Adaptive Enforcement with AI-Augmented Monitoring

Ginger Zhe Jin* D. Daniel Sokol† Liad Wagman‡

March 2026

Abstract

We study a novel dynamic inspection game in which a regulator commits to a detection technology, and a regulated firm chooses whether to engage in harmful conduct and, if detected and sanctioned, whether to incur a redesign cost that generates a modified violation requiring renewed detection. In this adaptive environment, bounded sanctions give rise to three Markov perfect equilibria: full compliance, harm until recognition, and persistent redesign.

We show that AI-augmented monitoring can shift the equilibrium from persistent redesign to harm until recognition, but only when regulatory investment crosses a regime-shifting threshold. Below this threshold, greater monitoring intensity may *increase* enforcement workload without reducing aggregate harm, as the regulator repeatedly detects adaptive violations while the firm continues to redesign. Thus, partial investments in AI monitoring can generate congestion rather than deterrence.

When the firm can also adopt AI to reduce its redesign cost, the regulator's deterrence threshold rises, reinforcing the strategic interaction between enforcement and evasion technologies. Moreover, congestion becomes particularly salient when AI-flagged violations require human review and regulatory review capacity is binding. In this case, the precision of AI triage—especially its false positive rate—matters as much as detection intensity. Enforcement effectiveness therefore depends not only on expanding detection, but also on allocating scarce human review resources efficiently.

Keywords: Adaptive enforcement; AI monitoring; inspection games; regulatory capacity; strategic obfuscation; regime shift.

JEL Codes: D82, K21, K42, L40.

*Department of Economics, University of Maryland, College Park. Email: jin@econ.umd.edu.

†USC Gould School of Law and USC Marshall School of Business (Marketing Department), University of Southern California; courtesy appointment in the Department of Economics. Email: dsokol@usc.edu.

‡Lally School of Management, Rensselaer Polytechnic Institute. Email: wagman@rpi.edu.

1 Introduction

The growing use of algorithms has made it increasingly difficult for law enforcement to keep pace with business practices that may harm consumers. Algorithmic systems now shape platform ranking and recommendation, targeted advertising, dynamic pricing, credit scoring, emissions reporting, and workplace management. These systems are updated through continuous experimentation and retraining, so small changes to code, prompts, data pipelines, or user interfaces can generate new variants of the same underlying conduct. When such conduct is opaque to outsiders, rapid iteration can outpace both market discipline and the evidentiary process required for legal action.

This gap is especially salient where harms scale quickly and are difficult to diagnose from outcomes alone. In antitrust, pricing algorithms may adjust at high frequency, potentially sustaining supracompetitive outcomes even without explicit agreement (Calvano et al., 2020; Assad et al., 2024). Asymmetric algorithm adoption can push prices toward monopoly levels (Brown and MacKay, 2023), and faster pricing algorithms may serve as coercive instruments in competitive interaction (Brown and MacKay, 2025). The FTC's 2023 monopolization complaint against Amazon alleged that search algorithms systematically prioritized sponsored products and Amazon's own brands in ways that entrenched market power, while the practices evolved through continuous experimentation (FTC, 2023).

In consumer protection, algorithmic ranking, disclosure design, or default rules can steer consumers in ways that shift surplus without generating an easily verifiable record. In financial supervision, credit-scoring and risk models are updated iteratively, potentially embedding discriminatory or unstable patterns. In environmental regulation, compliance systems and emissions reporting may rely on adaptive algorithms. In labor enforcement, algorithmic management tools influence wages, hours, and task allocation. As machine learning becomes a general-purpose input into firm decision-making, the enforcement problem becomes inherently dynamic: conduct evolves in response to what regulators monitor and litigate (Athey, 2018).

Public enforcement, however, remains structured around investigation, case selection, and retrospective remedies. Investigations take time, litigation is uncertain, and sanctions are bounded by statute and due process (Becker, 1968; Polinsky and Shavell, 2000). What distinguishes fast-moving algorithmic markets is that firms can modify conduct on a timescale short relative to the enforcement cycle.

Artificial intelligence (AI) offers regulators tools for adaptive enforcement: scalable monitoring, complaint triage, targeted audits, and post-remedy oversight that update as firms adjust tactics. AI expands the set of observable signals and may accelerate recognition of recurring patterns. Yet improved recognition alters strategic incentives. If regulators

become better at detecting a particular harmful pattern, firms may respond by redesigning the practice so that it again appears novel to the monitoring system.

To formalize this interaction, we study a dynamic inspection game in which a regulator commits to a detection technology and a firm chooses whether to deploy a harmful practice that yields private gain but imposes consumer harm. If detected and sanctioned, the firm may incur a redesign cost that generates a modified violation requiring renewed detection. A publicly observed recognition state captures whether the regulator has learned to detect the current pattern. Monitoring technology affects both immediate detection and the rate at which novel or previously missed patterns become recognizable. Sanctions are bounded.

Equilibrium behavior falls into one of three regimes: full compliance, harm until recognition, or persistent redesign. The economic intuition is that recognition creates an option value for the firm. Once recognized, the firm can comply or pay to redesign and return to the novel state. The value of that option depends on how quickly the novel state transitions back to recognition. When redesign is cheap and recognition slow, persistent redesign emerges: faster detection increases enforcement activity but does not reduce equilibrium harm. When redesign is costly, faster recognition shortens harm spells and induces compliance.

This framework yields several implications for AI monitoring investments. First, deterrence requires crossing a regime-shifting threshold in recognition speed; marginal improvements that do not eliminate persistent redesign are never optimal. Second, below-threshold AI adoption can increase enforcement workload without reducing harm, generating congestion rather than deterrence gains. Third, when firms can also use AI to lower redesign costs, the regulator's threshold rises, producing an arms-race dynamic. Fourth, when AI-flagged violations require human review and review capacity binds, triage precision and false-positive rates matter as much as detection intensity. Finally, when firms can shift activity across jurisdictions, monitoring incentives weaken, potentially trapping jurisdictions in under-enforcement.

The paper proceeds as follows. Section 3 presents the model. Section 4 characterizes Markov perfect equilibrium. Section 5 studies optimal monitoring investment under sanction caps. Section 6 provides numerical illustrations. Section 7 extends the model with heterogeneous redesign costs, cross-jurisdiction mobility, or capacity constraints in human review. Section 8 concludes. Proofs are in the Appendix.

2 Related Literature

Our analysis builds on the economics of public enforcement beginning with Becker (1968) and Stigler (1974), and the synthesis in Polinsky and Shavell (2000), which emphasizes the joint role of detection, adjudication, and sanctions under institutional constraints. We model

enforcement as a monitoring pipeline whose effectiveness evolves over time, following the spirit of Mookherjee and Png (1992), who distinguish broad monitoring from targeted investigation. Related work shows that sanctions alone may not achieve deterrence when penalties are socially costly or evidentiary standards bind (Shavell, 1987; Andreoni, 1991).

Our sanction cap captures real-world limits on penalties. In U.S. antitrust, the FTC's primary remedy under Section 5 is injunctive relief, and monetary penalties are statutorily constrained; criminal enforcement under the Sherman Act faces stringent evidentiary requirements and delay. Comparable limits arise in consumer protection, particularly following *AMG Capital Management v. FTC* (AMG Capital Management v. FTC, 2021). These constraints motivate environments in which monitoring technology, rather than sanctions alone, becomes central to deterrence.

The central friction in our model is a sanction cap combined with adaptive firm behavior. Enforcement under restricted penalties has been studied in repeated settings where reputational leverage substitutes for monetary sanctions (Harrington, 1988). The regulated party's information advantage also shapes enforcement design. Malik (1993) studies self-reporting under stochastic violations, and Kaplow and Shavell (1994) analyzes the social value of accuracy investments. Our triage extension builds on these insights by making review capacity a bottleneck: when enforcement is congestion-prone, false positives dilute deterrence.

Dynamic inspection regimes are surveyed in Avenhaus, von Stengel, and Zamir (2002). We depart from classic inspection games by introducing a firm-side redesign choice that resets detectability. This channel is motivated by algorithmic environments in which minor code, prompt, or interface changes can invalidate existing detectors without altering the underlying economic effect. For example, a platform investigated for a ranking algorithm can adjust weights or features while preserving similar demand effects, requiring regulators to rebuild an evidentiary record. Ortner et al. (2022) study adaptive responses in cartel screening; we instead characterize how recognition speed interacts dynamically with redesign incentives.

The strategic classification literature studies environments in which agents manipulate observed features to obtain favorable classification (Hardt et al., 2016). We share the focus on strategic responses to algorithmic rules but model equilibrium redesign of the harmful practice itself rather than feature manipulation around a fixed classifier. The recognition state in our model captures organizational learning by enforcement agencies updating detection procedures in response to evolving tactics.

Our motivating applications include algorithmic pricing and platform conduct. Calvano et al. (2020) show that learning algorithms can sustain supracompetitive prices in repeated oligopoly; Assad et al. (2024) provide large-scale evidence that margins increase primarily when all competitors adopt algorithmic pricing, underscoring strategic complementarities

relevant for enforcement. Real-time monitoring tools build on advances in text-based and NLP methods (Hoberg and Phillips, 2010, 2016; Duso et al., 2025), which demonstrate the feasibility of converting unstructured data into screening inputs. In our framework, the recognition hazard summarizes the effectiveness of such screening pipelines.

The model also connects to work on auditing algorithms (Sandvig et al., 2014) and accountability frameworks (Kroll et al., 2017). Evidence on public-sector IT adoption shows that technological improvements alone may not improve outcomes without complementary organizational capacity (Garicano and Heaton, 2010). This insight is consistent with our triage extension, in which recognition depends not only on statistical detection but also on the agency's ability to process signals into actionable enforcement. Lebovitz, Levina, and Lifshitz-Assaf (2021) document how AI tools evaluated against expert-labeled benchmarks can perform poorly in practice when tacit knowledge diverges from codified rules, reinforcing the importance of capacity constraints in AI-augmented enforcement.

Empirical detection of anticompetitive conduct provides additional motivation. Porter and Zona (1993) develop statistical tests for bid rigging in procurement auctions, establishing the feasibility of data-driven screening. Modern machine-learning extensions (Duso et al., 2025) expand these methods. In our model, the recognition hazard can be interpreted as summarizing the speed with which such screening tools convert raw data into actionable detection.

3 Model

3.1 Environment

Time is discrete and infinite, $t \in \{0, 1, 2, \dots\}$, with common discount factor $\delta \in (0, 1)$. There are two players: a regulated firm (the *firm*) and an enforcement agency (the *regulator*). The regulator represents consumer welfare and chooses monitoring technology, while the firm chooses whether to deploy a harmful algorithmic practice.

State variable. A publicly observed state $s_t \in \{0, 1\}$ summarizes whether the regulator's monitoring system currently *recognizes* the firm's harmful practice. State $s_t = 0$ means the current practice is *novel* to the monitoring pipeline (low detectability). State $s_t = 1$ means it is *recognized* (high detectability). The initial state is $s_0 = 0$.

We distinguish recognition from detection: while recognition simply refers to the state of $s_t = 1$, *detection* is a period-level event: after the regulator recognizes a conduct, it needs to take extra actions to litigate and sanction the firm in that period. As detailed below, we

assume detection occurs with a probability. So detection bundles recognizing, catching, and sanctioning. A practice can be recognized (the regulator understands the pattern) but not detected in a given period (the probabilistic detection draw fails).

This distinction captures the institutional reality that enforcement agencies must develop case theories, gather evidence, and establish investigative playbooks before they can effectively detect and prosecute a particular type of conduct. A novel algorithmic practice may cause harm from its first deployment, but until the agency understands the mechanism and develops detection protocols, prosecution is difficult. Once recognized, the same conduct becomes easier to identify and sanction. Below we will clarify the timing and resource differences between recognition and detection.

Regulator technology. At $t = 0$, the regulator chooses a monitoring technology $\tau \in \{L, A\}$, where L denotes legacy monitoring and A denotes AI-augmented monitoring. Technology affects detection probabilities and learning speed for novel patterns:

- If the firm deploys harmful conduct in state $s \in \{0, 1\}$, it can be immediately detected with probability $p_s^\tau \in [0, 1]$, where immediate detection refers to successful case development and sanction in the same period. We assume $p_1^\tau \geq p_0^\tau$ for each τ, and $p_s^A \geq p_s^L$ for each s.

- If the firm deploys harmful conduct in the state $s_t = 0$ and it is not immediately detected, the regulator can still learn to recognize the pattern at the end of period t with probability $\alpha^\tau \in [0, 1]$, but the detection (i.e., litigation and sanction) of this conduct must wait until the next period. We assume $\alpha^A \geq \alpha^L$.

- The regulator pays a per-period operational cost $c^\tau \geq 0$, with $c^A > c^L$. This fixed cost summarizes compute, data access, staffing, and governance overhead required to deploy trustworthy monitoring independent of caseload. In addition, there are two variable per-case costs: c_{rec} is the cost of opening and processing a case upon recognition (evidence gathering, legal review, building an investigative theory); c_{det} is the additional cost of completing a case with sanction (litigation, adjudication).

The distinction between immediate detection (p_s^τ) and learning (α^τ) captures two different margins of enforcement capacity. Immediate detection reflects whether the agency can identify and prosecute harmful conduct as it occurs, based on existing protocols and evidence. Learning reflects whether the agency can use near-misses, complaints, and partial evidence to obtain new detection capabilities for the future.

Firm actions and payoffs. In each period t, after observing s_t, the firm chooses an action a_t. If $s_t = 0$, available actions are *Comply* (C) or *Harm* (H). If $s_t = 1$, available actions are C, H, or *Adapt and Harm* (R). Action R represents redesigning the practice (or the code, prompt structure, data pipeline, or interface) to reset detection risk to the novel regime in the current period, at resource cost $k > 0$.

If the firm chooses C, it earns normalized payoff 0 and generates no consumer harm. If the firm chooses H or R, it earns private gain $g > 0$ in that period and generates consumer harm $h > 0$ in that period. If harmful conduct is detected, the firm pays a sanction $F \geq 0$ in that period. We interpret F as bounded by legal and institutional constraints.

The firm payoff in period t is

$$u_F(a_t, s_t) = \begin{cases} 0 & \text{if } a_t = C, \\ g - \mathbf{1}\{\text{detected}\}F & \text{if } a_t = H, \\ g - \mathbf{1}\{\text{detected}\}F - k & \text{if } a_t = R. \end{cases} \tag{1}$$

Given technology τ, the regulator payoff in period t is

$$u_R(a_t; \tau) = \begin{cases} -c^\tau & \text{if } a_t = C, \\ -h - c^\tau - c_{\text{rec}} \cdot \mathbf{1}\{s_{t+1} = 1\} - c_{\text{det}} \cdot \mathbf{1}\{\text{detected}\} & \text{if } a_t \in \{H, R\}. \end{cases} \tag{2}$$

This objective focuses on consumer harm and monitoring cost. The analysis can be extended to include weights on firm surplus or fine revenues.

3.2 Timing and State Transitions

Within each period t:

1. The state s_t is observed.

2. The firm chooses a_t. If $a_t = R$, the firm pays k and the harmful practice for detection purposes is treated as novel in that period.

3. If $a_t \in \{H, R\}$, detection occurs with probability $p^\tau_{s'_t}$, where $s'_t = s_t$ if $a_t = H$, and $s'_t = 0$ if $a_t = R$.

4. The next state s_{t+1} is realized.

   - If $a_t = C$, then $s_{t+1} = s_t$.[1]

[1]Note that we assume compliance does not change the state. This is natural when $s = 0$. If $s = 1$,

- If $a_t \in \{H, R\}$ and the effective state for recognition is 0, then

$$\Pr(s_{t+1} = 1 \mid a_t \in \{H, R\}, s'_t = 0) = p_0^\tau + (1 - p_0^\tau)\alpha^\tau. \tag{3}$$

  If immediate detection does not occur and learning does not occur, then $s_{t+1} = 0$.

- If $a_t = H$ and $s_t = 1$, then $s_{t+1} = 1$.

Define the recognition hazard for a novel pattern under technology τ as

$$q^\tau \equiv p_0^\tau + (1 - p_0^\tau)\alpha^\tau \in (0, 1]. \tag{4}$$

This quantity plays a central role in the analysis: it measures the probability that a novel harmful practice becomes recognized within one period of deployment. The recognition hazard combines two channels: immediate detection in the same period (with probability p_0^τ) and learning without immediate detection (with probability $(1 - p_0^\tau)\alpha^\tau$).

The distinction between these two channels is crucial for the firm's incentives in choosing R versus H: R can reset the state from 1 to 0 at the beginning of period t and thus reduce the risk of being recognized and sanctioned within t because we assume the probability of immediate detection is lower when the state is 0 ($p_0^\tau < p_1^\tau$). In addition, when immediate detection fails, the monitoring technology can also recognize the pattern via learning, but learning does not trigger detection until the next period. This delay gives the firm an option to choose R at the beginning of $t + 1$, effectively resetting the state to zero and requiring novel recognition before sanction. In doing so, the detection that could happen in $t + 1$ upon the recognition learned in t is lost in the process, because the recognized pattern no longer exists when the regulator attempts to litigate and sanction. Instead, the regulator must revent the wheel to recognize and sanction the novel pattern. By contrast, H does not carry this option value: if the regulator recognizes the firm's conduct in t, the state remains one in $t + 1$ under H and the firm faces a higher risk of being detected in $t + 1$ because we assume $p_1^\tau > p_0^\tau$ for any monitoring technology.

Remark 1 (AI's comparative advantage: learning speed α^τ)**.** *While AI-augmented monitoring could in principle improve both immediate detection (p_0^τ) and learning speed (α^τ), its comparative advantage likely lies in accelerating α^τ (the rate at which novel patterns become recognizable) rather than in one-shot detection of previously unseen tactics. Genuinely novel algorithmic conduct is difficult for* any *monitoring system to detect immediately, because*

compliance does not reset the state back to zero because the regulator often pays close attention to previously sanctioned firms even if the firm is in compliance with the consent decree it signed before.

the regulator has not yet developed the evidentiary framework, detection protocols, or legal theories needed to act. What AI changes is how quickly the agency can synthesize partial evidence, complaints, and near-misses into a reusable detection capability. This interpretation (that AI primarily raises α^τ) guides our numerical illustrations in Section 6, where we isolate the learning channel by varying α while holding p_0 fixed. In practice, improvements to p_0^τ (through better data integration, audit protocols, or real-time surveillance) and improvements to α^τ (through evidence synthesis, pattern recognition across matters, and codification of investigative playbooks) are both valuable, and different enforcement instruments map differently to these margins.

3.3 Equilibrium Concept

The model is a discounted stochastic game with perfect public monitoring. At $t = 0$ the regulator chooses a monitoring technology τ and commits. Thereafter, in each period the publicly observed state s_t and the realized immediate detection and learning draws are common knowledge.

A strategy for the regulator is a choice $\tau \in \{L, A\}$ at $t = 0$. A strategy for the firm is a mapping from public histories to actions. We focus on Markov perfect equilibrium (MPE), in which the firm's strategy depends only on the current state s_t and the chosen technology τ, not on the full history of play. Given the regulator's commitment to τ, the firm faces a discounted Markov decision problem (MDP) with states $\{0, 1\}$ and actions that depend on the state. The MPE of the firm's continuation game corresponds to an optimal stationary Markov policy for this MDP.

Remark 2 (On commitment and forward-looking behavior)**.** *The regulator's commitment assumption captures institutional frictions: adopting a monitoring technology requires procurement, staff training, and pipeline integration, all of which are difficult to reverse frequently. Both players are forward-looking. The regulator chooses technology anticipating the firm's equilibrium response; the firm chooses actions taking into account how current decisions affect future detection and recognition.*

4 Baseline Equilibrium Analysis

4.1 Value Functions and Bellman Equations

Fix a technology τ and define the expected flow payoffs from harmful conduct in each state:

$$a_0^\tau \equiv g - p_0^\tau F, \qquad a_1^\tau \equiv g - p_1^\tau F. \tag{5}$$

The quantity a_0^τ represents expected profit from harm when the practice is novel, and a_1^τ when recognized.

The Bellman equations characterizing the firm's optimization problem are:

$$V_0^\tau = \max\Big\{\delta V_0^\tau,\ a_0^\tau + \delta\big[(1-q^\tau)V_0^\tau + q^\tau V_1^\tau\big]\Big\}, \tag{6}$$

$$V_1^\tau = \max\Big\{\delta V_1^\tau,\ a_1^\tau + \delta V_1^\tau,\ -k + a_0^\tau + \delta\big[(1-q^\tau)V_0^\tau + q^\tau V_1^\tau\big]\Big\}. \tag{7}$$

In (6), the first term corresponds to choosing C (compliance), and the second term corresponds to choosing H (harm) when $s = 0$. In (7), the three terms correspond to C, H, and R (redesign) respectively when $s = 1$.

Lemma 1 (Existence of a stationary optimal policy)**.** *For each fixed technology τ, there exists a stationary deterministic policy for the firm that solves* (6)–(7)*. The associated value function (V_0^τ, V_1^τ) is the unique fixed point of the Bellman operator.*

Lemma 1 follows from standard results for discounted Markov decision problems with finite state and action sets; see Puterman (1994). We provide a complete proof in Appendix B for completeness.

4.2 A Focal Regime: Recognized Misconduct Is Deterred

The most relevant regime for adaptive enforcement is one in which, once a pattern is recognized, ongoing deployment of that same pattern is privately unattractive. This corresponds to settings where detection is high once the regulator understands what to look for, even if initial detection for novel variants is low.

Assumption 1 (Effective deterrence once recognized)**.** *For the monitoring technology under consideration, $a_1^\tau \le 0$, that is, $g \le p_1^\tau F$.*

Remark 3 (Economic interpretation of Assumption 1)**.** *Assumption 1 captures the empirical observation that enforcement agencies are typically much more effective against known threats than against novel ones. Once a pattern of conduct has been the subject of investigation,*

litigation, or rulemaking, the agency develops detection protocols, establishes legal precedent, and trains staff to identify similar violations. The assumption rules out the trivial case where even recognized misconduct is profitable, focusing attention on the more interesting dynamics between novel and recognized states. In terms of policy, the assumption corresponds to situations where remedies and consent decrees are effective at deterring the specific conduct that was litigated, but new variants can emerge that fall outside the scope of existing orders.

Under Assumption 1, action H in state $s = 1$ is weakly dominated by C because both lead to the same continuation state, while H yields weakly lower current payoff. Thus in state 1 the firm compares C and R.

Proposition 1 (Equilibrium characterization under Assumption 1)**.** *Fix a technology τ and suppose Assumption 1 holds. Then there exists a Markov perfect equilibrium in stationary pure strategies for the firm with the following structure. The value functions (V_0^τ, V_1^τ) are unique. The equilibrium policy is unique except on knife-edge parameter values where the firm is indifferent (in particular when $a_0^\tau = 0$, $a_1^\tau = 0$, or $k = k^*(\tau)$).*

1. ***Full compliance.*** *If $a_0^\tau \leq 0$ (equivalently $g \leq p_0^\tau F$), then $V_0^\tau = V_1^\tau = 0$ and the firm chooses C in both states.*

2. ***Harm until recognition.*** *If $a_0^\tau > 0$ and*

$$k \geq k^*(\tau) \equiv \frac{a_0^\tau}{1 - \delta(1 - q^\tau)}, \tag{8}$$

then the firm chooses H in state 0 and C in state 1. The value functions are

$$V_1^\tau = 0, \qquad V_0^\tau = \frac{a_0^\tau}{1 - \delta(1 - q^\tau)}. \tag{9}$$

3. ***Persistent redesign (cat-and-mouse).*** *If $a_0^\tau > 0$ and $k < k^*(\tau)$, then the firm chooses H in state 0 and R in state 1. The value functions are*

$$V_0^\tau = \frac{a_0^\tau - \delta q^\tau k}{1 - \delta}, \qquad V_1^\tau = V_0^\tau - k. \tag{10}$$

The proof appears in Appendix B. Proposition 1 shows that even when recognized patterns are deterred, bounded sanctions can still allow persistent misconduct through redesign when the redesign cost k is sufficiently small. We interpret AI-enabled rapid iteration as lowering k and increasing the feasibility of the cat-and-mouse regime.

Remark 4 (Economic intuition for the threshold $k^*(\tau)$)**.** *The threshold $k^*(\tau)$ balances the benefit of continued harmful activity against the cost of redesign. The numerator a_0^τ is the expected profit from one period of novel harm. The denominator $1-\delta(1-q^\tau)$ captures the effective discount rate for the novel state, accounting for the probability q^τ that recognition occurs. When $k < k^*(\tau)$, the present value of resuming harm in the novel state exceeds the redesign cost, so the firm redesigns upon recognition. When $k \geq k^*(\tau)$, the firm finds redesign too costly and instead accepts compliance once recognized. The threshold is decreasing in both detection probability p_0^τ and learning speed α^τ, as shown in Proposition 2.*

Corollary 1 (Expected discounted harm under equilibrium)**.** *Under the conditions of Proposition 1, let $H_t^\tau \in \{0,1\}$ denote the indicator that harmful conduct is deployed in period t under technology τ. Starting from $s_0 = 0$,*

$$\textit{Total Harm} = h\mathbb{E}\Big[\sum_{t=0}^{\infty}\delta^t H_t^\tau\Big] = \begin{cases} 0 & \textit{in full compliance,} \\ \dfrac{h}{1-\delta(1-q^\tau)} & \textit{in harm until recognition,} \\ \dfrac{h}{1-\delta} & \textit{in persistent redesign.} \end{cases} \tag{11}$$

The proof is in Appendix B. The key observation is that in the harm-until-recognition equilibrium, harm occurs only while the state is $s = 0$, and the state transitions from 0 to 1 with hazard q^τ each period. In the persistent-redesign equilibrium, harm occurs every period regardless of the state.

4.3 Comparative Statics: Recognition Speed and the Redesign Region

The threshold $k^*(\tau)$ in (8) summarizes how attractive redesign is. A higher $k^*(\tau)$ means redesign is privately optimal for a larger set of redesign costs.

Proposition 2 (Recognition speed shrinks the redesign region)**.** *Fix primitives (g, F, δ, k) and consider the regime of Proposition 1 with $a_0^\tau > 0$. Then:*

1. *$\partial k^*(\tau)/\partial p_0^\tau < 0$: The threshold is strictly decreasing in the detection probability for novel patterns.*
2. *$\partial k^*(\tau)/\partial \alpha^\tau \leq 0$: The threshold is weakly decreasing in the learning probability, with strict inequality when $p_0^\tau < 1$.*

3. *In the harm-until-recognition equilibrium, the expected discounted harm in Corollary 1 is strictly decreasing in q^τ.*

The proof is in Appendix B. Proposition 2 formalizes the core operational role of AI-augmented monitoring in the baseline model. Increasing detection for novel patterns p_0^τ or increasing recognition speed α^τ reduces the parameter region in which repeated redesign is privately optimal. This is the scenario in which faster surveillance and learning can restore deterrence even when formal sanctions are capped.

The economic mechanism is as follows. Immediate detection and learning both reduce the value of the novel state by increasing the probability that the firm will be recognized. A lower value of the novel state means a lower option value from redesigning to return to that state. When the option value falls below the redesign cost k, the firm stops redesigning and instead complies once recognized.

5 Design Monitoring Technology Under a Sanction Cap

This section studies the regulator's design problem when formal sanctions are capped. The baseline analysis isolates a simple logic: when the firm can cheaply redesign, marginal improvements in recognition that do not eliminate redesign have limited deterrence effects. We first restate this logic as a threshold condition on recognition speed, and then analyze adoption and continuous investment.

5.1 A minimal recognition speed to eliminate redesign

Maintain Assumption 1 and suppose $a_0^\tau > 0$, so novel harmful conduct is privately attractive. Proposition 1 implies that the redesign regime disappears if and only if the redesign cost exceeds the threshold $k^*(\tau)$.

Corollary 2 (Minimum recognition speed needed to deter redesign)**.** *Fix a technology τ and suppose Assumption 1 holds and $a_0^\tau > 0$. Define*

$$\underline{q}(\tau) \equiv \frac{a_0^\tau/k - (1-\delta)}{\delta}. \tag{12}$$

If $\underline{q}(\tau) \leq 0$, then redesign is deterred for any $q^\tau \in (0,1]$. If $\underline{q}(\tau) \in (0,1]$, then redesign is deterred if and only if $q^\tau \geq \underline{q}(\tau)$. If $\underline{q}(\tau) > 1$, redesign occurs for all feasible recognition speeds $q^\tau \in (0,1]$.

Corollary 2 identifies when monitoring improvements can make a difference. When $\underline{q}(\tau) > 1$, even perfect recognition cannot deter redesign because the sanction cap and the low

redesign cost jointly imply that the firm is willing to pay k repeatedly. In such cases, agencies must either raise sanctions (if legally possible) or raise the effective redesign cost through regulatory design (such as requiring pre-deployment review or imposing change-management obligations).

5.2 Sanctions and recognition as substitutes

The previous results hold fixed the sanction F. In many settings agencies face binding limits on F but can shift resources between adjudication and monitoring. The baseline model delivers two simple back-of-the-envelope substitution formulas that clarify how monitoring improvements can partially substitute for higher sanctions.

First, full compliance occurs whenever novel harmful conduct is privately unprofitable, i.e., $a_0^\tau = g - p_0^\tau F \leq 0$, which is equivalent to $F \geq g/p_0^\tau$. Second, even when $a_0^\tau > 0$, the redesign regime disappears when $k \geq k^*(\tau) = a_0^\tau/[1 - \delta(1 - q^\tau)]$, which can be rewritten as a lower bound on F.

Corollary 3 (Minimal sanctions to deter harm and deter redesign)**.** *Fix τ and suppose Assumption 1 holds. Define*

$$F^{\text{comp}}(\tau) \equiv \frac{g}{p_0^\tau}, \qquad F^{\text{no-redesign}}(\tau) \equiv \frac{g - k[1 - \delta(1 - q^\tau)]}{p_0^\tau}. \tag{13}$$

If $F \geq F^{\text{comp}}(\tau)$, full compliance is the unique equilibrium. If $F < F^{\text{comp}}(\tau)$ but $F \geq F^{\text{no-redesign}}(\tau)$, then harmful conduct occurs in state 0 but redesign is deterred, so equilibrium harm is transient. Increasing q^τ lowers $F^{\text{no-redesign}}(\tau)$, meaning that faster recognition weakly relaxes the sanction required to eliminate redesign.

Corollary 3 provides a substitution formula that is useful for policy analysis. If an agency faces a sanction cap $\bar{F}$ that is below the full-compliance threshold, it can still eliminate persistent redesign by investing in recognition speed until $F^{\text{no-redesign}}(\tau)$ falls below $\bar{F}$. The corollary also shows that the required recognition speed depends on the redesign cost k: firms with lower adaptation costs require faster recognition to deter redesign.

It is worth noting that "faster recognition" encompasses two distinct margins: immediate detection probability p_0^τ and the learning rate α^τ. Both contribute to the recognition hazard $q^\tau = p_0^\tau + (1 - p_0^\tau)\alpha^\tau$, and both appear in $F^{\text{no-redesign}}$. Improvements to p_0^τ are mechanically more deterrent because they increase the probability that harm is sanctioned immediately and thus directly reduce the firm's expected per-period gain $a_0^\tau = g - p_0^\tau F$ from novel conduct. In contrast, improvements to α^τ affect only the learning channel. However, the relevant policy comparison is cost-effectiveness, since different enforcement instruments map differently to

these margins. Data access and audit protocols may primarily affect p_0^τ, while analytic tooling and case-learning systems may primarily affect α^τ.

5.3 Adoption of AI-augmented monitoring

We now compare the regulator's welfare across monitoring technologies. Fix $\tau \in \{L, A\}$ and let the firm play the equilibrium characterized in Proposition 1 when Assumption 1 holds for that τ. Let W^τ denote the regulator's expected discounted payoff from $t = 0$, starting at $s_0 = 0$.

According to Equation 2, the regulator payoff is $-h - c^\tau - c_{\text{rec}} \cdot \mathbf{1}\{s_{t+1} = 1\} - c_{\text{det}} \cdot \mathbf{1}\{\text{detected}\}$ whenever harm is deployed and $-c^\tau$ otherwise. This implies:

$$W^\tau = -\frac{c^\tau}{1-\delta} \; - \; (h + c_{rec} q^\tau + c_{det} p_0^\tau)\, \mathbb{E}\Big[\sum_{t=0}^{\infty} \delta^t H_t\Big]. \tag{14}$$

or more specifically:

$$W^\tau = \begin{cases} -\dfrac{c^\tau}{1-\delta} & \text{in full compliance,} \\ -\dfrac{c^\tau}{1-\delta} - \dfrac{h + c_{rec} q^\tau + c_{det} p_0^\tau}{1 - \delta(1-q^\tau)} & \text{in harm until recognition,} \\ -\dfrac{c^\tau}{1-\delta} - \dfrac{h + c_{rec} q^\tau + c_{det} p_0^\tau}{1-\delta} & \text{in persistent redesign.} \end{cases}$$

Proposition 3 (When does the regulator adopt AI?)**.** *Suppose Assumption 1 holds under both $\tau = L$ and $\tau = A$. Then AI-augmented monitoring is strictly preferred by the regulator, $W^A > W^L$, if and only if*

$$\frac{c^A - c^L}{1-\delta} \; < \; \left(\mathbb{E}\Big[(h + c_{rec} q^L + c_{det} p_0^L) \sum_{t=0}^{\infty} \delta^t H_t^L\Big] \; - \; (h + c_{rec} q^A + c_{det} p_0^A) \mathbb{E}\Big[\sum_{t=0}^{\infty} \delta^t H_t^A\Big] \right), \tag{15}$$

where the harm terms are given by Corollary 1 under each technology's equilibrium.

The proof is immediate from (14). Proposition 3 provides a transparent adoption rule: the discounted operational cost of AI must be smaller than the discounted reduction in harm and corresponding recognition and detection costs. The right-hand side is largest when AI crosses the boundary in Corollary 2 and moves the equilibrium from persistent redesign to harm-until-recognition or full compliance.

5.4 Continuous recognition investment and lumpy value

To isolate the design implication suggested by Corollary 2, suppose the regulator starts from a baseline (legacy) recognition hazard $q^L > 0$ and can invest to raise recognition to any $q \in [q^L, 1]$ at a per-period operational cost $\kappa(q)$. Assume κ is increasing and $\kappa(q^L) = c^L$. For simplicity hold $a_0 > 0$ fixed and maintain Assumption 1. Define the redesign boundary

$$\bar{q} \equiv \max\left\{q^L,\ \frac{a_0/k - (1-\delta)}{\delta}\right\}. \tag{16}$$

Proposition 4 (No value of marginal recognition below the redesign boundary)**.** *Suppose $a_0 > 0$ and Assumption 1 holds. If $\bar{q} \in (q^L, 1)$, then any $q \in (q^L, \bar{q})$ is strictly dominated for the regulator by $q = q^L$: it yields the same equilibrium harm as in persistent redesign but strictly higher monitoring cost (including higher per-period fixed operational costs and more variable costs in recognition and detection driven by a heavier caseload under better monitoring technology.) Therefore any optimal recognition choice q^* either equals q^L or satisfies $q^* \geq \bar{q}$.*

Proposition 4 has important implications for how agencies should evaluate monitoring investments. Incremental improvements that raise detection rates or shorten learning times may look good on operational metrics but have no effect on deterrence if they leave the firm willing to redesign. The first welfare gains arrive only when recognition is fast enough to cross the redesign boundary. This suggests that agencies should evaluate investments with a "regime shift" mindset rather than by marginal improvements in detection and learning rates.

Corollary 4 (Congestion: marginal recognition is counterproductive below the boundary)**.** *Suppose the conditions of Proposition 4 hold and $c_{\text{rec}} > 0$. In the persistent-redesign regime, the regulator's expected discounted payoff under* (14) *is*

$$W_{\text{PR}}(q) = -\frac{h + c^\tau + c_{\text{det}}\, p_0^\tau + c_{\text{rec}}\, q^\tau}{1-\delta}, \tag{17}$$

which is strictly decreasing in q. Therefore any $q \in (q^L, \bar{q})$ is strictly worse for the regulator than $q = q^L$: incremental recognition investment below the regime-shifting boundary not only fails to reduce harm but actively increases enforcement costs through wasted case-opening effort on patterns the firm has already abandoned via redesign.

In the harm-until-recognition regime ($q \geq \bar{q}$), under the relatively mild condition $c_{\text{rec}}(1-\delta) < h + c_{\text{det}}\, p_0$, the regulator's payoff is strictly increasing in q.

Proof. In persistent redesign, harm occurs every period, detection events occur at rate p_0 per period, and recognition events occur at rate q per period (since $\Pr(s_t = 1) = q$ in steady

state). Substituting into the discounted sum of (14) yields (17). The derivative with respect to q is $-c_{\text{rec}}/(1-\delta) < 0$.

For harm until recognition, the regulator's payoff excluding the constant $c^\tau/(1-\delta)$ term is $W(q) = -(h + c_{\text{det}}\, p_0 + c_{\text{rec}}\, q)/(1-\delta(1-q))$. The derivative has numerator proportional to $(h + c_{\text{det}}\, p_0) - c_{\text{rec}}(1-\delta)$, which is positive under the stated condition. □

Remark 5 (Dependencies of the investment threshold)**.** *The redesign boundary $\bar{q}$ in* (16) *depends on both the firm's adaptation cost k and the sanction cap F (through $a_0 = g - p_0 F$). Lower k raises $\bar{q}$, meaning that as AI-enabled tools reduce the firm's cost of iterating around enforcement, agencies must achieve faster recognition to cross into the deterrent regime. Similarly, a lower effective sanction cap F, whether due to statutory limits, evidentiary burdens, or judicial doctrine, raises a_0 and thereby raises $\bar{q}$. These dependencies create a compound challenge: technological change may simultaneously lower firm adaptation costs and raise the monitoring investment threshold that agencies must clear. The recent experience of digital-markets enforcement, where sanctions have often been modest relative to firm profits and where algorithmic systems can be rapidly updated, illustrates this joint constraint.*

6 Numerical Illustration and Simulations

This section provides a numerical illustration of the baseline model. The goal is to highlight (i) how recognition speed q^τ determines whether the firm enters a persistent redesign regime, and (ii) how increases in q^τ can generate discrete reductions in expected harm by inducing a regime shift.

6.1 An illustrative parameterization and the regime shift

We consider the following parameter values:

$$(g, h, F, p_0, p_1, \delta, k) = (1,\ 1,\ 1.5,\ 0.1,\ 0.9,\ 0.95,\ 1).$$

These values satisfy Assumption 1 because $a_1 = g - p_1 F = 1 - 1.35 = -0.35 < 0$, while novel harmful conduct remains privately attractive because $a_0 = g - p_0 F = 1 - 0.15 = 0.85 > 0$. For illustration, we start with $c_{rec} = 0$ and $c_{det} = 0$ and will allow positive recognition and detection costs later in Section 6.3.

We compare two monitoring technologies that differ only in the learning parameter α^τ:

$$\alpha^L = 0.2, \qquad \alpha^A = 0.9,$$

Table 1: Illustrative parameterization and equilibrium regimes

Technology	α	q	k^*	Regime	Analytic harm	Simulated harm
Legacy (L)	0.20	0.28	2.69	persistent redesign	20.00	20.00
AI-augmented (A)	0.90	0.91	0.93	harm until recognition	1.09	1.09

which imply recognition hazards $q^L = p_0 + (1-p_0)\alpha^L = 0.28$ and $q^A = p_0 + (1-p_0)\alpha^A = 0.91$. This parameterization isolates the learning channel by holding immediate detection p_0 fixed, consistent with the interpretation in Remark 1. In practice, AI-augmented monitoring could improve both p_0 and α, and Section 5 discusses how both margins contribute to deterrence. The focus on α here reflects the view that genuinely novel algorithmic conduct is difficult to detect immediately regardless of monitoring technology; AI's comparative advantage may therefore lie primarily in accelerating the rate at which novel patterns become recognizable rather than in improving one-shot detection of previously unseen tactics.

Table 1 summarizes the implied redesign thresholds, equilibrium regimes, and expected discounted harm under each technology. The Monte Carlo estimates simulate the geometric time to first recognition under harm-until-recognition using 200,000 draws (standard error about 7×10^{-4} in this parameterization) and confirm the closed-form expressions in Corollary 1.

The parameterization in Table 1 illustrates the economic mechanism. Under the legacy technology, recognition is slow ($q^L = 0.28$), so a redesign buys the firm a long expected return to the novel state. When a novel pattern is privately profitable ($a_0 > 0$), the option value of resetting the state is high, and for $k = 1$ the firm finds redesign worthwhile. Equilibrium therefore features persistent redesign, which generates harm in every period despite the fact that recognized harms are deterred.

6.2 Arms-race illustration: varying the firm's adaptation cost

The above numerical illustration treats k as fixed, but AI-enabled tools may also lower the firm's cost of redesigning algorithmic practices. If AI reduces k for firms while simultaneously raising q for agencies, the net effect on equilibrium harm depends on the relative magnitudes. Proposition 2 shows that k^* is decreasing in q, so agency AI adoption can offset firm AI adoption by shrinking the redesign region. However, if firm-side cost reductions outpace agency-side recognition improvements, the redesign region expands and harm increases. This suggests an arms-race dynamic in which the welfare consequences of AI depend on the comparative rate of capability acquisition on each side. Empirically, understanding whether agencies can keep pace with firm-side innovation is therefore central to predicting enforcement outcomes.

Table 2: Equilibrium regimes across redesign costs and monitoring technologies

k	$k^*(q^L)$	Legacy regime	Legacy harm	$k^*(q^A)$	AI regime	AI harm	$\bar{q}$
3	2.69	harm until recog.	3.16	0.93	harm until recog.	1.09	n/a
2	2.69	persistent redesign	20.00	0.93	harm until recog.	1.09	0.395
1	2.69	persistent redesign	20.00	0.93	harm until recog.	1.09	0.842

To illustrate the arms-race dynamic, we compare three levels of the firm's redesign cost, $k \in \{3, 2, 1\}$, holding all other parameters at the baseline values $(g, h, F, p_0, p_1, \delta) = (1, 1, 1.5, 0.1, 0.9, 0.95)$. Table 2 summarizes the equilibrium regimes and harm levels under legacy and AI-augmented monitoring, along with the regime-shifting threshold $\bar{q}$ from (16).

The table reveals three important patterns:

Within-regime improvement ($k = 3$). When the firm's redesign cost is high ($k = 3 > k^*(0.28) = 2.69$), the equilibrium is "harm until recognition" even under legacy monitoring. In this case, AI monitoring does not cause a regime shift; the firm already stops redesigning because k exceeds the threshold. AI nonetheless reduces harm *within* the harm-until-recognition regime by speeding up recognition: discounted harm falls from $1/[1-\delta(1-q^L)] = 3.16$ to $1/[1-\delta(1-q^A)] = 1.09$.

Regime shift from firm-side AI ($k = 3 \to k = 2$). When AI technology reduces the firm's adaptation cost from $k = 3$ to $k = 2$, the firm crosses the redesign threshold under legacy monitoring ($k = 2 < k^*(0.28) = 2.69$) and switches from "harm until recognition" to "persistent redesign." Harm jumps from 3.16 to 20, a dramatic deterioration. The firm therefore has strong incentives to adopt AI that lowers its redesign cost. However, the regime-shifting threshold for the regulator is only $\bar{q} = 0.395$: a moderate investment in AI monitoring suffices to push q above this threshold and restore the "harm until recognition" regime.

Escalation ($k = 2 \to k = 1$). When the firm's adaptation cost drops further to $k = 1$, the regime-shifting threshold rises to $\bar{q} = 0.842$. The regulator must now invest substantially more in AI to cross this higher boundary. This escalation is consistent with the comparative statics in Proposition 2 and Remark 4: lower k raises $\bar{q}$, requiring faster recognition to achieve the same deterrence outcome.

The comparison between $k = 2$ ($\bar{q} = 0.395$) and $k = 1$ ($\bar{q} = 0.842$) has a direct policy implication for the timing of AI investment. When AI technology is still developing and the firm's adaptation cost has not yet fallen dramatically (e.g., from $k = 3$ to $k = 2$), the regime-shifting threshold is relatively low and may be politically and fiscally easier to reach. Waiting until AI matures and drastically reduces firm adaptation costs (e.g., to $k = 1$) raises the threshold and makes it harder for the regulator to catch up. With sufficient early

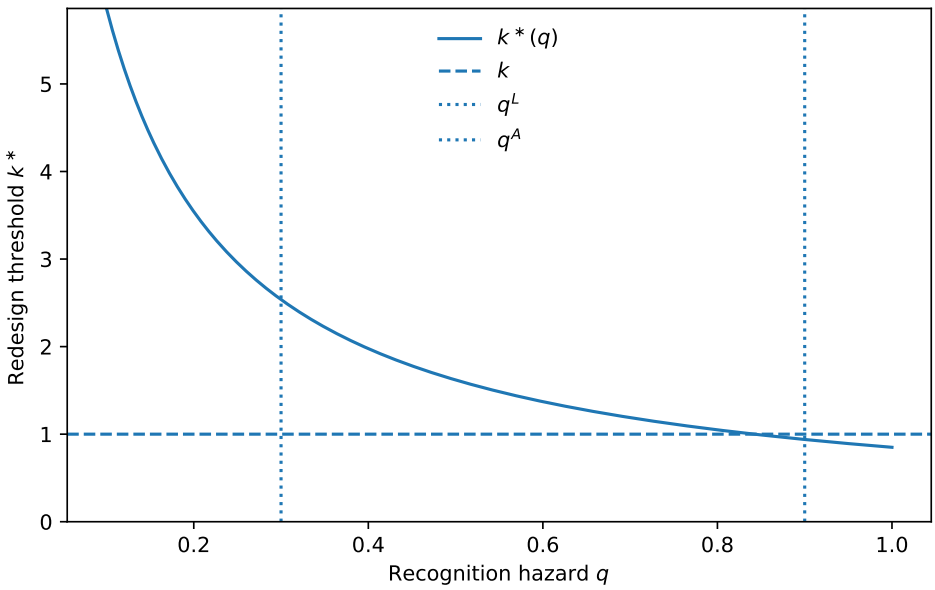


Figure 1: Redesign threshold as a function of recognition hazard. Notes: $(g, F, p_0, p_1, \delta, k) = (1, 1.5, 0.1, 0.9, 0.95, 1)$ and $a_0 = 0.85$. Vertical lines indicate $q^L = 0.28$ and $q^A = 0.91$.

investment, the regulator's preemptive adoption of AI pushes up the cost threshold above which the firm must invest to make persistent redesign profitable. Of course, these benefits of early adoption must be balanced against the possibility that the cost of quality-adjusted AI technology ($\kappa(q)$ in the current model) may decline over time, creating an option value of waiting.

Remark 6 (Budgetary constraints)**.** *In practice, agency monitoring investments are subject to budgetary constraints that may make it difficult to reach the regime-shifting threshold in a single budget cycle. The model's cost function $\kappa(q)$ can be reinterpreted as a budget-constrained optimization: the agency maximizes welfare subject to $\kappa(q) \leq B$, where B is the per-period budget. When $B < \kappa(\bar{q})$, the budget is insufficient to cross the regime-shifting threshold, and Proposition 4 implies that the agency is better off not investing in incremental monitoring at all (or investing in other enforcement margins such as raising the effective sanction F or raising the firm's redesign cost k through compliance mandates). Budget increases are most valuable when they cross the regime-shifting threshold, creating a case for concentrated rather than incremental investment. A fuller treatment of dynamic budgetary constraints with endogenous budget processes is left for future work.*

Under the AI-augmented technology, recognition is rapid ($q^A = 0.91$). A redesign now buys little time because the conduct is likely to be recognized again quickly. The continuation value from resetting falls below the redesign cost, and the firm instead chooses harm only

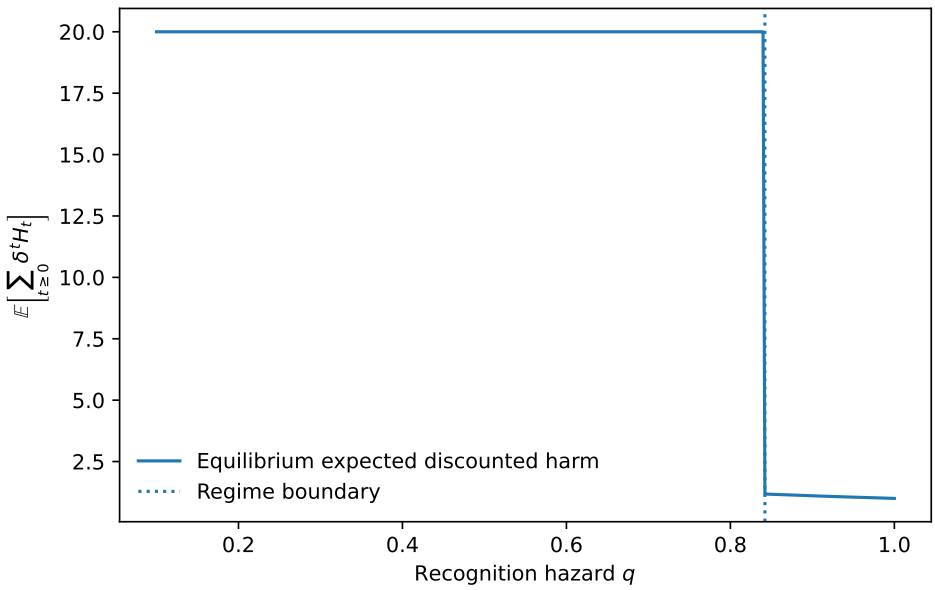


Figure 2: Equilibrium expected discounted harm as a function of recognition hazard. Notes: parameters as in Figure 1. The vertical line indicates the regime boundary where $k = k^*(q)$.

until recognition, then complies. In this parameterization, discounted harm collapses from $1/(1-\delta) = 20$ under persistent redesign to $1/[1-\delta(1-q^A)] = 1.09$ under harm-until-recognition.

Figure 1 plots the redesign threshold $k^*(q) = a_0/[1-\delta(1-q)]$. The threshold declines in q for a simple reason: faster recognition shortens the expected duration of the high-payoff phase after a reset. Equivalently, when recognition is fast, paying k to return to the novel state does not appreciably delay the moment at which the firm must stop to avoid sanctions.

Figure 2 converts this threshold into a piecewise harm schedule. In the redesign region ($k < k^*(q)$), equilibrium harm is pinned at $1/(1-\delta)$ and is locally insensitive to recognition speed because the firm responds to recognition by resetting the state. Once q crosses the boundary $k = k^*(q)$, a marginal increase in q reduces harm, because the firm switches to the harm-until-recognition regime.

6.3 Redesign frequency and workload implications

The baseline model also distinguishes between two conceptually different objects: the level of harm and the frequency with which recognition events trigger redesign. In the persistent-redesign equilibrium, harm occurs in every period, but recognition events still arrive with hazard q, and each recognition induces a redesign payment k. Thus, recognition improvements that do not deter redesign can raise the workload associated with processing novel variants

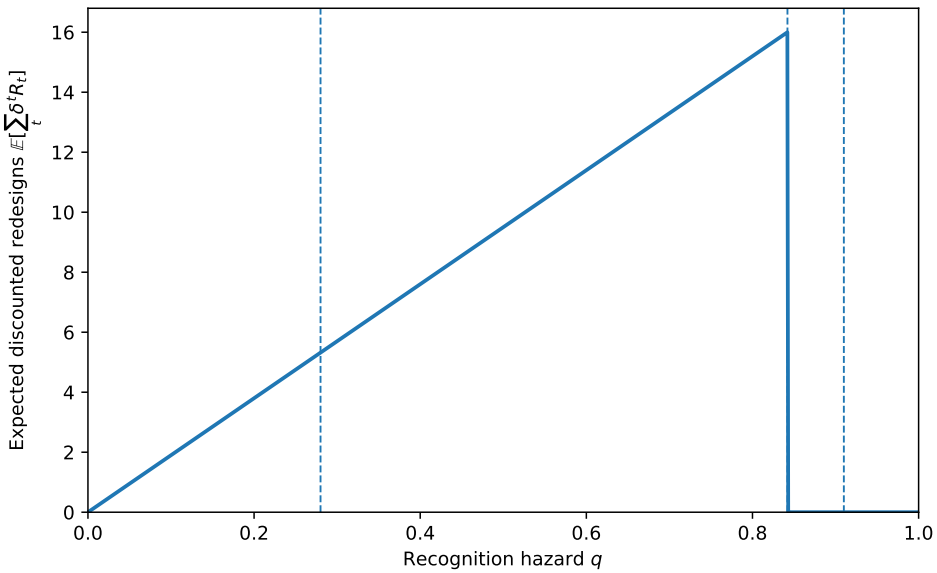


Figure 3: Expected discounted redesign frequency. Notes: The plot uses the parameterization from Table 1 and sets $R_t = 1$ when the firm chooses R in period t. For $q < \bar{q}$ the equilibrium is persistent redesign and (18) applies. For $q \geq \bar{q}$ redesign is deterred, so the expected discounted redesign frequency is zero. The vertical lines indicate $q^L = 0.28$ and $q^A = 0.91$.

even when they leave equilibrium harm unchanged. Corollary 4 formalizes this point: when the regulator incurs a variable cost c_{rec} per recognition event, the regulator's payoff is strictly decreasing in q within the persistent-redesign regime, making below-threshold monitoring investments actively counterproductive.

Let $R_t \in \{0, 1\}$ be the indicator that the firm chooses R in period t. In the persistent-redesign equilibrium, the firm redesigns if and only if $s_t = 1$. Because the transition from either state leads to $s_{t+1} = 1$ with probability q and to $s_{t+1} = 0$ with probability $1 - q$, we have $\Pr(s_t = 1) = q$ for all $t \geq 1$. Therefore, starting from $s_0 = 0$,

$$\mathbb{E}\Big[\sum_{t=0}^{\infty} \delta^t R_t\Big] = \sum_{t=1}^{\infty} \delta^t \Pr(s_t = 1) = \frac{\delta q}{1 - \delta} \qquad \text{(persistent redesign)}. \tag{18}$$

In the harm-until-recognition and full-compliance equilibria, $R_t = 0$ for all t, so the expression equals 0. Figure 3 plots (18) as a function of q under the parameterization in Table 1. The redesign workload rises linearly in q within the redesign region and then collapses to zero once recognition becomes fast enough to deter redesign. This discontinuity complements the harm discontinuity highlighted above: even if harm is flat below the boundary, the equilibrium volume of adaptations and recognition events is not.

Figure 4 illustrates the regulator's caseload burden implied by this nonlinearity. When

the regulator incurs variable recognition costs ($c_{\text{rec}} > 0$), the regulator's discounted loss is *increasing* in q below $\bar{q}$: each incremental improvement in recognition speed generates additional wasted case-opening effort without reducing harm. The loss then drops discontinuously at $\bar{q}$ and declines above it as faster recognition shortens harm spells. An agency that starts below the boundary should therefore either invest enough to cross it or not invest at all; intermediate upgrades are not merely ineffective but actively counterproductive.

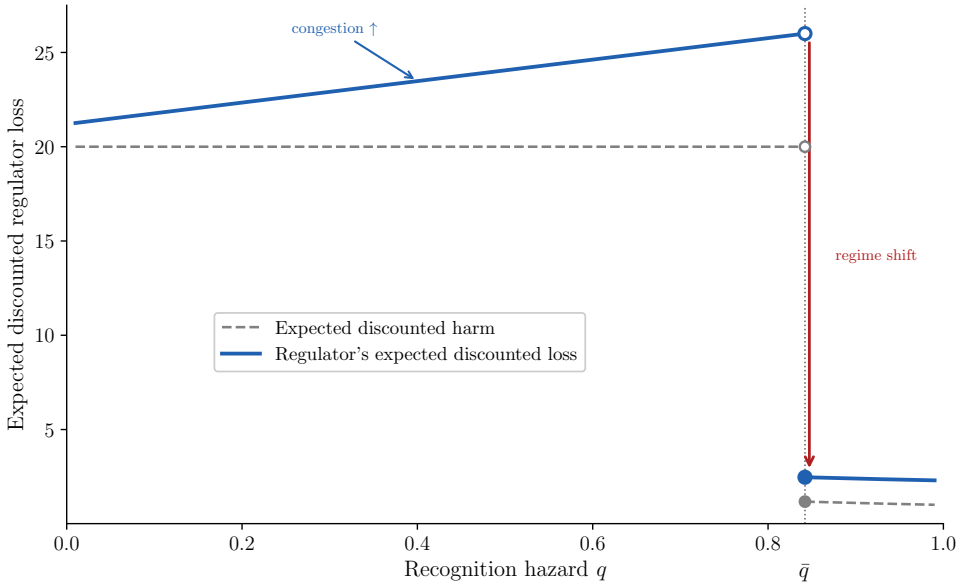


Figure 4: Regulator's discounted loss as a function of recognition hazard, with variable congestion costs. Notes: The plot uses the parameterization from Table 1, with $c_{\text{rec}} = 0.3$ and $c_{\text{det}} = 0.1$. The dashed line shows $\bar{q}$ from (16) with $a_0 = 0.85$, $k = 1$, and $\delta = 0.95$. For $q < \bar{q}$, the regulator's loss increases in q (Corollary 4): recognition events trigger costly case-opening on patterns the firm has already abandoned. Once $q \geq \bar{q}$, loss falls with q as faster recognition shortens harm spells.

Note that, in our simulation, the marginal caseload increase driven by better monitoring technology depends on c_{rec} but not c_{det}, because persistent redesign resets the state to zero every period upon recognition but keeps the risk of immediate detection at p_0, which we assume to be the same under legacy and AI-augmented monitoring technologies in our simulation. If AI-augmented monitoring also increases p_0, it will increase the regulator's caseload in detection as well.

7 Extensions

This section develops two extensions: the first introduces heterogeneity in adaptation costs, which generates a selection effect. We briefly discuss how the selection can result in a

multi-jurisdiction enforcement dilemma. The second extension assumes human review is required for AI-recognized violations but the regulator faces a capacity constraint in human review. We discuss how the AI triage reinforces the congestion effect.

7.1 Heterogeneous Redesign Costs and Selection Effects

The baseline model treats the redesign cost k as a fixed parameter, leaving no room for cross-sectional variation in how firms respond to enforcement. In practice, however, the ability to iterate around enforcement differs substantially across firms. A well-resourced technology company with dedicated engineering teams can redesign algorithmic practices at low marginal cost, whereas a smaller firm or one with less technical capacity may find redesign prohibitively expensive. This heterogeneity has consequences both for predicting which firms will engage in persistent cat-and-mouse behavior and for evaluating the aggregate welfare effects of improved monitoring.

This subsection introduces a distribution of redesign costs to explore these selection effects. The analysis reveals that the redesign threshold derived in Proposition 1 serves as a selection cutoff: firms with low adaptation costs persistently redesign, while those with high adaptation costs accept compliance once their conduct is recognized. Improvements in monitoring technology reduce aggregate harm through two channels. The direct channel shortens the expected duration of harmful conduct for firms that do not redesign. The indirect channel shrinks the mass of firms for which redesign is privately profitable, converting would-be redesigners into temporary offenders.

Consider a unit mass of firms indexed by their redesign cost k, drawn from a distribution G with support on $\mathbb{R}_+$. Each firm faces the dynamic problem analyzed in Section 4 and plays its type-contingent Markov perfect policy. We maintain Assumption 1 and the condition $a_0^\tau > 0$, so that novel harmful conduct is privately attractive regardless of type.

Proposition 5 (Selection into persistent redesign)**.** *Fix a monitoring technology τ and suppose Assumption 1 holds with $a_0^\tau > 0$. In Markov perfect equilibrium there exists a cutoff $\hat{k}(\tau) = k^*(\tau)$ such that:*

1. *Firms with $k < \hat{k}(\tau)$ choose H in state 0 and R in state 1 (persistent redesign).*
2. *Firms with $k \geq \hat{k}(\tau)$ choose H in state 0 and C in state 1 (harm until recognition).*

Aggregate expected discounted harm (per firm) starting from $s_0 = 0$ is

$$\bar{H}(\tau) = \frac{G(\hat{k}(\tau))}{1-\delta} + \frac{1 - G(\hat{k}(\tau))}{1 - \delta(1-q^\tau)}. \tag{19}$$

Moreover, $\hat{k}(\tau)$ is decreasing in both p_0^τ and α^τ, so improvements in recognition reduce aggregate harm both by shortening harm spells for high-k firms and by shrinking the mass of redesigning firms.

The proof is in Appendix B. The proposition has several implications worth emphasizing.

First, the selection prediction offers testable content. If one could observe firms' adaptation costs (or proxies such as engineering headcount, patenting intensity, or past product iteration speed), the model predicts that low-cost firms are more likely to exhibit repeated compliance-violation cycles, while high-cost firms are more likely to exhibit a single violation episode followed by sustained compliance. Evidence of such sorting would support the model's mechanism.

Second, the aggregate harm expression in (19) decomposes neatly into contributions from the two equilibrium regimes. The first term, $G(\hat{k}(\tau))/(1-\delta)$, is the discounted harm from persistent redesigners, who cause harm in every period. The second term, $[1-G(\hat{k}(\tau))]/[1-\delta(1-q^\tau)]$, is the discounted harm from firms in the harm-until-recognition regime, whose harm episodes are geometrically distributed with parameter q^τ. Monitoring improvements affect both terms: a higher q^τ directly reduces the second term by speeding recognition, and it indirectly reduces the first term by lowering $\hat{k}(\tau)$ and thereby shrinking the mass $G(\hat{k}(\tau))$ of persistent redesigners.

Third, the relative importance of the direct and indirect channels depends on the distribution G. When G places substantial mass near the cutoff $\hat{k}(\tau)$, small improvements in monitoring can generate large reductions in aggregate harm by converting marginal redesigners into temporary offenders. When G is concentrated at the extremes, the selection margin is less important, and harm reduction comes primarily from accelerating recognition for non-redesigning firms. This observation suggests that the value of monitoring investments depends on the industrial organization of the regulated population.

Remark 7 (Distributional assumptions and risk-based monitoring)**.** *For concreteness, one might assume G is log-normal or Pareto. The specific functional form affects the quantitative predictions but not the qualitative logic. A natural extension would endogenize the distribution by allowing firms to invest in lowering their own redesign costs, which would introduce an additional strategic margin and potentially generate multiplicity in the cross-sectional distribution of adaptation capabilities. We leave this richer interaction for future work.*

The heterogeneity result also suggests that regulators could tailor monitoring intensity to firm characteristics. Since the redesign threshold $\hat{k}(\tau)$ depends on recognition speed, firms with lower adaptation costs require faster recognition to deter persistent redesign. If agencies can observe proxies for k, such as engineering headcount, product iteration frequency, or

past compliance history, they could allocate monitoring resources accordingly, concentrating faster and more intensive oversight on firms with greater adaptation capability. This form of risk-based supervision would improve the efficiency of scarce enforcement resources by targeting investments where they are most likely to cross the regime boundary.

Remark 8 (Multi-jurisdiction enforcement dilemma)**.** *When firms can relocate across jurisdictions, an additional strategic tension arises: even if AI-augmented monitoring is locally welfare-improving, jurisdictions may rationally choose weaker monitoring to retain mobile economic activity. In the spirit of regulatory competition models (Tiebout, 1956; Oates and Schwab, 1988), suppose two jurisdictions simultaneously choose monitoring technologies and a partially mobile firm locates in the jurisdiction with weaker monitoring. If the hosting benefit from attracting the firm is large enough relative to the welfare gain from AI-augmented monitoring, each jurisdiction has a dominant strategy to choose legacy monitoring, even though both would be better off if both adopted AI. The formal result (Proposition 7 in Appendix B.11) shows that this enforcement dilemma is more severe when firm mobility is higher. AI-augmented tools can mitigate the dilemma by increasing the welfare gap between AI-augmented and legacy monitoring (for example, by reducing harm through faster recognition or by lowering operational costs through workflow automation), but the fundamental tension between local welfare maximization and competitive attraction of mobile firms remains.*

7.2 Human Review Capacity and AI Triage

The baseline model abstracts from the fact that, in practice, agency action is bottlenecked by human review capacity. AI-augmented monitoring generates learning signals that require human assessment before they translate into recognition and enforcement action. When the volume of signals exceeds the regulator's review capacity, the effective learning speed declines, slowing recognition and expanding the parameter region in which persistent redesign is an equilibrium. This subsection introduces a capacity constraint on human review that operates through the learning speed α of the baseline model, preserving the dynamic Markov structure, the redesign action, and the pure-strategy equilibrium characterization.[2] The extension builds on the heterogeneous-cost framework of Section 7.1 and captures a source of congestion that complements the variable recognition cost introduced in Corollary 4.

[2]We apply the capacity constraint to the learning channel α rather than to the immediate detection probability p_0. As discussed in Remark 1, p_0 captures detection through existing enforcement protocols, while α captures the slower process of developing recognition capability from accumulated evidence, which is the channel that AI augments and that depends on human review of AI-generated signals.

Consider the unit mass of firms with redesign costs drawn from G as in Section 7.1. The regulator deploys AI-augmented monitoring with detection probability $p_0 \equiv p_0^A$ and baseline learning speed $\alpha \equiv \alpha^A$. Each period, the AI monitoring tool generates a binary flag $z \in \{0, 1\}$ for each firm:

$$\Pr(z = 1 \mid \text{harmful conduct}) = \rho, \qquad \Pr(z = 1 \mid \text{compliant}) = \phi, \tag{20}$$

where $\rho \in (0, 1]$ is the sensitivity (true positive rate) and $\phi \in [0, 1)$ is the false positive rate. We interpret a more trustworthy tool as one with higher ρ for a given ϕ, or lower ϕ for a given ρ.

The regulator can review at most M flagged cases per period. Recognition of a novel harmful pattern requires that the learning signal generated by the firm's conduct receives human review, because human analysts must evaluate the evidence and determine whether flagged conduct constitutes a recognizable violation pattern.

Assumption 2 (Scarce review capacity)**.** $0 < \phi < M < \rho \leq 1$.

Assumption 2 captures the empirically relevant case in which the review pipeline can handle the flag volume when violations are rare ($m(0) = \phi < M$), but becomes congested if all firms violate ($m(1) = \rho > M$).

Definition 1 (Effective recognition hazard under capacity constraint)**.** *Given a fraction $x \in [0, 1]$ of firms deploying harmful conduct, the mass of AI-generated flags is $m(x) = x\rho + (1 - x)\phi$. The human confirmation probability is*

$$\pi(x) \equiv \min\left\{1,\ \frac{M}{x\rho + (1 - x)\phi}\right\}. \tag{21}$$

The effective learning speed is $\tilde{\alpha}(x) \equiv \alpha \cdot \pi(x)$, and the effective recognition hazard for a novel pattern is

$$\tilde{q}(x) \equiv p_0 + (1 - p_0)\,\alpha\,\pi(x). \tag{22}$$

When $m(x) \leq M$, all flags are reviewed, $\pi(x) = 1$, and $\tilde{q}(x) = q$, the unconstrained baseline. When $m(x) > M$, capacity binds, $\pi(x) < 1$, and $\tilde{q}(x) < q$. Because the capacity constraint modifies only the recognition hazard, replacing q with $\tilde{q}(x)$, while preserving the baseline's dynamic structure, the equilibrium characterization of Proposition 1 applies with $\tilde{q}(x)$ in place of q. Each firm, taking the aggregate violation rate x as given, plays a stationary Markov strategy determined by the capacity-constrained redesign threshold.

Proposition 6 (Capacity-constrained recognition)**.** *Maintain Assumption 1 and suppose $a_0 > 0$. Fix the capacity parameters (ρ, ϕ, M) satisfying Assumption 2. Given a fraction $x \in [0,1]$ of firms deploying harmful conduct, define the capacity-constrained redesign threshold*

$$\tilde{k}^*(x) \equiv \frac{a_0}{1 - \delta\big(1 - \tilde{q}(x)\big)}. \tag{23}$$

Then:

1. *$\tilde{k}^*(x)$ is weakly decreasing in $\pi(x)$, with strict inequality when $p_0 < 1$ and $\alpha > 0$.*
2. *When $m(x) \leq M$: $\tilde{k}^*(x) = k^*_{\text{baseline}}$.*
3. *When $m(x) > M$: $\partial\tilde{k}^*/\partial\phi \geq 0$ (strictly for $x < 1$), $\partial\tilde{k}^*/\partial M < 0$, $\partial\tilde{k}^*/\partial x > 0$.*
4. *Expected discounted harm in the harm-until-recognition regime, $1/[1 - \delta(1 - \tilde{q}(x))]$, is decreasing in $\pi(x)$.*

The proposition establishes that capacity constraints on human review amplify the conditions for persistent redesign. The mechanism operates through a simple chain: limited review capacity reduces the probability that AI-generated learning signals are confirmed by human analysts, which slows the effective recognition hazard, which raises the redesign threshold, which expands the set of firms for which persistent redesign is privately optimal. False positives contribute to this congestion: each compliant firm incorrectly flagged consumes review resources that would otherwise accelerate recognition of genuine violations. Even within the harm-until-recognition regime, where firms do not redesign, the capacity constraint lengthens expected harm spells by slowing the transition from the novel state to the recognized state.

The comparative statics in Part (3) highlight the role of triage quality. A higher false positive rate ϕ worsens congestion by inflating the flag pool, while greater review capacity M relieves it. A higher aggregate violation rate x intensifies congestion, creating a strategic externality: each firm's decision to deploy harmful conduct reduces recognition speed for all other firms.

Equilibrium violation rate with heterogeneous firms. Under the heterogeneous-cost framework of Section 7.1, each firm plays the pure-strategy Markov policy characterized in Proposition 1: a firm with redesign cost k chooses persistent redesign if $k < \tilde{k}^*(x)$ and harm until recognition otherwise. The steady-state fraction of firms deploying harmful conduct is therefore

$$x = G\big(\tilde{k}^*(x)\big), \tag{24}$$

where G is the distribution of redesign costs introduced in Section 7.1. An equilibrium is a fixed point x^* of the mapping $x \mapsto G(\tilde{k}^*(x))$.

Because $\tilde{k}^*(x)$ is weakly increasing in x when capacity binds (Proposition 6, Part 3), the composite mapping $x \mapsto G(\tilde{k}^*(x))$ is weakly increasing in x. Since it maps $[0,1]$ into $[0,1]$ and is continuous, at least one fixed point exists by the intermediate value theorem.[3] The next result shows that multiple fixed points can arise, reflecting the self-reinforcing nature of congestion.

Corollary 5 (Self-reinforcing congestion)**.** *Under the conditions of Proposition 6, define*

$$k_L^* \equiv \tilde{k}^*(0) = \frac{a_0}{1-\delta(1-q)}, \qquad k_H^* \equiv \tilde{k}^*(1) = \frac{a_0}{1-\delta\left(1-p_0-(1-p_0)\alpha\,\dfrac{M}{\rho}\right)}. \tag{25}$$

Then $k_H^* \geq k_L^*$, with strict inequality when $p_0 < 1$ and $\alpha > 0$.

1. *When capacity does not bind at $x_L \equiv G(k_L^*)$ (i.e., $m(x_L) \leq M$), $x_L^* = G(k_L^*)$ is a fixed point of* (24)*.*

2. *If $G(k_H^*) = 1$, then $x = 1$ is also a fixed point of* (24)*, with capacity binding. More generally, whenever a fixed point $x_H^* > x_L^*$ exists in the capacity-binding region, it satisfies $\tilde{k}^*(x_H^*) \in (k_L^*,\, k_H^*]$.*

Corollary 5 establishes that congestion is self-reinforcing: widespread violation reduces effective recognition, which raises the redesign threshold, making redesign attractive for a larger mass of firms and sustaining the high-violation state. The multiplicity arises from a strategic complementarity operating through the capacity constraint. When the aggregate violation rate is low, capacity is non-binding, recognition is fast, and the redesign threshold coincides with the baseline. When the violation rate is high, capacity binds, recognition slows, and the threshold rises, drawing additional firms into persistent redesign.[4] The regulator strictly prefers the low-violation fixed point: at x_L^*, fewer firms engage in persistent redesign, and recognition is faster for those in the harm-until-recognition regime.

The capacity constraint also affects the minimum recognition speed required to deter redesign. From Corollary 2, redesign is deterred for a firm with cost k if and only if the

[3]The mapping is continuous because $\pi(x)$ is continuous, $\tilde{q}$ is continuous in π, $\tilde{k}^*$ is continuous in $\tilde{q}$, and G is assumed continuous. When the mapping is weakly increasing, existence also follows from Tarski's fixed-point theorem.

[4]Starting from $s_0 = 0$ for all firms, the initial violation rate is $x = 1$, so capacity binds during the transient even in the harm-until-recognition regime. For firms with $k \geq k_H^*$, redesign is unprofitable at every level of congestion, so the transient does not affect regime selection. For firms with $k \in [k_L^*, k_H^*)$, expectations about the path of congestion determine which fixed point is selected, consistent with the self-reinforcing logic.

effective recognition hazard exceeds $\underline{q}(k) \equiv [a_0/k - (1-\delta)]/\delta$. Under the capacity constraint, achieving $\tilde{q}(x) \geq \underline{q}(k)$ requires a sufficient human confirmation probability.

Corollary 6 (Minimum confirmation probability for deterrence)**.** *Suppose $\underline{q}(k) > p_0$ (so that the learning channel is needed for deterrence) and $q \geq \underline{q}(k)$ (so that deterrence is feasible absent the capacity constraint). Redesign is deterred for a firm with cost k if and only if $\pi(x) \geq \underline{\pi}(k)$, where*

$$\underline{\pi}(k) \equiv \frac{\underline{q}(k) - p_0}{(1-p_0)\alpha} \in (0, 1]. \tag{26}$$

When capacity binds, this condition translates to $M \geq [x\rho + (1-x)\phi] \cdot \underline{\pi}(k)$.

Even when the regulator's AI monitoring achieves $q \geq \underline{q}(k)$ absent capacity constraints, the capacity bottleneck can prevent deterrence if the confirmation probability falls below $\underline{\pi}(k)$. False positives are costly here not because they generate sanctions against innocent firms (human review uncovers the truth), but because they consume scarce review resources that would otherwise accelerate learning about genuine violations.

Remark 9 (Two sources of enforcement congestion)**.** *The baseline model (Corollary 4) and this extension capture complementary sources of congestion. In the baseline, congestion arises from the firm's adaptive behavior: each recognition event in the persistent-redesign regime triggers costly case-opening effort (c_{rec}) on a pattern the firm has already abandoned via redesign, so the regulator's investigative resources are wasted on obsolete conduct. Here, congestion arises from the AI tool's imprecision: false positives consume scarce review capacity, crowding out investigation of genuine violations and slowing recognition. The two mechanisms reinforce each other: capacity-constrained recognition raises $\tilde{k}^*$, making persistent redesign more likely, which in turn generates more wasted case-opening effort.*

Remark 10 (Policy levers for capacity-driven congestion)**.** *Three levers mitigate capacity-driven congestion. First, expanding review capacity M directly raises π and, by narrowing the gap between k_L^* and k_H^*, reduces the mass of firms whose equilibrium regime depends on the aggregate violation rate. Second, reducing the false positive rate ϕ frees capacity for genuine violations by shrinking the flag volume from compliant firms. Third, improving sensitivity ρ has competing effects. Higher ρ does not enter the recognition hazard $\tilde{q}$ directly through the learning speed α, but it enters $\tilde{q}$ indirectly through $\pi(x)$: because ρ appears in the denominator of $\pi(x) = M/[x\rho + (1-x)\phi]$, a higher ρ increases total flag volume when violations are prevalent, reducing π and thereby slowing effective recognition. At the same time, higher ρ shifts the composition of flags toward true positives, which is valuable when the review pipeline can process the additional volume. The net effect on the equilibrium fixed point x^* is therefore*

ambiguous and depends on the level of M relative to flag volume. The practical implication is that deploying a more sensitive AI tool generates more actionable flags but also demands more review capacity; without commensurate investment in human reviewers, increased sensitivity can slow recognition by congesting the review pipeline.

8 Conclusion

This paper develops a novel, dynamic inspection game in which a firm can redesign harmful conduct to evade recognition, while a regulator can invest in monitoring technology that accelerates the rate at which novel patterns become detectable. The framework applies broadly across regulatory domains (antitrust, consumer protection, environmental regulation, financial supervision, labor enforcement) wherever firms can iteratively modify conduct and penalties are bounded. The core tension arises from a sanction cap: when formal penalties are bounded by legal or institutional constraints, traditional deterrence theory offers limited guidance, and enforcement efficacy depends critically on the speed of learning relative to the costs of firm adaptation.

The equilibrium analysis delivers a sharp trichotomy. When expected gains from novel harm are negative, the firm complies. When expected gains are positive but redesign is costly, the firm harms temporarily until its conduct is recognized, then stops. When expected gains are positive and redesign is cheap, the firm enters a persistent cat-and-mouse regime in which it redesigns upon each recognition event, generating harm in every period. The boundary between the second and third regimes is governed by a threshold that depends on the recognition hazard, the discount factor, and the expected flow payoff from novel harm. Faster recognition shrinks the redesign region by lowering this threshold.

We derive three main results with practical implications for enforcement agencies. First, agencies should evaluate AI monitoring investments with a *regime-shifting mindset*: monitoring investments have discontinuous returns. Marginal improvements in recognition speed that leave the firm in the redesign regime yield no deterrence benefit because equilibrium harm is flat in recognition speed throughout that region. Only investments that cross the redesign boundary generate welfare gains. Second, below-threshold AI adoption can be *counterproductive*: incremental monitoring improvements that fall short of the regime-shifting threshold increase enforcement workload (more frequent recognition events trigger more redesign cycles) without reducing equilibrium harm. Third, this problem can intensify as AI lowers the firm's adaptation cost, creating an arms-race dynamic that favors early investment by the regulator.

We examine extensions on heterogeneous adaptation costs and capacity-constrained triage.

The heterogeneity extension reveals that the redesign threshold operates as a selection cutoff across firms with different adaptation costs, generating testable predictions about which types of firms are most likely to engage in persistent evasion. If firms can shift activity across jurisdictions, monitoring incentives weaken, potentially trapping jurisdictions in under-enforcement. When human review is required for AI-recognized violations but human review capacity is the binding constraint, AI tools with high false-positive rates can flood the review pipeline and dilute effective detection, reinforcing the congestion effect identified in the baseline model.

Several avenues merit further exploration. The model treats the firm's redesign technology as exogenous, but in practice adaptation costs may be endogenous to prior investment in engineering capabilities or to the structure of the underlying algorithmic system. Allowing firms to invest in lowering their own k would introduce an additional strategic margin and could generate path dependence in the distribution of firm types. A related extension would model the regulator's own learning process more explicitly, allowing for feedback between enforcement actions and the accumulation of institutional knowledge. The model also abstracts from the dynamics of enforcement agency budgets, which may themselves respond to observed harm or to lobbying by regulated firms. Budgetary constraints can prevent agencies from reaching the regime-shifting threshold, creating a case for concentrated rather than incremental investment. A fuller treatment of dynamic budgetary constraints and multi-jurisdiction coordination could also connect more directly to the institutional realities that shape enforcement in practice.

Our central message is that adaptive enforcement is not simply a matter of deploying more technology. The value of AI-augmented monitoring depends on whether it is sufficient to cross the threshold that separates persistent misconduct from temporary misconduct. Investments that fall short of this threshold may be not merely ineffective but counterproductive, burning agency resources without reducing harm. For regulators facing fast-moving markets in which firms can iteratively redesign conduct, understanding where the threshold lies, and what it takes to cross it, is a step toward effective deterrence.

References

Mastrobuoni, Giovanni. 2020. "Crime is Terribly Revealing: Information Technology and Police Productivity", *The Review of Economic Studies* 87(6): 2727—2753.

Supreme Court of the United States. 2021. *AMG Capital Management, LLC v. Federal Trade Commission*, 593 U.S. ___ (2021), No. 19–508, decided April 22.

Andreoni, James. 1991. "Reasonable Doubt and the Optimal Magnitude of Fines: Should the Penalty Fit the Crime?" *RAND Journal of Economics* 22(3): 385–395.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. 2024. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Journal of Political Economy* 132(3): 723–771.

Athey, Susan. 2018. "The Impact of Machine Learning on Economics." In *The Economics of Artificial Intelligence: An Agenda*, ed. Ajay Agrawal, Joshua Gans, and Avi Goldfarb, 507–547. Chicago: University of Chicago Press.

Avenhaus, Rudolf, Bernhard von Stengel, and Shmuel Zamir. 2002. "Inspection Games." In *Handbook of Game Theory with Economic Applications*, ed. Robert J. Aumann and Sergiu Hart, vol. 3, 1947–1987. Amsterdam: Elsevier.

Axon. 2025. *2025 AI in Law Enforcement Trends Report*. Scottsdale, AZ: Axon Enterprise, Inc.

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

Brown, Zach Y., and Alexander MacKay. 2023. "Competition in Pricing Algorithms." *American Economic Journal: Microeconomics* 15(2): 109–156.

Brown, Zach Y., and Alexander MacKay. 2025. "Algorithmic Coercion with Faster Pricing." Working paper.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. 2020. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review* 110(10): 3267–3297.

Duso, Tomaso, Joseph E. Harrington Jr., Carl Kreuzberg, and Geza Sapi. 2025. "Public Communication and Collusion: New Screening Tools for Competition Authorities." CEPR Discussion Paper 20490 / CESifo Working Paper 12029.

The Economist. 2025. "Why Britain's Police Forces Are Taking to AI." December 6.

Federal Trade Commission. 2023. *FTC v. Amazon.com, Inc.: Complaint for Permanent Injunction and Other Relief*. Case No. 2:23-cv-01495, W.D. Wash., filed September 26.

Garicano, Luis, and Paul Heaton. 2010. "Information Technology, Organization, and Productivity in the Public Sector: Evidence from Police Departments." *Journal of Labor Economics* 28(1): 167–201.

Hardt, Moritz, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. 2016. "Strategic Classification." In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, 111–122. New York: Association for Computing Machinery.

Harrington, Winston. 1988. "Enforcement Leverage When Penalties Are Restricted." *Journal of Public Economics* 37(1): 29–53.

Hoberg, Gerard, and Gordon Phillips. 2010. "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis." *Review of Financial Studies* 23(10): 3773–3811.

Hoberg, Gerard, and Gordon Phillips. 2016. "Text-Based Network Industries and Endogenous Product Differentiation." *Journal of Political Economy* 124(5): 1423–1465.

Kaplow, Louis, and Steven Shavell. 1994. "Accuracy in the Determination of Liability." *Journal of Law and Economics* 37(1): 1–15.

Kroll, Joshua A., Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. 2017. "Accountable Algorithms." *University of Pennsylvania Law Review* 165(3): 633–705.

Larson, Heidi. 2025. "The Outlook for Vaccines Is Stormy." *The Economist: The World Ahead 2026*, November 12.

Lebovitz, Sarah, Natalia Levina, and Hila Lifshitz-Assaf. 2021. "Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts' Know-What." *MIS Quarterly* 45(3): 1501–1526.

Malik, Arun S. 1993. "Self-Reporting and the Design of Policies for Regulating Stochastic Pollution." *Journal of Environmental Economics and Management* 24(3): 241–257.

Mookherjee, Dilip, and I. P. L. Png. 1992. "Monitoring vis-a-vis Investigation in Enforcement of Law." *American Economic Review* 82(3): 556–565.

Oates, Wallace E., and Robert M. Schwab. 1988. "Economic Competition Among Jurisdictions: Efficiency Enhancing or Distortion Inducing?" *Journal of Public Economics* 35(3): 333–354.

Ortner, Juan, Sylvain Chassang, Kei Kawai, and Jun Nakabayashi. 2022. "Screening Adaptive Cartels." Working Paper 300, Princeton University, Center for Economic Policy Studies.

Polinsky, A. Mitchell, and Steven Shavell. 2000. "The Economic Theory of Public Enforcement of Law." *Journal of Economic Literature* 38(1): 45–76.

Porter, Robert H., and J. Douglas Zona. 1993. "Detection of Bid Rigging in Procurement Auctions." *Journal of Political Economy* 101(3): 518–538.

Public First and Center for Data Innovation (ITIF). 2026. *Global Public Sector AI Adoption Index*.

Puterman, Martin L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York: John Wiley & Sons.

Rochet, Jean-Charles, and Jean Tirole. 2006. "Two-Sided Markets: A Progress Report." *RAND Journal of Economics* 37(3): 645–667.

Sandvig, Christian, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms." Paper presented at the 64th Annual Meeting of the International Communication Association, Seattle, WA.

Shavell, Steven. 1987. "The Optimal Use of Nonmonetary Sanctions as a Deterrent." *American Economic Review* 77(4): 584–592.

Stigler, George J. 1974. "The Optimum Enforcement of Laws." In *Essays in the Economics of Crime and Punishment*, ed. Gary S. Becker and William M. Landes, 55–67. Cambridge, MA: National Bureau of Economic Research.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* 64(5): 416–424.

Appendix

A Notation

Table 3 summarizes the main notation used in the paper.

Symbol	Meaning
t	Period index, $t \in \{0, 1, 2, \dots\}$.
δ	Common discount factor, $\delta \in (0, 1)$.
s_t	Recognition state in period t; $s_t = 0$ means novel, $s_t = 1$ means recognized.
τ	Monitoring technology, $\tau \in \{L, A\}$ where L is legacy and A is AI-augmented.
p_s^τ	Detection probability in state $s \in \{0, 1\}$ under technology τ.
α^τ	Learning probability for a novel pattern under technology τ.
q^τ	Recognition hazard, $q^\tau = p_0^\tau + (1 - p_0^\tau)\alpha^\tau$.
c^τ	Per-period operational cost of technology τ, with $c^A > c^L$.
a_t	Firm action in period t; $a_t \in \{C, H\}$ if $s_t = 0$, $a_t \in \{C, H, R\}$ if $s_t = 1$.
C	Comply; generates payoff 0 and no consumer harm.
H	Harm; generates payoff $g - \mathbf{1}\{\text{detected}\}F$ and consumer harm h.
R	Adapt & Harm; leads to payoff $g - \mathbf{1}\{\text{detected}\}F - k$ and consumer harm h.
g	Firm's private gain from harmful conduct.
h	Consumer harm from one period of harmful conduct.
F	Sanction if detected, bounded by legal and institutional constraints.
k	Firm's redesign cost (resource cost of resetting recognition).
a_s^τ	Expected flow payoff from harm in state s: $a_s^\tau = g - p_s^\tau F$.
V_s^τ	Firm's equilibrium continuation value in state s under technology τ.
$k^*(\tau)$	Redesign threshold: $k^*(\tau) = a_0^\tau / [1 - \delta(1 - q^\tau)]$.
$\underline{q}(\tau)$	Minimum recognition speed to deter redesign: $\underline{q}(\tau) = (a_0^\tau/k - (1 - \delta))/\delta$.
$\bar{F}^{\text{comp}}(\tau)$	Minimal sanction for full compliance: $F^{\text{comp}}(\tau) = g/p_0^\tau$.
$F^{\text{no-redesign}}(\tau)$	Minimal sanction to deter redesign: $F^{\text{no-redesign}}(\tau) = (g - k[1-\delta(1-q^\tau)])/p_0^\tau$.
W^τ	Regulator's expected discounted payoff under technology τ.
$\bar{q}$	Redesign boundary in the continuous-investment problem.
$\kappa(q)$	Per-period operational cost as a function of recognition hazard q.
Triage, Heterogeneity, and Two-Jurisdiction Extensions (Section 7)	
ρ	True positive rate of the triage tool.
ϕ	False positive rate of the triage tool.
M	Human review capacity (maximum mass of firms that can be investigated).
x	Population violation rate.
$m(x)$	Mass of flagged firms: $m(x) = x\rho + (1 - x)\phi$.
$D(x)$	Detection probability for a violator given population rate x.
x^m	Interior mixed-strategy equilibrium violation rate.
G	Distribution of redesign costs across firms.
$\hat{k}(\tau)$	Selection cutoff for redesign: $\hat{k}(\tau) = k^*(\tau)$.
$\bar{H}(\tau)$	Aggregate expected discounted harm per firm.
j	Jurisdiction index, $j \in \{1, 2\}$.
μ	Probability that the firm is mobile.
b	Per-period hosting benefit (flow).
B	Present discounted value of hosting benefit: $B = b/(1 - \delta)$.
$w(\tau)$	Welfare from hosting under technology τ: $w(\tau) = B - c^\tau/(1 - \delta) - hH(\tau)$.

Table 3: Summary of notation

B Proofs

B.1 Proof of Lemma 1

The firm's optimization problem after the regulator commits to technology τ is a discounted Markov decision problem with finite state space $\{0, 1\}$ and finite action sets. For any discount factor $\delta \in (0, 1)$ and bounded per-period payoffs, the Bellman operator is a contraction with modulus δ under the sup-norm. By the Banach fixed-point theorem, there exists a unique fixed point (V_0^τ, V_1^τ). Standard results in Puterman (1994) guarantee that there exists a stationary deterministic policy that attains this fixed point. □

B.2 Proof of Proposition 1

Fix technology τ and suppose Assumption 1 holds, i.e., $a_1^\tau = g - p_1^\tau F \leq 0$.

Step 1: Dominated actions in state 1. In state $s = 1$, if the firm chooses H, it receives $a_1^\tau + \delta V_1^\tau \leq \delta V_1^\tau$ (with equality only when $a_1^\tau = 0$). Choosing C yields δV_1^τ. Thus H is weakly dominated by C in state 1, and we can restrict attention to $\{C, R\}$ in that state.

Step 2: Full compliance case. If $a_0^\tau \leq 0$, then in state 0, choosing H yields $a_0^\tau + \delta[(1 - q^\tau)V_0^\tau + q^\tau V_1^\tau]$. Choosing C yields δV_0^τ. Suppose by way of contradiction that $V_0^\tau > 0$ or $V_1^\tau > 0$. Then the firm must be choosing H or R in some state with positive value. But $a_0^\tau \leq 0$ and $a_1^\tau \leq 0$ imply that any harmful action yields non-positive flow payoff minus any adaptation cost. Hence $V_0^\tau = V_1^\tau = 0$, and C is optimal in both states.

Step 3: State-0 behavior when $a_0^\tau > 0$. In state 0, H dominates C if $a_0^\tau + \delta[(1 - q^\tau)V_0^\tau + q^\tau V_1^\tau] > \delta V_0^\tau$. Rearranging: $a_0^\tau > \delta q^\tau (V_0^\tau - V_1^\tau)$. When $a_0^\tau > 0$, we will show that in equilibrium, either $V_0^\tau \geq V_1^\tau$ (with $V_0^\tau - V_1^\tau \leq k$), or the inequality holds directly. In either case, H is strictly optimal in state 0 when $a_0^\tau > 0$.

Step 4: State-1 behavior and the threshold. In state 1, the firm compares:

$$\begin{aligned} \text{Comply (C):} \quad & \delta V_1^\tau, \\ \text{Redesign (R):} \quad & -k + a_0^\tau + \delta[(1 - q^\tau)V_0^\tau + q^\tau V_1^\tau]. \end{aligned}$$

Case A: Harm until recognition. Suppose the firm chooses C in state 1. Then $V_1^\tau = \delta V_1^\tau$, so $V_1^\tau = 0$. In state 0, the firm chooses H, so $V_0^\tau = a_0^\tau + \delta[(1 - q^\tau)V_0^\tau + q^\tau \cdot 0]$. Solving:

$V_0^\tau = a_0^\tau/[1-\delta(1-q^\tau)]$. For C to be optimal in state 1, we need $\delta \cdot 0 \geq -k + a_0^\tau + \delta(1-q^\tau)V_0^\tau$. Substituting V_0^τ:

$$\begin{aligned} 0 &\geq -k + a_0^\tau + \delta(1-q^\tau) \cdot \frac{a_0^\tau}{1-\delta(1-q^\tau)} \\ &= -k + a_0^\tau \cdot \frac{1-\delta(1-q^\tau)+\delta(1-q^\tau)}{1-\delta(1-q^\tau)} \\ &= -k + \frac{a_0^\tau}{1-\delta(1-q^\tau)}. \end{aligned}$$

This holds if and only if $k \geq a_0^\tau/[1-\delta(1-q^\tau)] = k^*(\tau)$.

Case B: Persistent redesign. Suppose the firm chooses R in state 1. Then the Bellman equations become:

$$\begin{aligned} V_0^\tau &= a_0^\tau + \delta[(1-q^\tau)V_0^\tau + q^\tau V_1^\tau], \\ V_1^\tau &= -k + a_0^\tau + \delta[(1-q^\tau)V_0^\tau + q^\tau V_1^\tau]. \end{aligned}$$

Subtracting: $V_0^\tau - V_1^\tau = k$, so $V_1^\tau = V_0^\tau - k$. Substituting into the first equation:

$$V_0^\tau = a_0^\tau + \delta[(1-q^\tau)V_0^\tau + q^\tau(V_0^\tau - k)] = a_0^\tau + \delta V_0^\tau - \delta q^\tau k.$$

Solving: $V_0^\tau(1-\delta) = a_0^\tau - \delta q^\tau k$, so $V_0^\tau = (a_0^\tau - \delta q^\tau k)/(1-\delta)$. For R to be optimal in state 1, we need:

$$-k + a_0^\tau + \delta[(1-q^\tau)V_0^\tau + q^\tau V_1^\tau] \geq \delta V_1^\tau.$$

This simplifies to $-k + a_0^\tau + \delta(1-q^\tau)V_0^\tau \geq 0$. Substituting V_0^τ and simplifying shows this holds when $k < k^*(\tau)$.

Step 5: Uniqueness. The value functions are determined by a contraction mapping, so they are unique. The policy is unique except when $a_0^\tau = 0$ (where both H and C are optimal in state 0), $a_1^\tau = 0$ (where both H and C yield identical payoffs in state 1), or $k = k^*(\tau)$ (where both C and R are optimal in state 1). □

B.3 Proof of Corollary 1

In full compliance, $H_t = 0$ for all t, so the sum is 0.

In harm-until-recognition, harm occurs in period t if and only if $s_t = 0$. Starting from $s_0 = 0$, the state transitions to 1 with probability q^τ each period that harm occurs. Let T be the first period in which $s_t = 1$. Then T is geometrically distributed with parameter q^τ. We

have:

$$\mathbb{E}\Big[\sum_{t=0}^{\infty}\delta^t H_t\Big] = \mathbb{E}\Big[\sum_{t=0}^{T-1}\delta^t\Big] = \mathbb{E}\Big[\frac{1-\delta^T}{1-\delta}\Big] = \frac{1}{1-\delta}\Big(1-\mathbb{E}[\delta^T]\Big).$$

For $T \sim \text{Geom}(q^\tau)$ (with support $\{1,2,\ldots\}$), we have $\mathbb{E}[\delta^T] = \sum_{t=1}^{\infty}\delta^t q^\tau(1-q^\tau)^{t-1} = \delta q^\tau/[1-\delta(1-q^\tau)]$. Therefore:

$$\mathbb{E}\Big[\sum_{t=0}^{\infty}\delta^t H_t\Big] = \frac{1}{1-\delta}\left(1-\frac{\delta q^\tau}{1-\delta(1-q^\tau)}\right) = \frac{1}{1-\delta(1-q^\tau)}.$$

In persistent redesign, $H_t = 1$ for all t, so the sum is $\sum_{t=0}^{\infty}\delta^t = 1/(1-\delta)$. □

B.4 Proof of Proposition 2

Recall $k^*(\tau) = a_0^\tau/[1-\delta(1-q^\tau)]$ with $a_0^\tau = g - p_0^\tau F$ and $q^\tau = p_0^\tau + (1-p_0^\tau)\alpha^\tau$. Define $D \equiv 1-\delta(1-q^\tau)$. Then $k^*(\tau) = a_0^\tau/D$.

Item 1. We have $\partial a_0^\tau/\partial p_0^\tau = -F < 0$ and $\partial D/\partial p_0^\tau = \delta(1-\alpha^\tau) \geq 0$. By the quotient rule: $\partial k^*/\partial p_0^\tau = [(-F)D - a_0^\tau\delta(1-\alpha^\tau)]/D^2 < 0$ when $a_0^\tau > 0$.

Item 2. We have $\partial a_0^\tau/\partial\alpha^\tau = 0$ and $\partial D/\partial\alpha^\tau = \delta(1-p_0^\tau) \geq 0$. Thus $\partial k^*/\partial\alpha^\tau = -a_0^\tau\delta(1-p_0^\tau)/D^2 \leq 0$, with strict inequality when $p_0^\tau < 1$.

Item 3. Expected harm in harm-until-recognition is $1/[1-\delta(1-q^\tau)]$, which is decreasing in q^τ since the denominator is increasing in q^τ. □

B.5 Proof of Proposition 3

From (14), $W^A - W^L = -(c^A - c^L)/(1-\delta) + h(H^L - H^A)$, where H^τ denotes expected discounted harm under technology τ. Thus $W^A > W^L$ if and only if $h(H^L - H^A) > (c^A - c^L)/(1-\delta)$. □

B.6 Proof of Corollary 2

Fix τ and suppose Assumption 1 holds and $a_0^\tau > 0$. By Proposition 1, redesign is deterred if and only if $k \geq k^*(\tau) = a_0^\tau/[1-\delta(1-q^\tau)]$. Rearrange:

$$k \geq \frac{a_0^\tau}{1-\delta(1-q^\tau)} \quad\Longleftrightarrow\quad 1-\delta(1-q^\tau) \geq \frac{a_0^\tau}{k} \quad\Longleftrightarrow\quad \delta q^\tau \geq \frac{a_0^\tau}{k} - (1-\delta).$$

Dividing by δ yields the threshold $q^\tau \geq \underline{q}(\tau)$ in (12). If $\underline{q}(\tau) \leq 0$, the inequality holds for all feasible $q^\tau \in (0,1]$. If $\underline{q}(\tau) > 1$, the inequality cannot hold for any feasible $q^\tau \in (0,1]$, so redesign occurs for all feasible recognition speeds. □

B.7 Proof of Corollary 3

The full-compliance threshold is immediate from $a_0^\tau = g - p_0^\tau F \leq 0$. For the no-redesign threshold, start from $k \geq k^*(\tau) = a_0^\tau/[1-\delta(1-q^\tau)]$ and substitute $a_0^\tau = g - p_0^\tau F$. Rearranging gives

$$p_0^\tau F \geq g - k[1-\delta(1-q^\tau)] \quad \Longleftrightarrow \quad F \geq \frac{g - k[1-\delta(1-q^\tau)]}{p_0^\tau} = F^{\text{no-redesign}}(\tau).$$

Finally, $F^{\text{no-redesign}}(\tau)$ is decreasing in q^τ because $1-\delta(1-q^\tau)$ is increasing in q^τ. □

B.8 Proof of Proposition 4

By definition of $\bar{q}$ in (16), $\bar{q} \geq q^L$ and $k < k^*(q)$ for all $q \in [q^L, \bar{q})$, while $k \geq k^*(q)$ for all $q \geq \bar{q}$, where $k^*(q) = a_0/[1-\delta(1-q)]$. For any $q \in (q^L, \bar{q})$, Proposition 1 therefore implies persistent redesign, which yields the same expected discounted harm $1/(1-\delta)$ as at $q = q^L$. Since κ is increasing and $\kappa(q) > \kappa(q^L)$ for all $q > q^L$, the regulator's discounted payoff is strictly lower at any $q \in (q^L, \bar{q})$ than at $q = q^L$. Hence no optimal choice lies in $(q^L, \bar{q})$, and any optimal q^* equals q^L or satisfies $q^* \geq \bar{q}$. □

B.9 Proof of Proposition 5

Fix τ and a type k. Under Assumption 1 and $a_0^\tau > 0$, Proposition 1 implies that the type chooses R in state 1 if and only if $k < k^*(\tau)$, and chooses C in state 1 if and only if $k \geq k^*(\tau)$. Thus the cutoff is $\hat{k}(\tau) = k^*(\tau)$. Aggregate expected discounted harm is a convex combination of the harm levels from Corollary 1 across types: types with $k < \hat{k}(\tau)$ generate $1/(1-\delta)$, while types with $k \geq \hat{k}(\tau)$ generate $1/[1-\delta(1-q^\tau)]$. Weighting by the population shares yields (19). Finally, Proposition 2 shows that $k^*(\tau)$ is decreasing in both p_0^τ and α^τ, so monitoring improvements weakly reduce the mass $G(\hat{k}(\tau))$ of redesigning types. □

B.10 Proof of Proposition 6

A firm's gain from violating is g and expected sanction is $D(x)F$, where $D(x)$ is defined in (**??**).

Under Assumption 2, $D(0) = \rho$ (since $m(0) = \phi < M$) and $D(1) = M$ (since $m(1) = \rho > M$). For x such that $m(x) > M$: $D(x) = \rho M/[\phi + x(\rho - \phi)]$, which is decreasing in x.

Item 1. If $g \leq MF$, then $g \leq D(x)F$ for all x (since $D(x) \geq M$), so compliance dominates. Unique equilibrium: $x^* = 0$.

Item 2. If $g \geq \rho F$, then $g \geq D(x)F$ for all x (since $D(x) \leq \rho$), so violation dominates. Unique equilibrium: $x^* = 1$.

Item 3. If $MF < g < \rho F$: At $x = 0$, $D(0)F = \rho F > g$, so compliance is optimal; $x^* = 0$ is an equilibrium. At $x = 1$, $D(1)F = MF < g$, so violation is optimal; $x^* = 1$ is an equilibrium.

For the interior equilibrium, firms must be indifferent: $g = D(x^m)F$. Since $g < \rho F$ and $D(x) = \rho$ when $m(x) \leq M$, indifference requires $m(x^m) > M$. Setting $g = \rho MF/[\phi + x^m(\rho - \phi)]$ and solving: $x^m = (M\rho F/g - \phi)/(\rho - \phi)$.

Comparative statics. Taking derivatives of x^m with respect to ϕ, ρ, and M: $\partial x^m/\partial \phi = \rho(MF/g - 1)/(\rho - \phi)^2 < 0$ since $MF < g$. $\partial x^m/\partial \rho = \phi(1 - MF/g)/(\rho - \phi)^2 > 0$ since $MF < g$. $\partial x^m/\partial M = \rho F/[g(\rho - \phi)] > 0$. □

B.11 Two-Jurisdiction Enforcement Game

This appendix provides the formal treatment of the two-jurisdiction enforcement dilemma summarized in Remark 8 in the main text.

Two jurisdictions $j \in \{1, 2\}$ simultaneously choose monitoring technology $\tau_j \in \{L, A\}$ at $t = 0$. A representative firm, mobile with probability $\mu \in [0, 1]$, locates in the jurisdiction that maximizes its equilibrium value $V_0^{\tau_j}$. Define per-jurisdiction welfare from hosting under technology τ as $w(\tau) \equiv B - c^\tau/(1 - \delta) - hH(\tau)$, where $B = b/(1 - \delta)$ is the PDV of hosting benefits and $H(\tau)$ is expected discounted harm.

Assumption 3 (Symmetric tie-breaking)**.** *If the firm is indifferent across jurisdictions, it chooses each with probability* 1/2.

Assumption 4 (Firm prefers lax monitoring)**.** $V_0^L > V_0^A$.

Proposition 7 (Partial-mobility enforcement dilemma)**.** *Suppose Assumptions 3 and 4 hold and* $w(A) > w(L) > 0$. *If*

$$(1 + \mu)w(L) > w(A), \tag{27}$$

then (L, L) is the unique Nash equilibrium of the jurisdiction game, even though (A, A) Pareto dominates (L, L). Moreover, the set of parameters satisfying (27) is increasing in μ.

Proof. Under Assumptions 3 and 4, the payoff matrix is:

	A	L
A	$\frac{w(A)}{2}$, $\frac{w(A)}{2}$	$\frac{1-\mu}{2}w(A)$, $\frac{1+\mu}{2}w(L)$
L	$\frac{1+\mu}{2}w(L)$, $\frac{1-\mu}{2}w(A)$	$\frac{w(L)}{2}$, $\frac{w(L)}{2}$

(28)

Strict dominance of L. If jurisdiction 2 chooses A: Jurisdiction 1 gets $w(A)/2$ from A and $(1+\mu)w(L)/2$ from L. Condition (27) implies $(1+\mu)w(L)/2 > w(A)/2$, so L is better.

If jurisdiction 2 chooses L: Jurisdiction 1 gets $(1-\mu)w(A)/2$ from A and $w(L)/2$ from L. From (27): $w(L) > w(A)/(1+\mu) \geq (1-\mu)w(A)$ (since $(1-\mu)(1+\mu) \leq 1$). Thus $w(L)/2 > (1-\mu)w(A)/2$, so L is better.

Hence L is strictly dominant; (L, L) is the unique Nash equilibrium.

Pareto dominance. At (A, A): payoff $w(A)/2$ each. At (L, L): payoff $w(L)/2$ each. Since $w(A) > w(L)$, (A, A) Pareto dominates.

Comparative static. Condition (27) becomes $\mu > w(A)/w(L) - 1$. This is easier to satisfy as μ increases. □

B.12 Proof of Proposition 6

All results follow from the chain $\pi \to \tilde{q} \to \tilde{k}^*$. Since $\tilde{q}(x) = p_0 + (1-p_0)\alpha\pi(x)$ is increasing in π (with derivative $(1-p_0)\alpha > 0$ when $p_0 < 1$ and $\alpha > 0$), and

$$\tilde{k}^*(x) = \frac{a_0}{1 - \delta + \delta\,\tilde{q}(x)}$$

is decreasing in $\tilde{q}$ (with derivative $-a_0\delta/[1-\delta+\delta\,\tilde{q}]^2 < 0$), it follows that $\tilde{k}^*$ is decreasing in π, establishing Part (1).

Part (2) is immediate: when $m(x) \leq M$, $\pi(x) = 1$ and $\tilde{q}(x) = q$.

For Part (3), when capacity binds, $\pi(x) = M/m(x)$. Differentiating:

$$\frac{\partial\pi}{\partial\phi} = -\frac{M(1-x)}{m(x)^2} \leq 0, \qquad \frac{\partial\pi}{\partial M} = \frac{1}{m(x)} > 0, \qquad \frac{\partial\pi}{\partial x} = -\frac{M(\rho-\phi)}{m(x)^2} < 0.$$

Since $\tilde{k}^*$ is decreasing in π, the signs reverse, with $\partial\tilde{k}^*/\partial\phi = 0$ when $x = 1$ (no compliant firms to generate false positives).

Part (4) follows from Corollary with $\tilde{q}(x)$ replacing q: the function $1/[1-\delta+\delta\,\tilde{q}]$ is decreasing in $\tilde{q}$ and hence in π. □

B.13 Proof of Corollary 5

At $x=0$: $m(0)=\phi<M$ by Assumption 2, so $\pi(0)=1$ and $\tilde{k}^*(0)=a_0/[1-\delta(1-q)]=k_L^*$. At $x=1$: $m(1)=\rho>M$ by Assumption 2, so $\pi(1)=M/\rho<1$ and $\tilde{q}(1)=p_0+(1-p_0)\alpha M/\rho<q$. Therefore $1-\delta(1-\tilde{q}(1))<1-\delta(1-q)$, and $k_H^*=a_0/[1-\delta(1-\tilde{q}(1))]>k_L^*$ when $p_0<1$ and $\alpha>0$.

For Item (1): if $m(G(k_L^*))\leq M$, then $\tilde{k}^*(G(k_L^*))=k_L^*$, and $x_L^*=G(k_L^*)$ satisfies (24).

For Item (2): if $G(k_H^*)=1$, then at $x=1$, $G(\tilde{k}^*(1))=G(k_H^*)=1=x$, so $x=1$ is a fixed point with capacity binding ($m(1)=\rho>M$). For the general bound, any fixed point in the capacity-binding region satisfies $x_H^*=G(\tilde{k}^*(x_H^*))$ with $x_H^*>x_L^*$ (since capacity does not bind at x_L^*). Because $\tilde{k}^*$ is increasing in x in the capacity-binding region (Proposition 6, Part 3) and $\tilde{k}^*(x)\leq\tilde{k}^*(1)=k_H^*$ for all $x\leq 1$, we have $\tilde{k}^*(x_H^*)\in(k_L^*,k_H^*]$. □

B.14 Proof of Corollary 6

$\tilde{q}(x)\geq\underline{q}(k)$ if and only if $p_0+(1-p_0)\alpha\pi(x)\geq\underline{q}(k)$, i.e., $\pi(x)\geq(\underline{q}(k)-p_0)/[(1-p_0)\alpha]=\underline{\pi}(k)$. Since $\underline{q}(k)>p_0$, we have $\underline{\pi}(k)>0$. Since $q\geq\underline{q}(k)$ (deterrence is feasible without the capacity constraint), $\underline{\pi}(k)\leq(q-p_0)/[(1-p_0)\alpha]=1$. When capacity binds, $\pi(x)=M/m(x)\geq\underline{\pi}(k)$ if and only if $M\geq m(x)\,\underline{\pi}(k)$. □