

STRATEGIC TAX EVASION THROUGH CUSTOMS OFFICE SELECTION*

EVIDENCE FROM 20 MILLION IMPORT TRANSACTIONS IN HAITI

Nikolaos Angelopoulos Kensley Blaise Rose Camille Vincent †

07.03.2026

Preliminary Draft. Please **DO NOT** cite.

Abstract

Trade taxation remains a central source of public revenue in low-capacity states, yet customs administrations often exhibit weak and uneven enforcement. In this paper, we investigate an under-explored margin of tax evasion: strategic selection among customs offices within a single country. We ask whether importers exploit spatial heterogeneity in enforcement by routing otherwise identical imports through offices where the expected cost of evasion is lower. We assembled comprehensive administrative data on more than 2.16 million import declarations in Haiti from 2017 to 2024. Because Haiti applies a uniform national tariff schedule, systematic differences in reporting behaviour across offices reflect differences in implementation rather than legal rules. Our empirical strategy exploits within-firm and within-product variation to estimate office-specific elasticities of declared import values with respect to ad-valorem tax rates. We find a large and economically meaningful response of declared values to taxes, with an elasticity close to minus one. Allowing the tax response to vary by office reveals substantial enforcement heterogeneity: secondary ports exhibit tax elasticities that are 2 to 3 times larger than those at the main port. These patterns persist under common-support restrictions and within-firm specifications that compare the same importer clearing the same product at the same time through different offices. We address concerns about endogenous office choice and measurement error with a re-centred exposure instrumental-variables strategy based on broker routing networks interacting with product-time tax intensity. The IV estimates closely mirror the baseline results. The findings show that decentralised customs administration generates exploitable enforcement heterogeneity, with important implications for revenue mobilisation and customs reform in fragile states.

JEL Codes: H26, H21, F13, O17

Keywords: trade taxation; tax evasion; customs enforcement; revenue mobilization; developing countries; Haiti

* Acknowledgment: TO BE ADDED. Usual disclaimers apply

† Nikolaos Angelopoulos is affiliated with the School of Economics, University of East Anglia; Kensley Blaise with the School of Economics, University of East Anglia; and Rose Camille Vincent with Utrecht School of Economics, Utrecht University.

1 Introduction

Trade taxation is often viewed as the last administratively viable instrument in low-capacity states. Unlike income and profit taxes, which rely on dense third-party reporting networks, reliable registries and credible audit capacity, *import duties* can be collected at a limited number of border gateways where goods physically cross the territory and where enforcement can, in principle, be concentrated. This administrative logic helps explain why trade taxes historically dominated public revenue and why they remain salient in many developing economies even after decades of trade liberalization and tax modernization efforts. When domestic tax bases are difficult to observe or enforce, border taxes remain an attractive source of revenue (see for e.g. [Baunsgaard and Keen, 2010](#); [Cagé and Gadenne, 2018](#)) and reductions in trade taxes have proven to be fiscally costly in settings with large informal sectors and limited enforcement of domestic indirect taxes (see for e.g. [Emran and Stiglitz, 2005](#)). For governments facing severe administrative constraints, customs duties are not merely a legacy instrument but a core pillar of fiscal capacity.

Yet the very features that make customs taxation feasible in low-capacity settings also make it vulnerable. Valuation is partly subjective, product classification is complex, and border posts are high-rent environments in which discretion, opacity, and repeated interactions between firms and officials can support corruption and collusion ([Yang, 2008](#); [Chalendard et al., 2023](#)). In fact, tax evasion through trade mis-invoicing has been documented as one of the largest and most persistent sources of revenue loss in developing countries. Estimates of illicit financial flows linked to trade manipulation suggest amounts on the order of \$1.6 trillion for 134 developing countries over the period 2009–2018, a magnitude that exceeds total development aid by a wide margin ([Global Financial Integrity, 2021](#)). For low-income countries, where trade taxes still account for roughly 15–30% of government revenue – compared to approximately 1–2% in high-income economies – failures in customs administration can pose an existential threat to fiscal capacity ([Keen and Ligthart, 2002](#); [Baunsgaard and Keen, 2010](#)). In such contexts, the central question is not whether mis-invoicing exists, but which mechanisms structure it and how institutional weaknesses shape both its incidence and organization.

This paper examines a previously under-researched margin of tax evasion: strategic selection among customs offices within a single country. We ask whether importers exploit spatial heterogeneity in enforcement by routing otherwise identical imports through offices where the expected cost of evasion is lower. This margin is conceptually distinct from the well-documented mechanisms of under-valuation, quantity misreporting, or mis-classification. It arises only when enforcement is uneven across locations, even though the legal framework governing trade taxation is uniform. When such heterogeneity exists, firms gain an additional degree of freedom: they can respond to tax incentives not only by adjusting how they declare imports, but also by choosing where they clear them.

We study this mechanism using comprehensive administrative data on 2.16 million import

transactions in Haiti from 2017 to 2024. For each transaction, we observe the importer identity, the product classification, country of origin, declared values and quantities, the applicable tariff schedule, and the customs office of clearance. Haiti’s tariff code and statutory rates are national, so – conditional on product, timing, and origin – differences in declared values across offices cannot be attributed to differences in the law.¹ Instead, they reflect differences in implementation, monitoring, and enforcement practices. This institutional feature allows us to treat the office of clearance as an enforcement environment embedded within a common legal and macroeconomic setting.

Our empirical strategy exploits within-importer and within-product variation in office choice. We compare the same importer clearing the same narrowly defined product – facing the same statutory tariff rate – across different customs offices, and we estimate office-specific semi-elasticities and elasticities of declared values with respect to the applicable tax rate. This design is motivated by a core identification challenge in the evasion literature: firms that evade may differ systematically from firms that comply, and enforcement may be targeted endogenously toward particular firms, products, or shipments. By conditioning on importer fixed effects and product fixed effects, we absorb time-invariant importer characteristics such as sophistication, political connections, compliance culture, and willingness to pay bribes, as well as product-specific attributes such as quality differentiation and valuation uncertainty. What remains is variation that operates through differences in the enforcement environment itself. This approach aligns with a broader methodological shift in public finance toward exploiting micro-level administrative variation and within-entity comparisons to disentangle behavioral responses from compositional differences (Kleven et al., 2011; Carrillo et al., 2017; Almunia and Lopez-Rodriguez, 2018).

The resulting pattern is consistent with spatial arbitrage in enforcement. The elasticity of declared values with respect to tax rates varies sharply across offices, ranging from approximately -0.9 at the main port to about -2.5 at smaller secondary offices. Crucially, this dispersion reflects genuine enforcement heterogeneity rather than product specialization or firm sorting. The same firm importing the same product declares systematically lower values when clearing through more permissive offices. These findings indicate that office choice itself constitutes a meaningful evasion margin and that decentralized enforcement structures can generate substantial differences in effective tax burdens within a single legal regime.

These findings place the paper at the intersection of three strands of literature. First, it relates to work on trade taxation and tariff evasion in developing countries, which shows that evasion responds strongly to tariff rates and is particularly pronounced for differentiated products, but typically infers evasion from cross-product or cross-country variation while treating enforcement as homogeneous (Fisman and Wei, 2004; Javorcik and Narciso, 2008). Second, it speaks to the literature on enforcement, discretion, and corruption in customs administrations, which emphasizes the role of organizational design and personnel incentives but has focused primarily on discretion

¹see robustness details

within offices rather than heterogeneity across them (Yang, 2008; Chalendard et al., 2023). Third, it contributes to the broader literature on state capacity and revenue mobilization in fragile settings, which highlights how persistent revenue leakage undermines the state’s ability to invest in enforcement capacity itself (Besley and Persson, 2011, 2013; Jensen and Weigel, 2025). By bringing these strands together, the paper highlights the organizational and spatial dimensions of enforcement as a first-order channel of revenue loss.

Haiti provides an unusually informative setting in which to study these issues. As one of the world’s most fragile states, Haiti exemplifies the challenges faced by governments with extremely limited fiscal capacity. The country ranks near the bottom of global governance and corruption indicators (Transparency International, 2024; World Bank, 2024b), and domestic revenue mobilization remains exceptionally low, at roughly 6–7% of GDP in recent fiscal years (World Bank, 2024a, 2025). In this environment, the budget constraint is binding in a literal sense: revenue shortfalls translate rapidly into failures of core state functions. External debt service also represents a growing fiscal burden, and limited domestic capital markets make revenue collection directly binding on expenditure. Every dollar lost to customs evasion translates immediately into reduced public services – including security provision and infrastructure maintenance – or increased dependence on foreign aid.²

Within this constrained fiscal structure, customs occupies a central role. In Haiti, the customs administration accounts for a majority share of total government revenue, yet a broad set of diagnostics documents pervasive undervaluation and corruption at the border.³ Despite sustained investment in modernization, enforcement outcomes remain weak.⁴ Understanding customs dynamics therefore carries direct policy relevance. Persistent revenue leakage undermines the state’s ability to invest in enforcement capacity itself, reinforcing low-capacity equilibria in which weak institutions and low compliance mutually sustain one another (Besley and Persson, 2011, 2013). More importantly, identifying the sources of revenue loss is central to reform design.

²The Inter-American Development Bank estimates that closing Haiti’s revenue gap — bringing collection to levels typical of countries at similar income levels — would require doubling customs revenue, an impossible target without addressing endemic evasion (Inter-American Development Bank, 2023).

³For instance, IMF technical assistance reports describe widespread undervaluation, noting that import values declared in Haiti are often 40–60% below values declared by exporting countries for the same shipments (International Monetary Fund, 2021). The Inter-American Development Bank estimates that Haiti forgoes \$200–300 million annually in customs revenue due to evasion, equivalent to 15–20% of total government revenue (Inter-American Development Bank, 2023). The World Bank investment climate assessment identifies customs corruption as the second-most-cited obstacle to doing business in Haiti, after electricity (World Bank, 2022). Evaluations of donor-supported modernization programs report mixed results: while computerization has expanded coverage, it has not translated into commensurate revenue gains, suggesting that technical interventions alone cannot overcome institutional resistance (USAID, 2020).

⁴International donors, including the IMF, World Bank, and USAID, have invested heavily in customs reform across developing countries, often with disappointing results (Cantens et al., 2024; World Bank, 2025). The standard reform package—computerization, risk-based inspections, performance contracts, and anti-corruption measures—presumes a baseline level of state capacity and political commitment that is frequently absent in fragile states. Although the rollout of ASYCUDA has expanded computerized coverage in Haiti, technology alone has proven insufficient to close enforcement gaps when personnel incentives, organizational oversight, and political constraints remain binding (Okunogbe and Santoro, 2023; Chalendard et al., 2023).

If leakage primarily reflects technological deficits—such as poor valuation databases or limited risk analytics—one would expect relatively uniform under-reporting across offices operating under the same systems. If, instead, evasion concentrates in specific offices where enforcement is weaker, this implies that decentralized administration generates exploitable heterogeneity even within a single institution. In that case, reforms focused solely on technology are likely to yield limited returns unless they are complemented by organizational reforms, monitoring, and targeted oversight that directly address personnel incentives and local enforcement practices (Okunogbe and Pouliquen, 2022; Chalendard et al., 2023).

Our contribution is fourfold. First, we document a spatial margin of tax evasion – strategic customs office selection – showing that decentralized enforcement generates exploitable heterogeneity in effective tax burdens within a single legal regime. While trade mis-invoicing is well documented, this spatial dimension of evasion within a country’s customs system has received little direct empirical attention as argued above. Second, we provide clean identification of enforcement effects using within-firm, within-product comparisons across customs offices operating under a uniform national tariff schedule. This approach allows us to separate behavioral responses to enforcement from compositional differences in firms and products that confound cross-sectional and aggregate analyses (Fisman and Wei, 2004; Javorcik and Narciso, 2008). Third, we quantify the magnitude and incidence of this mechanism. Office-specific elasticities of declared values with respect to tax rates differ nearly threefold, multi-office importers account for a disproportionate share of total trade value, and counterfactual equalization of enforcement stringency across offices would raise tariff revenue substantially in our estimates. These results indicate that spatial arbitrage is not a marginal phenomenon but a first-order channel of revenue loss in decentralized customs systems. Fourth, we contribute to debates on administrative design in low-capacity settings by showing that personnel incentives and monitoring, rather than technology alone, are central to explaining observed enforcement gaps. This evidence informs the sequencing of reforms in fragile states, highlighting the limits of purely technical solutions and the importance of organizational and incentive-compatible interventions (Besley and Persson, 2011; Okunogbe and Pouliquen, 2022; Okunogbe and Santoro, 2023; Chalendard et al., 2023).

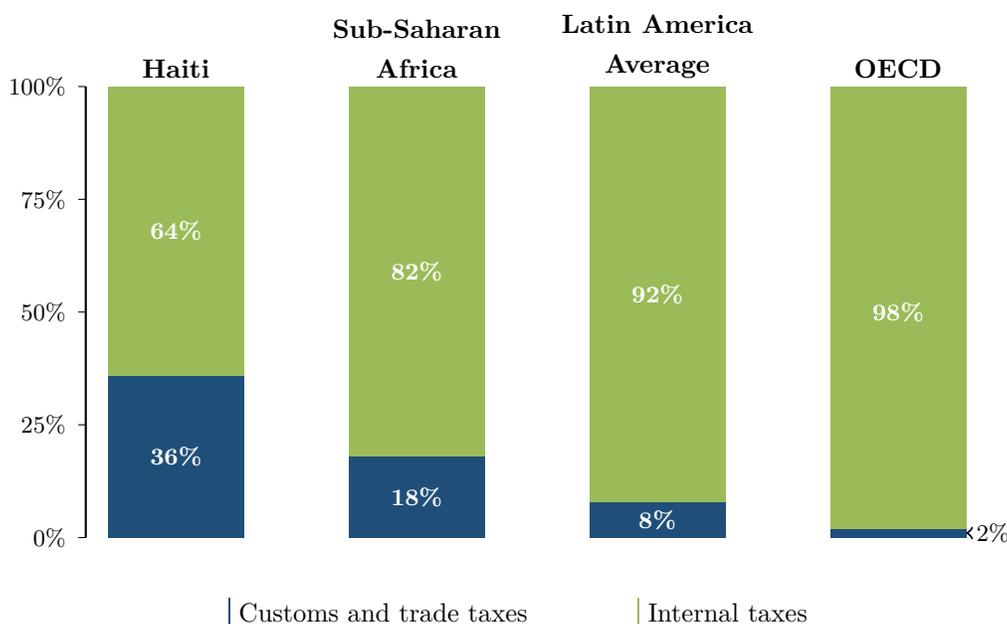
The remainder of this paper proceeds as follows. Section ?? reviews the related literature and institutional context. Section 2 describes the data sources. Section ?? presents the empirical framework and identification strategy. Section 3 reports the main results. Section 4 checks the robustness of our results and further examines underlying mechanisms. Section 5 concludes with policy implications for customs administration in fragile and low-capacity states.

2 Institutional Context and Data

Haiti’s fiscal structure and the primacy of import duties

Haiti presents an extreme case of fiscal dependence on trade taxation. Import duties and related border charges account for approximately 36% of total government revenue, making customs the single most important pillar of state fiscal capacity (International Monetary Fund, 2024; World Bank, 2025). As Figure 1 illustrates, this dependence far exceeds regional and global averages. It reflects fundamental constraints on alternative revenue instruments: a large informal sector and weak mechanisms for enforcing profit and income taxes, which require reliable accounting and third-party reporting (Besley and Persson, 2011). In contrast, import duties can be assessed at physical border crossings where goods are observable. The government’s direct price controls on petroleum products, calibrated explicitly to budgetary needs, illustrate how tightly customs revenue is bound to core state functions (International Monetary Fund, 2024).

FIGURE 1: Composition of government tax revenue: Haiti vs. comparators



Notes: The figure shows the share of total government tax revenue derived from customs and trade taxes (dark) versus internal taxes (light). Haiti’s customs dependence (36%) far exceeds regional and global averages. Data for Haiti from IMF Article IV (2024); comparator averages from World Bank World Development Indicators and IMF Government Finance Statistics.

Haiti is also highly centralized. The main port of Port-au-Prince operates under closer supervision from the customs administration (*Administration Générale des Douanes*), with greater staffing and tighter connections to central authority. Peripheral ports operate with less oversight and greater institutional distance from the capital. These features make Haiti informative for studying customs enforcement, but also suggest broader relevance: many low-income countries

share similar fiscal structures, including high dependence on trade taxes, weak internal revenue enforcement, and uneven territorial reach (Baunsgaard and Keen, 2010).

Customs declarations, coverage, and analysis sample

We use administrative microdata from the Haitian customs administration covering standard import declarations (procedure code IM4) over the period 2017–2024. The data originate from the Ministry of Economy and Finance of Haiti (*Ministère de l'Économie et des Finances*) and are compiled by the Central Tax Authority (*Direction Générale des Impôts*) and the Customs Administration (*Administration Générale des Douanes*), in collaboration with the Central Bank of Haiti (*Banque de la République d'Haïti*). The full administrative universe covers all import declarations recorded by the customs information system during this period.

For the purposes of this paper, we restrict attention to sea-cargo clearances under the standard import regime (IM4) and exclude observations processed under other import regimes and transport modes.⁵

Declarations are recorded at the product-line level. A single customs declaration may contain multiple product lines, each classified under the Harmonized System (HS).⁶ Each line reports a declared value, quantity measures, and net weight. Our unit of observation is therefore a declaration line—an importer–HS6–office–origin–month entry—corresponding to a specific product shipment cleared through a given customs office at a given time.

After cleaning and variable construction, the IM4 sea-cargo dataset contains approximately 2.41 million declaration-line observations spanning 92 months, 15 customs offices, and 189 origin countries. Appendix A reports detailed diagnostics on declaration structure, including the prevalence of multi-line declarations and the extent to which declaration-level values are repeated across lines. These diagnostics show that repeated declaration-level values account for a small share of total import value, indicating that the analysis exploits genuinely line-level variation.

For the econometric analysis, we focus on the largest importers by activity. Specifically, we restrict attention to the top 5,000 importers ranked by transaction count in the cleaned IM4 sea-cargo dataset (5,001 firms due to ties). This restriction captures the bulk of formal import activity while preserving substantial within-firm variation in clearance location. The resulting estimation sample contains approximately 1.58 million declaration-line observations, covers 4,184 HS6 products, and spans all 92 months of the sample period. These top importers account for 94.5% of total import value in the cleaned IM4 sea-cargo data.

⁵IM4 denotes the standard import procedure in the Haitian customs system.

⁶The Harmonized System is an internationally standardized product classification maintained by the World Customs Organization. We work at the six-digit level (HS6), which defines products consistently across countries.

Products and unit values

Products are classified at the six-digit level of the Harmonized System (HS6). We construct declared unit values in U.S. dollars per kilogram. Declared customs values are reported in Haitian gourdes and are converted to USD using a median monthly exchange rate inferred from the administrative records. Unit values are then obtained by dividing USD values by reported net weight.

To limit the influence of reporting errors and extreme observations, we trim the unit-value distribution by dropping the bottom and top 1% of unit values within each HS4×year cell. This trimming rule is applied uniformly across offices and time periods. By construction, it removes 2.35% of observations and 4.05% of total import value while preserving economically meaningful variation in declared values.

Taxes and effective tax wedges

For each declaration line, we summarize tax exposure using an effective ad-valorem tax wedge defined as the ratio of total taxes paid to declared customs value,

$$\tau \equiv \frac{\text{total taxes paid}}{\text{declared customs value}}.$$

We refer to τ as a wedge rather than a statutory tax rate because it is constructed from declared values and therefore reflects the combined outcome of the national tariff-and-fee schedule, exemptions, valuation practices, and enforcement at the point of clearance. It is an equilibrium object rather than a pure policy parameter, and understanding how it behaves in practice is one of the paper’s core concerns.

Because the denominator is the declared value itself, extreme ratios can arise when values are reported at very low levels. To limit undue leverage from such observations, we cap τ at 150%. Tax information is available for the large majority of declaration lines, though coverage varies across offices (Table 1); in the estimation sample, we restrict attention to observations with a positive, non-missing effective tax wedge.

Effective tax wedges are large and highly non-Gaussian. In the estimation sample, the unweighted mean wedge equals 30.1% and the median equals 28.7%, with an interquartile range of 22.5%–40.2% (10th–90th percentiles: 17.3%–44.8%). The value-weighted mean wedge is substantially lower at 22.0%, indicating that higher-value declaration lines face systematically lower effective wedges. The distribution is also markedly discrete: rounding τ to two decimal places, the ten most common values account for 88.7% of declaration lines and 76.9% of total import value. This concentration at a small number of values is consistent with the application of a rule-based tariff-and-fee schedule to declared values, potentially combined with exemptions and valuation discretion.

Haiti applies a nationally set tariff schedule. Statutory ad-valorem rates vary by product and

over time, but they are not intended to vary across customs offices; Appendix A presents evidence consistent with this institutional feature. Office-level differences in effective tax wedges for the same product and period are therefore unlikely to reflect office-specific tariff policy. Instead, they arise from differences in implementation and administration, including valuation scrutiny, inspection intensity, and the application of exemptions and fees.

Finally, because τ is constructed from declared values, it is mechanically related to reported unit values. For this reason, the analysis below does not treat τ as an exogenous policy variable. Instead, the empirical strategy emphasizes product–time variation in tax pressure and examines whether reporting responses to that variation differ systematically across customs offices.

Customs offices

Each declaration line is processed by a single customs office, which serves as the locus of valuation, inspection, and tax assessment. Import activity is highly concentrated across locations. In the cleaned IM4 sea-cargo data, Port-au-Prince (10DP) alone accounts for approximately 91.7% of total import value, while Cap-Haïtien (20CH) accounts for 6.7%; no other individual office exceeds 1%.

Table 1 reports descriptive statistics by customs office, including shipment counts, importer coverage, product breadth, median unit values, median effective tax wedges, and tax-wedge coverage. The table also indicates which offices are included in the regression analysis.

FIGURE 2: Ship Rerouting to Peripheral Ports in Haiti



Note: Arrow thickness represents the relative intensity of ship rerouting from Port-au-Prince to peripheral ports. Percentages indicate the average effective tax rate at each port. Lower effective tax rates at peripheral ports reflect reduced customs enforcement, correlating with higher rerouting activity.

Importers and access to multiple offices

Importers are identified by a stable firm identifier that allows activity to be tracked across products, offices, and time. Among the top-importer sample, 3,791 firms clear all shipments through a single office and account for 45.4% of total import value. By contrast, 1,209 firms clear shipments through two or more offices and account for 54.6% of total import value.

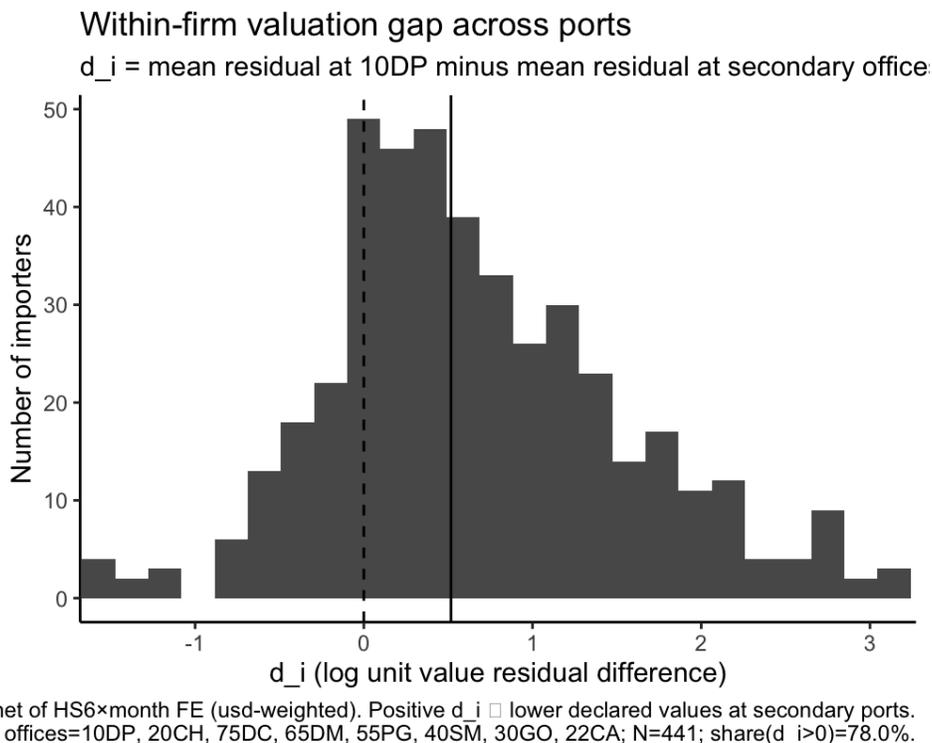
Summary statistics and descriptive patterns

Three descriptive facts motivate the empirical strategy. First, declared unit values and effective tax wedges display substantial dispersion even within narrowly defined HS6 product categories. Second, overlap across offices is extensive within the regression sample: 2,161 HS6 products are cleared through at least two offices, accounting for 96.9% of observations and 96.1% of import value. Third, a subset of economically dominant firms clears shipments through multiple offices, creating within-firm variation in clearance environments.

Figure 3 illustrates this variation. The figure plots the distribution of firm-level differences in residualized log unit values between Port-au-Prince and secondary offices, after absorbing

HS6×month fixed effects. For approximately 78% of multi-office firms, the same firm reports lower unit values when clearing through secondary offices than when clearing through Port-au-Prince.

FIGURE 3: Within-firm valuation gap across ports



Note: The figure shows the distribution of d_i , defined as the difference between a firm’s mean residualized log unit value at Port-au-Prince and at secondary offices. Residuals are obtained after absorbing HS6×month fixed effects and are weighted by shipment value within firm. Positive values indicate lower declared unit values at secondary offices.

These features imply that otherwise similar imports can face the same product–time tax environment but be cleared under different office environments. This structure motivates the within-firm, within-product empirical strategy developed in Section ??.

2.1 Baseline specification

Our unit of observation is a declaration line. Let ℓ index lines, and let $i(\ell)$ denote the importer, $j(\ell)$ the HS6 product, $k(\ell)$ the customs office, $o(\ell)$ the origin country, and $t(\ell)$ the month of clearance. We study how reported import unit values vary with tax pressure at the border.

For each line, we observe a declared unit value UV_ℓ in USD per kilogram and total border taxes paid. We construct an effective ad valorem tax wedge

$$\tau_\ell \equiv \frac{\text{total taxes paid}_\ell}{\text{declared customs value}_\ell},$$

TABLE 1: Descriptive statistics by customs office (IM4 sea imports, 2017–2024)

Customs office	In OLS	Share of value (%)	Shipments	Importers	HS6 products	Median unit value	Median tax wedge (%)	Tax coverage (%)
Port-au-Prince (10DP)	Yes	91.7	1,569,609	29,639	4,316	1.06	28.7	90.5
Cap-Haïtien (20CH)	Yes	6.7	527,815	19,220	2,439	0.29	28.7	98.2
Saint-Marc (40SM)	Yes	0.9	239,887	75,520	1,595	0.51	82.9	59.3
Miragoâne (65DM)	Yes	0.4	39,730	21,141	799	0.38	82.9	98.6
Gonaïves (30GO)	Yes	0.1	31,675	5,619	698	0.48	34.4	82.2
SONAPI (25BS)	Yes	0.1	609	23	283	7.16	23.0	9.9
Petit-Goâve (55PG)	Yes	0.1	3,098	1,418	191	0.51	40.2	98.2
Les Cayes (75DC)	Yes	0.1	66	13	11	0.13	28.7	100.0
Other offices	No	0.0	449	32	215	2.54	17.3	90.2

Notes: The table reports office-level descriptives for the IM4 sea-cargo analysis dataset (2017–2024), after constructing unit values (USD/kg) and trimming the 1st–99th percentiles within HS4×year cells. “Share of import value” is the office share of total declared import value in USD. “Median tax rate” is the median effective tax rate, defined as total taxes paid divided by declared customs value, capped at 150%. “Tax rate coverage” is the share of shipments with a positive, non-missing effective tax rate. “Other offices” aggregates the remaining offices outside the top eight by import value.

measured in units of one and capped at 1.5, as described in Section 2. We use the term wedge rather than statutory rate because τ_ℓ is an equilibrium outcome. It reflects the national tariff and fee schedule, exemptions, valuation outcomes, and implementation in practice.

Two features complicate interpretation. First, τ_ℓ uses declared value in the denominator, while $\log(\text{UV}_\ell)$ is itself a function of declared value. This mechanical linkage can induce co-movement even absent behavioral responses. Second, if enforcement differs across offices, firms and brokers may route shipments across clearance locations. Observed tax exposure may then be correlated with unobserved determinants of declared unit values.

Against this backdrop, we begin with a benchmark reduced-form specification in the spirit of the trade misinvoicing literature (Fisman and Wei, 2004; Javorcik and Narciso, 2008):

$$\log(\text{UV}_\ell) = \beta \tau_\ell + \alpha_{j(\ell)} + \gamma_{t(\ell)} + \delta_{k(\ell)} + \theta_{o(\ell)} + \mu_{i(\ell)} + \varepsilon_\ell. \quad (1)$$

HS6 fixed effects α_j absorb persistent differences in valuation levels and reporting difficulty across products. Month fixed effects γ_t capture common time shocks. Office fixed effects δ_k absorb level differences across clearance locations. Origin fixed effects θ_o control for systematic price differences by source country. Importer fixed effects μ_i absorb all time-invariant firm characteristics, including scale, product mix, and baseline reporting behavior. This fixed-effects structure follows the administrative-data approach in tax compliance research, which relies on within-entity comparisons under rich controls to separate behavior from composition (Kleven et al., 2011; Carrillo et al., 2017).

Since τ_ℓ enters in levels, β is a semi-elasticity. A 10 percentage-point increase in the wedge is associated with a $0.10 \times \beta$ change in log unit values.

We estimate (1) by value-weighted least squares, using declared customs value in USD as regression weights. Standard errors are clustered at the importer level (Wooldridge, 2010; Cameron and Trivedi, 2005). A negative $\hat{\beta}$ indicates that, conditional on the fixed-effects structure, higher effective wedges are associated with lower declared unit values.

The estimation sample is restricted to sea import declarations (procedure IM4) over 2017–2024, excluding office 10DA, and requiring positive declared value and a strictly positive wedge $\tau_\ell \in (0, 1.5]$. For computational tractability in the current implementation, we focus on the 5,000 importers with the largest number of declaration lines in the universe sample.

2.2 Office heterogeneity

To assess whether responsiveness differs across customs offices, we allow the slope on the effective wedge to vary by clearance location. Let k_0 denote the reference office (Port-au-Prince, 10DP). We estimate:

$$\log(\text{UV}_\ell) = \beta \tau_\ell + \sum_{k \neq k_0} \lambda_k (\tau_\ell \cdot \mathbb{1}\{k(\ell) = k\}) + \alpha_{j(\ell)} + \gamma_{t(\ell)} + \delta_{k(\ell)} + \theta_{o(\ell)} + \mu_{i(\ell)} + \varepsilon_\ell. \quad (2)$$

The semi-elasticity at the reference office is β . The implied semi-elasticity at office $k \neq k_0$ is $\beta + \lambda_k$. Office fixed effects δ_k remain included, so the λ_k terms capture differential sensitivity of declared unit values to tax pressure rather than mean valuation differences.

We test the null of homogeneous responsiveness across offices using Wald tests of the linear restrictions

$$H_0 : \lambda_k = 0 \quad \forall k \neq k_0.$$

Inference is based on standard errors clustered at the importer level (Wooldridge, 2010; Cameron and Trivedi, 2005).

A key institutional feature is that statutory tariff schedules are intended to be nationally uniform, so systematic slope differences across offices are unlikely to reflect office-specific statutory policy. Heterogeneity in the estimated semi-elasticities is consistent with differences in effective enforcement environments across clearance locations (Chalendard et al., 2023).

In the heterogeneity analysis we restrict attention to a fixed set of major offices. In the current implementation we retain the top offices by total declared customs value in the estimation sample and use 10DP as the reference office.

2.3 Design-based checks against compositional explanations

Office-specific slopes could reflect differences in the products and firms served by each location. We therefore implement three checks that tighten the comparison set.

First, we restrict to HS6 products observed in at least two of the selected offices and re-estimate (2). This common-support restriction ensures that office comparisons are made for identical HS6 codes.

Second, we split importers into single-office and multi-office firms, defined in the universe sample prior to tax conditioning, and re-estimate (1) and (2) separately by firm type. If office choice is an

evasion margin, heterogeneity should be more pronounced among firms with multi-office access.

Third, we restrict to importers that use at least two of the selected offices within the taxed estimation sample and replace separate HS6 and month fixed effects with HS6-by-month fixed effects. This tighter design compares shipments within narrowly defined product-time cells:

$$\log(\text{UV}_\ell) = \beta \tau_\ell + \sum_{k \neq k_0} \lambda_k \left(\tau_\ell \cdot \mathbb{1}\{k(\ell) = k\} \right) + \alpha_{j(\ell) \times t(\ell)} + \delta_{k(\ell)} + \theta_{o(\ell)} + \mu_{i(\ell)} + \varepsilon_\ell. \quad (3)$$

Estimator, weights, and clustering follow the baseline (Wooldridge, 2010; Cameron and Trivedi, 2005). In this specification, identification comes from within-importer variation in τ_ℓ across offices for shipments of the same HS6 product cleared in the same month.

2.4 Alternative identification: re-centered shift–share IV

The OLS designs compare declarations within importers and within product–time cells, but two concerns remain. The effective wedge τ_i is mechanically related to declared values through its denominator, and office choice may be endogenous. We therefore complement OLS with a re-centered shift–share IV that shifts identifying variation away from the importer’s own clearance choices and toward predetermined broker routing patterns interacted with product–time tax intensity, following Borusyak et al. (2022) and Borusyak and Hull (2023).

The outcome is $y_i = \log(\text{UV}_i)$. We estimate by 2SLS

$$y_i = \beta \tau_i + \sum_{k \neq k_0} \delta_k (\tau_i \cdot \mathbb{1}\{k(i) = k\}) + \gamma_1 \frac{1}{V_i} + \gamma_2 \frac{Q_i}{V_i} + \text{FE}_i + \varepsilon_i, \quad (4)$$

where V_i is declared customs value and Q_i is declared weight. The controls $1/V_i$ and Q_i/V_i absorb fixed-fee and specific-duty intensity implied by the accounting identity for the wedge. The fixed effects FE_i include HS6 and month effects, origin effects, broker effects, office×month effects, and either importer fixed effects or importer×month fixed effects. Estimation uses USD-value weights and two-way clustering by importer and broker.

In our baseline IV implementation we treat the overall wedge τ_i as an exogenous control and instrument only the office-specific interaction terms $\tau_i \cdot \mathbb{1}\{k(i) = k\}$. This targets the object of interest: differential tax responsiveness across enforcement environments within a uniform statutory schedule.

The instrument has a shift component and a share component. The shift is a product–time tax-intensity parameter. Within each HS6×month cell we estimate

$$T_i = \alpha_{h(i),m(i)} V_i + \beta_{h(i),m(i)} Q_i + u_i, \quad i \in (h, m),$$

where $T_i = \tau_i V_i$ is total taxes paid in USD. The coefficient $\alpha_{h,m}$ summarizes the marginal mapping

from declared value to taxes within a product–time environment, net of weight. We compute $\alpha_{h,m}$ by cross-fitting over importers: for each importer, the $\alpha_{h,m}$ applied to its observations is estimated using other importers only. This reduces mechanical feedback from using an importer’s own declarations to generate its instrument.

The share is broker-specific office exposure. Let $b(i)$ denote the broker. For each broker b and month m , define $z_{b,m,-i}^{(k)}$ as the leave-one-out share of that broker’s other importers’ declared value cleared through office k . This makes broker exposure predetermined with respect to importer i ’s own office choices in month m . We then re-center exposures within each HS6×month cell:

$$z_{rc,b,h,m}^{(k)} = z_{b,m,-i}^{(k)} - \text{wmean}_{i \in (h,m)}(z_{b(i),m,-i}^{(k)}; V_i),$$

so identification does not rely on exposure components common to all shipments within a product–time cell.

The office-specific instrument is

$$IV_i^{(k)} = \alpha_{h(i),m(i)} \times z_{rc,b(i),h(i),m(i)}^{(k)}.$$

Intuitively, the instrument increases office- k tax exposure for importers handled by brokers that systematically route other clients toward office k in a given month, and the strength of that induced exposure varies with product–time tax intensity. Re-centering ensures that the instrument captures differential exposure within product–time cells rather than level differences across cells.

In practice, re-centered exposures have negligible variance for very small offices, and the corresponding instruments are weak. We therefore screen instruments using first-stage diagnostics and retain only offices with adequate exposure support and first-stage relevance, using 10DP as the reference office. The IV estimation sample is the subset of the tax-conditioned analysis sample for which leave-one-out broker exposure is defined and the cross-fitted tax-intensity parameter is finite.

2.5 Interpretation and scope

The OLS specifications establish robust office heterogeneity under rich controls and within-firm comparisons. The re-centered shift–share IV complements this evidence by constructing variation in office-specific tax exposure from broker routing networks and product–time tax intensity, while removing mechanical own-observation contamination through leave-one-out shares and cross-fitting. The IV estimates are interpreted as causal evidence on differential office-specific tax responsiveness for the subset of offices with sufficient instrument support.

3 Results

3.1 Baseline results

Table 2 reports estimates of the baseline reduced-form relationship between declared unit values and tax pressure at the border. The dependent variable is $\log(\text{UV})$, where UV is the declared unit value (USD per kilogram). The key regressor is the effective border-tax wedge τ , defined as total taxes paid divided by declared customs value and capped at 1.5. All specifications are value-weighted (weights are declared customs value in USD) and standard errors are clustered at the importer level.

Column (1) estimates equation (1) with HS6 fixed effects, month fixed effects, customs office fixed effects, origin fixed effects, and importer fixed effects. The estimated semi-elasticity is large and negative: $\hat{\beta} = -0.703$ (s.e. 0.130). A 10 percentage point increase in the wedge is associated with an approximately 7 percent lower declared unit value. As discussed in Section ??, the wedge is mechanically related to declared values through its denominator, so we interpret this estimate as a disciplined reduced-form association rather than a clean policy elasticity.

3.2 Office-specific enforcement heterogeneity

Column (2) of Table 2 estimates equation (2) by allowing the semi-elasticity to vary across clearance locations. The coefficient on τ is the slope at the reference office, Port-au-Prince (10DP), and the interaction terms capture differential slopes relative to 10DP. A Wald test rejects the null of equal slopes across offices ($\chi^2(7) = 176.6$, $p < 0.001$).

The implied semi-elasticity at 10DP is -0.684 (s.e. 0.129). The implied slopes are substantially larger in magnitude at several other offices, including Cap-Haïtien (20CH) at -1.653 (s.e. 0.195) and Saint-Marc (40SM) at -2.642 (s.e. 0.252). Estimates for very small offices can be extreme. In the estimation sample, 10DP accounts for about 94 percent of declared value among the selected offices, while most other offices account for less than 0.1 percent, so precision varies sharply across locations. Since statutory tariff schedules are nationally uniform, cross-office differences in slopes are unlikely to reflect office-specific legal rates. They are consistent with heterogeneity in implementation, including inspection intensity, valuation scrutiny, and the application of exemptions.

3.3 Is office heterogeneity a product-composition artifact

A first concern is that office heterogeneity could reflect differences in product composition rather than differences in implementation. Columns (3) and (4) of Table 2 address this by restricting the sample to HS6 products that appear in at least two of the selected offices (common support). This restriction retains 1,526,690 declaration lines, or 96.9% of the estimation sample.

Two findings are immediate. First, the baseline semi-elasticity is essentially unchanged (column (3): -0.698 , s.e. 0.132). Second, heterogeneity remains strong (column (4): $\chi^2(7) = 174.1$,

$p < 0.001$). Product mix differences across offices therefore do not account for the office-specific pattern.

TABLE 2: Declared unit values and the effective tax wedge

	(1) Baseline	(2) Office slopes	(3) Baseline (CS)	(4) Office slopes (CS)
Effective tax wedge τ	-0.703*** (0.130)	-0.684*** (0.129)	-0.698*** (0.132)	-0.678*** (0.131)
$\tau \times \mathbb{1}\{20CH\}$		-0.969*** (0.187)		-0.974*** (0.187)
$\tau \times \mathbb{1}\{25BS\}$		-2.960 (2.117)		-3.026 (2.098)
$\tau \times \mathbb{1}\{30GO\}$		-1.373*** (0.314)		-1.376*** (0.316)
$\tau \times \mathbb{1}\{40SM\}$		-1.958*** (0.262)		-1.972*** (0.263)
$\tau \times \mathbb{1}\{55PG\}$		-3.644*** (0.572)		-3.654*** (0.573)
$\tau \times \mathbb{1}\{65DM\}$		-1.096 (0.811)		-1.106 (0.816)
$\tau \times \mathbb{1}\{75DC\}$		-16.170*** (2.156)		-16.440*** (2.194)
HS6 fixed effects	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes
Office fixed effects	Yes	Yes	Yes	Yes
Origin fixed effects	Yes	Yes	Yes	Yes
Importer fixed effects	Yes	Yes	Yes	Yes
Observations	1,575,976	1,575,976	1,526,690	1,526,690
R^2	0.778	0.778	0.773	0.774
Within R^2	0.007	0.008	0.007	0.008

Notes: Dependent variable is $\log(UV)$. τ is total border taxes paid divided by declared customs value, capped at 1.5. All regressions are weighted by declared customs value (USD). Standard errors clustered at the importer level are in parentheses. Columns (2) and (4) include $\tau \times$ office interactions with 10DP as the reference office, so the coefficient on τ is the slope at 10DP and interaction coefficients are slope differentials relative to 10DP. (CS) restricts to HS6 products observed in at least two selected offices. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

3.4 Is office heterogeneity driven by firm selection

A second concern is selection of firms into offices. We split importers into single-office and multi-office firms, defined in the pre-tax universe. In the estimation sample, single-office firms comprise 3,787 importers and account for 48.4% of declared import value, while multi-office firms comprise 1,209 importers and account for 51.6% of value.

Table 3 reports results separately for these groups. Columns (1) and (2) estimate the baseline specification within each group. The point estimates are similar: -0.690 (s.e. 0.136) for single-office firms and -0.659 (s.e. 0.234) for multi-office firms.

Columns (3) and (4) then allow office-specific slopes within each group. Equality of office slopes is rejected among both single-office firms ($\chi^2(4) = 25.7$, $p < 0.001$) and multi-office firms ($\chi^2(7) = 86.1$, $p < 0.001$). Because single-office firms do not switch offices, these comparisons

are necessarily across different firms. The within-firm design below is therefore the more stringent check.

3.5 Within-firm evidence among multi-office importers

Column (5) of Table 3 implements the tightest OLS design in equation (3). The sample is restricted to multi-office firms that use at least two of the selected offices. The specification includes importer fixed effects and HS6×month fixed effects, so identification comes from within-importer variation in τ across offices within the same HS6 product and month. Equality of office slopes is strongly rejected ($\chi^2(7) = 198.9$, $p < 0.001$).

The implied slope at 10DP is -0.552 (s.e. 0.102). Several secondary offices exhibit substantially larger semi-elasticities within the same firms, including 20CH and 40SM. This pattern is consistent with heterogeneous implementation environments across offices that are salient for firms with the ability to route shipments.

As a descriptive complement, we examine importer-office pairs for multi-office firms and compare the median line value (USD) in a firm’s main office and its secondary offices. The median of importer-office medians is 125 in the main office and 589 in secondary offices. The difference is not statistically precise in a clustered comparison ($t = 1.39$, $p = 0.166$).

3.6 HS6×month fixed effects robustness

Table 4 reports the baseline and office-heterogeneity estimates under HS6×month fixed effects, both in the full estimation sample and under the common-support restriction. The baseline semi-elasticity remains negative and precisely estimated (column (1): -0.492 , s.e. 0.060), and office heterogeneity remains strongly significant (column (2): $\chi^2(7) = 277.1$, $p < 0.001$). The common-support restriction leaves these conclusions unchanged (columns (3) and (4)).

Across specifications, declared unit values are strongly negatively associated with the effective border-tax wedge. The baseline semi-elasticity is between -0.49 and -0.70 depending on whether month and product effects are entered separately or as HS6×month fixed effects. This relationship varies systematically across customs offices. The null of homogeneous responsiveness across offices is rejected in the full sample, under common-support restrictions, within single-office and multi-office firm subsamples, and in within-firm designs that absorb importer fixed effects and HS6×month fixed effects. These patterns are consistent with meaningful cross-office differences in implementation and enforcement that shape the extent to which tax pressure is reflected in reported border values.

TABLE 3: Firm office usage, office heterogeneity, and within-firm evidence

	(1) Base, single	(2) Base, multi	(3) Slopes, single	(4) Slopes, multi	(5) Within-firm, multi
Effective tax wedge τ	-0.690*** (0.136)	-0.659** (0.234)	-0.679*** (0.135)	-0.632** (0.231)	-0.552*** (0.102)
$\tau \times \mathbb{1}\{20CH\}$			-0.999*** (0.266)	-0.886*** (0.231)	-0.979*** (0.199)
$\tau \times \mathbb{1}\{25BS\}$				-7.843** (2.941)	-7.043* (3.301)
$\tau \times \mathbb{1}\{30GO\}$			-1.271** (0.414)	-1.227** (0.433)	-0.350 (0.367)
$\tau \times \mathbb{1}\{40SM\}$			-1.489* (0.603)	-2.008*** (0.307)	-1.494*** (0.238)
$\tau \times \mathbb{1}\{55PG\}$				-3.280*** (0.563)	-3.301*** (0.711)
$\tau \times \mathbb{1}\{65DM\}$			-0.929 (1.061)	-2.087*** (0.464)	-1.953*** (0.482)
$\tau \times \mathbb{1}\{75DC\}$				-15.240*** (4.115)	-14.980*** (1.195)
HS6 fixed effects	Yes	Yes	Yes	Yes	No
Month fixed effects	Yes	Yes	Yes	Yes	No
HS6 \times month fixed effects	No	No	No	No	Yes
Office fixed effects	Yes	Yes	Yes	Yes	Yes
Origin fixed effects	Yes	Yes	Yes	Yes	Yes
Importer fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	1,109,738	466,238	1,109,738	466,238	460,877
R^2	0.742	0.804	0.742	0.804	0.899
Within R^2	0.008	0.006	0.008	0.007	0.007

Notes: Dependent variable is $\log(UV)$. All regressions are weighted by declared customs value (USD). Standard errors clustered at the importer level are in parentheses. Columns (3) to (5) include $\tau \times$ office interactions with 10DP as the reference office, so the coefficient on τ is the slope at 10DP and interaction coefficients are slope differentials relative to 10DP. Single-office and multi-office firm types are defined in the pre-tax universe. Column (5) restricts to multi-office firms that use at least two of the selected offices and includes HS6 \times month fixed effects. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

TABLE 4: Robustness: HS6×month fixed effects

	(1) Baseline	(2) Office slopes	(3) Baseline (CS)	(4) Office slopes (CS)
Effective tax wedge τ	-0.492*** (0.060)	-0.470*** (0.062)	-0.493*** (0.060)	-0.471*** (0.061)
$\tau \times \mathbb{1}\{20CH\}$		-1.110*** (0.147)		-1.109*** (0.145)
$\tau \times \mathbb{1}\{25BS\}$		-5.223 (3.479)		-5.184 (3.447)
$\tau \times \mathbb{1}\{30GO\}$		-1.126* (0.470)		-1.127* (0.466)
$\tau \times \mathbb{1}\{40SM\}$		-2.168*** (0.198)		-2.165*** (0.197)
$\tau \times \mathbb{1}\{55PG\}$		-3.991*** (0.591)		-3.995*** (0.586)
$\tau \times \mathbb{1}\{65DM\}$		-1.142** (0.409)		-1.141** (0.406)
$\tau \times \mathbb{1}\{75DC\}$		-10.990*** (1.168)		-10.990*** (1.162)
HS6×month fixed effects	Yes	Yes	Yes	Yes
Office fixed effects	Yes	Yes	Yes	Yes
Origin fixed effects	Yes	Yes	Yes	Yes
Importer fixed effects	Yes	Yes	Yes	Yes
Observations	1,575,976	1,575,976	1,526,690	1,526,690
R^2	0.883	0.883	0.879	0.879
Within R^2	0.004	0.006	0.004	0.006

Notes: Dependent variable is $\log(UV)$. All regressions are weighted by declared customs value (USD). Standard errors clustered at the importer level are in parentheses. Columns (2) and (4) include $\tau \times$ office interactions with 10DP as the reference office. (CS) restricts to HS6 products observed in at least two selected offices. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

4 Robustness Checks

4.1 Re-centered shift–share IV estimates

We complement the OLS evidence with an instrumental-variables design that shifts identification away from importers’ own clearance choices and toward quasi-experimental variation induced by broker routing networks interacting with product–time tax intensity. The design follows the re-centered exposure framework of [Borusyak et al. \(2022\)](#) and [Borusyak and Hull \(2023\)](#).

The construction has three ingredients. First, within each HS6×month cell we recover an ad-valorem tax-intensity parameter $\alpha_{h,m}$ by regressing total taxes paid on declared value and weight. We implement this step using cross-fitting by importer, so that an observation does not mechanically load on a tax-intensity estimate computed using the same importer. Second, for each broker and month we compute leave-one-out routing shares across offices using only the broker’s *other importers* in that month. Third, we re-center these broker routing shares within each HS6×month cell and interact them with $\alpha_{h,m}$ to form office-specific instruments. We then estimate 2SLS with USD-value weights and two-way clustering by importer and broker, conditioning on a rich fixed-effects structure that includes HS6, month, origin, broker, and office×month effects, and either importer fixed effects or importer×month effects.

A practical constraint is that the re-centered broker exposure has negligible variance for several small offices, so the corresponding instruments are uninformative. We therefore screen office instruments for relevance and retain only offices with nontrivial exposure support and strong first-stage relationships. [Table 5](#) reports the screening diagnostics for the top offices by import value in the estimation universe. The instruments are strong for Saint-Marc (40SM), Cap-Haïtien (20CH), and Gonaïves (30GO). By contrast, exposure variance is essentially zero for Bureau SONAPI (25BS) and Cayes (75DC), and is too limited for Miragoâne (65DM) and Petit-Goâve (55PG). The IV results below therefore compare Port-au-Prince (10DP) to the subset of secondary offices with usable exposure variation.

TABLE 5: First-stage screening for re-centered office instruments

Office	Kept	Own-IV t	sd(IV)	Share nonzero	Joint Wald F (approx.)	N
25BS (SONAPI)	No	<i>n.a.</i>	0.000	0.000	0.000	1,377,674
75DC (Cayes)	No	<i>n.a.</i>	0.000	0.000	0.174	1,377,674
40SM (Saint-Marc)	Yes	7.12	0.033	0.273	19.75	1,377,674
20CH (Cap-Haïtien)	Yes	4.93	0.130	0.694	7.90	1,377,674
30GO (Gonaïves)	Yes	3.96	0.025	0.206	5.93	1,377,674
65DM (Miragoâne)	No	1.65	0.016	0.077	1.42	1,377,674
55PG (Petit-Goâve)	No	0.73	0.002	0.004	0.39	1,377,674

Notes. “Own-IV t ” is the t -statistic on the office’s own re-centered instrument in a first-stage regression for the corresponding endogenous interaction, estimated with the baseline fixed effects (HS6, month, origin, importer, broker, office×month), USD-value weights, and two-way clustering by importer and broker. “Joint Wald F (approx.)” is the clustered Wald statistic on the full instrument set converted to $F = \chi^2/\text{df}$ for comparability.

Offices are retained if the instrument has nontrivial variation and $|t| \geq 3.5$.

Table 6 reports the second-stage estimates for the “heterogeneity-only” specification, which instruments the office-specific interactions $\tau_i \cdot \mathbb{1}\{k(i) = k\}$ for the retained offices while controlling flexibly for the baseline wedge τ_i . Both columns include mechanical controls implied by the accounting identity for the wedge, $1/V_i$ and Q_i/V_i , which absorb fixed-fee and specific-duty intensity. The estimated baseline semi-elasticity at the main port is negative and precisely estimated. Under importer fixed effects, a 10 percentage point increase in the wedge reduces declared unit values at Port-au-Prince by about 6.5 percent ($\hat{\beta} = -0.651$, s.e. 0.052). With importer \times month fixed effects, the corresponding effect is about 4.7 percent ($\hat{\beta} = -0.472$, s.e. 0.074). These magnitudes align closely with the OLS estimates under HS6 \times month fixed effects.

Crucially, the IV estimates preserve the core heterogeneity pattern. Secondary offices are substantially more responsive than the main port. The joint hypothesis of equal office slopes is rejected in both specifications. The Wald test for the instrumented office interactions yields $F \simeq 4.37$ ($p = 0.004$) under importer fixed effects and $F \simeq 7.19$ ($p < 0.001$) under importer \times month fixed effects. The implied office-specific semi-elasticities, reported in Table 7, show that the slope at Saint-Marc and Gonaïves is markedly steeper than at Port-au-Prince, and the slope at Cap-Haïtien is also materially more negative. In particular, under importer fixed effects the implied semi-elasticity is about -3.0 in Saint-Marc and -3.7 in Gonaïves, compared to -0.65 at Port-au-Prince.

TABLE 6: Re-centered shift–share IV estimates

	IV, importer FE	IV, importer \times month FE
Tax wedge τ_i (Port-au-Prince slope)	−0.651*** (0.052)	−0.472*** (0.074)
$\tau_i \times \mathbb{1}\{k = 40SM\}$	−2.371** (0.798)	−0.632 (0.763)
$\tau_i \times \mathbb{1}\{k = 20CH\}$	−0.633 (0.615)	−1.036*** (0.244)
$\tau_i \times \mathbb{1}\{k = 30GO\}$	−3.062 (1.586)	−1.659 (0.874)
$1/V_i$	7.312*** (1.666)	5.867*** (1.369)
Q_i/V_i	−0.074*** (0.017)	−0.060*** (0.014)
Observations	1,375,987	1,375,987
R^2	0.768	0.879
Within R^2	0.092	0.078

Notes. The endogenous regressors are the office-specific interactions $\tau_i \cdot \mathbb{1}\{k(i) = k\}$ for $k \in \{20CH, 40SM, 30GO\}$, instrumented with re-centered broker exposures interacted with cross-fitted HS6 \times month tax intensity. Port-au-Prince (10DP) is the reference office. Both columns include HS6, month, origin, broker, and office \times month fixed effects; the second column replaces importer fixed effects with importer \times month fixed effects. Regressions are USD-value weighted. Standard errors are two-way clustered by importer and broker. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

The IV results sharpen the interpretation of the OLS evidence. OLS documents steep tax responsiveness and substantial slope heterogeneity across offices under rich fixed effects and within-firm comparisons. The shift–share design adds an orthogonal source of identifying variation that

TABLE 7: Implied IV semi-elasticities by office

Office	Importer FE	Importer×month FE
10DP (Port-au-Prince)	−0.651 (0.052)	−0.472 (0.074)
20CH (Cap-Haïtien)	−1.285 (0.639)	−1.508 (0.232)
40SM (Saint-Marc)	−3.022 (0.762)	−1.104 (0.769)
30GO (Gonaïves)	−3.714 (1.557)	−2.130 (0.869)

Notes. Implied semi-elasticities equal $\hat{\beta}$ for 10DP and $\hat{\beta} + \hat{\delta}_k$ for office k . Standard errors are computed using the estimated variance-covariance matrix with two-way clustering by importer and broker.

is tied to broker routing networks and constructed to be leave-one-out and re-centered within product–time cells. Under this alternative variation, the estimated response of declared unit values to tax pressure remains negative and economically large, and the office heterogeneity persists. The agreement across OLS and IV is difficult to reconcile with a purely mechanical relationship between τ_i and declared values, or with a story in which heterogeneity is driven only by importer sorting across offices.

4.2 Other Robustness Checks & Mechanisms

We complement the main OLS and IV evidence with additional checks designed to rule out three practical concerns: sensitivity to extreme unit values, confounding from origin-specific time shocks, and the mechanical structure of unit values as a ratio. Throughout, we keep the estimator and inference aligned with the main analysis: value-weighted least squares (weights equal declared USD value) and standard errors clustered at the importer level (Cameron and Trivedi, 2005; Wooldridge, 2010).

We first re-estimate the office-slope specification on a trimmed sample that removes extreme unit values within HS4-by-year cells (dropping observations outside the 1st and 99th percentiles of $\log(UV_\ell)$). Under the tight HS6-by-month fixed-effects structure, the estimated semi-elasticity at the reference office remains negative and precisely estimated, and the hypothesis of homogeneous responsiveness across offices is rejected at conventional levels by wide margins. We then strengthen the design by adding origin-by-month fixed effects, which absorb any origin-specific price movements or composition changes that vary over time and could otherwise contaminate the relationship between the wedge and declared values. The estimates remain similar in sign and magnitude, and office heterogeneity remains strongly present (Table 8).

Declared unit values are ratios, $UV_\ell = V_\ell/Q_\ell$, and the effective tax wedge τ_ℓ also uses declared value in its denominator. To ensure that the baseline relationship is not driven mechanically by this accounting structure, we re-specify the outcome in levels.

Specifically, we replace $\log(UV_\ell)$ with $\log(V_\ell)$, where V_ℓ is declared customs value in USD, and include $\log(Q_\ell)$ as a control, where Q_ℓ is declared weight in kilograms. This specification removes ratios from the left-hand side while holding physical quantities fixed. The coefficient on τ_ℓ therefore

TABLE 8: Robustness to outliers and origin-specific time shocks

	(1) HS6×month FE	(2) + Origin×month FE
Tax wedge τ_ℓ (10DP slope)	-0.383*** (0.056)	-0.436*** (0.075)
Wald test: equal office slopes (p -value)	3.4×10^{-50}	2.1×10^{-37}
Observations	1,537,117	1,537,117
Weights	USD value	USD value
Fixed effects	HS6×month, importer, origin, office	HS6×month, importer, origin×month, office
SE clustering	importer	importer

Notes: Both columns use the office-slope specification with Port-au-Prince (10DP) as the reference office, so the reported coefficient on τ_ℓ is the semi-elasticity at 10DP. The sample trims $\log(UV_\ell)$ within HS4-by-year cells at the 1st and 99th percentiles. The Wald test reports the p -value for the null that all office-specific slope differentials are zero. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

captures how tax pressure is reflected in reported values for shipments of a given declared weight.

Formally, we estimate

$$\log(V_\ell) = \beta \tau_\ell + \sum_{k \neq k_0} \lambda_k (\tau_\ell \cdot \mathbb{1}\{k(\ell) = k\}) + \rho \log(Q_\ell) + FE_\ell + \varepsilon_\ell,$$

where FE_ℓ includes HS6×month fixed effects, importer fixed effects, origin fixed effects, and office fixed effects. Estimation uses USD-value weights and standard errors clustered at the importer level.

The estimated semi-elasticity remains negative and precisely estimated, indicating that higher tax pressure is associated with lower declared customs values even when conditioning on declared quantities. Moreover, the pattern of office-level heterogeneity closely mirrors the baseline results: several secondary offices exhibit substantially larger responses than the reference office. These findings confirm that the main results are not an artifact of unit-value arithmetic but reflect systematic adjustment in declared values in response to tax pressure.

5 Conclusion

This paper studies a margin of tax evasion that has received little systematic attention: strategic selection among enforcement locations within a single national customs system. Using comprehensive administrative data on more than two million import declarations in Haiti, we document large and systematic differences in how declared import values respond to ad-valorem tax pressure across customs offices operating under a uniform tariff schedule.

Baseline estimates show that declared unit values are highly tax-responsive, with elasticities consistent with substantial under-invoicing. Allowing this response to vary by clearance location reveals pronounced enforcement heterogeneity. Secondary ports exhibit tax elasticities that are two to nearly three times larger than those at the main port of Port-au-Prince. A sequence of increasingly demanding designs shows that these differences do not reflect product composition or

TABLE 9: Accounting identity robustness: dependent variable $\log(V_\ell)$ with $\log(Q_\ell)$ control

	(1) Baseline	(2) Office slopes
Tax wedge τ_ℓ	-0.517*** (0.052)	-0.501*** (0.050)
$\log(Q_\ell)$	0.703*** (0.047)	0.703*** (0.048)
$\tau_\ell \times \mathbb{1}\{k(\ell) = 20\text{CH}\}$		-0.938*** (0.196)
$\tau_\ell \times \mathbb{1}\{k(\ell) = 25\text{BS}\}$		-6.528*** (1.596)
$\tau_\ell \times \mathbb{1}\{k(\ell) = 30\text{GO}\}$		-0.239 (0.748)
$\tau_\ell \times \mathbb{1}\{k(\ell) = 40\text{SM}\}$		-1.432*** (0.273)
$\tau_\ell \times \mathbb{1}\{k(\ell) = 55\text{PG}\}$		-2.010*** (0.465)
$\tau_\ell \times \mathbb{1}\{k(\ell) = 65\text{DM}\}$		-0.426 (0.459)
$\tau_\ell \times \mathbb{1}\{k(\ell) = 75\text{DC}\}$		3.681 (3.041)
Observations	1,537,117	1,537,117
Weights	USD value	USD value
Fixed effects	HS6×month, importer, origin, office	HS6×month, importer, origin, office
SE clustering	importer	importer

Notes: Dependent variable is $\log(V_\ell)$, where V_ℓ is declared customs value. Column (2) allows the semi-elasticity on τ_ℓ to vary by office, with 10DP as the reference. The sample is the same trimmed sample as in Table 8. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

firm sorting. Office-specific responses persist when restricting attention to identical HS6 products traded across multiple offices and when comparing behavior within the same firm importing the same product in the same month.

To address remaining concerns about endogenous office choice and mechanical linkage between declared values and effective tax wedges, we complement the OLS analysis with a re-centered exposure instrumental-variables strategy. This design exploits predetermined broker routing networks interacted with product–time tax intensity. The IV estimates closely mirror the baseline patterns, both in sign and relative magnitude across offices, reinforcing the interpretation that the observed heterogeneity reflects differences in effective enforcement rather than residual selection.

The results point to a general mechanism through which decentralized enforcement undermines revenue collection even under uniform statutory rules. When enforcement intensity varies across locations, firms can arbitrage these differences by reallocating declarations, amplifying the fiscal consequences of limited state capacity. More broadly, the findings highlight that enforcement heterogeneity is not merely a nuisance for identification but a central object for understanding compliance behavior and revenue performance in low-capacity settings.

References

- Almunia, M. and Lopez-Rodriguez, D. (2018). Under the radar: The effects of monitoring firms on tax compliance. *American Economic Journal: Economic Policy*, 10(1):1–38.
- Baunsgaard, T. and Keen, M. (2010). Tax revenue and (or?) trade liberalization. *Journal of Public Economics*, 94(9-10):563–577.
- Besley, T. and Persson, T. (2011). *Pillars of Prosperity: The Political Economics of Development Clusters*. Princeton University Press.
- Besley, T. and Persson, T. (2013). Taxation and development. In Auerbach, A. J., Chetty, R., Feldstein, M., and Saez, E., editors, *Handbook of Public Economics*, volume 5, pages 51–110. Elsevier.
- Borusyak, K. and Hull, P. (2023). Nonrandom exposure to exogenous shocks. *Econometrica*, 91(6):2155–2185.
- Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift–share research designs. *Review of Economic Studies*, 89(1):181–213.
- Cagé, J. and Gadenne, L. (2018). Tax revenues and the fiscal cost of trade liberalization, 1792–2006. *Economic Journal*, 128(614):1037–1079.
- Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: Methods and Applications*. Cambridge University Press.
- Cantens, T., Jadotte, E., MacWilliam, D. C., and Suarez, J. (2024). Haiti: Strengthening customs administration in an insecure environment.
- Carrillo, P. E., Pomeranz, D., and Singhal, M. (2017). Dodging the taxman: Firm misreporting and limits to tax enforcement. *American Economic Journal: Applied Economics*, 9(2):144–164.
- Chalendard, C., Fernandes, A. M., Raballand, G., and Rijkers, B. (2023). Corruption in customs. *The Quarterly Journal of Economics*, 138(1):575–636.
- Emran, M. S. and Stiglitz, J. E. (2005). On selective indirect tax reform in developing countries. *Journal of Public Economics*, 89(4):599–623.
- Fisman, R. and Wei, S.-J. (2004). Tax rates and tax evasion: Evidence from “missing imports” in china. *Journal of Political Economy*, 112(2):471–500.
- Global Financial Integrity (2021). Trade-related illicit financial flows in 134 developing countries: 2009–2018.

- Inter-American Development Bank (2023). Haiti country strategy 2023–2026. Estimates Haiti loses \$200–300 million annually in customs revenue.
- International Monetary Fund (2021). Haiti: Technical assistance report on customs administration (selected issues). Technical Report IMF Country Report (details to be completed if you have the exact TA module), IMF.
- International Monetary Fund (2024). Haiti: 2024 article iv consultation. IMF Country Report 24/85, International Monetary Fund, Washington, DC.
- Javorcik, B. S. and Narciso, G. (2008). Differentiated products and evasion of import tariffs. *Journal of International Economics*, 76(2):208–222.
- Jensen, A. D. and Weigel, J. L. (2025). No Taxation without Administration: Bringing the State Back into the Public Finance of Developing Countries. *Journal of Economic Literature*.
- Keen, M. and Ligthart, J. E. (2002). Coordinating tariff reduction and domestic tax reform. *Journal of Economic Theory*, 102(2):385–407.
- Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S., and Saez, E. (2011). Unwilling or unable to cheat? evidence from a tax audit experiment in denmark. *Econometrica*, 79(3):651–692.
- Okunogbe, O. and Pouliquen, V. (2022). Technology, taxation, and corruption: Evidence from the introduction of electronic tax filing. *American Economic Journal: Economic Policy*, 14(1):341–372.
- Okunogbe, O. and Santoro, F. (2023). The promise and limitations of information technology for tax mobilization. *World Bank Research Observer*, 38(2):295–324.
- Transparency International (2024). Corruption perceptions index 2024. <https://www.transparency.org/en/countries/haiti>. Haiti score 16/100, rank 168/180.
- USAID (2020). Customs modernization program: Final evaluation – haiti.
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, 2 edition.
- World Bank (2022). Haiti: Country private sector diagnostic / investment climate assessment.
- World Bank (2024a). Haiti country overview. <https://www.worldbank.org/en/country/haiti/overview>. GDP per capita current US\$2,142.6 in 2024.
- World Bank (2024b). Worldwide governance indicators 2024: Government effectiveness. <https://info.worldbank.org/governance/wgi/>.

World Bank (2025). Haiti: Strengthening public financial management project — program document.

Yang, D. (2008). Can enforcement backfire? crime displacement in the context of customs reform in the philippines. *The Review of Economics and Statistics*, 90(1):1–14.

A Additional Background Information

B Additional Robustness Checks

Appendix References