ISSN No. 2454 - 1427

CDE
December 2025

Discretionary Enforcement and Strategic Compliance

Omer F. Baris
Email: omer.baris@nu.edu.kz
Nazarbayev University,
Graduate School of Public Policy

Shreekant Gupta
Email: shreekantgupta@csep.org
Centre for Development Economics,
Centre for Social and Economic Progress

Eduardo Araral Jr
Email: ed_araral@nus.edu.sg.
National University of Singapore,
Lee Kuan Yew School of Public Policy

Working Paper No. 355

Centre for Development Economics
Delhi School of Economics
Delhi- 110007

Discretionary Enforcement and Strategic Compliance

Omer F. Baris Shreekant Gupta Eduardo Araral Jr.

December 7, 2025

Abstract

We extend the canonical 2 × 2 inspection game to a 3 × 3 framework that allows strategic non-compliance and an inspection–quality choice, and test it in nine laboratory sessions (164 subjects) in India, Singapore, and Kazakhstan. Static evidence confirms the Harrington paradox: firms comply far more than risk-neutral theory predicts, yet Boards choose the cheap, low-detection audit in over 60% of rounds. Dynamic logit and multinomial specifications reveal four mechanisms. First, strong inertia: current compliance is the strongest predictor of future compliance. Second, inspection quality matters but only briefly. Thorough inspections raise compliance on impact and the effect dissipates within approximately three rounds. Third, cursory audits have no *immediate* effect; in longer-lag models their positive association with compliance is consistent with behavioral inertia rather than direct deterrence. Fourth, punishment works slowly: individual fines are insignificant on impact, while a history of fines has a modest cumulative influence, indicating graduated deterrence. Moral framing partially counteracts the compliance drop caused by strategic complexity, and country/demographic controls leave core results unchanged. Together, these findings reconcile the Harrington paradox: sticky firm behavior sustains *excess* compliance, whereas fading inspection salience and the prevalence of cheap cursory audits generate under-enforcement. Policy-wise, improving inspection *quality*, rather than increasing frequency or fine size, yields the highest marginal deterrence when regulators face capacity constraints and firms can game detection technologies.

Keywords: Enforcement · Inspection · Compliance · Corruption

JEL Codes: C91 · D64 · D73 · J24

Omer F. Baris: Nazarbayev University, Graduate School of Public Policy omer.baris@nu.edu.kz. Shreekant Gupta: Centre for Development Economics and Centre for Social and Economic Progress shreekantgupta@csep.org. Eduardo Araral Jr.: National University of Singapore, Lee Kuan Yew School of Public Policy ed_araral@nus.edu.sg. We thank colleagues for helpful comments and discussions. This work was partially supported by a research grant from National University of Singapore.

1. Introduction

How do firms' *strategic non-compliance* and regulators' *discretionary inspection choices* jointly shape compliance? We extend the standard 2×2 inspection–compliance game to a 3×3 framework that lets firms invest in concealment and regulators choose inspection quality (cursory vs. thorough), and we test it in the lab. We offer a behavioral reconciliation of the Harrington (1988) paradox: *sticky* compliance, short-lived salience of thorough inspections, and *placebo-like* cursory audits jointly generate excess compliance alongside under-inspection. Unlike studies that vary audit *probability* or taxpayers' *beliefs* about audits (e.g., Slemrod and Yitzhaki 2002; Kleven et al. 2011; Iyer, Reckers, and Sanders 2010; Anderson and Stafford 2003), we separate *frequency* from *informativeness*, showing that reallocating effort toward thorough inspections delivers the highest marginal deterrence even when fines are fixed. Policy-wise, *quality* dominates frequency when regulators face capacity constraints.

Traditional enforcement models treat compliance as a static trade-off between costs and expected penalties, typically via random auditing in the Allingham-Sandmo/Becker tradition (Allingham and Sandmo 1972; Becker 1968) or conditional/trigger audit rules (Landsberger and Meilijson 1982; Greenberg 1984). Despite standard mixed-strategy logic (Harsanyi 1973), enforcement data often show *over*-compliance with *under*-inspection (Harrington 1988), compounded by measurement/detection bias in official compliance statistics (Harrington and Morgenstern 2004; Nyborg and Telle 2006). In practice, firms divert resources toward evasion technologies (*non-inspectabilty* in the sense of Heyes 1994) while regulators, operating under capacity constraints, choose not only *whether* to inspect but also *how* (quality), creating margins the binary (2×2) cannot capture.

Because enforcement is a strategic interaction with asymmetric information, a game-theoretic lens (rather than one-sided decision theory) better captures how inspection choices and concealment co-evolve (Bier and Lin 2013). Our 3×3 design formalizes the margins of concealment and inspection quality: the firm chooses among compliance, strategic non-compliance (SNC), and pure non-compliance (PNC); the regulator chooses no inspection (NI), cursory/low-detection inspection (CI), or thorough/high-detection inspection (TI). Two empirical regularities motivate our contribution. First, compliance remains high relative to risk-neutral predictions (the Harrington paradox). Second, enforcement boards disproportionately select cheap, low-detection audits, even when high-quality inspections have larger immediate deterrent effects. SNC operationalizes non-inspectability (Heyes 1994) and parallels blame-distancing mechanisms (plausible deniability), where actors strategically reduce traceable evidence so wrongdoing is hard to prove (e.g., Tergiman and Villeval 2023). In our setting, SNC lowers detectability without bearing the full compliance cost, making low-quality inspections observationally consistent with "compliance on paper" even when violations persist.

Laboratory data have drawbacks, but they offer a distinct advantage where field data are noisy or partially observed: undetected violators are often (mis-)coded as compliant because status is inferred from audit outcomes, not observed directly. Misclassification arises from (i) firms' non-inspectability investments and (ii) regulators' endogenous inspection technology. Audit integrity and informed discretion matter for real abatement (Duflo et al. 2013, 2018). The lab lets us observe underlying choices without this noise: we benchmark the canonical 2×2 and confirm the paradox, then move to 3×3 to study how concealment and inspection quality jointly reshape incentives and outcomes.

Prior experiments related to Harrington's dynamic enforcement logic provide qualified validation (Cason and Gangadharan 2006) or test alternative audit schemes (Clark, Friesen, and Muller 2004), but they do not study the *interaction* of inspection quality and strategic evasion. Much empirical and experimental work varies audit *probability* or beliefs. Belief shocks raise self-reported income (Slemrod and Yitzhaki 2002; Kleven et al. 2011), whereas changing audit *rates* without stronger detection often has mixed effects (e.g., Lindeboom, van der Klaauw, and Vriend 2016), and in VAT systems impacts depend on third-party reporting and chain propagation (Pomeranz 2015). We instead separate *frequency* from *informativeness*: the Board chooses between thorough (TI) and cursory (CI) inspections, so only TI bites on SNC (via high detectability), whereas both TI and CI affect blatant PNC. Experimental work also shows that structured punitive mechanisms can raise compliance in regulatory analogues of public-goods environments (Anderson and Stafford 2003), consistent with our finding of graduated deterrence.

Our experiment introduces two explicit strategic dimensions–SNC and CI–and estimates their dynamics. Methodologically, we separate firm concealment (pure vs. strategic non-compliance) from inspection quality (cursory vs. thorough) in a dynamic lab setting. Using lag-block logit and multinomial models with average marginal effects and robust clustering, we find:

- **Inertia:** current compliance best predicts future compliance;
- **Inspection-quality salience:** thorough inspections raise compliance on impact, but effects dissipate within roughly three rounds;
- **Cursory audits as placebos:** CI shows no immediate deterrence; any longer-lag association consistent with inertia rather than deterrence;
- **Graduated deterrence:** single fines rarely move behavior contemporaneously, but fine histories have modest cumulative bite;

- **Framing and other controls:** moral framing partially offsets the compliance drop from strategic complexity; country and demographics leave core results unchanged.

Substantively, we reconcile Harrington (1988) and offer a policy rule-of-thumb: in capacity-constrained environments, reallocating effort toward *inspection quality* delivers the highest marginal deterrence.

We next present the canonical 2 × 2 model and a concise equilibrium summary of the 3 × 3 extension (Section 2), describe the experimental design (Section 3), and report results 4. Section 6 concludes. Full derivations and additional robustness (frame interactions, country splits, alternative lag structures) appear in the appendices.

2. Theoretical Framework: Inspection, Compliance, and the Harrington Paradox

2.1. The Canonical Model

Inspection-compliance models characterize enforcement as a strategic interaction between a firm (Firm) and a regulator (Board), typically following the crime-and-punishment tradition (Becker 1968). Agents weigh expected costs and benefits in a binary (2 × 2) game: Firm either complies (C) or not (NC), and Board either inspects (I) or does not inspect (NI) (Figure 1).

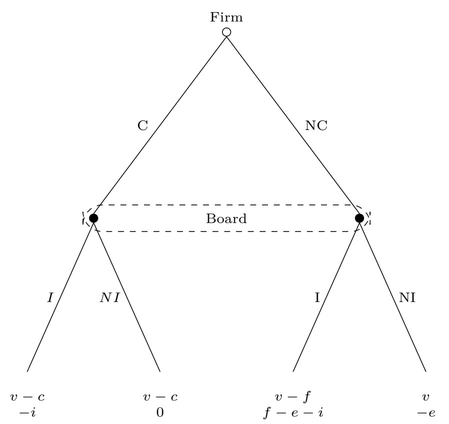

A. Extensive-form.

		Board	
		I	NI
Firm	C	$(v-c, -i)$	$(v-c, 0)$
	NC	$(v-f, f-e-i)$	$(v, -e)$

B. Normal form.

FIGURE 1. Canonical 2 × 2 inspection–compliance game tree (Panel A) and normal-form representation (Panel B). Both players choose their actions *without observing* the orther's. Our notation follows Germani et al. (2017).

The game is simultaneous: each player chooses an action without observing the other's choice. Payoffs depend on compliance cost $c > 0$, inspection cost $i > 0$, fine $f > 0$ upon detected noncompliance, and environmental harm $e > 0$ internalized by Board. The term v denotes the firm's gross value from operating; it appears in all outcomes, so it cancels out in all payoff comparisons and does not affect equilibrium behavior. We nevertheless retain v in the payoff notation to remain fully consistent with Germani et al. (2017).

The Mixed-Strategy Nash Equilibrium and the Harrington Paradox. When $f > c$ and $f > i$, no pure-strategy Nash equilibrium exists. In the unique mixed-strategy Nash equilibrium, Board inspects with probability c/f and Firm complies with probability $1 - i/f$. Compliance is increasing in the fine and decreasing in inspection cost, while inspection is decreasing in the fine and increasing in compliance cost (these comparative statics are standard; see, e.g., Harsanyi 1973; Heyes 2000; Franckx 2011; Cochard, Le Gallo, and Franckx 2014; Germani et al. 2017).

Empirically, however, this canonical model struggles with two robust stylized facts: firms often *over*-comply while boards *under*-inspectrelative to the mixed-strategy benchmark—the **Harrington paradox** (Harrington 1988). Proposed explanations include targeted enforcement (Harrington 1988), regulatory dealing (Heyes and Rickman 1999), voluntary disclosure and leniency (Livernois and McKenna 1999; Sterner and Coria 2012), reputational motives (Lyon and Maxwell 2008; Sterner and Coria 2012), 'data illusion' in measured compliance (Nyborg and Telle 2006) and asymmetric information (Raymond 1999).

Identifying audit effects in field data is notoriously hard because non-compliance is hidden and measurement further clouds inference. Undetected violations appear as compliance in administrative data (Harrington and Morgenstern 2004), and audit rates are endogenous to risk (see Section 6.3 Andreoni, Erard, and Feinstein 1998).

These theoretical and empirical limitations motivate us to move beyond the binary framework and explicitly model inspection *quality* alongside firm *concealment*, consistent with risk-informed regulatory games (Bier and Lin 2013, arguing for strategic analysis under private information).

2.2. The Extended Game

Because real-world regulatory environments are far more complex than the canonical 2×2 game, we extend it to a 3×3 framework. A key limitation of the binary model is that it (i) treats all inspections as equally effective and (ii) assumes detected violations closely track true non-compliance. In practice, inspection *quality* varies, and recorded compliance reflects both detection technology and enforcement choices.

Two implications follow. First, *detection bias*: official compliance rates can be biased upward when a nontrivial share of inspections are low-quality. Undetected violators are recorded as compliant, potentially inducing regulatory complacency and under-inspection. This mechanical misclassification overstates compliance and helps sustain the Harrington paradox of high measured compliance alongside low observed inspection. Second, *selective compliance*: noncompliant firms often choose *how* to violate, not just *whether* to violate

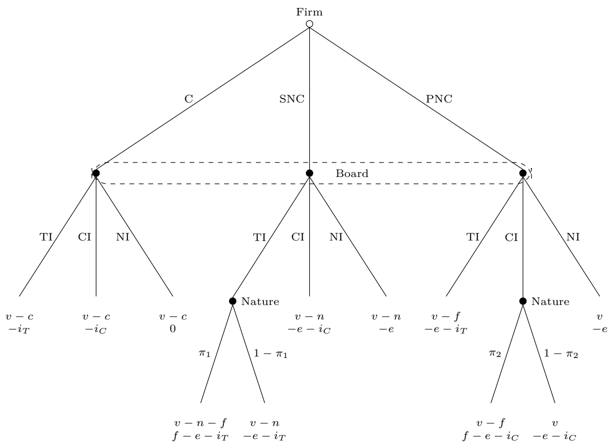

A. Extensive-form representation.

		Board		
		TI	CI	NI
	C	$(v - c, -i_T)$	$(v - c, -i_C)$	$(v - c, 0)$
Firm	SNC	$(v - n - \pi_1 f, \pi_1 f - e - i_T)$	$(v - n, -e - i_C)$	$(v - n, -e)$
	PNC	$(v - f, f - e - i_T)$	$(v - \pi_2 f, \pi_2 f - e - i_C)$	$(v, -e)$

B. Normal form representation.

FIGURE 2. **Extended inspection-compliance game.** Firm chooses whether to comply or, if not, whether to conceal (SNC) or violate blatantly (PNC). Board chooses whether to inspect and at what intensity (cursory CI vs. thorough TI). Payoffs and detection hinge jointly on Firm's non-compliance type and Board's inspection choice.

(e.g., by investing in *non-inspectability* to lower detection probability even when inspected Heyes 1994). Field evidence documents such margins: Gupta, Saksena, and Baris (2019) show polluting firms in Punjab, India adopt partial compliance while regulators vary intensity across no, cursory, and thorough checks given capacity constraints and signals. Because detected violations proxy for true behavior in the field, standard datasets suffer both under-detection and endogenous auditing (Andreoni, Erard, and Feinstein 1998, Section 6.3); laboratory experiments, by contrast, let us observe underlying choices directly while holding primitives fixed.

From a theoretical standpoint, the most natural way to capture heterogeneous inspections and strategic compliance is to model inspection effort and concealment on continuous scales, with detection probabilities determined by an underlying technology. Continuous-type models of disclosure and optimal enforcement exist (Bier and Lin 2013) and can encode finer gradations of effort and intensity (e.g., marginal deterrence, Stigler 1970). Such formulations are attractive for policy analysis, but they quickly become analytically intractable and unwieldy in the lab: they require subjects to reason about continuous probabilities, induce complicated best responses, and make equilibrium predictions less transparent.

We therefore adopt a parsimonious 3×3 abstraction that preserves the two policy-relevant margins—*concealment* (SNC vs. PNC) and *inspection quality* (CI vs. TI)—while yielding closed-form equilibria and transparent comparative statics that map directly into testable hypotheses. As with most lab designs, this abstraction trades some external realism for internal validity; we therefore treat field heterogeneity as a target for future work rather than something the current experiment can identify.

The resulting 3×3 game (Figure 2) preserves the key economic trade-offs of effort, concealment, and detection, while remaining experimentally implementable and analytically transparent:

- **Firm:** Chooses between Full Compliance (C), Strategic Non-Compliance (SNC; concealment reduces detectability), and Pure Non-Compliance (PNC; blatant violation, no concealment investment).
- **Board:** Chooses between No Inspection (NI), Cursory Inspection (CI, low-quality, low-cost), and Thorough Inspection (TI, high-quality, high-cost).
- **Detection technology:** Let $0 < \pi_1, \pi_2 < 1$ with

$$\pi_1 \equiv \mathbb{P}(\text{detect} \mid \text{SNC}, \text{TI}), \qquad \pi_2 \equiv \mathbb{P}(\text{detect} \mid \text{PNC}, \text{CI}).$$

Thorough inspections always uncover blatant violations and can detect strategic violations with probability π_1; cursory inspections detect blatant violations with probability π_2 but never detect SNC. Formally, the detection probabilities are

$$\mathbb{P}(\text{detect} \mid s_F, s_B) = \begin{array}{c|ccc} & \text{NI} & \text{CI} & \text{TI} \\ \hline \text{C} & 0 & 0 & 0 \\ \text{SNC} & 0 & 0 & \pi_1 \\ \text{PNC} & 0 & \pi_2 & 1 \end{array}$$

Notation and assumptions. As earlier, Firm pays a fine f when caught; Board bears inspection costs (i_C for cursory and i_T for thorough inspection with $i_T > i_C$), and internalizes environmental harm e if Firm does not comply. Firm bears compliance ($c > 0$) or strategic evasion ($n > 0$, non-inspectability investment) costs, with $0 < n < c$ (strategic evasion makes sense only if the concealment cost lies below full compliance cost, cf. Heyes 1994). Pure non-compliance (i.e. blatant violation) entails no such investment ($n = 0$). Finally, we assume $f > i_T$ so that thorough inspections can be optimal; otherwise TI would be weakly dominated by NI.

Discussion. Our 3 × 3 framework is not a cosmetic extension: it separates *frequency* from *quality*. This is novel relative to much of the compliance literature, which varies the *probability of inspection* in a 2 × 2 setting and studies how $\mathbb{P}(C)$ responds. In a 2 × 2, the inspection choice primarily controls *frequency*: the regulator either inspects or not, typically assuming perfect detection upon inspection (or a fixed exogenous detection rate), so there is no action-based quality margin. By contrast, in our 3 × 3 framework, Board chooses between *Thorough* (TI) and *Cursory* (CI): overall frequency is $q_{TI} + q_{CI}$ while *quality* is the TI–CI composition. Shifting mass between TI and CI generates different best responses for SNC versus PNC—a composition effect the binary game cannot capture. Audit integrity matters for both measured and true compliance (Duflo et al. 2013); discretion over intensity can deliver real abatement when targeting is informative (Duflo et al. 2018), while predictable monitoring can be gamed (Cochard, Le Gallo, and Franckx 2014). In line with field evidence, raising audit *rates* without a credible increase in informativeness may do little (Lindeboom, van der Klaauw, and Vriend 2016); our design isolates that 'bite' via TI versus CI.

We distinguish *choice probabilities* (q_{TI}, q_{CI}, q_{NI}) from *detection parameters* (π_1, π_2): the former are strategies; the latter are technological primitives. Reallocating mass from CI to TI, holding the overall inspection rate $(q_{TI} + q_{CI})$ fixed, increases detection risk for both SNC and PNC. However, because SNC is only exposed to TI while PNC is exposed to both TI and CI, the relative attractiveness of the two non-compliance modes depends sensitively on the TI–CI composition—a margin the binary game cannot capture. Thus, our framework highlights enforcement complexities beyond binary compliance—especially the role of *Strategic Non-Compliance* (concealment lowers detectability without incurring the full compliance cost) and *Cursory Inspections* (low-cost, low-salience checks). Together, these mechanisms can generate "symbolic enforcement": apparent compliance on paper alongside undetected violations. The canonical 2 × 2 nests as a special case (collapse PNC with SNC and CI with TI).

Heterogeneity further shapes outcomes. Larger or politically connected firms may face different cost maps—lower average compliance cost c (scale) yet greater ability to fund concealment n; inspection costs i_C and i_T can rise with facility complexity; and external pressure (complaints, prior violations) can shift enforcement toward TI. Where bureaucratic discretion or rent-seeking is salient, inspections may target firms with deep pockets without commensurate penalties—another path to symbolic enforcement. These are *testable hypotheses* rather than claims the experiment identifies.

Most 2 × 2 studies manipulate a single audit probability (e.g. Germani et al. 2017); sometimes even via belief updates rather than realized inspections (Slemrod and Yitzhaki 2002; Kleven et al. 2011). Field evidence shows that raising audit *rates* without credible, informative detection may do little (Lindeboom, van der Klaauw, and Vriend 2016). Our 3 × 3 separates the *rate* $q_{TI} + q_{CI}$ from the *composition* (q_{TI}, q_{CI}), which differentially disciplines SNC vs. PNC and cannot be mimicked in a binary design. We therefore implement Board as a strategic player rather than a fixed lottery, so that inspection frequencies $(q_{\mathrm{TI}}, q_{\mathrm{CI}}, q_{\mathrm{NI}})$ are equilibrium *outcomes*, not imposed parameters. In principle, this allows us to study how enforcement responds to firm behavior, whether inspectors endogenously under-inspect when observed compliance is high, and how changes in inspection *quality* feed back into the Harrington paradox. In this paper, however, we focus on how these endogenous inspection choices shape Firm's compliance behavior and do not attempt a full analysis of inspectors' strategy selection; Board behavior is modeled explicitly to provide a coherent strategic environment and a basis for future work.

In sum, our framework accommodates both *selective compliance*—as firms adjust behavior to expected inspection *quality*, not just frequency—and *detection bias*—as low-quality inspections inflate recorded compliance. In the lab, we observe firms' true compliance and concealment choices and Boards' inspection quality directly, rather than inferring behavior from noisy violation records. This allows us to separate actual compliance from measurement and auditing distortions in a way that real-world audit data, by construction, cannot.

2.3. Equilibrium summary and comparative statics

We now characterize equilibrium play in the 3×3 game and show how it extends the canonical mixed-strategy benchmark of the 2×2 model. Throughout this subsection we maintain $f > i_T > i_C > 0$ and $0 < n < c$.

- **Pure-strategy Nash equilibria (regions).** There are two regions in (π_1, π_2)-space with pure Nash equilibria:
  - $(\mathrm{SNC}, \mathrm{TI})$ is a Nash equilibrium when

$$\frac{i_T}{f} \leq \pi_1 \leq \frac{c-n}{f};$$

  - $(\mathrm{PNC}, \mathrm{CI})$ is a Nash equilibrium when

$$1 - \frac{i_T - i_C}{f} \leq \pi_2 \leq \frac{n}{f}.$$

  Outside these regions, no pure-strategy equilibrium exists.
- **Mixed-strategy Nash equilibrium.** Let Firm's mixed strategy be $\boldsymbol{p} = (p_{\mathrm{SNC}}, p_{\mathrm{PNC}}, p_{\mathrm{C}})$ over $(\mathrm{SNC}, \ \mathrm{PNC}, \ \mathrm{C})$ and Board's be $\boldsymbol{q} = (q_{\mathrm{TI}}, q_{\mathrm{CI}}, q_{\mathrm{NI}})$ over $(\mathrm{TI}, \ \mathrm{CI}, \ \mathrm{NI})$. When neither pure region obtains, there is a unique mixed equilibrium given by

$$p^*_{\mathrm{SNC}} = \frac{i_T}{\pi_1 f} - \frac{i_C}{\pi_1 \pi_2 f}, \qquad p^*_{\mathrm{PNC}} = \frac{i_C}{\pi_2 f}, \qquad q^*_{\mathrm{TI}} = \frac{c-n}{\pi_1 f}, \qquad q^*_{\mathrm{CI}} = \frac{n - c(1-\pi_1)}{\pi_1 \pi_2 f},$$

with $p_{\mathrm{C}}^* = 1 - p_{\mathrm{SNC}}^* - p_{\mathrm{PNC}}^*$ and $q_{\mathrm{NI}}^* = 1 - q_{\mathrm{TI}}^* - q_{\mathrm{CI}}^*$. For interior mixing (all probabilities in $(0, 1)$), the following must hold:

$$c > n, \quad \frac{n}{c} > 1 - \pi_1, \text{ and } \frac{i_C}{i_T} < \pi_2, \qquad \text{with } \pi_1, \pi_2 \in (0, 1)\,, \tag{1}$$

ensuring all probabilities lie strictly between 0 and 1 (full derivation in App. A).

On the interior mixed region (under (1)) the Firm's compliance probability simplifies to

$$p_{\mathrm{C}}^* = 1 - \frac{i_T}{\pi_1 f} + \frac{(1 - \pi_1)\, i_C}{\pi_1 \pi_2 f}. \tag{2}$$

Pure compliance as non-equilibrium. Notice that neither pure-strategy equilibrium involves full compliance. In fact, C can never be part of a pure-strategy Nash equilibrium in the 3×3 game. If Firm were to choose C, Board's best reply is always NI, since inspections are strictly costly ($i_T, i_C > 0$) and never produce additional benefits when the Firm complies. Given NI, however, Firm strictly prefers PNC to both SNC and C. Thus (C, s_B) cannot be a Nash equilibrium for any Board strategy s_B. As in the canonical 2×2 model, full compliance arises only as part of a mixed-strategy equilibrium, never as a strict best response, underscoring why the observed prevalence of high compliance in the field is hard to rationalize within a static one-shot framework.

Let ϵ_x denote the elasticity of p_{C}^* with respect to x. Using (2) and the partial derivatives under conditions (1) the elasticities are:

$$\epsilon_{i_T} = -\frac{i_T \pi_2}{\pi_1 \pi_2 f - \pi_2 i_T + (1 - \pi_1)\, i_C} < 0, \tag{3}$$

$$\epsilon_{i_C} = \frac{(1 - \pi_1)\, i_C}{\pi_1 \pi_2 f - \pi_2 i_T + (1 - \pi_1)\, i_C} > 0, \tag{4}$$

$$\epsilon_f = \frac{\pi_2 i_T - (1 - \pi_1)\, i_C}{\pi_1 \pi_2 f - \pi_2 i_T + (1 - \pi_1)\, i_C} > 0, \tag{5}$$

$$\epsilon_{\pi_1} = \frac{\pi_2 i_T - i_C}{\pi_1 \pi_2 f - \pi_2 i_T + (1 - \pi_1)\, i_C} > 0, \tag{6}$$

$$\epsilon_{\pi_2} = -\frac{(1 - \pi_1) i_C}{\pi_1 \pi_2 f - \pi_2 i_T + (1 - \pi_1)\, i_C} < 0. \tag{7}$$

Thus, on the interior mixed region, equilibrium compliance increases in the fine f and detection parameter π_1, and decreases in the cost of thorough inspections i_T and in π_2; a higher cost of cursory inspections i_C raises compliance by making low-quality inspection relatively less attractive.

In comparison, the equilibrium compliance in the canonical 2×2 game is a function of fine and inspection cost: $1 - i/f$ with $f > i > 0$. The relevant elasticities are

$$\epsilon_f^{2\times 2} = -\epsilon_i^{2\times 2} = \frac{i}{f-i}.$$

Both frameworks predict that higher fines increase compliance and higher inspection costs reduce it, but the 3×3 model additionally separates the effects of thorough versus cursory inspections and of inspection *quality* via (π_1, π_2).

3. Experimental Design

Participants and setting. We ran in-person laboratory sessions at three universities in India, Singapore, Kazakhstan, respectively, recruiting university students from standard subject pools. In total, 164 participants took part in nine controlled sessions (India: 72, Singapore: 38, Kazakhstan: 54). In each session, randomly paired participants played both the 2×2 (Figure 1) and 3×3 (Figure 2) versions of the inspection-compliance game.

We chose three locations to diversify institutional backgrounds (environmental enforcement salience, bureaucratic capacity, and language), while keeping comparable lab protocols and subject pools. Our cross-country analysis is descriptive/robustness-oriented and our main estimates pool countries with country fixed effects (and interactions in robustness checks). As is standard in lab experiments, student subject pools trade external heterogeneity for internal control; we therefore report country-by-country robustness to guard against pool-specific idiosyncrasies.

Roles (Firm/Board) were randomly assigned at the start and held fixed. Pairs were randomly rematched every round to prevent reputation/collusion and to mimic environments where firms face different inspectors over time.

We implemented two labeling framesin both games: neutral ("left, center, right"; "up, middle, down") and explicit ("inspect/comply"). Framing effects are modest and does not alter core results; we therefore pool across frames but report frame-by-quality interactions for robustness checks.

Each subject completed four 35-round blocks (two 2×2, two 3×3; neutral and labeled), for 140 rounds total. Sessions lasted approximately 70 minutes (mean 68, SD 7.1). To mitigate early learning, we discarded the first five rounds of each block and paid 3 randomly selected rounds *per block*. Thus our analysis uses 120 rounds per subject, focusing on stable strategic behavior rather than trial-and-error.

We fixed block order (neutral 2×2, neutral 3×3, labeled 2×2, labeled 3×3) to preserve cognitive continuity when moving from binary to ternary actions. This creates potential order confounds (learning, fatigue, drift). We address these by checking robustness of our analysis by (i) dropping the last 10 rounds of each block; (ii) dropping the first 10 rounds, and (iii) including round and block×round trends. Core coefficients are unchanged.

3.1. Calibration and theoretical mixed-strategy equilibria

In all rounds each participant began with 60 points. We kept primitives fixed across rounds to hold incentives constant and let behavior adjust endogenously:

$$c = 40,\; f = 50,\; e = 50,\; i = i_T = 40,\; i_C = 10,\; n = 10,\; \pi_1 = .90,\; \pi_2 = .50. \tag{8}$$

Thus the primary sources of variation are game structure (2×2 vs. 3×3) and framing (neutral vs. explicit). Full payoff matrices are reported in Table 1.

TABLE 1. Payoff structures used in the lab.

a. Payoff structure in the 2×2 game.

		Board I	Board NI
Firm	C	$(20, 20)$	$(20, 60)$
	NC	$(10, 20)$	$(60, 10)$

b. Payoff structure in the 3×3 game.

		Board TI	Board CI	Board NI
Firm	C	$(20, 20)$	$(20, 50)$	$(20, 60)$
	SNC	$90\% : (0, 20)$ $10\% : (50, -30)$	$(50, 0)$	$(50, 10)$
	PNC	$(10, 20)$	$50\% : (60, 0)$ $50\% : (10, 50)$	$(60, 10)$

Notes: We calibrate detection probabilities ($\pi_1 = 0.90, \pi_2 = 0.50$) to reflect asymmetries between thorough inspections detecting strategic non-compliance and cursory inspections detecting blatant violations.

Detection is perfect in the 2×2 game ($\mathbb{P}(\text{detect} \mid \text{NC}, I) = 1$). In the 3×3 π_1 is the probability of detecting SNC under TI, and π_2 is the probability of detecting PNC under CI. Under our calibration $i_T/f = 0.8$, $(c - n)/f = 0.6$, $1 - (i_T - i_C)/f = 0.4$, $n/f = 0.2$, both pure-strategy equilibrium regions (Section 2.3) are empty, so both pure-strategy regions are empty and the game lies in the mixed-equilibrium case.

The 2×2 mixed equilibrium implies $\mathbb{P}(\text{Inspect}) = 0.80$; in 3×3 the mixed equilibrium implies $q^*_{\text{TI}} = 0.67$ and $q^*_{\text{CI}} = 0.27$ (total inspection $= 0.94$, hence $q^*_{\text{NI}} = 0.07$). Likewise, in 2×2 the MSE yields $\mathbb{P}(\text{Not Comply}) = 0.80$, while in 3×3 it yields $p^*_{\text{SNC}} = 0.44$ and

$p^*_{\mathrm{PNC}} = 0.40$ (total 0.84). Frequency and quality are not equivalent: expected sanctions differ by violation type,

$$\mathbb{E}(\text{fine} \mid \text{SNC}) = \pi_1\, q_{\mathrm{TI}}\, f, \qquad \mathbb{E}(\text{fine} \mid \text{PNC}) = (q_{\mathrm{TI}} \cdot 1 + q_{\mathrm{CI}}\, \pi_2)\, f.$$

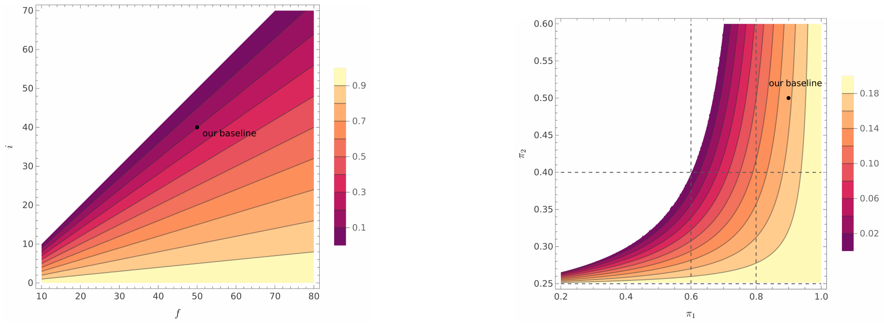

A. 2×2 game. Compliance is a function of the fine f and inspection cost i. Contours plot p^*_{C} over the feasible region $f > i > 0$. Iso-compliance lines are straight rays, showing that sensitivity collapses to the single ratio i/f (slope)—no action-based quality margin. The black dot marks our baseline $(f, i) = (50, 40)$, where $p^*_{\mathrm{C}} = 0.20$.

B. 3×3 game. Contours plot p^*_{C} over (π_1, π_2) with parameters $(c, n, f, i_T, i_C) = (40, 10, 50, 40, 10)$. Vertical and horizontal dashed lines at $\pi_1 = 0.6$, $\pi_1 = 0.8$, and $\pi_2 = 0.25$ and $\pi_2 = 0.40$ indicate pure Nash equilibrium boundaries/feasibility edges discussed in Sec. 2.3. The black dot marks our baseline $(\pi_1, \pi_2) = (0.90, 0.50)$, where $p^*_{\mathrm{C}} \approx 0.156$.

FIGURE 3. Compliance maps in the canonical 2×2 game (Panel A) and the extended 3×3 game (Panel B). Warmer colors indicate higher compliance.

In the 2×2, as compliance depends only on the ratio i/f, iso-compliance lines are linear in i/f (Figure 3A). At our baseline, $\epsilon_f^{2\times2} = 4$ and $\epsilon_i^{2\times2} = -4$ (Figure 4), and sensitivity collapses to this single ratio: there is no action-based inspection-quality margin. This reinforces why separating thorough vs. cursory inspections in our design is substantively new.

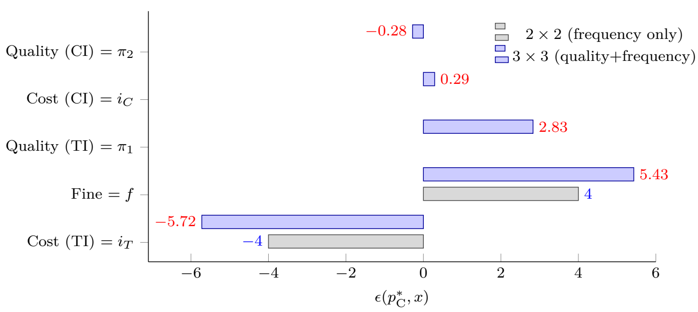

FIGURE 4. Elasticities ϵ_x of mixed-equilibrium compliance p^*_{C} at the baseline calibration $(\pi_1, \pi_2, i_T, i_C, f) = (0.90, 0.50, 40, 10, 50)$. In 3×3 games compliance is most sensitive to the cost of thorough inspections and fines; thorough-inspection quality has a sizable effect, while cursory tweaks are modest. Static sensitivities emphasize inspection quality over cursory improvements.

In 3×3 games, compliance is strongly increasing in thorough inspection quality (π_1) and weakly decreasing in cursory inspection quality (π_2). Thus, improving cursory inspections mainly shifts non-compliance to strategic non-compliance rather than into compliance (Figure 3B). The elasticities in (3)-(7) evaluated at the baseline calibration (Figure 4) are:

$$\epsilon_{i_T} \approx -5.72, \quad \epsilon_{i_C} \approx +0.29, \quad \epsilon_f \approx +5.43, \quad \epsilon_{\pi_1} \approx +2.83, \quad \epsilon_{\pi_2} \approx -0.28.$$

These indicate (i) large leverage from lowering the *cost of thorough inspections* and raising *fines*, (ii) improving the quality of thorough inspections (π_1) matters far more than improving quality of cursory inspections (π_2), which can even slightly reduce p_C^* by shifting pure non-compliance to strategic non-compliance and (iii) raising the cost of cursory inspections, combined with cheaper thorough checks, raises compliance.

We use these results to organize our empirical tests (cf. Germani et al. 2017); full statements, derivations, and proofs are relegated to Appendix A. In what follows, we compare observed play in the lab to these mixed-strategy benchmarks, focusing on overall compliance, and the composition of non-compliance between SNC and PNC.

4. Results

We examine empirical compliance patterns in three steps, moving from round-level analysis to heterogeneity across subjects, and finally to dynamic regression analyses. To quantify deviations from the MSE equilibrium benchmark, we define *over-compliance* (oc) as the difference between observed compliance and the equilibrium prediction in each environment.

4.1. Overall Compliance and Overcompliance

TABLE 2. Compliance frequencies by Firm players in 2×2 and 3×3 games: observed vs. theoretical MSE (%).

	2×2 **games**				3×3 **games**		
	Obs	**MSE**	Δ		**Obs**	**MSE**	Δ
Comply	39.6	20	+19.6	Comply	26.0	16	+10.0
Not Comply	60.4	80	−19.6	Strategic NC	32.6	44	−11.4
				Pure NC	41.4	40	+1.4

We observe four robust empirical patterns. First, compliance is remarkably stable over time in both treatments (s.d. 0.037 in 2×2 and 0.032 in 3×3 games): there is no downward drift toward equilibrium and no clear sign of strategic learning or fatigue, and the same patterns obtain if we drop either the first or the last ten rounds of each block (see Appendix B).

Second, compliance is consistently higher in the 2×2 games than in the 3×3 games at every round (30 out of 30) and and for the vast majority of participants (69 out of 82).

Third, both environments exhibit compliance well above the MSE equilibrium benchmarks, indicating that participants engage in substantial over-compliance relative to the noisy best-response prediction.

Finally, the gap between observed and predicted compliance is larger in the 2×2 game than in 3×3, foreshadowing the Harrington-style paradox: in the simpler binary environment, compliance lies far above the mixed-strategy benchmark, whereas introducing a strategic non-compliance option in the 3×3 game substantially attenuates (but does not fully eliminate) this surplus. We return to this point, and to the country-specific nuance in Kazakhstan, later in this section.

4.1.1. Compliance by Round

Table 2 reports overall compliance frequencies by game type. Figure 5 plots the fraction of compliant choices by round, pooling all sessions, separately for the 2×2 and 3×3 games.Country-specific trajectories and exact counts for each game is reported in the Appendix B; they display the same qualitative patterns.

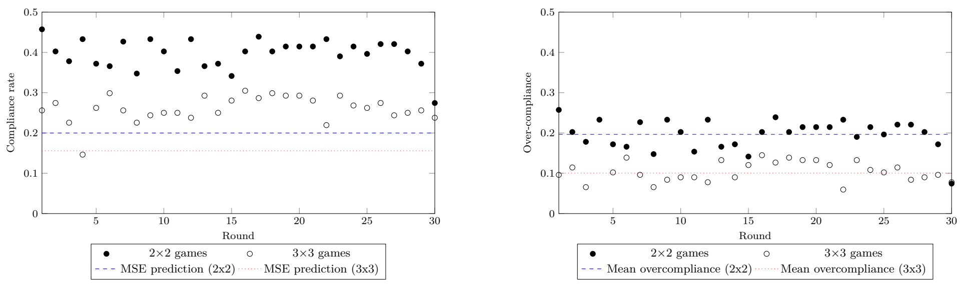

A. Average compliance per round. B. Average over-compliance by round.

FIGURE 5. Compliance (Panel A) and over-compliance (Panel B) in 2×2 and 3×3 games relative to MSE predictions. Over-compliance is defined as the difference between observed compliance and the MSE equilibrium prediction. Dots show the compliance and over-compliance rates in each of the 30 rounds.

Pooling all countries, mean compliance is 0.40 in 2×2 and 0.26 in 3×3 games. One-sample tests against the mixed-strategy benchmarks yield $t = 28.8$ and $t = 17.0$, respectively, and a paired round-level test shows a large and precise gap between games (mean difference 0.136, $t = 14.7$, Wilcoxon $p < 0.001$).

Country-by-country breakdowns confirm (i) significantly higher compliance in both games than the Nash prediction and (ii) substantially higher compliance in 2×2 than in 3×3 games (paired $t > 9.8$, Wilcoxon $p < 0.001$ in each case). Round-level country statistics are reported in Appendix Tables A2–A3.

Turning to over-compliance, round-level averages are roughly twice as large in the 2×2 game: $\overline{oc}_{2\times2} = 0.197$ (s.d. 0.037) versus $\overline{oc}_{3\times3} = 0.100$ (s.d. 0.032). The mean gap is 0.096 ($t = 10.4$, Wilcoxon $z = 4.76$, both $p < 0.001$).

4.1.2. Compliance by Participant

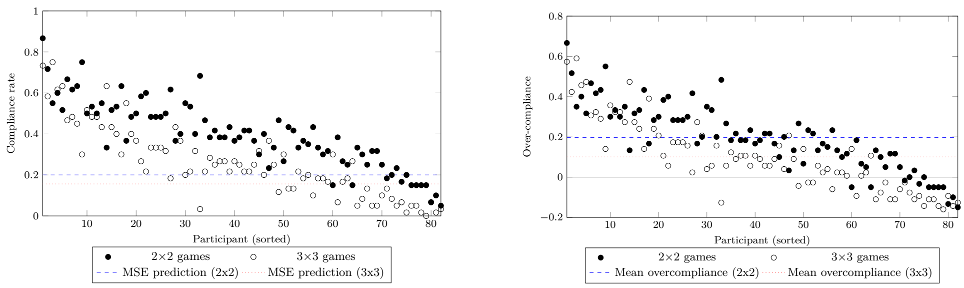

A. Average compliance by participants. B. Average over-compliance by participants.

FIGURE 6. Compliance (Panel A) and over-compliance (Panel B) in 2×2 and 3×3 games relative to MSE predictions. Dots show the compliance and over-compliance rates by each player.

Participant-level overcompliance reveals an even sharper distinction. Mean over-compliance ($N = 82$) is again significantly positive in both treatments: 0.197 in 2×2 and 0.100 in 3×3 games (standard errors 0.018 and 0.020, respectively). The within-subject difference remains large (0.096; standard error 0.017; paired $t = 5.73$; Wilcoxon $z = 5.30$, both $p < 0.001$). Subjects systematically over-comply relative to the MSE benchmark in both games, but the simpler 2×2 environment generates roughly twice as much over-compliance as the more complex 3×3 environment.

Figure 6 shows that for majority of participants (66 of 82), the 3×3 over-compliance value lies below the 2×2 value. High surplus types reduce their compliance substantially, and marginal types move toward zero or slightly negative values. Complexity thus reduces not only mean overcompliance but also its dispersion and upper tail, further reinforcing the conclusion that overcompliance is systematically smaller in more complex environments.

Confirming this, the distributional patterns in Figure 7 differ both in location and shape. The 2×2 environment produces a long right tail of individuals with very large overcompliance (up to 0.6–0.7 above equilibrium prediction), whereas the 3×3distribution is compressed downward and loses this tail entirely. Thus, the difference between the 2×2 and 3×3 environments is not driven by a few extreme individuals but by a systematic shift of the entire distribution.

Country breakdowns reveal one nuance. In Kazakhstan ($N = 27$), over-compliance in the 3×3 game is positive at the round level but not statistically different from zero once we average at the individual level, whereas over-compliance in the 2×2 game remains large

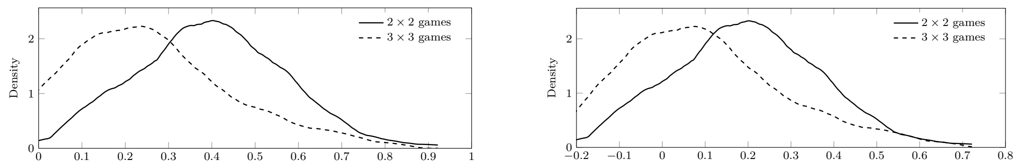

A. Compliance by participants.

B. Over-compliance by participants.

FIGURE 7. Distributions of compliance and over-compliance by participant in 2 × 2 and 3 × 3 games. Panel A shows the distribution of individual compliance rates; Panel (B) shows the distribution of individual over-compliance relative to the MSE benchmark. Subjects over-comply in both treatments, but the 2 × 2 games exhibit roughly twice as much over-compliance as the 3 × 3 games.

and highly significant. Thus, in this setting the Harrington-style paradox of persistently high compliance relative to the mixed equilibrium virtually disappears once strategic non-compliance is available.

TABLE 3. Classification of Firm behaviour relative to MSE benchmark by game type.

	1%-level				5%-level			
	Count		Percentage		Count		Percentage	
	2×2	3×3	2×2	3×3	2×2	3×3	2×2	3×3
Over-compliance	89	53	27.1	16.2	96	61	29.3	18.6
Statistically equal	234	275	71.3	83.8	224	251	68.3	76.5
Under-compliance	5	0	1.5	0.0	8	16	2.4	4.9
Total id–games	328	328	100	100	328	328	100	100

One observation is an id–game pair (30 rounds). Over-compliance (under-compliance) is defined as a sample proportion of compliance significantly above (below) the mixed-strategy equilibrium benchmark at the stated significance level (1% or 5%) using a normal-approximation z-test: $H_0 : \widehat{p} = p_0$ vs. $H_1 : \widehat{p} \neq p_0$.

As a complementary classification, we test for each id–game pair whether the observed compliance rate differs significantly from the MSE benchmark. Using a conservative 1% two-sided z-test, about 27% of Firm players in the 2×2 games can be classified as significant over-compliers, declining to 16% in the 3×3 games, while virtually no one under-complies (5 out of 328 id-games, all in 2×2, and none in 3×3; Table 3). The majority of players are statistically indistinguishable from the MSE benchmark, but among those who deviate, the mass lies overwhelmingly on the side of over-compliance, especially in the simpler 2×2 environment. This reinforces our earlier result that complexity sharply reduces the extent of surplus compliance without generating systematic under-compliance.

As a robustness check, we relax the significance threshold from 1% to 5% when classifying participant-game cells as over-, under-, or near-equilibrium. The shares move mechanically but the pattern is unchanged. At the 5% level, 29.3% of id-game pairs are classified as over-compliant in the 2×2 game, declining to 18.6% in 3×3. The share that is statistically indistinguishable from the MSE benchmark is 68.3% in 2×2 and 76.5% in 3×3, while under-compliance remains rare (2.4% vs. 4.9%). Overall, deviations from the MSE benchmark are dominated by over-compliance, and the incidence of such deviations is substantially lower in the more complex 3×3 environment.

Summary. Across all metrics, the data show a clear and systematic pattern: compliance is higher in 2×2 than in 3×3 games, and over-compliance declines sharply in 3×3 games. Statistically, these patterns are very strong. At both the round and participant level, over-compliance is significantly positive in each treatment, and the difference in over-compliance between 2×2 and 3×3 games is large and highly significant in paired comparisons ($p < 0.001$

in all cases). Thus, the reduction in over-compliance under complexity is not a visual artifact of a few influential observations; it appears as a robust, within-sample shift across rounds and across participants. In simple environments, subjects regularly generate substantial over-compliance—complying far more than equilibrium requires. In more complex environments, this surplus is sharply reduced, even though the underlying incentive structure is unchanged.

4.2. Dynamic Determinants of Compliance

We now turn to dynamic behavior: how do firms adjust compliance from one round to the next in response to inspections and detection, and does this depend on inspection *quality* as well as *frequency*?

To answer this we analyze round-by-round compliance choices of Firm players. We estimate a battery of dynamic models in which the dependent variable is the binary indicator $C_{it} \in \{0, 1\}$ for full compliance by Firm i in round t (training rounds removed). Our aim is to separate (i) immediate responsiveness to what happened in the previous round from (ii) medium-run persistence and feedback effects operating through inspection quality/frequency and sanction history.

We estimate linear probability models (LPM) with individual fixed effects and clustered standard errors, and pooled logistic models reported as average marginal effects (AME). All specifications include round fixed effects, a round×3×3trend and session fixed effects; because sessions are nested within countries in our design, session dummies absorb country differences in LPM. As a robustness check we later replace session FE with country FE and obtain essentially identical dynamic estimates.

The baseline dynamic equation is

$$\begin{aligned} C_{it} = \rho\, C_{i,t-1} \;+\; \beta_{TI}\, TI_{j,t-1} \;+\; \beta_{CI}\, CI_{j,t-1} \;+\; \phi\, \text{Caught}_{i,t-1} \\ +\; \alpha_i \;+\; \gamma_k \;+\; \lambda_s \;+\; \delta_t \;+\; \eta_{g\times t} \;+\; u_{it}, \end{aligned} \tag{9}$$

where $TI_{j,t-1}$ and $CI_{j,t-1}$ indicate the opponent Board's thorough and cursory inspections in $t-1$, and $\text{Caught}_{i,t-1}$ indicates whether Firm i was detected and fined in $t-1$. The fixed effects are individual (α_i), country (γ_k), session (λ_s), and round (δ_t); $\eta_{g\times t}$ is a full interaction between the 3×3 indicator and round fixed effects with $g \in \{2\times2, 3\times3\}$. Session effects are included in all specifications to absorb lab-day shocks.

Lag structure and discounted histories. For each specification (LPM and Logit) we first estimate the player's immediate reaction to what happened in the previous round $(t-1)$ without controlling for the history. But this specification limits our ability to interpret longer lags (1 through T) for each predictor, which would potentially allow us to trace behavioral reactions better over time, detecting short-run responsiveness, delayed reactions, and temporal decay. While multi-lag models provide granular insight into behavioral timing– revealing, for example, that thorough inspections (TI) exert immediate influence, while the effects of punishment (being caught) may emerge more slowly– they come with analytical trade-offs. Lag coefficients are often collinear and noisy, making it difficult to interpret individual effects and reducing overall model precision.

To capture medium and long-run feedback without an unwieldy set of lags, we add *exponentially discounted* histories from lags $t-2$ through $t-5$:

$$\text{past}_x^{(\delta)} = \sum_{k=2}^{5} \delta^{k-1} L_k x, \qquad x \in \{\text{TI}, \text{CI}, \text{Caught}, \text{C}\},$$

where $0 < \delta < 1$ governs decay. By construction these histories exclude L_1 (the model already contains the immediate lag). This weighting scheme gives greater importance to more recent experiences–e.g., round $t-2$ has weight 0.5, while round $t-5$ contributes only 0.0625– capturing a realistic behavioral assumption that memory and strategic salience diminish over time. In our main specification we set $\delta = 0.5$; Appendix Table **??** and Figure 8 show robustness to alternative values. Table A4 also report sensitivity to other lag lengths.

Augmenting (9) with *discounted histories* built from lags $t-2$ through $t-5$ allows us to retain dynamic structure while gaining both statistical efficiency and interpretability, especially for testing persistence, accumulation, and behavioral inertia. This parsimonious summary increases precision and captures empirically plausible bounded memory / reinforcement dynamics. It is commonly used in models of reinforcement learning and bounded memory (see Erev and Roth 1998; Camerer and Hua Ho 1999), and aligns with evidence from behavioral economics suggesting declining influence of past outcomes over time (e.g. Rabin 2002; Fudenberg and Levine 2006). Figure 8 and Appendix Table **??** show that results are stable across δ.

Empirical hypotheses. Guided by the mixed-strategy benchmarks and elasticities, we test:

H1. Quality vs. frequency. Holding frequency constant, reallocating inspections toward TI (quality) raises compliance more than expanding CI (frequency at low quality).

H2. SNC targeting. Thorough inspections deter strategic noncompliance on impact; cursory inspections mainly reshuffles blatant violations to strategic noncompliance (PNC→SNC) with little net compliance gain.

H3. Inertia. Compliance exhibits strong persistence (state dependence).

H4. Short-run fines. Per-round fines have weak or null immediate effects; cumulative fine history has a positive (gradual) effect.

H5. Framing. Moral labels modestly increase compliance but do not overturn H1–H4.

H6. Complexity. Compliance is lower in 3×3 than 2×2 (strategic complexity), even with constant primitives.

4.2.1. Main findings

Table 4 reports LPM estimates of (9) with α_i absorbed via `reghdfe` and cluster-robust SEs at the subject level. To aid probability interpretation, we also estimate pooled logits with the same covariates (session and country dummies included) and cluster-robust SEs. We omit subject fixed effects in logit to avoid incidental-parameter bias in short panels; we therefore report AMEs rather than coefficients. For the nonlinear models, the *average marginal effect* for covariate x is

$$\text{AME}_x = \frac{1}{N}\sum_{i,t}\frac{\partial\,\mathbb{P}(C_{it}=1\mid X_{it})x}{\partial - \mathit{NoValue}-},$$

evaluated at each observation and averaged over the sample. When discounted histories are included, $t-1$ is excluded from the history by design.

(i) Inspections: Quality dominates frequency. A thorough inspection in $t-1$ raises compliance in t by about 0.15–0.18 (Logit AME = 0.175^{***}; LPM = 0.167^{***}) probability points, while a cursory inspection has a small, imprecise effect. Across specifications we decisively reject $H_0 : \beta_{TI} = \beta_{CI}$ (bottom row), confirming that improving *quality* is substantially more effective than increasing *frequency at low quality*. Past thorough inspections also matter in the medium run (past_ti AME = 0.095^{***}), indicating that quality investments pay off beyond the current round.

(ii) Compliance inertia and accumulation. Compliance is state-dependent (one-lag AME = 0.140^{***}) and accumulates through experience. With discounted histories, the one-lag effect attenuates (to AME = 0.091^{***}) and the history of own compliance is large (past_C AME = 0.342^{***}), indicating genuine persistence beyond serial correlation.

(iii) Sanctions work through experience, not one-shot shocks. Being caught and fined in $t-1$ has no detectable immediate effect (AME ≈ 0). In contrast, a richer history of being caught raises current compliance (AME = 0.094), consistent with deterrence operating through accumulated experience rather than single shocks.

(iv) Game complexity works through these channels.. In reduced form (cols. 1-2), the 3×3 environment is associated with lower compliance—even controlling for inspection *frequency*, detection, and individual fixed effects—consistent with the aggregate evidence from the previous section that greater strategic complexity coincides with lower compliance and less over-compliance under equal incentives. When we additionally condition on lagged outcomes and discounted histories (cols. 3–4)—the main mechanisms through which complexity plausibly operates—the residual 3×3 coefficient attenuates and is statistically indistinguishable from zero at conventional levels. This pattern is *consistent* with 3×3 affecting compliance primarily

TABLE 4. Regression results.

	(1) LPM	(2) Logit	(3) LPM	(4) Logit
Past compliance				
($t-1$)	0.066**	0.140***	0.020	0.091***
	(0.027)	(0.026)	(0.030)	(0.025)
History ($L2-L5$)			0.189***	0.342***
			(0.045)	(0.032)
Past fines				
($t-1$)	-0.015	-0.005	-0.008	0.021
	(0.020)	(0.020)	(0.023)	(0.022)
History ($L2-L5$)			0.060**	0.094***
			(0.030)	(0.030)
Past inspections–Thorough				
($t-1$)	0.167***	0.175***	0.164***	0.153***
	(0.029)	(0.027)	(0.030)	(0.027)
History ($L2-L5$)			0.134***	0.095***
			(0.039)	(0.035)
Past inspections–Cursory				
($t-1$)	-0.029*	-0.019	-0.024	-0.009
	(0.016)	(0.022)	(0.016)	(0.020)
History ($L2-L5$)			-0.028	0.005
			(0.031)	(0.032)
Explicit labels/moral framing	0.058***	0.052***	0.062***	0.046***
	(0.017)	(0.016)	(0.018)	(0.012)
Game is 3×3	-0.156***	-0.137***	-0.098	-0.072
	(0.028)	(0.023)	(0.116)	(0.075)
Observations	6910	6910	4197	4197
R-squared / Pseudo-R2	0.176		0.244	
Test: TI=CI (p)	0.000	0.000	0.000	0.010

All models cluster SEs at subject (id). Absorb: session and country FE; round FE and game×round included. Discounted history aggregates L2–L5 (excludes L1); Logit columns report AMEs.

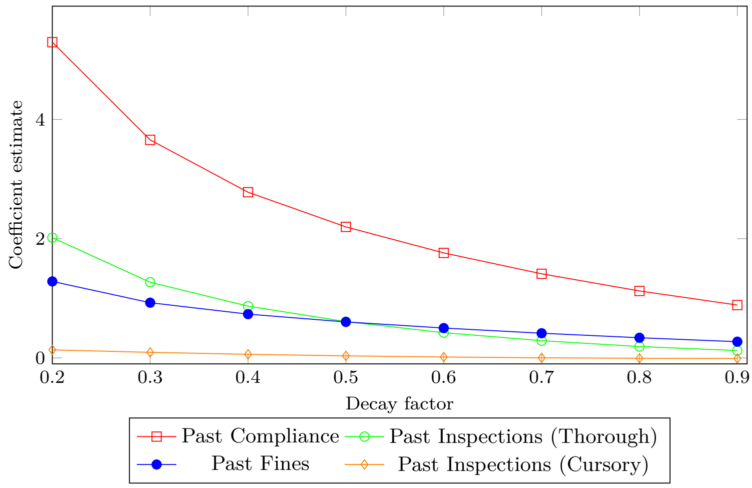

FIGURE 8. The Effect of decay factor on discounted history coefficients

via its impact on inspection composition and dynamic feedback, rather than via a separate, direct penalty. We emphasize that this is a mediation-style interpretation of reduced-form estimates, not a causal decomposition. Conditioning on L_1 and histories treats inspection choices and past outcomes as mediators; the attenuation therefore indicates accounted-for pathways rather than proof of full mediation. Results are similar replacing session FE with country FE.

Robustness.. Results are stable to (i) alternative decay factors $\delta \in \{0.3, 0.7\}$, (ii) trimming early/late rounds, and (iii) replacing discounted sums with simple L2–L5 counts (Appendix Tables **??**–**??**). In all cases, TI $\gg$ CI, and most of the 3×3 impact is accounted for by inspection composition and enforcement feedback.

For alternative decay factor specifications and their effects on coefficient estimates see Appendix Table **??** and Figure 8.

Cross-country stability. We replace session FE with country FE; estimates of the dynamic terms (lagged compliance, opponent TI/CI and the penalty) are virtually unchanged, while standard errors are somewhat tighter owing to fewer fixed-effect parameters (Table 6). A multilevel logit with sessions nested in countries yields comparable marginal effects and confirms that most cross-sectional variation is captured at the session level rather than at the country level.

Country-specific AMEs are tightly clustered: persistence and the impact of thorough inspections are positive and of similar magnitude in all three countries; cursory inspections and one-shot punishment remain small and imprecise. This cross-country stability strengthens external validity and suggests our dynamic mechanisms are not country-specific artifacts.

TABLE 5. Estimated elasticities of compliance

Regressor	ε
Past compliance	
($t-1$)	0.092^{***} (0.025)
History ($L2-L5$)	0.319^{***} (0.030)
Past fines	
($t-1$)	0.026 (0.028)
History ($L2-L5$)	0.113^{***} (0.036)
Past inspections–Thorough	
(–1)	0.242^{***} (0.042)
History ($L2-L5$)	0.141^{***} (0.052)
Past inspections–Cursory	
(–1)	−0.012 (0.026)
History ($L2-L5$)	0.007 (0.039)
Observations	4197
Model	Logit (AME); discounted histories L2–L5

Elasticity $\varepsilon_x = \text{AME}_x \times \bar{x}/\bar{C}$. Means $\bar{x}$ and $\bar{C}$ are computed on the Specification (4) estimation sample saved prior to `margins`. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Summary Moving from no thorough inspection to a thorough inspection increases the probability of full compliance by roughly 15-18 percentage points. One additional unit in the discounted compliance history raises current compliance by about 34 percentage points. These economically large effects rationalize policy that prioritizes inspection *quality* and on sustaining high-quality enforcement long enough for history to accumulate.

4.3. Composition of non-compliance.

We now look into non-compliance behavior through multinomial logit analysis which lets us ask a richer question than “did the firm comply?”: which of the three strategies it chooses in round t. Table **??** uses the most policy-relevant game, the 3×3 environment, and includes two blocks of dynamics: First we specify the model with one lag only, the most recent round, capturing short-run reactions. Then, the second specification captures the medium-run memory without the multicollinearity problems of separate lags with an exponentially-weighted average of five lags with $\delta = 0.5$.

We classify $Y_{it} \in \{\text{FC}, \text{SNC}, \text{PNC}\}$ and estimate a multinomial logit with FC as the base outcome:

$$\mathbb{P}(Y_{it} = c \mid X_{it}) = \frac{\exp\left(X_{it}'\beta_c\right)}{1 + \sum_{d \in \{\text{SNC},\text{PNC}\}} \exp\left(X_{it}'\beta_d\right)}, \quad c \in \{\text{SNC}, \text{PNC}\},$$

TABLE 6. Country-specific average marginal effects (AMEs) on $\mathbb{P}(C = 1)$

	Country		
	Kazakhstan	Singapore	India
Panel A: One-lag model (no history, N=6,910)			
Past compliance ($t-1$)	0.115***	0.157***	0.150***
	(0.023)	(0.030)	(0.028)
Past fines ($t-1$)	-0.004	-0.006	-0.006
	(0.017)	(0.023)	(0.022)
Past inspections–Thorough ($t-1$)	0.145***	0.197***	0.189***
	(0.025)	(0.032)	(0.029)
Past inspections–Cursory ($t-1$)	-0.016	-0.021	-0.020
	(0.018)	(0.025)	(0.024)
Panel B: One-lag + discounted history (L2–L5), N=4,197			
Past compliance			
($t-1$)	0.073***	0.100***	0.101***
	(0.020)	(0.027)	(0.027)
History ($L2-L5$)	0.273***	0.374***	0.378***
	(0.031)	(0.036)	(0.035)
Past fines			
($t-1$)	0.017	0.023	0.023
	(0.018)	(0.024)	(0.024)
History ($L2-L5$)	0.075***	0.103***	0.104***
	(0.024)	(0.033)	(0.033)
Past inspections–Thorough			
($t-1$)	0.122***	0.168***	0.170***
	(0.022)	(0.030)	(0.030)
History ($L2-L5$)	0.076***	0.104***	0.105***
	(0.028)	(0.039)	(0.038)
Past inspections–Cursory			
($t-1$)	-0.007	-0.010	-0.010
	(0.016)	(0.022)	(0.022)
History ($L2-L5$)	0.004	0.006	0.006
	(0.026)	(0.035)	(0.036)
Controls & FE: Round FE; 3×3×Round trend; Label; Session FE			
Estimator & SE: Logit; AMEs; SEs clustered by subject (id)			

Panel A uses the one-lag dynamic logit; Panel B augments with exponentially discounted histories over lags 2–5 (constructed from $L2$–$L5$). Entries are AMEs within each country's covariate distribution (Stata: `margins, dydx(...) over(country_id)`). Stars denote * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

where X_{it} includes past compliance (FC), fines, inspections (TI and CI) in one-lag and discounted histories ($\delta = 0.5$) with round FE, session FE, and round$\times 3 \times 3$. We report average marginal effects (AMEs) on the probabilities of SNC and PNC, which sum to the negative of the AME on FC by construction.

The multinomial logit results (Table 7) synthesizes the evidence from the binary models in Table 4, the country-specific AMEs in Table 6, the elasticity summary in Table 5, and the outcome-specific multinomial AMEs (Table 7). To organize the findings from most policy-relevant to most diagnostic, and to make explicit how the dynamic patterns corroborate the reduced-form patterns documented earlier we emphasize the following findings.

(1) Inspection quality *is the workhorse driver of compliance (confirms Section 4.1)..* Across specifications, a thorough inspection in $t-1$ substantially increases full compliance in t (binary AME ≈ 0.15–0.18, Table 4, cols. 1–4). The multinomial decomposition shows *how* this happens: TI_{t-1} shifts probability mass out of SNC and into FC (e.g. $+0.156$ to FC and -0.129 to SNC in the discounted-history model; Table 7). By contrast, cursory checks are small and imprecise for FC in the binary models, and in the multinomial they mainly *reshuffle* between non-compliant modes (reducing SNC, increasing PNC with little net gain in FC). This "reshuffle, not resolve" pattern for CI is precisely the frequency-without-quality critique emphasized in the design. These dynamic results are fully consistent with the earlier static patterns: environments or sessions with higher *TI* intensity sustain higher compliance and fewer strategic deviations, even when overall inspection *frequency* is similar.

(2) Compliance is state-dependent and accumulates (builds on Section 4.2).. Lagged own compliance is predictive (one-lag AME ≈ 0.14), indicating inertia/state dependence in the short run. When we add discounted histories (constructed from $L2$–$L5$), the one-lag effect shrinks (to ≈ 0.09) and the stock of own past compliance becomes large (AME ≈ 0.34). The multinomial AMEs align: past compliance raises *FC* and compresses PNC. In other words, *good compliance begets more compliance*, providing a dynamic counterpart to the aggregate over-compliance patterns documented earlier.

(3) Sanctions operate through experience*, not one-shot shocks..* Being caught in $t-1$ has little immediate effect on t (binary AME ≈ 0), but a richer history of detection increases *FC* (binary AME ≈ 0.09) and lowers SNC in the multinomial. This supports a learning/deterrence channel that requires *sustained* detection rather than isolated hits—again consistent with the earlier observation that sporadic enforcement does not undo strategic noncompliance.

(4) Game complexity's direct *effect is mediated by these channels..* In reduced form, 3×3 is associated with lower compliance (cols. 1–2 of Table 4), echoing our aggregate evidence

TABLE 7. Average marginal effects by outcome (multinomial logit)

	FC (full)	SNC	PNC
Panel A: One-lag model (no history), N = 4,756			
Last = SNC	-0.194***	-0.096***	0.255***
	(0.040)	(0.029)	(0.042)
Last = PNC	-0.169***	-0.082***	0.226***
	(0.042)	(0.031)	(0.037)
Caught (t−1)	0.019	0.014	-0.041*
	(0.023)	(0.022)	(0.024)
Opp. TI (t−1)	0.153***	0.156***	-0.026
	(0.026)	(0.027)	(0.024)
Opp. CI (t−1)	-0.018	-0.011	0.073***
	(0.021)	(0.020)	(0.019)
Framed (explicit labels)	0.076***	0.052***	-0.148***
	(0.018)	(0.014)	(0.022)
Panel B: One-lag + discounted histories (L2–L5, δ = 0.5**),** N = 4,100			
Last = SNC	-0.096***	0.136***	-0.039
	(0.029)	(0.026)	(0.024)
Last = PNC	-0.082***	-0.047**	0.129***
	(0.031)	(0.023)	(0.028)
History compliance (L2–L5)	0.328***	-0.044	-0.284***
	(0.039)	(0.037)	(0.040)
History SNC (L2–L5)	0.024	0.324***	-0.348***
	(0.029)	(0.031)	(0.028)
History caught (L2–L5)	0.088***	-0.066*	-0.022
	(0.031)	(0.037)	(0.038)
History TI (L2–L5)	0.089**	-0.039	-0.050
	(0.035)	(0.034)	(0.034)
History CI (L2–L5)	0.001	-0.051*	0.050
	(0.032)	(0.030)	(0.031)
Caught (t−1)	0.014	0.023	-0.036
	(0.022)	(0.022)	(0.022)
Opp. TI (t−1)	0.156***	-0.129***	-0.027
	(0.027)	(0.024)	(0.023)
Opp. CI (t−1)	-0.011	-0.061***	0.072***
	(0.020)	(0.017)	(0.018)
Framed (explicit labels)	0.052***	0.050***	-0.102***
	(0.014)	(0.013)	(0.016)

Notes: Multinomial logit with outcomes FC, SNC, PNC; base outcome immaterial for AMEs. Entries are average marginal effects on outcome probabilities; SEs (clustered by subject id) in parentheses. Controls in all specs: round FE, 3×3×round trend, session FE, and moral framing. Discounted histories are constructed from lags $t-2$–$t-5$ with decay $\delta = 0.5$ and exclude $t-1$. For any regressor, the three AMEs sum to ≈ 0 by the probability simplex. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

on lower FC and less over-compliance in the complex environment. Once we control for lagged outcomes and discounted histories (cols. 3–4), the residual 3×3 effect attenuates to statistical insignificance. The natural interpretation is mediation: complexity affects the *composition* of inspections and the dynamic feedback path (fewer TI, different detection histories), and it is through these levers that 3×3 lowers compliance. We thus reconcile the strong reduced-form gap with a mechanism-consistent story.

(5) Effects are remarkably stable across countries (supports external validity). Country-specific AMEs (Table 6) show tightly clustered estimates for the central levers (lagged C, TI), while CI and one-shot fines remain small everywhere. The cross-country consistency suggests that our dynamic channels (TI and accumulated experience) are not artifacts of any single site.

(6) Elasticities confirm policy salience and ordering of levers.. Elasticities from the preferred specification (Table 5) preserve the same ranking: TI (impact and history) and own compliance history deliver the largest proportional effects on $\mathbb{P}(FC=1)$ at observed means, detection history is meaningful but smaller, and CI is near zero. These scale-free summaries complement AMEs and are useful for cost-effectiveness comparisons when agencies trade off quality vs. frequency.

Identification. The identifying variation comes from *within-subject* changes across rounds, after saturating round effects and a round$\times 3 \times 3$ trend, and absorbing session shocks (with robustness to replacing session FE by country FE). The discounting approach summarizes $L2$–$L5$ without multicollinearity from a full lag stack, improving precision while remaining theory-consistent (bounded memory / reinforcement). The results are robust to alternative decay factors, trimming, and to simple-count histories (Appendix Tables **??**–**??**). Taken together, the dynamic models *strengthen* the core story from earlier sections: **quality** inspections and the **accumulation of experience** are the central policy levers; frequency without quality mostly rearranges non-compliance; complexity harms compliance largely by shifting those very channels.

Limitations and scope.. We report pooled-logit AMEs (not FE-logit) to avoid incidental-parameter bias in short panels; LPM with individual FE corroborates the signs and ordering. The models are reduced-form (not structural) and focus on short- to medium-run dynamics ($L1$ plus discounted $L2$–$L5$). These choices are conservative and bias against finding large dynamic effects—yet the quality and history channels remain strong and stable.

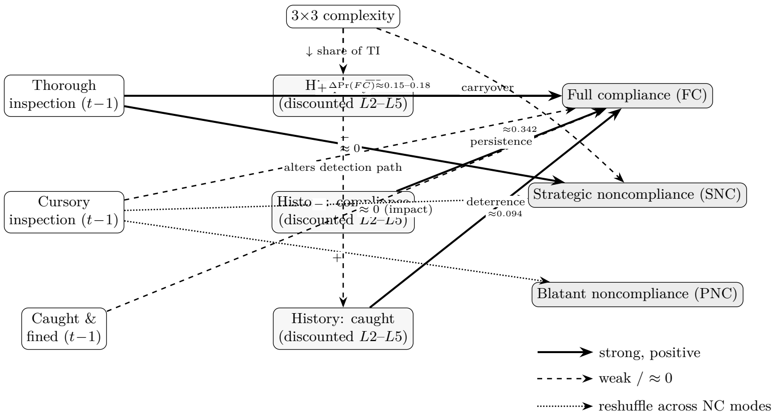

FIGURE 9. Dynamic channels: inspection *quality* boosts FC on impact and via history; cursory checks mainly reshuffle between SNC and PNC; detection deters through accumulated experience. The 3×3 environment lowers compliance largely by shifting these channels (mediation).

5. Discussion

The central implication of our results is straightforward: *tilt the inspection portfolio toward quality (TI)* and sustain it long enough for *compliance history* and *detection experience* to accumulate. The dynamic estimates show that doing so neutralizes much of the raw 3×3 disadvantage and raises full compliance (FC) rather than cosmetically reshuffling across non-compliant modes.

The dynamic analysis clarifies *how* the system moves from incentives to behavior:

1. **Quality dominates frequency.** A thorough inspection in $t-1$ increases $P(FC)$ by about 0.15–0.18 (Logit AME / LPM), whereas a cursory inspection (CI) has at best small and imprecise effects on FC. Discounted TI histories (L2–L5) further raise FC, indicating carryover beyond the impact round. Across specifications we decisively reject $H_0 : \beta_{TI} = \beta_{CI}$ (Table 4).
2. **State dependence and accumulation.** Compliance is persistent on impact (lag-1 AME $\approx$ 0.14), and *strongly* so through accumulated experience: the discounted five-round compliance history (excluding L_1) has a large marginal effect. This pattern—one-lag attenuation once histories enter, alongside large history coefficients—is hard to reconcile with pure serial correlation and is consistent with genuine state dependence.
3. **Sanctions deter through experience, not one-shot shocks.** Being caught and fined in $t-1$ is near zero on impact, but richer *fine histories* increase $P(FC)$. This explains why static elasticities can appear sizable for f while short-run effects of a single fine are weak: punishment accumulates rather than bites immediately.
4. **Complexity acts through these channels.** In reduced form, 3×3 is associated with lower FC and less over-compliance. Once we condition on lagged outcomes and discounted

histories—the very channels complexity plausibly shifts—the residual direct 3×3 effect attenuates and is statistically indistinguishable from zero. We interpret this as *mediation*: complexity lowers compliance *via* its influence on the composition of inspections and reinforcement dynamics, rather than through an autonomous penalty.[1]

5.1. Revisiting the Harrington paradox

By construction, the MSE benchmark isolates what noisy best-response can explain. Deviations above it are our measure of overcompliance. Two facts now connect the static and dynamic results:

- In the simpler 2×2 environment, overcompliance is large and persistent; enforcement variables have limited additional bite. This mirrors the "Harrington paradox" pattern: high compliance despite low inspection intensity.
- In 3×3, overcompliance is roughly halved and the distribution compresses downward. Precisely where surplus erodes, dynamic levers (TI, accumulated fines, accumulated own C) become more salient. Thus the static gap in surplus is *explained* by dynamic channels: when cognition and beliefs are taxed, quality inspections and reinforcement histories matter more.

Our evidence qualifies, rather than contradicts, the paradox. High compliance without heavy enforcement appears feasible in *simple* strategic settings (our 2×2), where norms can dominate behavior. As complexity rises (our 3×3 or many real-world regimes), overcompliance shrinks, the system becomes reliant on *precision inspections* and *consistent signals*, and fines contribute mainly through accumulated experience. In short: the paradox travels well only when the strategic environment is simple enough to support intrinsic motivation.

5.2. Ruling out design-only explanations

Could mechanics of the task, not behavior, drive the differences?

(a) **Sign prediction.** The MSE benchmark is *lower* in 3×3 than 2×2; a design-artifact story would predict the opposite of what we observe (higher raw C and surplus in 3×3).

(b) **Replication across levels.** The 2×2 surplus premium replicates at both round and participant levels with the same sign and large magnitude.

(c) **Distributional structure.** 2×2 exhibits a long right tail of surplus; 3×3 is compressed. Mechanical explanations rarely generate such structured distributional shifts.

(d) **Within-subject identification.** Each subject plays both environments; differences are within-person under matched incentives, pointing away from purely mechanical drivers.

[1] Formally, we show attenuation to statistical zero under rich dynamic controls; we do not claim a full causal mediation decomposition.

We thus view design features as a *conduit* through which complexity raises cognitive demands; they cannot account for the direction, magnitude, or distribution of the effects on their own.

Using *real* opponents (not scripted automata) strengthens the interpretation of overcompliance: sustaining compliance above MSE under belief uncertainty is costly. At the same time, human opponents introduce heterogeneity in beliefs and cognitive load. We therefore interpret "overcompliance" as a composite of norm adherence and responses to strategic uncertainty. The dynamic results show how inspection *quality* and accumulated experience restore compliance when intrinsic motivation is attenuated by complexity.

5.3. Limitations and robustness priorities

We cannot separately identify cognitive and normative channels; our dynamic regressions are reduced-form and do not estimate a structural learning model. Belief updating is unobserved, and binary compliance simplifies multi-dimensional real-world choices. Robustness already implemented shows: (i) results are stable to alternative decay factors ($\delta \in \{0.3, 0.7\}$), (ii) trimming early/late rounds, and (iii) replacing discounted sums with L2–L5 counts; in all cases TI $\gg$ CI and 3×3's direct effect attenuates once dynamic channels are included. Priority extensions include placebo histories, subject-level random slopes for TI, and explicit learning specifications.

6. Conclusion

This paper studied compliance in inspection games under matched incentives but differing strategic complexity. Two robust facts emerge. First, relative to the mixed-strategy equilibrium (MSE) benchmark, subjects display substantial *moral surplus*—excess compliance not mandated by payoffs—in the simpler $2{\times}2$ environment; this surplus is roughly halved and its distribution compressed in the more complex $3{\times}3$ environment. Second, dynamic reduced-form evidence shows *how* compliance is sustained: (i) *quality* inspections (TI) have large, immediate effects and persistent medium-run carryover; (ii) compliance is strongly state-dependent and accumulates through experience; and (iii) punishment operates mainly through *histories* of detection rather than one-shot fines. Conditioning on these channels, the residual direct effect of complexity attenuates and becomes statistically indistinguishable from zero, consistent with complexity acting *via* inspection composition and reinforcement dynamics rather than through an autonomous penalty.

Policy implications.. The regulator's portfolio should tilt toward *thorough* inspections and be sustained long enough for compliance and detection histories to accumulate. This strategy

(a) raises full compliance rather than merely reshuffling non-compliance modes, and (b) offsets much of the raw disadvantage of complex environments without requiring drastic increases in inspection frequency or fine levels. In practice, this translates into (i) minimum TI share targets, (ii) monitoring discounted compliance/detection histories as operational KPIs, and (iii) designing schedules that preserve streaks of TI rather than sporadic bursts.

Conceptual contribution.. Our results qualify the Harrington paradox: high compliance with light enforcement is feasible primarily in *simple* strategic settings where moral surplus is abundant. As complexity rises, intrinsic motivation becomes fragile and enforcement must work through *precision* (quality) and *consistency* (history). The moral-surplus lens thus explains when and why compliance exceeds equilibrium predictions, and when regulators must rely more heavily on targeted, high-quality enforcement.

Limits and external validity.. We provide reduced-form dynamics rather than a structural learning model; belief updating and cognitive frictions are not separately identified, and laboratory compliance is binary relative to multi-dimensional real-world decisions. Nonetheless, the cross-country stability of average marginal effects and the within-subject design mitigate concerns that results hinge on locale or composition.

Future work.. Promising directions include: (i) structural models that nest bounded memory and belief formation; (ii) optimal dynamic inspection design with TI capacity constraints; (iii) field pilots that implement TI-tilts and discounted-history dashboards; and (iv) heterogeneity in moral surplus and TI responsiveness, to target quality where it is most productive. In sum, *quality over quantity*, reinforced over time, is both behaviorally effective and institutionally actionable. Building compliance capital through thorough inspections and accumulated detection experience offers a scalable path to durable regulatory performance in complex environments.

References

Allingham, Michael G. and Agnar Sandmo. 1972. "Income Tax Evasion: A Theoretical Analysis." *Journal of Public Economics* 1 (3–4): 323–338. 10.1016/0047-2727(72)90010-2.

Anderson, Lisa R. and Sarah L. Stafford. 2003. "Punishment in a Regulatory Setting: Experimental Evidence from the VCM." *Journal of Regulatory Economics* 24 (1): 91–110. 10.1023/a:1023952115422.

Andreoni, James, Brian Erard, and Jonathan Feinstein. 1998. "Tax Compliance." *Journal of Economic Literature* 36 (2): 818–860.

Becker, Gary S. 1968. "Crime and Punishment: An Economic Approach." *Journal of Political Economy* 76 (2): 169–217. 10.1086/259394.

Bier, Vicki M. and Shi-Woei Lin. 2013. "Should the Model for Risk-Informed Regulation Be Game Theory Rather than Decision Theory?" *Risk Analysis* 33 (2): 281–291. 10.1111/j.1539-6924.2012.01866.x.

Camerer, Colin and Teck Hua Ho. 1999. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica : journal of the Econometric Society* 67 (4): 827–874. 10.1111/1468-0262.00054.

Cason, Timothy N. and Lata Gangadharan. 2006. "An Experimental Study of Compliance and Leverage in Auditing and Regulatory Enforcement." *Economic Inquiry* 44 (2): 352–366. 10.1093/ei/cbj019.

Clark, Jeremy, Lana Friesen, and Andrew Muller. 2004. "The Good, the Bad, and the Regulator: An Experimental Test of Two Conditional Audit Schemes." *Economic Inquiry* 42 (1): 69–87. 10.1093/ei/cbh045.

Cochard, Francois, Julie Le Gallo, and Laurent Franckx. 2014. "Regulation of Pollution in the Laboratory: Random Inspections, Ambient Inspections, and Commitment Problems." *Bulletin of Economic Research* 67 (S1). 10.1111/boer.12035.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2013. "Truth-Telling by Third-Party Auditors and the Response of Polluting Firms: Experimental Evidence from India." *The Quarterly Journal of Economics* 128 (4): 1499–1545. 10.1093/qje/qjt024.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2018. "The Value of Regulatory Discretion: Estimates from Environmental Inspections in India." *Econometrica : journal of the Econometric Society* 86 (6): 2123–2160. 10.3982/ecta12876.

Erev, Ido and Alvin E. Roth. 1998. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88 (4): 848–881.

Franckx, Laurent. 2011. "Ambient Environmental Inspections in Repeated Enforcement Games."

Fudenberg, Drew and David K Levine. 2006. "A Dual-Self Model of Impulse Control." *American Economic Review* 96 (5): 1449–1476. 10.1257/aer.96.5.1449.

Germani, Anna Rita, Pasquale Scaramozzino, Andrea Morone, and Piergiuseppe Morone. 2017. "Discretionary Enforcement and Strategic Interactions between Enforcement Agencies and Firms: A Theoretical and Laboratory Investigation." *Journal of Regulatory Economics* 52 (3): 255–284. 10.1007/s11149-017-9341-y.

Greenberg, Joseph. 1984. "Avoiding Tax Avoidance: A (Repeated) Game-Theoretic Approach." *Journal of Economic Theory* 32 (1): 1–13. 10.1016/0022-0531(84)90071-1.

Gupta, Shreekant, Shalini Saksena, and Omer F. Baris. 2019. "Environmental Enforcement and Compliance in Developing Countries: Evidence from India." *World Development* 117: 313–327. 10.1016/j.worlddev.2019.02.001.

Harrington, Winston. 1988. "Enforcement Leverage When Penalties Are Restricted." *Journal of Public Economics* 37 (1): 29–53. 10.1016/0047-2727(88)90003-5.

Harrington, Winston and Richard D. Morgenstern. 2004. *Choosing Environmental Policy: Comparing Instruments and Outcomes in the United States and Europe*. Rff Press Series. Routledge. 10.

4324/9781936331468.

Harsanyi, John C. 1973. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *International Journal of Game Theory* 2 (1): 1–23. 10.1007/bf01737554.

Heyes, Anthony G. 1994. "Environmental Enforcement When Inspectability Is Endogenous: A Model with Overshooting Properties." *Environmental & Resource Economics* 4 (5): 479–494. 10.1007/bf00691924.

Heyes, Anthony G. 2000. "Implementing Environmental Regulation: Enforcement and Compliance." *Journal of Regulatory Economics* 17 (2): 107–129. 10.1023/a:1008157410380.

Heyes, Anthony G. and Neil Rickman. 1999. "Regulatory Dealing—Revisiting the Harrington Paradox." *Journal of Public Economics* 72 (3): 361–378. 10.1016/s0047-2727(98)00098-x.

Iyer, Govind S., Philip M. J. Reckers, and Debra L. Sanders. 2010. "Increasing Tax Compliance in Washington State: A Field Experiment." *National Tax Journal* 63 (1): 7–32. 10.17310/ntj.2010.1.01.

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Soren Pedersen, and Emmanuel Saez. 2011. "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark." *Econometrica : journal of the Econometric Society* 79 (3): 651–692. 10.3982/ecta9113.

Landsberger, Michael and Isaac Meilijson. 1982. "Incentive Generating State Dependent Penalty System." *Journal of Public Economics* 19 (3): 333–352. 10.1016/0047-2727(82)90060-3.

Lindeboom, Maarten, Bas van der Klaauw, and Sandra Vriend. 2016. "Audit Rates and Compliance: A Field Experiment in Care Provision." *Journal of Economic Behavior & Organization* 131: 160–173. 10.1016/j.jebo.2015.08.016.

Livernois, John and C.J. McKenna. 1999. "Truth or Consequences: Enforcing Pollution Standards with Self-Reporting." *Journal of Public Economics* 71 (3): 415–440. 10.1016/s0047-2727(98)00082-6.

Lyon, Thomas P. and John W. Maxwell. 2008. "Corporate Social Responsibility and the Environment: A Theoretical Perspective." *Review of Environmental Economics and Policy* 2 (2): 240–260. 10.1093/reep/ren004.

Nyborg, Karine and Kjetil Telle. 2006. "Firms' Compliance to Environmental Regulation: Is There Really a Paradox?" *Environmental and Resource Economics* 35 (1): 1–18. 10.1007/s10640-006-9001-7.

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

Rabin, Matthew. 2002. "A Perspective on Psychology and Economics." *European Economic Review* 46 (4–5): 657–685. 10.1016/s0014-2921(01)00207-0.

Raymond, Mark. 1999. "Enforcement Leverage When Penalties Are Restricted: A Reconsideration under Asymmetric Information." *Journal of Public Economics* 73 (2): 289–295. 10.1016/

s0047-2727(98)00106-6.

Slemrod, Joel and Shlomo Yitzhaki. 2002. "Tax Avoidance, Evasion, and Administration." In *Handbook of Public Economics*, pp. 1423–1470. Elsevier. 10.1016/s1573-4420(02)80026-x.

Sterner, Thomas and Jessica Coria. 2012. *Policy Instruments for Environmental and Natural Resource Management*. An RFF Press Book. Resources for the Future, 2 ed.

Stigler, George J. 1970. "The Optimum Enforcement of Laws." *Journal of Political Economy* 78 (3): 526–536. 10.1086/259646.

Tergiman, Chloe and Marie Claire Villeval. 2023. "The Way People Lie in Markets: Detectable vs. Deniable Lies." *Management Science* 69 (6): 3340–3357. 10.1287/mnsc.2022.4526.

Appendix A. Proofs

Proof of Proposition **??**. From the Firm's perspective:

- PNC is a best-response to NI, since $0 > -n > -c$.
- When $p_2 > \frac{n}{f}$, SNC is a best-response to *CI* since $-n > -p_2 f$ and $-n > -c$.
- When $p_1 > \frac{c-n}{f}$, C is a best-response to TI since $-c > -n - p_1 f$ and $-c > -f$ (given $c < f$).

  From the Board's perspective

- NI is a best-response to C since $0 > -i > -t$.
- When $p_1 > \frac{t}{f}$, TI is a best-response to SNC since $p_1 f - e - t > -e > -e - i$.
- When $p_2 > 1 - \frac{t-i}{f}$, *CI* is a best-response to PNC since $p_2 f - e - i > f - e - t > -e$.

□

Proof of Proposition **??**. First note that when the Firm plays C, the Board's best response to NI since $0 > -i > -t$. But when the Board plays NI the Firm's best response is PNC since $0 > -n > -c$. The Board's best response to PNC cannot be NI because TI is a better response ($f - e - t > -e$) as we assumed $f > t$. Moreover, the Firm's best response to TI cannot be PNC since $f > c$ and the Board's best response to SNC cannot be *CI* since $-e - i < -e$. Therefore, We have two candidates for a pure strategy Nash equilibrium: (SNC, TI) and (PNC, CI).

a. (SNC, TI) : SNC is the Firm's best response to TI when

$$-n - p_1 f \geq -c \Rightarrow p_1 \leq \frac{c-n}{f}$$

   Otherwise, the Firm's best response to TI is C.

   TI is the Board's best response to SNC when

$$p_1 f - e - t \geq -e \Rightarrow p_1 \geq \frac{t}{f}$$

Otherwise, the Board's best response to SNC is NI.

b. (PNC, CI) : PNC is the Firm's best response to CI when

$$-p_2 f \geq -n \Rightarrow p_2 \leq \frac{n}{f}$$

Otherwise, the Firm's best response to CI is SNC.

CI is the Board's best response to PNC when

$$p_2 f - e - i \geq f - e - t \Rightarrow p_2 \geq 1 - \frac{t-i}{f}$$

Otherwise, the Board's best response to PNC is TI.

Note that when $p_1 \geq \frac{t}{f}$ and $p_2 > \frac{n}{f}$ there will be no pure strategy Nash equilibria. □

Proof of Proposition ??. We check the expected payoffs for the Board ($E\pi_B$) and the Firm ($E\pi_F$), respectively.

For the Board:

$$\begin{aligned}
E\pi_B(NI) &= -(\alpha_1 + \alpha_2)e \\
E\pi_B(TI) &= \alpha_1(p_1 f - e) + \alpha_2(f - e) - t \\
E\pi_B(CI) &= -\alpha_1 e + \alpha_2(p_2 f - e) - i
\end{aligned}$$

Making $E\pi_B(TI) = E\pi_B(CI) = E\pi_B(NI)$ and simplifying we get:

$$\begin{aligned}
E\pi_B(NI) = E\pi_B(CI) &\Rightarrow -(\alpha_1 + \alpha_2)e = -\alpha_1 e + \alpha_2(p_2 f - e) - i \Rightarrow \alpha_2^* = \frac{i}{p_2 f} \\
E\pi_B(NI) = E\pi_B(TI) &\Rightarrow -(\alpha_1 + \alpha_2)e = \alpha_1(p_1 f - e) + \alpha_2(f - e) - t \\
&\Rightarrow \alpha_1 p_1 f + \alpha_2 f - t = \alpha_1 p_1 f + \frac{i}{p_2 f} f - t = 0 \Rightarrow \alpha_1^* = \frac{p_2 t - i}{p_1 p_2 f}
\end{aligned}$$

Similarly, for the Firm:

$$\begin{aligned}
E\pi_F(C) &= -c \\
E\pi_F(SNC) &= -\beta_1 p_1 f - n \\
E\pi_F(PNC) &= -\beta_1 f - \beta_2 p_2 f
\end{aligned}$$

Making $E\pi_F(C) = E\pi_F(SNC) = E\pi_F(PNC)$ and simplifying we get:

$$\begin{aligned}
E\pi_F(C) = E\pi_F(SNC) &\Rightarrow -c = -\beta_1 p_1 f - n \\
&\Rightarrow \beta_1^* = \frac{c-n}{p_1 f}
\end{aligned}$$

$$
\begin{aligned}
E\pi_F(C) = E\pi_F(PNC) \quad &\Rightarrow \quad -c = -\beta_1 f - \beta_2 p_2 f \\
&\Rightarrow \quad -c = -\frac{c-n}{p_1 f} f - \beta_2 p_2 f \\
&\Rightarrow \quad \beta_2^* = \frac{n - c(1+p_1)}{p_1 p_2 f}
\end{aligned}
$$

□

Appendix B. Supplementary Tables

TABLE A1. Outcomes and strategy frequencies.

Panel A. Neutral

Game 1 (2 × 2)

Firm	Board: Left	Board: Right	Total
Up	504 (0.204)	422 (0.171)	926 (0.376)
Down	845 (0.343)	689 (0.280)	1534 (0.623)
Total	1349 (0.548)	1111 (0.451)	2460 (1.000)

Game 2 (3 × 3)

Firm	Board: Left	Board: Center	Board: Right	Total
Up	250 (0.101)	112 (0.045)	168 (0.068)	530 (0.215)
Middle	299 (0.121)	175 (0.071)	199 (0.080)	673 (0.273)
Down	478 (0.194)	338 (0.137)	441 (0.179)	1257 (0.510)
Total	1027 (0.417)	625 (0.254)	808 (0.328)	2460 (1.000)

Panel B. Labeled

Game 3 (2 × 2)

Firm	Board: I	Board: NI	Total
C	560 (0.227)	465 (0.189)	1025 (0.416)
NC	781 (0.317)	654 (0.265)	1435 (0.583)
Total	1341 (0.545)	1119 (0.454)	2460 (1.000)

Game 4 (3 × 3)

Firm	Board: TI	Board: CI	Board: NI	Total
C	352 (0.143)	153 (0.062)	246 (0.100)	751 (0.305)
SNC	400 (0.162)	241 (0.097)	290 (0.117)	931 (0.378)
NC	334 (0.135)	184 (0.074)	260 (0.105)	778 (0.316)
Total	1086 (0.441)	578 (0.234)	796 (0.323)	2460 (1.000)

Absolute and relative frequencies of strategies played. Each cell shows raw counts with relative frequencies in parentheses. Each game includes 30 rounds played by 82 randomly matched pairs.

TABLE A2. Compliance and over-compliance by country and game (per round)

Country	Compliance $comp$			Over-compliance		
	2×2 (mean)	3×3 (mean)	$\Delta_c omp$ (2x2−3x3)	2×2 (mean)	3×3 (mean)	$\Delta_o c$ (2x2−3x3)
Kazakhstan	0.310	0.189	0.122***	0.110***	0.029***	0.082***
Singapore	0.492	0.318	0.175***	0.292***	0.158***	0.135***
India	0.411	0.284	0.127***	0.211***	0.124***	0.087***
All pooled	0.397	0.260	0.136***	0.197***	0.100***	0.096***

Means are computed per round (30 rounds per game block). Over-compliance is defined as $oc_{2\times2} = comp_{2\times2} - 0.20$ and $oc_{3\times3} = comp_{3\times3} - 0.156$, where 0.20 and 0.156 are the mixed-strategy equilibrium predictions. $\Delta_c omp$ and $\Delta_o c$ are within-country differences (2×2 minus 3×3). Significance stars refer to one-sample t-tests against the MSE prediction for $comp$ and against zero for oc, and to paired t-tests for differences. *** $p < 0.01$. Wilcoxon signed-rank tests yield the same significance pattern (all $p < 0.01$).

TABLE A3. Over-compliance by country and game (per participant averages)

Country	Mean over-compliance (per subject)			N subjects
	2×2	3×3	$\Delta_o c$ (2x2−3x3)	
Kazakhstan	0.110***	0.029	0.082**	54
Singapore	0.292***	0.158***	0.135***	38
India	0.211***	0.124***	0.087***	72
All pooled	0.197***	0.100***	0.096***	164

For each subject and game we compute the average compliance over the 30 paid rounds and subtract the corresponding MSE benchmark (0.20 in 2×2, 0.156 in 3×3). Entries report the mean of these subject-level deviations. Stars on the 2×2 and 3×3 columns are from one-sample t-tests of $oc = 0$; stars on $\Delta_o c$ are from paired t-tests on subject-level differences. *** $p < 0.01$, ** $p < 0.05$. Wilcoxon signed-rank tests give the same significance pattern; in particular, Kazakh 3×3 over-compliance is small and not significantly different from zero, but the within-subject gap $\Delta_o c$ remains significant.

TABLE A4. Dynamic determinants of full compliance (3 × 3 game only)

	$\sum$ Lags 1-2	$\sum$ Lags 3-5	$\sum$ Lags 6-10
Own behavior			
Full compliance (FC)	0.98***	2.36***	1.70***
Strategic non-compliance (SNC)	0.04	0.41*	-0.23
Board's decisions			
Thorough inspection (TI)	1.33***	-0.08	-0.10
Cursory inspection (CI)	-0.28	0.02	0.75***
Caught & fined	0.47**	0.92***	0.39
Framing			
Game 4 (moral framing)	0.16*		
Observations	3 280		
Pseudo-R^2	0.22		

Robust s.e. clustered at the participant level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Country, gender, age and round-trend controls are included but omitted from the table. "Lag 1-Lag 10" are the ten period-specific coefficients; the right-hand columns report Wald sums of the indicated blocks. See Table A6 for binary-game estimates and average marginal effects.

Appendix C. Experiment Instructions

C.1. Introduction

C.1.1. Screen 1

You have volunteered to participate in a decision making experiment. This experiment is a study of individual behavior. The instructions are simple. If you follow them carefully and make good decisions, you may earn a considerable amount of money, which will be paid to you privately in cash at the end of the experiment today.

There are 24 people in the room who are participating in this experiment. You must not communicate with any other participant in any way during the experiment. You will not be asked for your name, student number or any other personal information that might reveal your identity.

If you have read and understood these policies please click Next to read about experiment instructions.

C.1.2. Screen 2

In this experiment there will be four games. In each game there will be 20 rounds. Therefore, there will be 80 rounds in total. Each round will take about one minute.

You will make a decision in each round against a randomly picked opponent who is also participating in the experiment. At the end of each game, we will randomly pick one of the rounds to determine your earnings. All rounds are equally likely to be chosen for your

TABLE A5. Determinants of Compliance and Strategic Behavior (Full model)

	(1)	(2)	(3)	(4)	(5)
	Base	All	Games 1&3	Games 2&4	Games 2& 4
VARIABLES		Dependent variable: FC_t			SNC_t
Full compliance (FC) $t-1$	0.829***	0.537***	0.365***	0.569***	0.308***
	(0.071)	(0.075)	(0.081)	(0.189)	(0.111)
Lags $t-2$ to $t-5$		1.710***	1.220***	2.259***	-0.238
		(0.081)	(0.161)	(0.120)	(0.207)
Strategic non-compliance $t-1$				-0.086	1.030***
				(0.110)	(0.074)
Lags $t-2$ to $t-5$				0.006	2.033***
				(0.127)	(0.135)
Thoroughly inspected (TI) $t-1$	0.795***	0.852***	0.781***	0.969***	-0.839***
	(0.068)	(0.063)	(0.102)	(0.106)	(0.090)
Lags $t-2$ to $t-5$		0.518***	0.654***	0.646***	-0.282
		(0.138)	(0.135)	(0.222)	(0.237)
Cursorily Inspected (CI) $t-1$				-0.076	-0.368***
				(0.144)	(0.099)
Lags $t-2$ to $t-5$				0.218	-0.356***
				(0.180)	(0.098)
Caught (penalty) $t-1$	-0.028	-0.036	-0.261**	0.096	0.215**
	(0.091)	(0.079)	(0.113)	(0.163)	(0.090)
Lags $t-2$ to $t-5$		0.553***	0.285	0.611***	-0.291
		(0.167)	(0.285)	(0.228)	(0.225)
Game 2	-0.599***	-0.405***			
	(0.058)	(0.070)			
Game 3	0.166**	0.159*	0.143*	1.522***	
	(0.073)	(0.084)	(0.079)	(0.391)	
Game 4	-0.193***	-0.067		0.292***	0.294***
	(0.064)	(0.078)		(0.112)	(0.082)
Singapore	0.529***	0.314***	0.446***	0.200	0.075
	(0.121)	(0.116)	(0.092)	(0.139)	(0.085)
India	0.362***	0.204***	0.231***	0.231***	0.105
	(0.048)	(0.053)	(0.051)	(0.088)	(0.099)
Female	-0.002	0.002	0.012	-0.007	0.179*
	(0.058)	(0.056)	(0.067)	(0.060)	(0.100)
Age	0.009	0.012	0.004	0.028**	-0.004
	(0.013)	(0.011)	(0.014)	(0.012)	(0.016)
Round	0.000	-0.000	-0.007*	-0.005	-0.007
	(0.003)	(0.003)	(0.004)	(0.007)	(0.006)
Constant	-1.782***	-2.732***	-2.192***	-3.792***	-1.377***
	(0.283)	(0.273)	(0.395)	(0.246)	(0.347)
Observations	9,758	9,430	4,510	4,182	4,100

This table reports logit regressions of full compliance across all games (columns 1–2), and disaggregated by game structure (columns 3-4). The specification includes one-period lags for recent behavior and collapsed cumulative indicators for inspection, compliance, and being caught. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

TABLE A6. Lagged Enforcement and Compliance History: Expanded Multi-Lag Models

This appendix table presents models with disaggregated lag structures ($L1$–$L3$ or up to $L5$), estimated across game structures. These models allow precise identification of temporal dynamics, including decay in inspection effects and delayed punishment. While individual lag coefficients are occasionally noisy, they support the findings in Table X: inspections deter in the short run, punishment has delayed and limited effect, and behavioral inertia–especially for strategic non-compliance–is strong and persistent. These models motivate the use of collapsed lag variables in the main text for interpretability and robustness.

earnings, so you should think carefully about each decision in each round. Your total earnings will be the sum of your points in randomly picked rounds in four games. Your earnings will depend on the decisions that you and the other players make. Your total earnings will be converted to real money at the end of the experiment if you complete all four parts.

You will receive the instructions for each game at the beginning of that game.

C.1.3. Screen 3

Before we begin the decision making rounds, you need to answer the following questions so that we can verify that you understand how your total earnings are calculated. These answers will not affect your actual earnings in any way.

C.2. Game One

C.2.1. Screen 1

In this game the computer will randomly assign you a role in the game: Player A or Player B. Half of the participants will be assigned as Player A and the other Player B. Your assigned role (Player A or B) will be communicated to you on the computer screen.

If you are assigned as *Player A*, you will stay as *Player A* for the entire experiment. Similarly, if you are assigned as *Player B*, you will stay as *Player B* for the entire experiment

Your task in this game is to make decisions as per your assigned role for 20 rounds. In each round, each Player A will be randomly paired with a Player B by the computer. You will play 20 rounds and in each round the computer will randomize the other player in the game. You will not know the identity of the person with whom you are paired each round.

In sum, your role (Player A or B) will not change during the rounds of the game, but the players with whom you will be paired will randomly change each round.

C.2.2. Screen 2

In this game you both players have two strategies. You will not be able to see the other player's choice until you submit your choice and the other player will not see your choice. You will earn points in each round according to the following rules:

A If you are player A you will move UP or DOWN.
- If you move UP you will earn 40 points.
- If you move DOWN, you will earn 30 or 70 points depending on Player B's move.
  - You will get 30 points if Player B moves LEFT
  - You will get 70 points if Player B moves RIGHT

B If you are Player B you will move LEFT or RIGHT.
- If you move LEFT you will earn 40 points.

- If you move RIGHT, you will earn 30 or 70 points depending on Player A's move.
  - You will get 70 points if Player A moves UP
  - You will get 30 points if Player A moves DOWN.

In the following table we summarize these rules.

In each cell the number on the left indicates points for Player A and the number on the right indicates points for Player B.

		Player B	
		LEFT	RIGHT
Player A	UP	(20,20)	(20,60)
	DOWN	(10,20)	(60,10)

You will now play 20 rounds of this game. In each round you will see the payoff table and decide about your choice. If you have understood the game and the rules please click I AM READY to begin playing.

C.3. Game Two

This game is similar to Game One. You will keep your role (Player A or Player B) assigned to you previously and you will again be paired with another player randomly by the computer. You will not know the identity of the person with whom you are paired each round. Your task is again to make decisions as per your assigned role for 30 rounds.

In this game, you will have three options, instead of two. Player A can now move UP or MIDDLE or DOWN and Player B can move LEFT or CENTER or RIGHT. You will not be able to see the other player's choice until you submit your choice and the other player will not see your choice. You will earn points in each round according to the following rules:

A If you are Player A you can move UP, MIDDLE or DOWN.

- If you move UP you will earn 20 points.
- If you move MIDDLE, You will get 0 or 50 points depending on Player B's move:
  - If you move MIDDLE and Player B moves LEFT you will get 50 points (with 10% probability) or 0 points (with 90% probability),
  - If you move MIDDLE and Player B moves CENTER or RIGHT, you will get 50 points for sure.
- If you move DOWN, you will earn 10, or 60 points depending on Player B's move:
  - If you move DOWN and Player B moves LEFT you will get 10 points,
  - If you move DOWN and Player B moves CENTER you will get you will get 60 points (with 50% probability) or 10 points (with 50% probability),
  - If you move DOWN and Player B moves RIGHT you will get 60 points,

B If you are Player B you can move LEFT, CENTER or RIGHT.

- If you move LEFT you will earn 20 or loose 30 points depending on Player A's move.
  - If you move LEFT and Player A moves UP or DOWN you will get 20 points,
  - If you move LEFT and Player A moves MIDDLE you will get 20 points (with 90% probability) or loose 30 points (with 10% probability).
- If you move CENTER, you will earn 50 or 0 points depending on Player A's move.
  - If you move CENTER and Player A moves UP you will get 50 points,
  - If you move CENTER and Player A moves MIDDLE you will get 0 points,
  - If you move CENTER and Player A moves DOWN you will get 50 points (with 50% probability) or 0 points (with 50% probability),
- If you move RIGHT, you will earn 10 or 60 points depending on Player A's move.
  - If you move RIGHT and Player A moves UP you will get 60 points,
  - If you move RIGHT and Player A moves MIDDLE or DOWN you will get 10 points.

In the following table we summarize these rules.

In each cell the number on the left indicates points for Player A and the number on the right indicates points for Player B.

		Player B		
		LEFT	CENTER	RIGHT
Player A	UP	(20,20)	(20,50)	(20,60)
	MIDDLE	(50,-30) 10% prob. (0,20) 90% prob.	(50,0)	(50,10)
	DOWN	(10,20)	(60,0) 50% prob. (10,50) 50% prob.	(60,10)

You will now play 30 rounds of this game. In each round you will see the payoff table and decide about your choice. If you have understood the game and the rules please click I AM READY to begin playing.

C.4. Game Three

C.4.1. Screen 1 for the Firm

In this game you are assigned the role of a Firm. In each round, you will start with 50 points. You will be paired with another player who is assigned the role of the environmental regulation agency. You have two options: COMPLY with the environmental regulations or NOT COMPLY. The Agency will also decide between two options: to INSPECT or to NOT INSPECT. If you decide to COMPLY, the cost is 30 points. Your final score in this round will be 20 points. If you decide to NOT COMPLY, and the Agency decides to INSPECT you will get caught and pay a penalty of 40 points to the Agency. In this case, your final score will be 10 points for this round. If you decide to NOT COMPLY, and the Agency decides to NOT INSPECT you will not pay any penalty. For your information, the Agency also starts with 50 points, and the cost of inspection is 30 points. If you choose NOT COMPLY, the Agency will loose 40 points due to environmental costs.

The player (Agency) you will be paired will change randomly from round to round. You will not be able to see the other player's choice before you submit your decision. The Agency will not be able to see your decision unless submitting his/her decision to INSPECT. The points you earn in a round will depend on the decisions made by you and the person you are paired with in that round.

C.4.2. Screen 1 for the Agency

In this game you are assigned the role of the environmental regulation agency. You will be paired with another player who is assigned the role of a Firm. In each round, you will start with 50 points. You have two options: INSPECT the firm or NOT INSPECT the Firm. The Firm also has two options: to COMPLY or to NOT COMPLY with environmental regulations. If the Firm is not complying it will cost you 40 points due to environmental damage. To INSPECT, you will pay 30 points, as cost of inspection. If you INSPECT and the FIRM is NOT COMPLYING, you will collect 40 points from the Firm.

For your information, the Firm also starts the round with 50 points, and the cost of compliance paid by the Firm is 30 points.

The player (Firm) you will be paired will change randomly from round to round. You will not be able to see the other player's choice before you submit your decision. The Firm will not be able to see your decision unless submitting his/her decision. The points you earn in a round will depend on the decisions made by you and the person you are paired with in that round.

C.5. Game Four

In this game you will keep the same role as you had in Game Three. Again, you will be matched with a different player in the other role in each round. And again, your decisions together with the decisions of the people that you will be matched with will determine your earnings.

The strategies available to each player are similar to the previous game. But in this game both players have one more option available before they decide.

The Agency can now hire two types of inspectors. Expert inspectors cost 30 points but they can conduct a FULL INSPECTION, which increases the chance of detection if the Firm is not complying. Inexperienced inspectors cost just 5 points but they can conduct only a PARTIAL INSPECTION which decreases the chance of detection if the Firm is not complying. The new alternative option for The Firm is PARTIAL COMPLIANCE. FULL COMPLIANCE costs 30 points to the Firm and eliminates the possibility of a fine. PARTIAL COMPLIANCE costs 5 points. It reduces the probability of detection but it does not prevent environmental damage.

The probability of the fine depends on the Firm and Inspectors mutual choices:

- If The Firm decides to NOT COMPLY and The Agency chooses FULL INSPECTION, the Firm will be caught.
- If The Firm decides to NOT COMPLY and The Agency chooses PARTIAL INSPECTION, the Firm will be caught 50% of the time.
- If The Firm chooses PARTIAL COMPLIANCE and The Agency chooses FULL INSPECTION, the Firm will be caught 90% of the time.
- If The Firm chooses PARTIAL COMPLIANCE and The Agency chooses PARTIAL INSPECTION, the Firm will not be caught.

Ending The Experiment

You have played four games in a total of 120 rounds. The computer will now randomly pick three rounds from each game to calculate your final earnings. Your total earnings will be converted to cash at a rate of 10 cents per point and you will be paid this amount privately in cash.

After taking the following short survey, you will be given a 5 digit code. Please write down and remember this code to collect your envelope.

Thank you for your participation.