspector assignment was prevalent.

Based on the findings of an early incarnation of this article, a number of customs inspectors were sanctioned for corruption and removed from their posts. The assignment of declarations was delegated to GasyNet, which agreed to randomize the assignment of declarations to inspectors. This delegated randomization provides an opportunity to assess whether we are indeed identifying IT manipulation that we will exploit in Section IX.

V. Do Deviant Declarations Exhibit a Higher Risk of Tax Evasion?

If excess interactions were the product of accidental deviations from official rules in inspector assignment, then the characteristics of these declarations should not systematically differ from those of other declarations. In contrast, if excess interactions were the product of deliberate IT manipulation to assign a specific declaration to a preferred inspector with whom the broker has a corruption agreement, then a higher risk of customs fraud, which would indicate higher susceptibility to tax evasion, would be expected for such declarations.

On average, declarations characterized by higher excess interaction shares have higher risk scores and are subject to higher

corruption scheme. Unfortunately, it is not possible to test this possibility with our data. More generally, we are not able to ascertain who initiates the scheme.

tax rates, as is shown in Figure IV, Panels A and B, which present polynomial plots of these risk characteristics against the excess interaction share. By contrast, initial unit prices relative to median import unit prices tend to fall with the excess interaction share, as shown in Figure IV, Panel C, suggesting that declarations of brokers that interact excessively with some inspectors are more likely to be undervalued. The excess interaction share is indeed positively correlated with hypothetical tax revenue losses calculated on the basis of the initial registration of the declaration by the broker (that is, before the inspector assesses the declaration and carries out any adjustment), as shown in Figure IV, Panel D.

Table II presents estimates of unconditional bivariate ordinary least squares (OLS) regressions of declaration characteristics commonly associated with tax evasion on the excess interaction share. The standard errors are two-way clustered by inspector and by broker. A 10% (0.10) increase in the excess interaction share is associated with an increase in the risk score of half a point, a 3.1% higher tax rate, a 7.8% increase in the probability that the declaration contains multiple HS6 products (i.e., is mixed), a 5% increase in the share of the declaration's value accounted for by differentiated products, a 9.4% increase in the probability of valuation advice being issued, and a 5.9% decrease in the initial price relative to median import price. These significantly lower initial prices may explain why the excess interaction share is not significantly correlated with the initially declared value, despite being associated with a higher initially declared weight. Consistent with this interpretation of undervaluation, a 10% increase in the excess interaction share is associated with a 6.3% increase in initial (i.e., before any adjustment made by customs) hypothetical tax revenue losses.

Note that because declarations subject to excess interaction are subject to higher taxes and tend to be larger (in weight), they are subject to a substantially higher theoretical tax liability. Figure IV, Panel E shows that declarations subject to excess interaction are significantly more likely to be "high potential tax yield" declarations—defined as those for which the hypothetical tax yield based on external reference prices exceeds \$20,000, for which the incentives to evade are largest.[36] Although only one in four declarations not subject to significant excess interaction are

36. The cutoff of US\$20,000 corresponds roughly to the top quartile of the hypothetical tax revenue yield distribution.

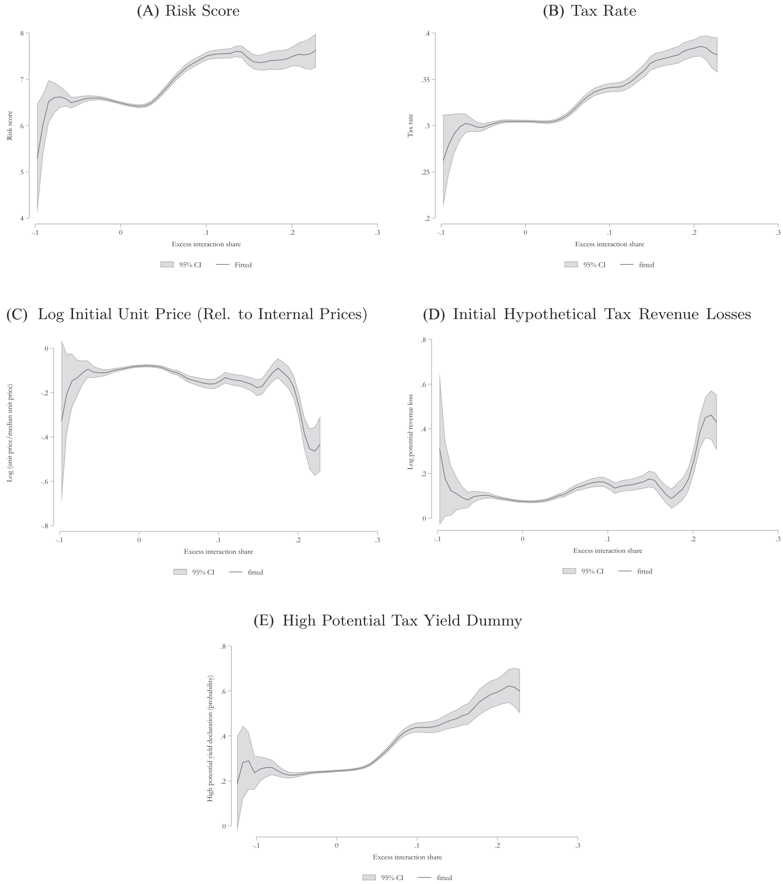

FIGURE IV

Tax Evasion Risk and Excess Interaction

The graphs, generated using STATA 17, show weighted local polynomial plots (using the Epanechnikov kernel function) of a selected number of declaration characteristics capturing tax evasion risk on calibrated excess interaction shares for the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention). The sample used is the regression sample used to generate baseline results (see e.g. Table III). Excess interaction share is the difference between the share of given broker's declarations handled by an inspector in a given semester and the hypothetical share she would be expected to handle if the allocation of declarations to inspectors were random conditional on their productivity, as prescribed by official rules, calculated using calibration methods (as explained in Section IV). Initial hypothetical tax revenue losses refer to the tax revenue losses estimated based on internal reference prices. High potential tax yield declarations are those for which the hypothetical tax yield (if the declaration was valued using external reference prices) exceeds $20,000. CI stands for confidence interval.

TABLE II

TAX EVASION RISK AND EXCESS INTERACTION

Before Delegated Randomization of Inspector Assignment

Dependent variable	Risk score (1)	Tax rate (2)	Red channel dummy (3)	Mixed shipment dummy (4)	Differentiated share (5)	Valuation advice dummy (6)
Excess interaction share	5.178***	0.306***	0.088	0.775**	0.503**	0.937***
	(1.051)	(0.064)	(0.240)	(0.336)	(0.186)	(0.330)
Observations	44,522	45,058	45,058	45,058	45,058	45,058
R-squared	0.006	0.010	0.000	0.005	0.002	0.019
Dependent variable		Log initial value (7)	Log initial weight (8)	Log initial unit price (relative to internal prices) (9)	Initial hyp. tax rev. losses (10)	High potential tax yield dummy (11)
Excess interaction share		0.256	1.457***	−0.591**	0.632**	1.540***
		(0.358)	(0.503)	(0.243)	(0.240)	(0.235)
Observations		45,058	45,058	45,033	45,033	31,402
R-squared		0.000	0.001	0.002	0.003	0.020

Notes. Standard errors clustered two-way by inspector and by broker are presented in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively. *Excess interaction share* is the difference between the share of a given broker's declarations handled by an inspector in a given semester and the hypothetical share she would be expected to handle if the allocation of declarations to inspectors were random conditional on their productivity, as prescribed by official rules, calculated using calibration methods (as explained in Section IV). "Observations" refers to the number of nonsingleton observations. OLS estimation is used. The sample covers the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention).

high potential yield declarations, half of all declarations subject to significant excess interaction are (see Online Appendix Table A3).

In short, declarations characterized by excess interaction have characteristics commonly associated with an elevated risk of tax evasion.[37] Online Appendix Table A7 presents regressions examining the determinants of the excess interaction share. The risk score and issuance of valuation advice (a proxy for undervaluation) are the most salient predictors of deviations from conditional random assignment of inspectors to declarations. The evidence of declarations with higher excess interaction shares being at a higher risk of tax evasion is confirmed using excess-interaction measures based on binomial logit models (see Online Appendix Table A9).

VI. Are Deviant Declarations Treated Differently?

This section assesses whether inspectors treat the deviant declarations differently—in a preferential manner—from other declarations. If excess interactions were accidental, then inspectors should provide no differential treatment to deviant declarations, beyond the increased scrutiny that may be legitimately expected because these declarations were shown to be at a higher risk for tax evasion in Section V. Similarly, if IT department staff were simply bribed to assign certain declarations to the least competent inspector, we would not necessarily expect the chosen inspector to treat manipulated declarations any differently from the way she handles other declarations. Inspectors complicit in a corruption agreement, by contrast, would plausibly provide, in exchange for a bribe, preferential treatment to manipulated declarations. To assess whether inspectors treat deviant declarations—those of brokers with whom they interact excessively—differently than other declarations, the following specification is estimated by OLS:

$$Y_d = \beta_E ES_{ibt} + \beta_X X_d + \mu_i + \nu_b + \kappa_c + \pi_p + \tau_m + \epsilon, \quad (5)$$

37. An alternative explanation for these findings is that inspectors offer discounts to certain importers to maximize future tariff revenue. However, Online Appendix Table A8 shows that excess interaction is not correlated with proxies for the average trade elasticity of products included in the declaration based on Broda and Weinstein (2006) and Fontagné, Guimbard, and Orefice (2022), nor with the average durability or “stickiness” of trade relationships based on Martin, Mejean, and Parenti (2020) of the products included in the declaration. This suggests inspectors are not targeting importers with the highest sensitivity to tariff discounts.

where Y_d is one of the declaration-level customs outcomes described in Section III (clearance time, fraud records, value and tax adjustments, hypothetical tax revenue losses). The main regressor of interest is the excess interaction share ES_{ibt} defined in Section IV. The vector of declaration characteristics X_d includes the tax rate, the risk score, a dummy for the red channel, a dummy for being a mixed shipment, the share of differentiated products, and a dummy for GasyNet's valuation advice. Inspector fixed effects μ_i, broker fixed effects ν_b, HS two-digit product fixed effects π_p, source country fixed effects κ_c, and month-year fixed effects τ_m are also controlled for. The independent and identically distributed (i.i.d.) error is ϵ.

The inclusion of inspector fixed effects accounts for heterogeneity across inspectors in their average productivity, ability, work ethic, and other time-invariant characteristics that may affect their performance. Broker fixed effects account for heterogeneity in their import patterns, efficacy, record-keeping, and other characteristics that may affect customs clearance. Since average differences in inspector and broker characteristics are accounted for, the specification is stringent in that identification of the coefficient on the excess interaction share is based on the interaction between the inspector and broker relative to other pairings of inspectors and brokers. Standard errors are clustered two-way by inspector and by broker.[38]

VI.A. Main Findings

The results from estimating equation (5) are shown in Table III. Inspectors assess declarations registered by brokers with whom they interact excessively frequently significantly faster than other declarations. Column (1) implies that a 10% increase in the excess interaction share is associated with a 20% (or approximately a four-hour) reduction in clearance times. Declarations characterized by excess interaction are also less likely to be deemed fraudulent: column (2) shows that a 10% increase in

38. Due to the inclusion of a large set of fixed effects, our estimates are obtained using the reghdfe Stata command drawing on Guimarães and Portugal (2010). The current version of the command eliminates from the number of observations singletons and adjusts standard errors for their exclusion. A singleton is an observation unique in the sample in having a given fixed effect equal to one: for example, a declaration with imports from source country A if no other declaration reports importing from country A.

TABLE III
DIFFERENTIAL TREATMENT BY INSPECTORS
Before Delegated Randomization of Inspector Assignment

Dependent variable	Time (1)	Fraud (2)	Δlog value (3)	Δlog tax (4)	Hyp. tax revenue losses (5)
Excess interaction share	−2.008***	−0.275**	−0.079***	−0.086***	0.389**
	(0.361)	(0.101)	(0.022)	(0.031)	(0.175)
Declaration characteristics	Yes	Yes	Yes	Yes	Yes
Inspector fixed effects	Yes	Yes	Yes	Yes	Yes
Broker fixed effects	Yes	Yes	Yes	Yes	Yes
Source country fixed effects	Yes	Yes	Yes	Yes	Yes
HS2-product fixed effects	Yes	Yes	Yes	Yes	Yes
Month-year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	41,121	44,522	44,434	40,471	44,497
R-squared	0.318	0.214	0.152	0.132	0.211
p-value joint significance of broker fixed effects	.000	.000	.000	.000	.000

Notes. Standard errors clustered two-way by inspector and by broker are presented in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively. *Excess interaction share* is the difference between the share of a given broker's declarations handled by an inspector in a given semester and the hypothetical share she would be expected to handle if the allocation of declarations to inspectors were random conditional on their productivity, as prescribed by official rules, calculated using calibration methods (as explained in Section IV). Declaration characteristics include the tax rate, the risk score, a dummy for the red channel, the share of value accounted for by differentiated products, a dummy indicating whether the declaration was mixed, and a dummy indicating that the declaration was subject to valuation advice. "Observations" refers to the number of nonsingleton observations. OLS estimation is used. The sample covers the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention).

the excess interaction share is associated with a 2.8% reduction in the likelihood of fraud being recorded. This is a large effect given that the unconditional probability of fraud being recorded is 8% (see Online Appendix Table A3).

In the same vein, columns (3) and (4) show that value and tax adjustments are significantly lower for declarations characterized by excess interaction. A 10% increase in the excess interaction share is linked to a 0.8% lower increase in value and a 0.9% lower increase in tax yield. These are again sizable effects given that the unconditional averages of value and tax adjustment are 2%. The significantly lower likelihood of the tax burden being revised upward is perturbing because declarations characterized by excess interaction are more likely to be undervalued to start with, as shown in Section V. Inspectors thus seem to exacerbate, rather than reduce, the disparities between declarations characterized by excess interaction and other declarations. As a result, excess interaction is associated with sizeable tax revenue losses. Column (5) implies that a 10% increase in the excess interaction share is associated with a tax revenue loss of 3.9%.

In summary, inspectors treat the declarations of brokers with whom they interact excessively frequently preferentially: they clear these declarations more quickly and subject them to significantly laxer tax enforcement. If inspectors were honest, no preferential treatment should be observed.

VI.B. Robustness Checks

We subject these findings of preferential treatment by inspectors to several robustness checks. First, we estimate equation (5) using the measures of excess interaction based on inspector-specific logit models. The results, presented in Online Appendix Table A10, provide clear evidence of preferential treatment given by inspectors to the declarations of brokers with whom they interact excessively frequently. Second, to address potential selection bias, we rely on a propensity score matching approach described in Online Appendix B to identify a set of control declarations that are most similar to those that are treated, that is, have significant excess interaction, based on observable risk characteristics.[39] The results from estimating equation (5) for the matched sample (Panel A) or using propensity score weighted least squares as proposed by Hirano, Imbens, and Ridder (2003) (Panel B) are shown in Online Appendix Table A12 and are consistent with our main findings displayed in Table III.[40] Third, we estimate variants of equation (5) that progressively add the different types of fixed effects (Panels A–D) instead of including them all at once, and control for all risk characteristics considered in Table II (Panel E), and find the

39. The balance tests for this propensity matching approach shown in Online Appendix Table A11 indicate no significant differences on average across treated and control declarations on all but two of those risk characteristics: the probability of receiving valuation advice and the tax rate (the latter at a 10% significance level only).

40. The number of observations in Online Appendix Table A12, Panel A is substantially smaller than in Table III since our matching approach uses the nearest-neighbor matching algorithm that selects for each treated declaration a single control declaration.

patterns of preferential treatment by inspectors to declarations with excess interaction to be maintained in Online Appendix Table A13.[41] Fourth, we add more stringent types of fixed effects to equation (5): inspector-semester and broker-semester; inspector-month and broker-month; inspector-semester, broker-semester, and importer-semester; or importer-broker. Online Appendix Table A15 shows that these fixed effects do not affect the qualitative pattern of results. Fifth, we estimate equation (5) using either the indicator for significant excess interaction defined in Section IV (instead of the excess interaction share) or two other indicators based on significance levels of 95% or 99.9%. The findings in Panels A–C of Online Appendix Table A16 are qualitatively similar to those in Table III. Finally, we construct measures of excess interaction for three alternative samples that modify the restrictions described in Section III: a sample excluding only brokers registering fewer than 20 declarations per semester (versus 50 in our main sample), a sample excluding brokers registering fewer than 100 declarations per semester, and a sample not excluding any brokers. The estimates of equation (5) for these three samples shown in Panels D–F of Online Appendix Table A16 are qualitatively unchanged relative to those in Table III.

VI.C. Alternative Explanations

This section evaluates salient alternative explanations for the findings of differential preferential treatment of deviant declarations by inspectors by running a set of additional tests. First, one possibility is that our excess interaction share merely reflects "familiarity" between inspector and broker, whereby the fact that certain brokers interact very frequently with an inspector reduces fixed inspection costs. Alternatively, inspectors may update their prior beliefs about brokers' likely compliance based on their past interactions with them and consequently be less likely to scrutinize brokers with whom they interact frequently for which they have a sizable pool of past interactions to base their inferences on. To assess the validity of these explanations for our results, we add to equation (5) a measure of "familiarity": the total number of prior transactions of that same broker cleared by the same inspector over the preceding semester.[42] The results in Table IV,

41. We also show that the results are robust to different types of clustering of standard errors in Online Appendix Table A14.

42. Our excess interaction share measure is based on identifying deviations in the share of a given broker's declarations handled by a given inspector. By

TABLE IV

Alternative Explanations for Differential Treatment Before Delegated Randomization of Inspector Assignment

Dependent variable	Time (1)	Fraud (2)	Δlog value (3)	Δlog tax (4)	Hyp. tax revenue losses (5)
Panel A: Controlling for familiarity					
Excess interaction share	−1.797***	−0.278***	−0.081***	−0.082***	0.323*
	(0.429)	(0.096)	(0.023)	(0.028)	(0.176)
Familiarity	−0.042	−0.000	0.000	−0.001	0.012*
	(0.027)	(0.005)	(0.002)	(0.001)	(0.006)
Observations	40,990	44,359	44,273	40,324	44,335
R-squared	0.321	0.214	0.153	0.133	0.211
Panel B: Controlling for congestion					
Excess interaction share	−2.002***	−0.275**	−0.079***	−0.085***	0.389**
	(0.362)	(0.100)	(0.022)	(0.031)	(0.175)
Congestion	0.098**	−0.004	0.001	0.000	−0.001
	(0.038)	(0.004)	(0.001)	(0.001)	(0.005)
Observations	41,121	44,522	44,434	40,471	44,497
R-squared	0.318	0.214	0.152	0.132	0.211
Panel C: Excluding declarations registered outside regular business hours					
Excess interaction share	−2.033***	−0.270***	−0.079***	−0.088***	0.385**
	(0.358)	(0.093)	(0.021)	(0.028)	(0.177)
Observations	40,285	43,497	43,409	39,534	43,473
R-squared	0.316	0.220	0.156	0.136	0.210
Panel D: Adding importer fixed effects					
Excess interaction share	−2.051***	−0.170**	−0.069***	−0.081***	0.242***
	(0.287)	(0.071)	(0.021)	(0.024)	(0.087)
Observations	40,311	43,691	43,601	39,678	43,669
R-squared	0.393	0.327	0.292	0.297	0.429
Panel E: Adding importer fixed effects and excess interaction share with importers					
Excess interaction share	−2.056**	−0.221*	−0.047	−0.075*	0.092
	(0.837)	(0.122)	(0.035)	(0.040)	(0.166)
Excess interaction share with importer	−0.172	0.016	−0.003	0.006	0.099
	(0.474)	(0.104)	(0.018)	(0.027)	(0.105)
Observations	9,537	10,281	10,263	9,184	10,278
R-squared	0.371	0.308	0.238	0.226	0.240
Panel F: Dropping top three inspectors with the largest share of declarations with excess interaction each semester					
Excess interaction share	−2.207***	−0.258***	−0.085***	−0.087***	0.367*
	(0.432)	(0.082)	(0.022)	(0.029)	(0.180)
Observations	32,542	35,222	35,152	31,935	35,203
R-squared	0.322	0.222	0.161	0.141	0.201

Notes. Standard errors clustered two-way by inspector and by broker are presented in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively. *Excess interaction share* is the difference between the share of a given broker's declarations handled by an inspector in a given semester and the hypothetical share she would be expected to handle if the allocation of declarations to inspectors were random conditional on their productivity, as prescribed by official rules, calculated using calibration methods (as explained in Section IV). All specifications control for the tax rate, the risk score, a dummy for the red channel, the share of value accounted for by differentiated products, a dummy indicating whether the declaration was mixed, a dummy indicating that the declaration was subject to valuation advice, inspector fixed effects, broker fixed effects, source country fixed effects, and month-year fixed effects. "Observations" refers to the number of nonsingleton observations. OLS estimation is used. The sample covers the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention).

Panel A show that the familiarity measure itself has some predictive power: it is linked to slightly higher tax revenue losses but does not significantly predict the incidence of fraud or value adjustment. More important for our purposes, controlling for familiarity only marginally reduces the impact of the excess interaction share, which remains strongly statistically significant in all specifications. Put differently, the results do not appear to be driven by familiarity or learning, which, in any case, cannot explain why deviant declarations would be more risky to start with.

Second, a possible explanation for differential treatment is that it reflects congestion and fluctuations in inspectors' workload. Specifically, when inspectors get very busy they may be tempted to exert less scrutiny and speed up clearance merely to be able to manage increased traffic. If this increase in their workload is generated by absenteeism of other inspectors, we might see a simultaneous increase in the excess interaction share and a decrease in scrutiny and clearance times. To control for such congestion, we add to equation (5) the number of declarations assigned to a given inspector over the course of the calendar month as a proxy for their workload. While Table IV, Panel B shows that this measure of congestion is clearly positively correlated with clearance time, the impact of the excess interaction share on the other customs outcomes is hardly affected by its inclusion.[43]

Third, one may worry that the patterns documented are an artifact of dubious declarations being more likely to be registered outside of regular business hours, that is, late in the evening, at night, or during the weekend. This could help explain excess interaction because there are typically many fewer inspectors active and they may monitor incoming declarations less aggressively because they are fatigued or want to go home. However, Table IV, Panel C shows that the results are robust to excluding declarations registered outside of regular business hours, which account for less than 3% of all declarations.

---

contrast, the familiarity measure is based on the absolute number of interactions between the broker and the inspector. Whereas inspectors will interact more with large brokers, and hence be more "familiar" with them, they will not necessarily interact excessively with large brokers, since our excess interaction share is a relative measure.

43. Similarly, as mentioned already, Online Appendix Table A15 shows that the results are robust to controlling for inspector-month and broker-month fixed effects, which can also proxy for workload and congestion.

Fourth, one may be concerned that the results are driven by (excess) interaction between inspectors and importers themselves rather than brokers, who are supposed to represent the interests of importers. We address this possibility in two ways. We augment equation (5) with importer fixed effects in Table IV, Panel D, and this hardly impacts the qualitative pattern of results.[44] In Table IV, Panel E we add to equation (5) the excess interaction share between importers and inspectors. The measure is defined analogously to the calibrated excess interaction share between brokers and inspectors, for importers that registered at least 50 declarations during the semester, which leads to a reduction in sample size. The excess interaction share between inspectors and importers does not significantly predict fraud, value, or tax adjustment and does not seem correlated with tax revenue losses. By contrast, the excess interaction share between inspectors and brokers remains robustly significant. These results justify our focus on brokers rather than importers. The fact that brokers seem to be the primary protagonists of the specific corruption scheme we document may be because they have more to gain from it; there are far fewer brokers than importers, and brokers interact more frequently with inspectors than importers do. Moreover, lobbying customs is the core business of brokers in many developing countries.

It is worth noting that while most importers work exclusively with one broker, Online Appendix Table A17 furnishes evidence that importers who use multiple brokers systematically steer their most risky declarations to brokers with the highest propensity to have excess interaction. It is difficult to ascertain, however, whether they do so because they know about the corruption scheme or whether they were simply offered favorable terms by brokers engaging in excess interaction.

Fifth, given the limited number of inspectors working in Toamasina, one may worry that our results are driven by a few individuals, rather than reflecting widespread corruption. Table IV, Panel F replicates our baseline results but excluding for each semester the top three inspectors with the greatest share of declarations with significant excess interaction. The coefficients remain statistically significant and of similar magnitude as our baseline results. The results are thus not driven by a select few inspectors

44. Online Appendix Table A15 shows that including importer-broker fixed effects does not qualitatively change results either.

(even though the tax revenue losses associated with the scheme are very concentrated, as we show in Section VIII). Evidence that the results are also not driven by a select few brokers is provided in Online Appendix Table A14, where we exclude for each semester the top five brokers with the greatest share of declarations with significant excess interaction.

Sixth, another potential concern is that results might be driven by inspectors specializing in clearing different goods. This concern is mitigated by the fact that formally, there is no specialization across different inspectors: they all clear the same set of goods. However, one may nonetheless wonder whether the IT department staff who are manipulating assignment are systematically assigning declarations containing certain products to unwitting inspectors that do not have the requisite expertise to adequately evaluate them; they may be seeking out inspectors that are the worst at detecting fraud for particular sets of products. To address this concern, Online Appendix Table A18 presents regressions where the unit of observation is an item (recall that a declaration can contain multiple items). The dependent variables are the log of the initially declared unit price, adjustments in that unit price, the finally registered unit price, the adjustment in weight (finally registered − initially declared) and an item-specific measure of the hypothetical tax revenue loss. The main regressor of interest is still the excess interaction share and the set of controls now includes HS eight-digit product-inspector fixed effects, broker fixed effects, source country fixed effects, month-year fixed effects, and a vector of declaration characteristics (the risk score, a dummy for the red channel, a dummy for being a mixed shipment, a dummy for GasyNet's valuation advice) and the item-specific tax rate. The HS eight-digit product-inspector fixed effects capture the comparative advantage of the inspector in detecting fraud in different types of products. The item-level initially declared unit price is significantly negatively correlated with excess interaction (column (1)). Excess interaction is associated with a lower item-level initial unit price but also with significantly lower adjustments to the unit price. As a result, the final unit price is even more negatively correlated with excess interaction. Excess interaction is also associated with lower weight adjustment and higher potential tax revenue losses, but these associations are not statistically significant at conventional significance levels.

Seventh, evidence of heterogeneity in the differential treatment of deviant declarations is hard to reconcile with explanations other than corruption for our main results. We estimate equation (5) allowing the excess interaction share to be interacted with the tax rate. Differential treatment by inspectors that interact excessively frequently with a given broker appears particularly pronounced for declarations subject to higher taxes: these are especially less likely to be deemed fraudulent and exhibit significantly higher tax revenue losses, as seen in Online Appendix Table A19.

Some final evidence consistent with corruption is provided by the analysis of inspector reassignments made by the customs port manager. Such reassignments are substantially more likely when declarations are initially assigned to an inspector with whom the broker is not interacting excessively frequently (see Online Appendix Table A20). This is inconsistent with reassignments being random. Moreover reassigned declarations typically yield higher fraud findings, value adjustments, and tax adjustments. This is especially the case if they are taken away from inspectors with initial excess interaction, suggesting that these nonrandomly assigned declarations were more risky to start with. By contrast, reassigned declarations from inspectors without excess interaction toward inspectors with excess interaction do not yield increased fraud findings or tax adjustments, as is shown in Online Appendix Table A21.

VII. How Costly Is Corruption?

How much tax revenue is lost because of the corruption scheme we document? To answer this question, we calculate how much more tax revenue would have been collected if there was no significant excess interaction between inspectors and brokers.[45] The key input into this back-of-the-envelope calculation are estimates of the impact of excess interaction between inspectors and brokers—β_E in equation (5)—on tax revenue losses. We focus on a measure of hypothetical tax revenue losses described in Section III that accounts for underreporting of quantities and is based on prices reported by countries exporting to Madagascar, which are arguably less likely to be endogenous to underinvoicing

45. We abstract from dynamic effects of offering tariff discounts today on future tariff revenues and from uncertainty about tariff rates. Importantly, note that we do not evaluate the social welfare effects of the prevailing tariff structure nor those of the corruption scheme we unveil.

in Madagascar.[46] In addition, to rely on an unbiased estimate of the overall impact of corruption on tax revenue losses, we estimate equation (5) including only controls that are plausibly exogenous to corruption: the tax rate, the dummy for mixed shipment, the share of differentiated products, source country fixed effects, HS two-digit product fixed effects, and month-year fixed effects.[47] Our β_E estimate (presented in Online Appendix Table A22, Panel B, column (11)) indicates that a 10% increase in the excess interaction share is associated with a 21% increase in tax revenue losses.

Using this estimate, we calculate for each declaration the counterfactual tax revenue that would have been collected in the absence of significant excess interaction between inspectors and brokers as $T^{NC} = T * exp(\widehat{\beta_E * ES})$, where T is the actual tax yield.[48] We are effectively asking how much more tax revenue would have been collected if declarations subject to excess interaction had been treated by inspectors like declarations that were not. The results of this exercise are presented in Online Appendix Table A24, Panel A for declarations characterized by significant excess interaction (in the first two columns) and for all declarations (in the last two columns). Interestingly, declarations with significant excess interaction yield more tax revenue, \$11,423 on average, despite being undervalued, than the average declaration with \$10,446. This finding is consistent with declarations with significant excess interaction being subject to a higher tax liability, as was shown in Section V. In the absence of corruption, the average declaration with significant excess interaction would have yielded an additional \$2,962 in tax revenue. Put differently, the tax yield on declarations likely to be the object of corruption agreements would have been 26% higher. This number is a lower bound on total tax revenue losses per declaration associated with corruption since the set of declarations characterized by significant excess interaction likely includes some that were randomly assigned and not the object of corruption schemes. Across all declarations, average (and hence aggregate) tax yield would have been 3% higher in the period before the delegated randomization intervention. These estimates do not reflect the gains

46. Results for all other measures of hypothetical tax revenue losses described in Section III are reported in Online Appendix C.

47. The controls in equation (5) that are potentially endogenous to corruption are inspector and broker fixed effects and the risk score.

48. The details of this calculation are provided in Online Appendix C.

associated with eliminating tax evasion altogether, but only the gains from eliminating tax evasion due to the specific corruption scheme we uncover. Our methodology does not address the (rather plausible) possibility that tax evasion can also result from deals made between randomly assigned inspectors, brokers, and/or importers.[49]

VIII. Who Benefits from Corruption?

This section analyses which types of inspectors and brokers participated in the corruption scheme and tries to shed light on how much they gained by doing so.

VIII.A. Which Inspectors and Brokers Participated?

We start by assessing which types of inspectors participated in the scheme by regressing the average share of declarations handled by the inspector subject to significant excess interaction on a number of inspector characteristics. The results are presented in Online Appendix Table A25, Panel A. Tenure is by far the strongest predictor of handling declarations characterized by excess interaction, but the effect is highly nonlinear: new inspectors (the omitted category) have a significantly lower propensity to handle declarations subject to excess interaction than more established inspectors, and this association is robust to controlling for inspector age (column (2)). There is some evidence that male inspectors handle a larger share of declarations subject to excess interaction, but the difference with their female colleagues is only borderline statistically significant at the 10% level and loses significance when controlling for age. For the subset of inspectors for whom we have information on educational attainment, we find that those with a management degree have a higher propensity to handle declarations subject to excess interaction (column (3)).

Comparable estimates for brokers are presented in Online Appendix Table A25, Panel B. Brokers based in Toamasina handle a significantly higher share of declarations assessed by

49. These calculations do not take into account potential beneficial effects of tariff discounts on future trade volumes: lower levels of taxes may encourage imports in subsequent periods. More generally, customs administrations have the dual objective of facilitating trade and collecting tax revenues and these objectives may conflict with one another both in the short- and in the long-run because of such dynamic effects.

inspectors with whom they interact excessively. On average, the share of declarations they handle that is subject to excess interaction is 5.9% higher than the average share of non-Toamasina based brokers, and the difference is significant at the 10% level. However, this association loses significance once we include a dummy identifying brokers that import on behalf of only one importer (column (2)) and when we control for the broker's average share of all declarations and for broker tenure (column (3)). Though none of the results are significant at conventional significance levels, they suggest that brokers who serve only one importer exhibit lower excess interaction. Brokers who have been active for a longer period of time tend to have more declarations subject to excess interaction.

Finally, for 20 inspectors who participated in the survey of inspectors that we implemented in 2017, we are able to correlate their views with the share of declarations they handled that were subject to excess interaction. Online Appendix Table A26 shows that inspectors with a higher share of declarations with excess interaction report on average significantly higher overall job (but not pay) satisfaction, report higher esprit de corps, and are much more likely to claim that they know the most fraudulent firms.[50]

Taken together, the fact that excess interaction increases with inspector tenure and that brokers based in Toamasina have a higher propensity to handle declarations with excess interaction points to the importance of establishing personal relationships and private information acquisition.

VIII.B. How Are Tax Revenue Losses Distributed?

Unfortunately, our data do not allow us to identify how participants in the corruption scheme divided the surplus generated by tax savings associated with the scheme; it is thus not possible to precisely pinpoint how much each participant gained, which is furthermore complicated by the fact that officials in the IT department, the port manager, and importers likely have taken a cut. Instead, we present estimates of the distribution of tax revenue losses across inspectors and brokers by semester, ranking inspectors and brokers in terms of their contribution to overall

50. Excess interaction is not significantly correlated with pay satisfaction, views about the adequacy of training, discretion, perceived corruption (among brokers, colleagues, and supervisors), punishment for unethical behavior, reporting of threats by brokers, or meritocracy.

revenue losses (with rank 1 assigned to the inspector or broker with the highest revenue losses in a given semester), using our preferred measure of counterfactual additional tax yield (calculated using external reference prices and correcting for potential underreporting of quantities) if there had not been significant excess interaction.

Table V reports that on average an inspector collects \$4.8 million worth of tax revenue per semester (6% of total taxes collected in the port per semester). Average tax revenue losses associated with the unveiled corruption scheme equal \$140,000 per inspector per semester (3% of the revenues they collect). Yet this number masks large heterogeneity across inspectors. In a typical semester, the inspector accountable for the largest tax revenue losses incurs \$677,000 worth of losses—roughly four times the average loss. Yet she also collects \$6.5 million worth of taxes (or 8% of total taxes collected in the port per semester). Hence, over the period considered, tax revenues collected by the "top" inspector would have been 11% higher without the corruption scheme.

Despite fairly widespread participation of inspectors in the scheme (documented in Table I), the tax losses associated with participation in the corruption scheme are highly concentrated: the "top" inspector accounts for roughly one-third of the total revenue losses associated with the scheme in a given semester, the "top" two inspectors jointly account for more than half (55%) of all revenue losses, and the top three inspectors jointly account for more than two-thirds of all revenue losses. These statistics attest to the granularity of tax evasion associated with the scheme; if each semester the top three most corrupt inspectors had not participated in the scheme, overall tax revenue collection in the port would have been almost 2% higher (as opposed to 3% if none of the inspectors had participated). The behavior of a select few inspectors thus has macro-fiscal implications.

The concentration of tax revenue losses in the hands of a select few inspectors prompts the question as to why corruption was neither detected nor sanctioned sooner. The answer to this question may partly lie in the targeting of declarations subject to high tax liability noted in Section V. Figure V, Panel A plots the tax yield per declaration against the excess interaction share. If anything, the relationship between tax yield and excess interaction is positive. Inspectors with more excess interaction have higher average (unconditional) tax yields per declaration, as is shown in Online Appendix Figure A3. This helps explain why conventional inspector performance metrics—such as total tax yield or average

TABLE V

CONCENTRATION OF TAX REVENUE LOSSES BY SEMESTER

Rank per semester	Total taxes collected average per semester (US$)	Tax losses average per semester (US$)	% Total taxes collected per semester	% Total tax losses per semester
Panel A: By inspector				
(ranked in terms of tax revenue losses, from largest to smallest, by semester)				
Inspector rank (in a given semester)				
1	6,491,683	677,109	8.2	32.8
2	6,552,940	499,588	8.0	22.4
3	5,901,505	324,983	7.2	13.4
4	5,663,769	265,376	7.3	10.5
5	4,890,071	180,883	6.1	7.4
Rank 6–10 (combined)	25,181,080	347,773	31.2	13.4
Rank 11 and higher (combined)	23,764,220	5,242	31.9	0.2
Average per inspector per semester	4,802,772	140,875	6.1	6.1
Panel B: By broker				
(ranked in terms of tax revenue losses, from largest to smallest, by semester)				
Broker rank (in a given semester)				
1	1,734,389	514,420	2.3	24.2
2	3,294,547	353,191	4.0	16.2
3	2,237,922	288,914	2.7	12.7
4	3,059,864	239,310	4.0	10.6
5	3,671,714	186,521	4.8	7.7
Rank 6–10 combined	11,782,514	535,779	15.1	22.0
Rank 10–20 (combined)	13,552,180	182,819	17.1	6.5
Rank 21 and higher (combined)	39,112,140	0	50.1	0.0
Average by broker per semester	1,749,709	51,322	2.2	2.2
Panel C: Overall				
Total per semester	78,445,268	2,300,954	100	2.9

Notes. Tax losses are calculated as the difference between the counterfactual tax yield collected in the absence of significant excess interaction and actual tax yield. Counterfactual tax yield is calculated using external reference prices (see Section III) taking into consideration underreporting of quantities (see Section VII for details). Inspectors (brokers) are ranked each semester on the basis of their total tax revenue losses (with rank 1 denoting the inspector with the highest tax losses), with ties arbitrarily split in the case of nonparticipation in the scheme (we assume that inspectors (brokers) that do not participate in the scheme do not contribute to tax losses associated with the scheme). To avoid having these arbitrary splits affect the rankings, we assign to each of the nonparticipating inspectors the average tax yield of inspectors (brokers) that did not participate in the scheme that semester. This effectively amounts to calculating the average over all possible permutations of randomly assigned splits in the case of tie breaks. The statistics in this table reflect averages across semesters (note that it is possible for a different inspector or broker to assume rank 1 in different semesters). The sample covers the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention).

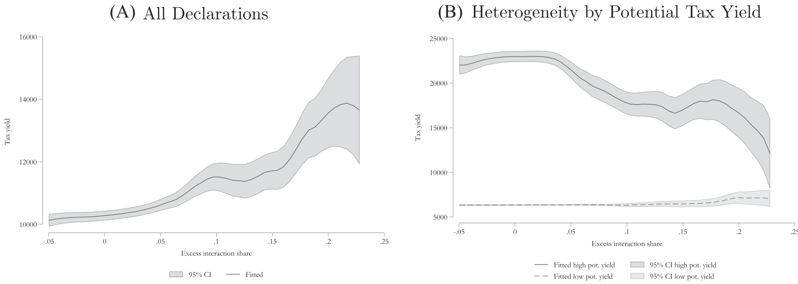

FIGURE V

Tax Yield and Excess Interaction Share

The graphs show weighted local polynomial plots (using the Epanechnikov kernel function) of the tax yield per declaration (in US$) on the share of declarations per inspector and semester that were subject to significant excess interaction (see Section IV). Panel A combines all declarations, whereas Panel B distinguishes across high potential yield and low potential yield declarations. High potential yield (low potential yield) declarations are those for which the hypothetical tax yield (if the declaration was valued using external reference prices) exceeds \$20,000 (is less than \$20,000). The sample covers the period January 1, 2015 to November 17, 2017 (i.e., before the delegated randomization intervention).

tax yield per declaration—may not obviously point to corruption. In fact, based on tax revenue collection numbers alone, one might be tempted to conclude that many of the inspectors with the highest excess interaction are top performers.

Figure V, Panel B reveals that this conclusion would be driven by selection which masks important performance differences. Dividing declarations into "high potential yield" declarations (with a hypothetical tax yield based on external reference prices exceeding \$20,000) and "low potential yield" (all other declarations) reveals that the association between the excess interaction share and tax yield is clearly negative for high potential yield declarations and nonexistent for other declarations. Inspectors with more excess interaction collect substantially less tax revenue on these high potential yield declarations. Yet their average tax yield across all declarations is higher despite their inferior performance simply because they attract a higher share of such high potential yield declarations, as was shown in Figure IV, Panel E (recall that declarations subject to excess interaction are significantly more likely to be high potential yield declarations).[51] The ability of corrupt inspectors to appropriate lucrative declarations thus

51. Online Appendix Figure A4 plots the average inspection share of each inspector per semester against the share of declarations subject to excess interaction

helps explain why they manage to collect more taxes on average despite turning a blind eye to undervaluation among some of the most valuable declarations. Perversely, the inspectors who are most implicated in the corruption scheme and responsible for the largest revenue losses, presumably pocketing the largest illegal bribes, also exhibit nominally superior revenue collection performance.

Tax revenue losses associated with the corruption scheme are also very concentrated among a fairly limited set of brokers, as is shown in Table V, Panel B. The broker accountable for the largest revenue losses in a given semester on average pays $1.7 million worth of taxes (or roughly 2.3% of the total taxes collected in our sample in a given semester) but at the same time evades $514,000 worth of taxes. In other words, their total tax liability would be 29% higher without the corruption scheme. The top three brokers in terms of their contribution to overall tax revenue losses account for half of all revenue losses, the top five brokers account for 71.5% of revenue losses associated with the scheme but only for 17% of overall tax revenue.

IX. Did Delegated Randomization of Inspector Assignment Curb Corruption?

After presenting a preliminary version of this article to the director general (DG) of customs, internal audits were launched and a number of inspectors were either sanctioned for corruption or strongly encouraged to opt for voluntary retirement, and the head of the IT department was suspended. The DG also decided to reform the assignment of declarations to inspectors, by delegating it to GasyNet. Using its own software, GasyNet randomly assigned declarations to active inspectors. This delegated randomization intervention provides us with a unique opportunity to assess whether the excess interactions we document are indeed the product of IT manipulation and hence to validate our methodology to detect corruption. It simultaneously offers a case study of the effectiveness of IT interventions to curb corruption and reduce fraud.[52]

and shows that inspectors with more excess interaction assess a higher share of high potential yield declarations.

52. However, note that the reform exploited in our analysis is not a natural experiment and thus causal claims from its impact need to be taken with caution.

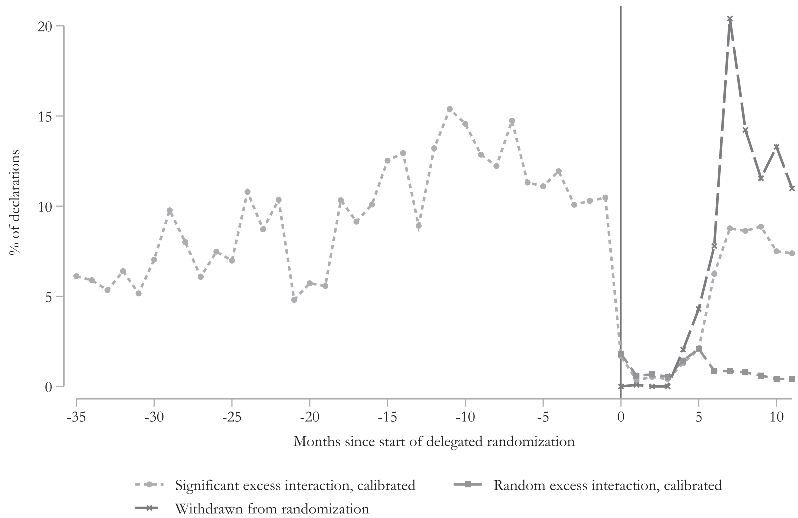

FIGURE VI

Evolution of Nonrandom Assignment

The line with circles "Significant excess interaction, calibrated" depicts the share of all import declarations that are characterized by significant excess interaction, calculated using calibration methods (as explained in Section IV). The vertical bar denotes the start of the delegated randomization intervention in which the assignment of declarations to inspectors was delegated to GasyNet. Soon after this start, the customs IT department managed to withhold several declarations from the randomization process. The prevalence of these declarations is shown by the line with crosses "Withdrawn from randomization." The line with squares "Random excess interaction, calibrated" refers to the share of randomized declarations that are characterized by significant excess interaction, calculated using calibration methods. The sample covers the period January 1, 2015 to November 17, 2018.

IX.A. Prevalence of Excess Interaction during Delegated Randomization Period

The delegated randomization of inspector assignment started on November 18, 2017 and led to the virtual disappearance of excess interaction, as shown in Figure VI, which plots the evolution of the share of declarations characterized by significant excess interaction after automatic assignment. While the prevalence of excess interaction trended upward in the period preceding the delegated randomization intervention, it suddenly and precipitously fell to nearly zero after the start of delegated randomization, indicated by the vertical bar in the graph. The delegated randomization intervention thus effectively eliminated excess interaction between inspectors and brokers.

However, approximately four months after the start of the delegated randomization intervention excess interaction resurfaced, plausibly driven by a new form of IT manipulation: the withholding of certain declarations from the delegated randomization. IT department staff complicit in the corruption scheme figured out a way to temporarily shut down the automatic notification that GasyNet receives when a declaration is registered, thus preventing GasyNet from randomizing the inspector assignment of these declarations. Approximately 7% of declarations (1,275 declarations out of 17,736 declarations registered in the delegated randomization period) were withheld from delegated randomization. These declarations were readily identified by comparing the set of declarations randomized by GasyNet to the set of declarations that cleared customs daily. The set of declarations withheld from delegated randomization likely includes declarations that were not deliberately "targeted" to bypass the randomization. Disabling the automatic notifications for some period implied that none of the declarations registered during that period were randomized by GasyNet, whether or not they were part of a corruption agreement.[53]

The evolution of the withholding of declarations from delegated randomization is depicted by the line with crosses in Figure VI and is remarkably similar to the evolution of significant excess interaction. In fact, 36% of the declarations that were withheld are characterized by significant excess interaction. Conversely, 63% of the declarations characterized by significant excess interaction in the delegated randomization period were withheld from randomization. Interestingly, nonrandom assignment is persistent: for a given pairing of a broker with a particular inspector the share of withheld declarations is correlated with past deviations from random assignment, as shown in Online Appendix Table A5, suggesting that withholding declarations from random assignment reflects a continuation of corruption agreements.

53. The withholding of declarations subject to corruption agreements likely involves coordination between brokers and customs IT department staff: they are likely to agree on a particular time slot during which the delegated randomization is temporarily shut down and the declaration is registered. However, other brokers, who are not part of corruption agreements may also register declarations during these time slots, which implies that not all declarations that are withheld from delegated randomization are part of corruption agreements.

To ascertain that IT manipulation is driving the excess interaction we conduct a simple placebo test: we calculate the prevalence of excess interaction for the subsample of declarations that were randomized by GasyNet. Any excess interaction in this subsample should be purely accidental. Indeed, there is hardly any excess interaction in this subsample, as is shown by the line with rectangles for "random excess interaction" in Figure VI. The only period with some excess interaction is five to seven months after the start of the delegated randomization intervention, when a number of inspectors went on repeated strikes (resulting in a higher average workload, and possibly higher excess interaction shares, for the remaining inspectors). Put differently, without bypassing the delegated randomization, there would not have been a resurgence of excess interaction between inspectors and brokers. The patterns in Figure VI are very similar when based on measures of excess interaction based on inspector-specific logit models, as shown in Online Appendix Figure A5.

IX.B. Excess Interaction and Evasion Risk during the Delegated Randomization Period

Declarations withheld from delegated randomization are not only characterized by significantly higher excess interaction shares but are also significantly more at risk of tax evasion on average than declarations that were randomized by GasyNet, as is shown in Table VI, Panel A, which replicates some of the key specifications presented in Table II for the sample of withheld declarations. They are subject to tax rates that are 8.8% higher, have risk scores that are 1.2 points higher, are significantly heavier, and exhibit 19.7% lower initial unit prices relative to median import unit prices. These declarations exhibit 19.9% higher tax revenue losses than similar declarations whose assignment to inspectors was randomized by GasyNet. Random excess interaction (i.e., excess interaction in the sample of declarations whose assignment was randomized by GasyNet) is not correlated with declaration characteristics commonly associated with tax evasion, as is shown in Panel B; all the R^2's are zero and none of the coefficients are significant.

Even in the delegated randomization period, the excess interaction share is significantly correlated with declaration characteristics associated with tax evasion risk, as is shown in Panel C, which replicates Table II using the entire sample of declarations

TABLE VI

Tax Evasion Risk and Excess Interaction During Delegated Randomization of Inspector Assignment

Dependent variable:	Excess interaction share (1)	Tax rate (2)	Risk score (3)	Log initial weight (4)	Log initial unit price (relative to internal prices) (5)	Initial hyp. tax revenue losses (6)
Panel A: Withheld from randomization						
Withheld from randomization (WFR)	0.064**	0.088***	1.173***	0.160	−0.197***	0.199***
	(0.026)	(0.010)	(0.222)	(0.117)	(0.061)	(0.061)
Observations	17,736	17,738	17,169	17,738	17,728	17,728
R-squared	0.153	0.026	0.011	0.001	0.011	0.012
Panel B: Excess interaction—delegated randomized declarations only						
Random excess interaction share		0.083	−0.024	−0.302	0.076	−0.078
		(0.173)	(2.316)	(1.251)	(0.169)	(0.158)
Observations		16,461	15,925	16,461	16,454	16,454
R-squared		0.000	0.000	0.000	0.000	0.000
Panel C: Excess interaction						
Excess interaction share		0.307***	4.212**	0.392	−0.840**	0.836**
		(0.084)	(1.360)	(0.844)	(0.311)	(0.312)
Observations		17,736	17,167	17,736	17,726	17,726
R-squared		0.008	0.004	0.000	0.005	0.006

TABLE VI
CONTINUED

Dependent variable:	Excess interaction share (1)	Tax rate (2)	Risk score (3)	Log initial weight (4)	Log initial unit price (relative to internal prices) (5)	Initial hyp. tax revenue losses (6)
Panel D: Combined measures						
Withheld from randomization (WFR)		0.079***	1.015***	0.177	−0.110***	0.114***
		(0.009)	(0.234)	(0.139)	(0.032)	(0.033)
Excess interaction share		0.093	1.196	0.147	0.032	−0.032
		(0.094)	(1.609)	(1.251)	(0.104)	(0.098)
WFR×Excess interaction share		0.051	1.169	−0.386	−1.286***	1.262***
		(0.067)	(2.000)	(1.531)	(0.327)	(0.325)
Observations		17,736	17,167	17,736	17,726	17,726
R-squared		0.027	0.011	0.001	0.015	0.016

Notes. Standard errors clustered two-way by inspector and by broker are presented in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively. WFR stands for withheld from randomization. *Excess interaction share* is the difference between the share of a given broker's declarations handled by an inspector in a given semester and the hypothetical share she would be expected to handle if the allocation of declarations to inspectors were random conditional on their productivity, as prescribed by official rules, calculated using calibration methods (as explained in Section IV). *Random excess interaction share* is the excess interaction share calculated using only the set of declarations that were not withheld from randomization. "Observations" refers to the number of nonsingleton observations. OLS estimation is used. The sample covers the period November 18, 2017 to November 17, 2018 (i.e. the period of the delegated randomization intervention).

(randomized and withheld from randomization by GasyNet) in this period. However, these correlations are entirely driven by declarations withheld from randomization by GasyNet, as is shown in Panel D, in which we interact the excess interaction share with the dummy for being withheld from randomization. While being withheld from randomization continues to significantly predict tax evasion risk, the excess interaction share only has predictive power when interacted with being withheld from randomization (consistent with the results in Panel A). The declarations withheld from randomization and cleared by inspectors with a higher excess interaction share have significantly lower initial unit prices and significantly higher initial tax revenue losses (Panel D, columns (5) and (6)). This suggests the declarations withheld from randomization by GasyNet that were targeted by corruption schemes were assigned to certain "preferred" inspectors.

IX.C. Differential Treatment during the Delegated Randomization Period

To evaluate the extent to which the IT manipulation during the delegated randomization period reflects a continuation of corruption, Table VII examines whether inspectors treat the manipulated declarations differently. The table replicates the specifications in Table III but uses different proxies for excess interaction.

Declarations that were withheld from delegated randomization are cleared significantly faster than declarations that were not, as shown in Panel A. The estimates also point to a reduced likelihood of being reported fraudulent and lower value and tax adjustments but these effects are not statistically significant. Declarations withheld from delegated randomization exhibit significant and substantial tax revenue losses of 17.5% on average, relative to other declarations, ceteris paribus.

Panel B shows that for the subsample of declarations randomized by GasyNet, random excess interaction does not predict how long inspectors take to clear goods, whether they will report the declaration as being fraudulent, or change the value or the tax yield. Random excess interaction is negatively correlated with tax revenue losses, suggesting that it is linked to lower, not higher, tax losses, in this sample of declarations randomized by GasyNet.

When we extend the sample by including withheld declarations, excess interaction is again associated with significantly accelerated clearance, significantly reduced fraud, lower value and tax adjustments, and significantly higher tax revenue losses, as

TABLE VII

Differential Treatment during Delegated Randomization of Inspector Assignment

Dependent variable:	Time (1)	Fraud (2)	Δ log value (3)	Δ log tax (4)	Hyp. tax revenue losses (5)
Panel A: Withheld from randomization					
Withheld from randomization (WFR)	−0.853***	−0.015	−0.003	−0.003	0.175***
	(0.114)	(0.020)	(0.005)	(0.005)	(0.032)
Observations	16,455	17,169	17,147	15,188	17,159
R-squared	0.239	0.389	0.271	0.250	0.191
Panel B: Random excess interaction share—delegated randomized declarations only					
Random excess interaction share	−0.193	−0.031	−0.002	0.005	−0.287*
	(0.716)	(0.138)	(0.031)	(0.039)	(0.151)
Observations	15,899	15,925	15,907	13,692	15,918
R-squared	0.227	0.394	0.275	0.259	0.164
Panel C: Excess interaction share					
Excess interaction share	−2.352***	−0.187**	−0.051*	−0.038	0.421**
	(0.600)	(0.069)	(0.025)	(0.027)	(0.150)
Observations	16,453	17,167	17,145	15,186	17,157
R-squared	0.232	0.390	0.271	0.250	0.185

TABLE VII
CONTINUED

Dependent variable:	Time (1)	Fraud (2)	Δ log value (3)	Δ log tax (4)	Hyp. tax revenue losses (5)
Panel D: Excess interaction share and declarations withheld from randomization (and their interaction)					
Excess interaction share	−0.639***	0.007	0.004	0.002	0.129***
	(0.113)	(0.019)	(0.005)	(0.006)	(0.027)
Withheld from randomization (WFR)	−1.060*	−0.075	−0.014	−0.012	−0.097
	(0.522)	(0.078)	(0.021)	(0.023)	(0.101)
WFR×Excess interaction share	−2.542**	−0.370**	−0.128**	−0.083*	0.956**
	(0.899)	(0.124)	(0.043)	(0.046)	(0.303)
Observations	16,453	17,167	17,145	15,186	17,157
R-squared	0.241	0.390	0.272	0.250	0.192

Notes. Standard errors clustered two-way by inspector and by broker are presented in parentheses. ***, **, and * indicate significance at 1%, 5%, and 10% levels, respectively. WFR stands for withheld from randomization. *Excess interaction share* is the difference between the share of a given broker's declarations handled by an inspector in a given semester and the hypothetical share she would be expected to handle if the allocation of declarations to inspectors were random conditional on their productivity, as prescribed by official rules, calculated using calibration methods (as explained in Section IV). *Random excess interaction share* is the excess interaction share calculated using only the set of declarations that were not withheld from randomization. "Observations" refers to the number of nonsingleton observations. All specifications control for the tax rate, the risk score, a dummy for the red channel, the share of value accounted for by differentiated products, a dummy indicating whether the declaration was mixed, and a dummy indicating that the declaration was subject to valuation advice, inspector fixed effects, broker fixed effects, source country fixed effects, and month-year fixed effects. OLS estimation is used. The sample covers the period November 18, 2017 to November 17, 2018 (i.e. the period of the delegated randomization intervention).

shown in Panel C. However, this preferential treatment is driven by the declarations withheld from delegated randomization since we did not observe these correlations in the sample of declarations randomized by GasyNet analyzed in Panel B.

In Panel D we consider the entire sample of declarations and include the excess interaction share, a dummy for being withheld from delegated randomization, and the interaction between these two measures. The coefficients on the interaction term are consistently highly statistically significant. Preferential treatment is most pronounced for declarations that are withheld from delegated randomization and handled by inspectors who interact excessively frequently with a given broker. Such declarations are especially rapidly cleared, especially less likely to be deemed fraudulent, are subject to significantly lower value and tax adjustments, and as a result, exhibit higher tax revenue losses.

The preferential treatment by inspectors of declarations characterized by excess interaction was thus enabled by manipulation of the IT system. Our placebo tests show clearly that when declarations are truly randomly assigned, there is hardly any excess interaction. Whatever accidental excess interaction nonetheless arises is not correlated with customs outcomes. By contrast, declarations withheld from delegated randomization are associated with excess interaction and an increased risk of tax evasion. They receive privileged treatment from inspectors, especially when such inspectors are handling a significantly larger share of a given broker's declarations than would be expected had the assignment of declarations followed official rules. All in all, these results corroborate our methodology to detect corruption and also attest to the difficulties associated with dislodging systemic corruption.

An event study of the effect of the delegated randomization on tax yield per declaration, shown in Online Appendix Figure A6, is consistent with this interpretation. Tax revenues increased significantly in the first few months of delegated randomization, but these gains were not sustained.[54]

54. For the event study, we use a sample including six months before and after the start of the delegated randomization and estimate an OLS regression of log 1 plus tax yield per declaration on dummies that define the position of the month relative to November 2017 as well as inspector, broker, source country, HS2-product, and calendar month fixed effects.

Online Appendix Table A24 (Panel B) presents estimates of the costs of corruption during the delegated randomization period, following a similar approach to that described in Section VII.[55] According to our preferred estimates which calculate hypothetical tax yield using prices reported by exporters, declarations that were likely the object of corruption—notably those withheld from randomization cleared by an inspector who interacted excessively frequently with the broker registering them—would have yielded an additional \$11,223 in tax revenue. This represents a 129.8% increase over actual tax yield. Aggregate tax yield in this period would have been 2.6% higher had the randomization not been undermined by a new form of IT manipulation. These back-of-the-envelope estimates are crude and must be interpreted with caution given the difficulties inherent in measuring hypothetical tax yield and identifying deviant declarations.

X. Conclusion

Corrupt governance and limited state capacity to raise tax revenue constrain development, yet surprisingly little is known about the extent to which tax evasion is facilitated by (which) bureaucrats. Evidence on effectiveness of reforms to remedy institutionalized corruption is also limited. These questions are especially pertinent for customs agencies in low-income countries, which tend to be more reliant on tax revenues collected at the border than developed countries, despite suffering higher levels of evasion.

This article presents a new methodology to detect a specific form of corruption between customs inspectors and customs brokers, which we believe can be readily replicated in other contexts in which random assignment is used to deter corruption. Our approach is based on identifying deviations from random assignment of import declarations to inspectors, which is prescribed by official rules. Such deviations result in excessively frequent pairing of brokers with the inspector(s) they are conspiring with.

55. We use estimates from regressions of hypothetical tax revenue losses on the excess interaction share, a dummy for declarations being withheld, and their interaction shown in Online Appendix Table A23. Tax revenue in the absence of corruption is now calculated as $T^{NC} = T * exp(\widehat{\beta_E} * ES + \widehat{\beta_P} * WFR + \widehat{\beta_{EP}} * ES * WFR)$, where *WFR* is a dummy for declarations whithheld from randomization. The details on this calculation are shown in Online Appendix C.

Applying this methodology to Madagascar's main port of Toamasina unveiled that 10% of declarations were handled by inspectors who were not randomly assigned, plausibly because of manipulation of the IT system that assigns them. Nonrandomly assigned declarations were shown to be subject to higher tax rates, have higher potential tax yield, higher risk scores, and lower unit prices than those reported for declarations containing the same goods. Nonrandom assignment is thus associated with higher tax revenue losses. Customs inspectors are shown to provide preferential treatment to these deviant declarations by clearing them faster; being less likely to require value, weight, and tax adjustments; and failing to identify fraud. Such corruption is costly; tax yield for nonrandomly assigned declarations would have been 26% higher in the absence of excess interaction between inspectors and brokers. Overall tax revenues collected in Toamasina would have been 3% higher in the absence of the corruption scheme unveiled in this study. These tax losses are very concentrated among a select few inspectors and brokers, whose propensity to engage in corruption increases with tenure in the port. Paradoxically, inspectors responsible for the largest tax revenue losses tend to collect more tax revenue per declaration, because they manage to control the assessment of the most lucrative declarations. Corruption is thus positively correlated with (naive) measures of tax revenue yield.

An intervention to curb corruption by having a third party randomize inspector assignment validates our methodology as it led to the temporary disappearance of excess interaction between inspectors and brokers. It also triggered a novel form of IT manipulation. While manipulation of inspector assignment was eventually weeded out with the help of improved IT infrastructure, our results serve as a reminder that technology is not a panacea in the fight against corruption. Rather, our results illustrate how IT solutions can be captured by bureaucrats and economic operators and serve as a conduit to corruption.

International Trade Centre, Switzerland
World Bank, USA
World Bank, USA, and Université Clermont Auvergne, Centre National de la Recherche Scientifique, and Centre d'Études et de Recherches sur le Développement International Clermont-Ferrand, France
World Bank, USA, and Utrecht University, the Netherlands

Supplementary Material

Supplementary material is available at the *Quarterly Journal of Economics* online.

Data Availability

Code replicating the tables and figures in this article can be found in Chalendard et al. (2022) in the Harvard Dataverse, https://doi.org/10.7910/DVN/DPO4BS.

References

Atkin, David, and Amit K. Khandelwal, "How Distortions Alter the Impacts of International Trade in Developing Countries," *Annual Review of Economics*, 12 (2020), 213–238. https://doi.org/10.1146/annurev-economics-081919-041554

Banerjee, Abhijit, Rema Hanna, and Sendhil Mullainathan, "Corruption," in "The Handbook of Organizational Economics," Robert Gibbons and John Roberts, eds. (Princeton, NJ: Princeton University Press, 2012), Ch. 27, 1109–1147. https://doi.org/10.3386/w17968

Bardhan, Pranab, "Corruption and Development: A Review of Issues," *Journal of Economic Literature*, 35 (1997), 1320–1346. https://www.jstor.org/stable/2729979

Baunsgaard, Thomas, and Michael Keen, "Tax Revenue and (or?) Trade Liberalization," *Journal of Public Economics*, 94 (2010), 563–577. https://doi.org/10.1016/j.jpubeco.2009.11.007

Besley, Timothy, and Torsten Persson, "The Origins of State Capacity: Property Rights, Taxation, and Politics," *American Economic Review*, 99 (2009), 1218–1244. https://doi.org/10.1257/aer.99.4.1218

Bhagwati, Jagdish, "On the Underinvoicing of Imports," *Bulletin of the Oxford University Institute of Economics & Statistics*, 27 (1964), 389–397. https://doi.org/10.1016/B978-0-444-10581-3.50016-1

Broda, Christian, and David E. Weinstein, "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121 (2006), 541–585. https://doi.org/10.1162/qjec.2006.121.2.541

Casaburi, Lorenzo, Michael Kremer, and Ravindra Ramrattan, "Crony Capitalism, Collective Action, and ICT: Evidence from Kenyan Contract Farming," University of Zurich Working Paper, 2019.

Chalendard, Cyril, Ana M. Fernandes, Gael Raballand, and Bob Rijkers, "Replication Data for: 'Corruption in Customs'," (2022), Harvard Dataverse, https://doi.org/10.7910/DVN/DPO4BS.

Chalendard, Cyril, Gaël Raballand, and Antsa Rakotoarisoa, "The Use of Detailed Statistical Data in Customs Reforms: The Case of Madagascar," *Development Policy Review*, 37 (2019), 546–563. https://doi.org/10.1111/dpr.12352

Chalendard, Cyril, Alice Duhaut, Ana M. Fernandes, Aaditya Mattoo, Gaël Raballand, and Bob Rijkers, "Does Better Information Curb Customs Fraud?," CESifo Working Papers 8371, 2020. https://dx.doi.org/10.2139/ssrn.3633656

Dincecco, Mark, and Nico Ravanilla, "The Importance of Bureaucrats in a Weak State: Evidence from the Philippines," available at SSRN 2773884 (2017). http://dx.doi.org/10.2139/ssrn.2773884

Dutt, Pushan, and Daniel Traca, "Corruption and Bilateral Trade Flows: Extortion or Evasion?," *Review of Economics and Statistics*, 92 (2010), 843–860. https://doi.org/10.1162/REST_a_00034

Ferraz, Claudio, and Frederico Finan, "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes," *Quarterly Journal of Economics*, 123 (2008), 703–745. https://doi.org/10.1162/qjec.2008.123.2.703

Finan, Frederico, Benjamin A. Olken, and Rohini Pande, "The Personnel Economics of the Developing State," in *Handbook of Economic Field Experiments,* vol. 2, Abhijit Banerjee and Esther Duflo, eds. (Amsterdam: Elsevier, 2017), Ch. 6, 467–514. https://doi.org/10.1016/bs.hefe.2016.08.001

Fisman, Raymond, and Shang-Jin Wei, "Tax Rates and Tax Evasion: Evidence from 'Missing Imports' in China," *Journal of Political Economy*, 112 (2004), 471–496. https://doi.org/10.1086/381476

Fontagné, Lionel, Houssein Guimbard, and Gianluca Orefice, "Tariff-Based Product-Level Trade Elasticities," *Journal of International Economics*, 137 (2022), 103593. https://doi.org/10.1016/j.jinteco.2022.103593

Freund, Caroline, Mary Hallward-Driemeier, and Bob Rijkers, "Deals and Delays: Firm-Level Evidence on Corruption and Policy Implementation Times," *World Bank Economic Review*, 30 (2016), 354–382. https://doi.org/10.1093/wber/lhv001

Gordon, Roger, and Wei Li, "Tax Structures in Developing Countries: Many Puzzles and a Possible Explanation," *Journal of Public Economics*, 93 (2009), 855–866. https://doi.org/10.1016/j.jpubeco.2009.04.001

Guimarães, Paulo, and Pedro Portugal, "A Simple Feasible Procedure to Fit Models with High-Dimensional Fixed Effects," *The Stata Journal*, 10 (2010), 628–649. https://doi.org/10.1177%2F1536867X1101000406

Hirano, Keisuke, Guido W. Imbens, and Geert Ridder, "Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score," *Econometrica*, 71 (2003), 1161–1189. https://doi.org/10.1111/1468-0262.00442

Kaufmann, Daniel, and Shang-Jin Wei, "Does 'Grease Money' Speed up the Wheels of Commerce?," NBER Working Paper no. 7093, 1999. https://doi.org/10.3386/w7093

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

———, "Making Moves Matter: Experimental Evidence on Incentivizing Bureaucrats through Performance-Based Postings," *American Economic Review*, 109 (2019), 237–270. https://doi.org/10.1257/aer.20180277

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*, 79 (2011), 651–692. https://doi.org/10.3982/ECTA9113

Krinsky, Itzhak, and A. Leslie Robb, "On Approximating the Statistical Properties of Elasticities," *Review of Economics and Statistics*, 68 (1986), 715–719. https://doi.org/10.2307/1924536

Laajaj, Rachid, Marcela Eslava, and Tidiane Kinda, "The Costs of Bureaucracy and Corruption at Customs: Evidence from the Computerization of Imports in Colombia," Documento CEDE, 2019. https://dx.doi.org/10.2139/ssrn.3334529

Leff, Nathaniel H., "Economic Development through Bureaucratic Corruption," *American Behavioral Scientist*, 8 (1964), 8–14. https://doi.org/10.1177%2F000276426400800303

Lichand, Guillaume, and Gustavo Fernandes, "The Dark Side of the Contract: Do Government Audits Reduce Corruption in the Presence of Displacement by Vendors?" University of Zurich working paper, 2019.

Martin, Julien, Isabelle Mejean, and Mathieu Parenti, "Relationship Stickiness and Economic Uncertainty," CEPR Discussion Paper DP15609, 2020.

Niehaus, Paul, and Sandip Sukhtankar, "Corruption Dynamics: The Golden Goose Effect," *American Economic Journal: Economic Policy*, 5 (2013), 230–269. https://doi.org/10.1257/pol.5.4.230

Olken, Benjamin A., and Rohini Pande, "Corruption in Developing Countries," *Annual Review of Economics*, 4 (2012), 479–509. https://doi.org/10.1146/annurev-economics-080511-110917

Pepinsky, Thomas B., Jan H. Pierskalla, and Audrey Sacks, "Bureaucracy and Service Delivery," *Annual Review of Political Science*, 20 (2017), 249–268. https://doi.org/10.1146/annurev-polisci-051215-022705
Pomeranz, Dina, and José Vila-Belda, "Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities," *Annual Review of Economics*, 11 (2019), 755–781. https://doi.org/10.1146/annurev-economics-080218-030312
Rauch, James E., "Networks versus Markets in International Trade," *Journal of International Economics*, 48 (1999), 7–35. https://doi.org/10.1016/S0022-1996(98)00009-9
Rijkers, Bob, Leila Baghdadi, and Gael Raballand, "Political Connections and Tariff Evasion Evidence from Tunisia," *World Bank Economic Review*, 31 (2017), 459–482. https://doi.org/10.1093/wber/lhv061
Sequeira, Sandra, "Corruption, Trade Costs, and Gains from Tariff Liberalization: Evidence from Southern Africa," *American Economic Review*, 106 (2016), 3029–3063. https://doi.org/10.1257/aer.20150313
Sequeira, Sandra, and Simeon Djankov, "Corruption and Firm Behavior: Evidence from African Ports," *Journal of International Economics*, 94 (2014), 277–294. https://doi.org/10.1016/j.jinteco.2014.08.010
Shleifer, Andrei, and Robert W. Vishny, "Corruption," *Quarterly Journal of Economics*, 108 (1993), 599–617. https://doi.org/10.2307/2118402
———, *The Grabbing Hand: Government Pathologies and Their Cures* (Cambridge, MA: Harvard University Press, 2002).
Sison, Cristina P., and Joseph Glaz, "Simultaneous Confidence Intervals and Sample Size Determination for Multinomial Proportions," *Journal of the American Statistical Association*, 90 (1995), 366–369. https://doi.org/10.1080/01621459.1995.10476521
Slemrod, Joel, "Tax Compliance and Enforcement," *Journal of Economic Literature*, 57 (2019), 904–954. https://doi.org/10.1257/jel.20181437
Wier, Ludvig, "Tax-Motivated Transfer Mispricing in South Africa: Direct Evidence Using Transaction Data," *Journal of Public Economics*, 184 (2020), 104153. https://doi.org/10.1016/j.jpubeco.2020.104153
World Bank, *Doing Business 2020: Comparing Business Regulation in 190 Economies*. Washington, DC: Work Bank 2020. https://openknowledge.worldbank.org/handle/10986/32436
Xu, Guo, "The Costs of Patronage: Evidence from the British Empire," *American Economic Review*, 108 (2018), 3170–3198. https://doi.org/10.1257/aer.20171339
Xu, Guo, Marianne Bertrand, and Robin Burgess, "Social Proximity and Bureaucrat Performance: Evidence from India," NBER Working Paper no. 25389, 2018. https://doi.org/10.3386/w25389
Yang, Dean, "Can Enforcement Backfire? Crime Displacement in the Context of Customs Reform in the Philippines," *Review of Economics and Statistics*, 90 (2008a), 1–14. https://doi.org/10.1162/rest.90.1.1
———, "Integrity for Hire: An Analysis of a Widespread Customs Reform," *Journal of Law and Economics*, 51 (2008b), 25–57.
Zitzewitz, Eric, "Forensic Economics," *Journal of Economic Literature*, 50 (2012), 731–769. https://doi.org/10.1257/jel.50.3.731