Enforcement as Message:

A Dynamic Theory of Legitimacy and Compliance*

Ke Liang†

March 2026

Abstract

Why do formally identical legal rules generate sharply different patterns of compliance across jurisdictions and over time? This paper argues that the answer lies not only in sanctions and state capacity, but also in the *congruence* between law on the books and law in action. I develop a dynamic reputation model in which a government privately varies in its commitment to enforcing announced rules, while citizens differ in their willingness to comply voluntarily when they perceive the state as legitimate. The government's enforcement choice leaves a noisy public trace. Citizens update beliefs about governmental fidelity, and those beliefs shape both opportunistic behavior and the share of citizens who comply for normative reasons. The model yields three main results. First, in a static setting, hidden enforcement can improve deterrence by making citizens uncertain about actual permissiveness, but it also encourages under-enforcement by a strategic government. Second, once enforcement choices affect future legitimacy, lax enforcement becomes dynamically costly because it erodes trust, expands future opportunism, and raises the future burden of control. Third, in the infinite horizon, legitimacy becomes an endogenous state variable: enforcement today manages not only current violations but also the future composition of compliance. The theory formalizes a Fullerian intuition that legality depends on congruence between declared rule and official action, and it explains how jurisdictions operating under similar legal codes can diverge into high-trust, largely self-enforcing regimes or low-trust, coercion-dependent ones.

*Work in progress; please do not share without permission.

†Princeton University. Email: ke.liang@princeton.edu.

1 Introduction

Why do some legal rules become social norms while others remain merely formal commands? Across jurisdictions and policy domains, laws that appear similar on paper often produce sharply different patterns of compliance in practice. In some settings, citizens obey even when monitoring is imperfect and sanctions are infrequent. In others, formally identical rules are routinely ignored unless backed by visible and continuous coercion. This puzzle extends far beyond any single domain of governance. Speed limits, environmental regulations, and tax statutes often share comparable legal design, clarity, and formal codification, yet yield markedly different behavioral norms across places and over time. These differences suggest that compliance depends not only on the substantive content of law, but also on how the state enforces it and on whether those enforcement choices sustain the belief that the rule is legitimate and normatively binding.

This paper develops a theory of norm formation grounded in the politics of enforcement. Its central claim is that enforcement is both an instrument and a message. It is an instrument because it changes the expected cost of violation. It is also a message because it conveys whether the government is faithful to the law it publicly proclaims. When state practice appears congruent with formal law, citizens are more likely to regard legal rules as meaningful and worthy of obedience. When violations are visibly tolerated, by contrast, they may infer that the law is weakly binding, selectively applied, or not intended to govern everyday conduct. Enforcement therefore shapes compliance through two channels at once: by altering material incentives and by shaping beliefs about the legitimacy of legal authority. Those beliefs, in turn, influence both voluntary compliance and opportunistic behavior.

The model formalizes this mechanism in a dynamic reputation model. A long-lived government is privately either a committed type, which always enforces the law as written, or a strategic type, which chooses an enforcement gap between *de jure* and *de facto* law. Citizens are short-lived and heterogeneous in their normative willingness to comply. Each citizen has a legitimacy threshold: if perceived governmental commitment is high enough, the citizen complies voluntarily; otherwise, she behaves opportunistically and chooses violation by trading off private benefits against expected sanctions. Citizens do not directly observe the government's enforcement gap. Instead, they observe noisy public signals of enforcement activity and update beliefs about whether the government is truly committed to the law. These beliefs matter because they determine the composition of the population: higher legitimacy increases the share of voluntary compliers and reduces the share of citizens who enter the deterrence problem at all.

This setup produces a dynamic feedback between enforcement, legitimacy, and compliance. A more permissive enforcement choice may save current enforcement costs, but it also generates signals that make the government appear less committed in the future. Lower legitimacy then expands the future pool of opportunistic citizens, increases aggregate violations, and raises the future burden of control. Conversely, strict enforcement can sustain trust, induce more voluntary compliance, and make future enforcement less costly. Norm formation, in this account, is not an exogenous cultural

background condition. It is an equilibrium outcome shaped by how state behavior influences public beliefs and, through those beliefs, the social composition of compliance.

The analysis proceeds in three steps. I begin with a static benchmark that isolates the one-period trade-off between enforcement strictness and enforcement cost. In that environment, hidden enforcement has an asymmetric effect. Because citizens cannot perfectly observe the government's actual permissiveness, they remain exposed to the possibility that the government is committed. This opacity improves deterrence for any given enforcement gap. But the same opacity also makes laxity attractive to the strategic government, because sufficiently large increases in tolerated violations may cease to produce additional behavioral responses. The static model therefore reveals a fundamental tension: hidden enforcement can discipline citizens while undermining governmental discipline.

The dynamic model shows why this static logic is incomplete. In a two-period setting, the strategic government recognizes that current permissiveness affects future legitimacy. Lax enforcement today lowers tomorrow's posterior belief that the government is committed, which reduces future voluntary compliance and increases future enforcement burdens. The strategic type therefore internalizes a reputational cost of under-enforcement that is absent in the static benchmark. This dynamic incentive can induce stricter enforcement than in the one-shot game and, for sufficiently patient governments, can even make it optimal for the strategic type to mimic the committed type in order to preserve trust.

The infinite-horizon model sharpens this mechanism. Legitimacy becomes an endogenous state variable and a form of political capital. The strategic government no longer chooses enforcement merely to manage current violations; it also chooses enforcement to manage the future social foundations of compliance. This yields three qualitatively distinct equilibrium regimes. At high legitimacy, voluntary compliance is widespread and strict enforcement becomes self-enforcing: the strategic type pools with the committed type because preserving trust is more valuable than saving on current enforcement costs. At low legitimacy, by contrast, voluntary compliance is scarce and reputation has little value to protect. In that region, the strategic type tolerates violations to economize on enforcement, and legal order remains coercion-dependent and fragile. Between these extremes lies an intermediate regime in which the strategic type partially invests in credibility and may semi-separate, balancing short-run enforcement savings against the long-run value of sustaining legitimacy.

These dynamics imply that small early differences in enforcement can generate sharply divergent long-run trajectories. Jurisdictions governed by the same formal law may end up in very different equilibria because early enforcement choices alter beliefs, beliefs alter compliance, and compliance alters the future cost of governing. In one setting, law becomes self-enforcing because citizens come to treat obedience as normal and violations as deviant. In another, the same law remains brittle because visible tolerance prevents legitimacy from taking hold. The model thus explains how similar legal rules can give rise to different compliance norms not because citizens are inherently different, but because state practice places them on different belief paths.

The paper contributes to three broader debates. First, it contributes to research on compliance by integrating deterrence and legitimacy within a single dynamic framework. Rather than treating legitimacy as an exogenous attitude or inherited cultural disposition, the model shows how it is produced, sustained, and eroded through ordinary enforcement decisions. Second, it contributes to work on state capacity and implementation by showing that enforcement gaps matter not only because they alter immediate behavior, but also because they shape whether citizens believe that the state stands behind its own rules. Third, it contributes to scholarship on norm formation by identifying a mechanism through which legal order becomes self-reinforcing. In this account, norms do not arise from repeated behavior or socialization alone; they emerge from repeated public inferences about whether legal rules are genuinely meant to bind.

More fundamentally, the paper formalizes a jurisprudential insight first articulated by Lon Fuller: that authority of law rests on congruence between declared rules and official actions. When such congruence is visible, law governs through reasons rather than threats; when it is persistently violated, citizens perceive coercion as arbitrary and withdraw their allegiance. Legitimacy, in this account, is not an exogenous sentiment but an equilibrium outcome sustained by credible signals of fidelity. The model thus reinterprets legitimacy, and more broadly institutional reputation, as a dynamic equilibrium sustained by credible, repeated signaling of fidelity to law. It offers a unified explanation for variation in compliance under a common code, from traffic enforcement to environmental regulation and constitutional policing, where gaps between law on the books and law in action determine whether citizens experience the legal order as principled or performative. Its broader implication is that preserving legitimacy is not ancillary to governance, but one of its central tasks. Legitimacy invites compliance: it lowers the cost of rule by transforming coercion into consent.

The remainder of the paper proceeds as follows. Section 2 situates the argument within the literature on jurisprudence, legitimacy, enforcement, and reputation. Section 3 presents the model. Section 4 analyzes the static benchmark and shows how hidden enforcement creates a tension between citizen deterrence and governmental discipline. Section 5 and section 6 study the two-period and infinite-horizon games, characterizing the reputational value of congruent enforcement and the long-run dynamics of legitimacy. Section 7 discusses the broader normative and empirical implications for legality, trust, and the governance of rules.

2 Related Literature

This paper contributes to three broad literature: the philosophy of law, the study of legitimacy and procedural justice, and formal and empirical work on enforcement, compliance, and reputation.

Philosophy of Law and the Morality of Legality. The paper is directly inspired by Lon Fuller's *The Morality of Law*, which defines law as "the purposive human enterprise of subjecting human conduct to the governance and control of general rules," emphasizing the communal and interactive nature of law. Fuller's principle that "there must be congruence between declared rule

and official action" (Fuller, 1964, p.106), which forms the core focus of this paper. Fuller insists that a regime that routinely issues secret, retroactive, or self-contradictory directives cannot claim fidelity to law at all.

This view contrasts with H.L.A. Hart's positivist account, which grounds legality in social recognition rather than moral content (Hart & Green, 2012). Following Waldron (1994), who emphasizes law as a network of institutions bound by publicly knowable norms, this paper treats congruence as an institutional condition for legitimacy: the state must not depart from the rules it declares, even when deviation might advance instrumental goals.

Legitimacy and Procedural Justice. The paper also relates to an extensive literature on legitimacy in political psychology and sociology. One strand emphasizes procedural fairness as the foundation of perceived legitimacy (Tyler & Mitchell, 1993). Tyler (2001) and Tyler (2006) show that people's willingness to obey the law depends less on fear of punishment than on their judgment of whether legal authorities act fairly, respectfully, and consistently. Recent work extends this argument to policing, demonstrating that compliance depends on whether citizens perceive authorities as legitimate rather than merely powerful (Tyler, 2025). Similarly, Pettit (2023) proposes the "tough-luck test," suggesting that citizens accept adverse outcomes when they trust that the law was guided by public reason. Together, these studies imply that legitimacy invites compliance: people obey when they view the system as fair and law-abiding.

Formal approaches to legitimacy connect these insights to collective choice. Patty & Penn (2011) model legitimacy as the procedural consistency of decision rules rather than the optimality of outcomes. Penn & Patty (2023) extend this reasoning to enforcement, showing that citizens respond differently when authorities are perceived to act for public safety rather than for rent extraction. In their framework, as in this paper, legitimacy depends on how enforcement conveys purpose. Benabou & Tirole (2006) models the need to sustain belief in a just world, paralleling the role of institutional trust in maintaining compliance. Empirically, Stokes & Hunnicutt (2018) and Dickson et al. (2022) show that authorities' actions and messaging can directly shape behavior by enhancing perceived legitimacy, even when material incentives remain unchanged.

Law, Enforcement, and Reputation. From a law-and-economics perspective, this paper connects to the classic deterrence literature (Becker, 1968; Becker & Stigler, 1974) and its extensions in institutional design (Shavell, 2004). However, standard deterrence models typically abstract from how enforcement practices affect beliefs about institutional integrity and future enforcement. Instead, this paper draws on the reputation literature in economics, which studies how repeated interaction and signaling discipline behavior when commitment is imperfect. Seminal work by Kreps & Wilson (1982) and Milgrom & Roberts (1982) shows how reputational incentives can sustain cooperative behavior even when short-run incentives favor deviation. Milgrom et al. (1990) further emphasize how past actions serve as a bond for future conduct. In this paper, enforcement choices play an analogous role: they are publicly observed, costly signals that shape beliefs about governmental fidelity and discipline a strategic state.

The paper is close in spirit to Kartik et al. (2025), which likewise studies how long-run governance outcomes depend on belief updating under imperfect monitoring. The connection lies in the shared emphasis on path dependence and the way current behavior shapes future governance through its effects on belief. The key difference is that the object of reputation here is not effort provision or replacement risk, but fidelity to law. The paper is also related to Chen et al. (2025), which shows how current enforcement choices can build a reputation for strictness and thereby influence future behavior. The central distinction is that this paper studies not reputation for toughness per se, but reputation for legal fidelity and the way it sustains legitimacy, voluntary compliance, and the emergence of legal norms.

The paper also connects to work on discrepancies between law on the books and law in action. Holland (2016) conceptualizes forbearance as intentional non-enforcement; Fenske et al. (2023) document how such gaps shape regulated actors. Holland (2016) conceptualize forbearance as intentional non-enforcement, while Fenske et al. (2023) document how enforcement gaps shape regulated actors' behavior. Related work explores when insincere or symbolic rules can improve compliance (Gilbert, 2015), or how bureaucratic capacity constrains legal congruence (Huber & McCarty, 2004). This paper complements these accounts by modeling enforcement gaps as reputational signals rather than as purely distributive or capacity-driven choices. Furthermore, Stone (2016) critiques the dominant economic model of legal design, which assumes rational and self-interested "bad man" comply with laws solely to avoid sanctions. She argues this overlooks a significant portion of legal subjects who comply because they believe compliance has intrinsic or epistemic value. Such "good persons" can't be glossed over in policy design.

Finally, the paper speaks to recent work linking social norms, trust, and law-abidingness. Acemoglu & Jackson (2017) show that laws conflicting with prevailing norms can backfire; Benabou & Tirole (2011) model how laws interact with moral signaling. Empirically, Besley & Dray (2024) demonstrate that trust in government reduces the cost of policy implementation by increasing voluntary compliance. A related line of work highlights the role of beliefs and expectations in sustaining compliance:Bassetto & Phelan (2008) shows that enforcement institutions can support multiple behavioral equilibria, with expectations of widespread evasion making noncompliance self-reinforcing; Dilmé & Garrett (2019) shows that publicized enforcement successes can generate residual deterrence by shaping beliefs about future punishment. Further, Ma et al. (2025) highlight how shared responsibility between citizens and officials can blur accountability and shape institutional performance. This paper connects these insights by modeling legality as a reciprocal commitment between state and citizens: enforcement signals fidelity, beliefs generate trust, and trust sustains compliance.

In sum, this paper bridges the philosophical idea of law's internal morality with formal theories of enforcement and reputation. It contributes to the literature by providing a dynamic reputation model in which compliance norms arise endogenously from the interaction between enforcement choices and belief updating, explaining how jurisdictions operating under identical legal codes can diverge sharply in legitimacy, compliance, and the long-run cost of governance.

3 Model

3.1 Environment

Time is discrete and infinite, $t = 1, 2, \cdots$, with discount factor is $\beta \in [0, 1)$. A statute specifies a de jure upper bound $r \in [0, 1]$ on some activity $x \in [0, 1]$. Any choice above the statutory limit constitutes a *violation*,

$$v \equiv x - r \ \in \ V \equiv [0, \bar{v}], \qquad \bar{v} = 1 - r.$$

In each period, the government chooses an *enforcement gap* $\Delta_t \in V$. The enforcement gap is the distance between the *de jure* limit and the *de facto* tolerance boundary: violations $v \leq \Delta_t$ are tolerated, while violations $v > \Delta_t$ are sanction-eligible. A sanction-eligible violation is detected with probability $p(v) \in [0, 1)$ and, if detected, the citizen pays a fine $w > 0$. The detection probability $p(v)$ is strictly increasing in v, capturing the idea that larger violations are more visible, easier to verify, or more likely to trigger official action.

3.2 Citizens

Citizens are short-lived: a unit mass enters each period, makes exactly one decision, and exits. Citizens differ in their *normative commitment* to the law. Each citizen draws a legitimacy standard $\gamma \in [0, 1]$ independently from a continuous distribution F with density $f(\gamma) > 0$ on $[0, 1]$. The parameter γ measures how much perceived governmental legitimacy the citizen requires in order to comply voluntarily, independent of expected sanction. This heterogeneity of legitimacy standard reflects the empirical finding that citizens vary markedly in their intrinsic motivation to comply: some comply almost unconditionally, some are nearly always opportunistic, and many lie in between (Tyler, 2006). The distribution F therefore summarizes the heterogeneity in intrinsic motivation to obey the law.

Let $\mu_t \in [0, 1]$ denote *perceived legitimacy* at the end of period t, interpreted as the public posterior belief that the government is intrinsically committed to enforcing the statute (type C, defined below). At the beginning of period t, given public belief μ_{t-1}, a citizen with legitimacy standard γ then behaves as follows:

- **Voluntary compliance:** if $\mu_{t-1} \geq \gamma$, the citizen complies voluntarily and choose $v = 0$.
- **Opportunism:** if $\mu_{t-1} < \gamma$, the citizen behave opportunistically and choose $v \in V$ to maximize expected utility.

The threshold rule captures that a citizen complies voluntarily when the perceived probability of governmental commitment is high enough to satisfy her own normative standard. Hence, the *mass of opportunists* at belief μ is

$$\lambda(\mu) \equiv \Pr(\gamma > \mu) = 1 - F(\mu),$$

which is strictly decreasing and continuously differentiable. Higher legitimacy therefore reduces the share of citizens who enter the deterrence problem at all.

Opportunist payoff. An opportunist who chooses violation v obtains private benefit $b(v)$. Since the current enforcement gap Δ_t is not directly observed, the opportunist conditions on public belief and the government's equilibrium policy. Let $\sigma(\mu) \in \Delta(V)$ denote the strategic type's (possibly mixed) period-t enforcement strategy at belief μ. For any $v \in V$, define the *strategic exceedance probability*,

$$\sigma(v; \mu) \equiv \Pr(v > \Delta_t \mid \mu) = \sigma(\mu)\big([0, v)\big),$$

which represents the probability that a violation of size v exceeds the strategic type's chosen enforcement gap and is therefore sanction-eligible. Then, the *total sanction-eligibility probability* for violation $v > 0$ is

$$q(v \mid \mu, \sigma) = \mu \cdot \mathbf{1}\{v > 0\} + (1 - \mu) \cdot \sigma(v; \mu).$$

combining the certain enforcement of the committed type (probability μ) with the probabilistic enforcement of the strategic type (probability $1 - \mu$). For any fixed strategic-type policy σ, the total sanction probability $q(v \mid \mu, \sigma)$ is strictly increasing in μ, making violations more costly even holding the government's actual behavior constant. The opportunist thus solves

$$\max_{v \in V} \; u(v; \mu, \sigma) \equiv b(v) - wp(v)q(v \mid \mu, \sigma).$$

Assumption 1 (Citizen Primitives)**.** The benefit function $b : V \to \mathbb{R}_+$ and detection function $p : V \to [0, 1)$ satisfy: (i) b is C^2, strictly increasing, and strictly concave, with $b(0) = 0$; (ii) p is C^2, strictly increasing, and strictly convex, with $p(0) = 0$.

This assumptions on b and p, concavity of benefit and convexity of detection, imply that the opportunist's objective is strictly concave in v, ensuring a unique best response.

3.3 Government

At $t = 0$, Nature draws the government's type $\theta \in \{C, S\}$, which is fixed over all periods and privately observed. The public prior is

$$\mu_0 \equiv \Pr(\theta = C) \in (0, 1).$$

Government types differ only in enforcement feasibility:

- **Committed type** ($\theta = C$): sets $\Delta_t = 0$ for all $t \geq 1$.
- **Strategic type** ($\theta = S$): chooses $\Delta_t \in V$ to minimize discounted enforcement costs.

The committed type captures intrinsic adherence to legal principle and the rule of law: enforcement choices are independent from cost considerations. By contrast, the strategic type faces enforcement constraints and trade offs enforcement strictness against its associated costs.

Let $v^*(\Delta)$ denote the opportunist best response under observed enforcements, and let $v^*(\mu, \sigma)$ denote an opportunist best response under hidden enforcement. Aggregate violations are: $X(\mu, \Delta) \equiv \lambda(\mu)\, v^*(\Delta)$ and $X(\mu, \sigma) \equiv \lambda(\mu)\, v^*(\mu, \sigma)$. The strategic government's per-period payoff is

$$\Pi(\Delta; \mu) = -K(X, \Delta),$$

where $K : \mathbb{R}_+ \times V \to \mathbb{R}_+$ is an enforcement/capacity cost.

Assumption 2 (Government Cost)**.** The cost function $K : \mathbb{R}_+ \times V \to \mathbb{R}_+$ is C^2 and satisfies: (i) $K_X > 0$: higher violations increase cost; (ii) $K_\Delta < 0$: tolerance saves enforcement resources; (iii) $K_{X\Delta} < 0$: strict increasing differences in $(X, -\Delta)$; (iv) $K_{\Delta\Delta} > 0$: strict convexity in Δ; (v) $K_\Delta(X, \bar{v}) = 0$ and $K_X(X, \bar{v}) > 0$ for all $X > 0$: rule out full tolerance.

This assumption gives the strategic type a standard trade-off. Condition (i) says that aggregate violations are socially, administratively, or politically costly. Condition (ii) says that strict law enforcement is more costly. Condition (iii) imposes complementarity: tightening (lower Δ) is especially costly when violations are widespread, because each unit of enforcement gap now tolerates many more violations. Condition (iv) rules out pathological curvature. Condition (v) is the no-full-tolerance boundary condition: at the boundary $\Delta = \bar{v} = 1 - r$, the government is already tolerating the maximum feasible violation, so the direct marginal savings from further relaxation are exhausted: $K_\Delta(X, \bar{v}) = 0$. But if aggregate violations remain positive, their marginal cost is strictly positive: $K_X(X, \bar{v}) > 0$. Since a smaller Δ reduces violations, a government at full tolerance can strictly lower total cost by moving slightly inward from $\bar{v}$. This guarantees that $\Delta = \bar{v}$ cannot be optimal whenever $X > 0$.

3.4 Information and Timing

Static benchmark. In the static benchmark, public belief is fixed at μ_0, there is no intertemporal updating, and the government chooses enforcement once. Both *observed-enforcement*, in which citizens observe Δ before choosing v, and *hidden-enforcement*, in which citizens do not observe Δ and instead act based on prior belief, are considered.

Dynamic learning. Citizens do not observe Δ_t directly but observe a noisy public signal

$$s_t = \Delta_t + \epsilon_t,$$

where $\{\epsilon_t\}$ are i.i.d. with density g and full support on $\mathbb{R}$, independent of type and history. The signal captures imperfect public monitoring: enforcement actions generate observable traces (prosecution statistics, news coverage, administrative records), but these traces are noisy and do not perfectly reveal the government's enforcement gap.

Let $\mathcal{H}_{t-1} = \{s_1, \ldots, s_{t-1}\}$ denote public history, summarized by posterior μ_{t-1}. At $t = 0$, Nature draws $\theta \in \{C, S\}$, the public holds prior $\mu_0 = \Pr(\theta = C)$, and $\mathcal{H}_0 = \varnothing$. Each period $t \geq 1$:

1. Public belief μ_{t-1} is commonly observed.
2. The government chooses Δ_t (conditional on θ and μ_{t-1}); simultaneously, each citizen chooses whether to comply voluntarily (if $\mu_{t-1} \geq \gamma$) or act opportunistically, and opportunists choose v_t given μ_{t-1} and the equilibrium strategy.
3. Signal $s_t = \Delta_t + \epsilon_t$ is realized and observed publicly; beliefs update via Bayes' rule to μ_t.

Assumption 3 (Signal structure)**.** The noise $\{\epsilon_t\}_{t\geq 1}$ is i.i.d. across periods, independent of type and history, with continuous density g that is strictly positive on $\mathbb{R}$. The family of shifted densities

$\{g(\cdot - \Delta) : \Delta \in V\}$ satisfies the monotone likelihood ratio property ($MLRP$): for any $\Delta' > \Delta$, the likelihood ratio

$$\frac{g(s - \Delta')}{g(s - \Delta)}$$

is weakly increasing in s.

The signal structure formalizes imperfect monitoring of enforcement. In each period, citizens do not observe the government's enforcement gap Δ_t directly. Instead, they observe a public signal $s_t = \Delta_t + \epsilon_t$, which is informative but noisy. The signal should be understood broadly as a summary of the publicly visible traces of enforcement activity: prosecution rates, arrest patterns, administrative outputs, press coverage, court fillings, or other publicly available indicators. The signal structure assumption implies that higher signals are "bad news" for legitimacy: a government with a large enforcement gap ($\Delta' > \Delta$) generates higher signals in expectation. The Bayes' rule implies that, after observing signal s, the posterior probability that the government is committed is

$$\mu'(s; \mu, \Delta^S) = \frac{\mu\, g(s)}{\mu\, g(s) + (1 - \mu)\, g(s - \Delta^S)}.$$

By MLRP, $\mu'(s; \mu, \Delta^S)$ is strictly decreasing in s whenever $\Delta^S > 0$: citizens lower their assessment of governmental commitment upon observing a high signal, as the committed government always chooses $\Delta_t = 0$.

Equilibrium. The solution concept is Markov Perfect Bayesian Equilibrium. A strategy for the strategic type maps current public belief μ into an enforcement choice (or distribution over choices) $\sigma(\mu)$. Citizens best respond given current belief and the government's equilibrium strategy, and beliefs are updated by Bayes' rule wherever possible.

4 Static Benchmark

I begin with a static benchmark in which public belief $\mu \in (0, 1)$ is fixed and there is no intertemporal learning. This benchmark is useful for two reasons. First, it isolates the *pure contemporaneous tradeoff* facing the strategic government: relaxing enforcement saves direct enforcement costs, but it also changes the amount of violation that must be governed. Second, it provides the natural terminal-period benchmark for the dynamic model. In this static environment, legitimacy affects behavior only through current citizen decisions; it does not yet function as a state variable that the government may wish to preserve.

Fix belief $\mu \in (0, 1)$. The strategic type chooses an enforcement gap Δ once and minimizes

$$K(X, \Delta),$$

where X denotes aggregate violations induced by the chosen policy. Two citizen-side benchmarks will be useful throughout:

$$v^s \equiv \arg\max_{v \in V} \big[b(v) - w\, p(v)\big], \qquad \hat{v}(\mu) \equiv \arg\max_{v \in V} \big[b(v) - \mu w\, p(v)\big].$$

The quantity v^s is the opportunist's optimal violation when enforcement is effectively known to be strategic, so sanctions apply only when the violation exceeds the realized tolerance threshold. By contrast, $\hat{v}(\mu)$ is the opportunist's optimal violation when she expects to face sanction with effective probability scaled by the current legitimacy belief μ. Under Assumption 1, both maximizers are unique and interior.

Lemma 1. *$\hat{v}(\mu)$ is strictly decreasing in μ, and $\hat{v}(\mu) \geq v^s$ for all $\mu \in [0,1]$, with strict inequality for $\mu < 1$ whenever both optima are interior.*

Lemma 1 formalizes the most basic citizen-side force in the model: *legitimacy deters*. A higher belief that the government is committed makes opportunistic citizens behave more cautiously, because even violations that would be tolerated by the strategic type may still be punished by the committed type. Thus, legitimacy compresses the violation margin before we even introduce any dynamic incentive for the government to protect it. Legitimacy thus operates like a decentralized compliance technology.

4.1 Observed Enforcement

I first consider the case in which citizens observe the government's enforcement gap Δ directly before choosing v. An opportunist then solves

$$\max\Big\{ \max_{v\in[0,\Delta]} b(v),\ \max_{v\in(\Delta,\bar{v}]} \big[b(v) - w\,p(v)\big] \Big\}.$$

Because enforcement is observed, the citizen knows exactly which violations will be tolerated and which will be sanction-eligible.

Lemma 2 (Optimal Violation, Observed)**.** *There exists a unique threshold $\Delta^O \in (0, v^s)$ satisfying $b(\Delta^O) = b(v^s) - w\,p(v^s)$, such that*

$$v_O^*(\Delta) = \begin{cases} v^s, & \Delta \leq \Delta^O, \\ \Delta, & \Delta > \Delta^O. \end{cases}$$

Lemma 2 yields a simple but important cutoff structure. When the enforcement gap is small, tolerated violations are too limited to compete with the risky but more profitable choice v^s, so the opportunist continues to choose v^s. Once the gap becomes sufficiently large, however, the citizen finds it optimal to move exactly onto the tolerance boundary. Intuitively, observed laxity creates a *safe harbor*: if the government openly tolerates a broader range of violations, opportunists exploit that slack one-for-one.

This observation sharply disciplines the government's problem. Small positive gaps do not reduce violations at all relative to stricter policies, while sufficiently large gaps mechanically enlarge the tolerated violation level.

Proposition 1 (Enforcement Gap, Observed)**.** *Fix $\mu \in (0,1)$. In the observed static benchmark:*

(i) No optimum lies in $(0, \Delta^O)$.

(ii) On $(\Delta^O, \bar{v})$, aggregate violations are $X_O(\mu, \Delta) = \lambda(\mu)\Delta$, and any interior optimum $\Delta_O^* \in (\Delta^O, \bar{v})$ satisfies

$$\lambda(\mu)\, K_X(\lambda(\mu)\Delta_O^*, \Delta_O^*) + K_\Delta(\lambda(\mu)\Delta_O^*, \Delta_O^*) = 0.$$

(iii) If $\lambda(\mu) > 0$, then $\Delta = \bar{v}$ is not optimal.

Proposition 1 shows that under observed enforcement the government cannot relax policy without paying a contemporaneous behavioral price. Part (i) rules out the region $(0, \Delta^O)$ because moving from full enforcement to a small amount of laxity does not change citizen behavior, so it only weakly raises direct laxity without generating any offsetting reduction in the violation burden. Part (ii) describes the economically relevant interior region: once $\Delta > \Delta^O$, tolerated violations move exactly with the enforcement gap, so the government balances the marginal saving from relaxing strictness against the induced increase in aggregate violations. Part (iii) then rules out full tolerance. If any opportunists exist, setting $\Delta = \bar{v}$ maximally enlarges tolerated violations and is therefore too costly on the citizen side.

The logic is intuitive. Under observability, laxity is immediately and fully exploited. The government may still choose some positive gap if strict enforcement is costly enough, but it can never "hide" behind citizens' uncertainty. Observed laxity directly becomes higher tolerated noncompliance.

4.2 Hidden Enforcement

Now turn to the case in which citizens do *not* observe the enforcement gap before acting. This is the key informational environment for the rest of the paper. Hidden enforcement changes the citizen's problem because a violation that is tolerated by the strategic type may still be punished by the committed type. For $v \le \Delta$, the citizen is safe only under the strategic type, not under the committed type. The sanction-eligibility probability is therefore

$$q(v \mid \mu, \Delta) = \begin{cases} \mu, & 0 < v \le \Delta, \\ 1, & v > \Delta. \end{cases}$$

Lemma 3 (Optimal Violation, Hidden)**.** *Fix $\mu \in (0,1)$. There exists a unique threshold $\Delta^H(\mu) \in (0, v^s)$ satisfying*

$$b\big(\Delta^H(\mu)\big) - \mu w\, p\big(\Delta^H(\mu)\big) = b(v^s) - w\, p(v^s),$$

such that

$$v_H^*(\mu, \Delta) = \begin{cases} v^s, & \Delta < \Delta^H(\mu), \\ \Delta, & \Delta^H(\mu) \le \Delta < \hat{v}(\mu), \\ \hat{v}(\mu), & \Delta \ge \hat{v}(\mu). \end{cases}$$

Moreover, $\Delta^H(\mu) < \Delta^O$ and $\frac{d\Delta^H(\mu)}{d\mu} > 0$.

Lemma 3 captures the central informational advantage of opacity. Because citizens cannot observe the realized enforcement gap, even tolerated violations remain risky: they may still be

punished if the government is the committed type. Relative to the observed case, this makes opportunists more cautious and shifts the cutoff at which tolerated violations become attractive downward from Δ^O to $\Delta^H(\mu)$. In other words, hidden enforcement creates *residual deterrence* inside the tolerance region.

The piecewise structure is also revealing. For small gaps, citizens still choose the risky benchmark v^s. For intermediate gaps, they move to the hidden tolerance boundary Δ, as in the observed case, but only after a smaller threshold. For very large gaps, however, the citizen no longer follows the enforcement gap upward: she stops at $\hat{v}(\mu)$. This upper region is where hidden enforcement most starkly differs from observed enforcement. Once the gap is large enough, additional government laxity no longer induces more violation, because citizens remain disciplined by the residual probability of facing the committed type.

Proposition 2 (Enforcement Gap, Hidden)**.** *Fix $\mu \in (0,1)$. In the hidden static benchmark:*

(i) No optimum lies in $(0, \Delta^H(\mu))$.

(ii) On $(\Delta^H(\mu), \hat{v}(\mu))$, aggregate violations $X_H = \lambda(\mu)\Delta$, and any interior optimum Δ_H^* satisfies

$$\lambda(\mu)\, K_X(\lambda(\mu)\Delta_H^*, \Delta_H^*) + K_\Delta(\lambda(\mu)\Delta_H^*, \Delta_H^*) = 0.$$

(iii) On $[\hat{v}(\mu), \bar{v}]$, aggregate violations are $X_H = \lambda(\mu)\hat{v}(\mu)$ (constant), so conditional on choosing from this region the government optimally sets $\Delta = \bar{v}$.

Proposition 2 shows why opacity creates a temptation toward laxity. As in the observed case, part (i) rules out very small positive gaps because they do not change behavior. Part (ii) describes the intermediate region in which aggregate violations still move one-for-one with the gap, so the government solves a familiar interior tradeoff. But part (iii) is the distinctive result: once Δ is large enough, further increases in laxity do *not* increase violations. Citizen behavior is capped at $\hat{v}(\mu)$, so the government can continue relaxing enforcement without worsening aggregate noncompliance. In that region, stricter enforcement is pure cost from the government's point of view, and the optimal choice is maximal tolerance $\Delta = \bar{v}$.

This is the static model's central tension in its cleanest form. Opacity improves deterrence on the citizen side, but precisely because it improves deterrence, it can make large enforcement gaps privately attractive to the government. Hidden enforcement therefore contains both a compliance benefit and a governance hazard.

Proposition 3 (Static Comparison)**.** *Fix $\mu \in (0,1)$. Hidden enforcement improves citizen-side deterrence and increases the government's incentive to relax enforcement:*

(i) For any enforcement gap Δ, $X_H(\mu, \Delta) \leq X_O(\mu, \Delta)$, with strict inequality for $\Delta \in (\Delta^H(\mu), \Delta^O)$.

(ii) Full tolerance $\Delta = \bar{v}$ may be optimal under hidden enforcement even though it is ruled out under observed enforcement.

Proposition 3 puts the two informational environments side by side. Hidden enforcement dominates observed enforcement from the perspective of contemporaneous deterrence: for any given

policy, it generates weakly lower aggregate violations. Yet that very feature makes it easier for the government to relax. Under observability, laxity is immediately converted into higher tolerated violations; under opacity, citizens remain partially disciplined even when enforcement is weak. The result is a wedge between *citizen-side deterrence* and *government-side discipline*.

The static benchmark therefore delivers the paper's motivating puzzle. Opacity can make citizens behave better while making the government behave worse. In a one-shot world, nothing counters that temptation: whenever a sufficiently large "cost-free laxity" region exists, the strategic government exploits it. The dynamic model matters because it endogenizes the missing force. Once current enforcement choices shape future legitimacy, laxity is no longer costless.

5 Two-period Model

The two-period model is the minimal environment in which reputation matters. Unlike the static benchmark, period-1 enforcement now affects not only current violations but also future beliefs, and hence future compliance costs. This setting is analytically useful because it isolates the reputational mechanism in its simplest form: the government chooses current policy knowing that today's signal changes tomorrow's legitimacy, but there is only one future period to protect. The two-period model therefore sits conceptually between the static benchmark and the full infinite-horizon problem.

There are two periods, $t \in \{1, 2\}$, and the government discounts the future at rate $\beta \in (0, 1)$.

5.1 Period-2 Subgame

I begin from the end of the game. In period 2 there is no continuation value, so the strategic type solves exactly the static hidden-enforcement problem given belief μ_1 entering the terminal period. Define the period-2 value function as the minimized cost

$$C_2(\mu_1) \equiv \min_{\Delta_2 \in V} K\big(X_H(\mu_1, \Delta_2),\ \Delta_2\big).$$

Lemma 4 (Period-2 Value Function)**.** *C_2 is continuously differentiable on $(0, 1)$, and $C_2'(\mu_1) < 0$: the period-2 cost is strictly decreasing in legitimacy.*

Lemma 4 establishes the value of legitimacy in the simplest possible way. A government that enters the terminal period with higher legitimacy faces fewer opportunists, lower violations, and therefore lower enforcement costs. Legitimacy is thus not merely a symbolic asset; it is an operational resource that reduces the future burden of governance. This observation is what turns the static tension into a dynamic tradeoff. Once legitimacy has continuation value, current laxity becomes costly because it threatens a valuable future state.

5.2 Period-1 Problem

In period 1, the strategic type chooses $\Delta_1 \in V$ to minimize current cost plus discounted expected period-2 cost:

$$\mathcal{C}_1(\Delta_1; \mu_0) \equiv K\big(X_H(\mu_0, \Delta_1),\, \Delta_1\big) + \beta\, \mathbb{E}_{s_1}\big[C_2\big(\mu'(s_1; \mu_0, \Delta_1)\big)\big]\,. \tag{1}$$

The expectation is taken over the period-1 signal $s_1 = \Delta_1 + \epsilon_1$, where ϵ_1 is drawn from $g(\cdot - \Delta_1)$. Beliefs are updated by Bayes' rule:

$$\mu'(s; \mu_0, \Delta_1) = \frac{\mu_0\, g(s)}{\mu_0\, g(s) + (1 - \mu_0)\, g(s - \Delta_1)}.$$

It is convenient to isolate the continuation component:

$$\Phi(\Delta_1; \mu_0) \equiv \mathbb{E}_{s_1}\big[C_2\big(\mu'(s_1; \mu_0, \Delta_1)\big)\big] = \int_{-\infty}^{\infty} C_2\bigg(\frac{\mu_0\, g(s)}{\mu_0\, g(s) + (1 - \mu_0)\, g(s - \Delta_1)}\bigg)\, g(s - \Delta_1)\, ds.$$

This is the expected period-2 cost induced by a period-1 enforcement choice.

Lemma 5 (Reputational Cost)**.** *Under Assumptions 1–3:*

(i) $\Phi(0; \mu_0) = C_2(\mu_0)$*: if the strategic type pools with the committed type, expected future cost equals* C_2 *evaluated at the current belief.*

(ii) $\Phi(\Delta_1; \mu_0)$ *is strictly increasing in* Δ_1 *for* $\Delta_1 > 0$*: laxity today raises expected future cost.*

(iii) $\frac{\partial \Phi}{\partial \Delta_1}\big|_{\Delta_1 = 0^+} > 0$*: the marginal reputational cost of deviating from full enforcement is strictly positive.*

Lemma 5 is the key bridge from the static benchmark to the dynamic model. Part (i) says that if the strategic type perfectly mimics the committed type by choosing $\Delta_1 = 0$, then beliefs are preserved on average, and the government carries its current legitimacy into period 2. Part (ii) shows that relaxing enforcement today damages expected future legitimacy and thereby raises future costs. Part (iii) strengthens the point: the reputational penalty appears immediately at the margin. The first unit of laxity is not free. Even a small departure from full enforcement leaves a detectable trace in expected beliefs.

Economically, Φ is the dynamic analogue of the "cost-free laxity" region in the static hidden-enforcement problem. In the static model, sufficiently large hidden gaps may become attractive because they no longer worsen citizen behavior. In the two-period model, that logic is no longer complete: even if current violations stop rising, laxity still worsens future beliefs. Reputation restores discipline precisely where the static model loses it.

5.3 Equilibrium

Definition 1 (Two-Period Equilibrium)**.** A *two-period MPBE* consists of: (i) a period-1 enforcement gap Δ_1^{**} minimizing $\mathcal{C}_1(\Delta_1; \mu_0)$; (ii) a period-2 policy $\Delta_2^*(\mu_1)$ minimizing $C_2(\mu_1)$; (iii) belief updating via Bayes' rule.

Proposition 4 (Reputational Discipline in Two Periods)**.** *Let $\Delta_H^*(\mu_0)$ denote the static hidden-enforcement optimum at belief μ_0. Under Assumptions 1–3, any interior two-period equilibrium gap Δ_1^{**} satisfies:*

(i) $\Delta_1^{**} < \Delta_H^*(\mu_0)$: the two-period gap is strictly smaller than the static optimum.

(ii) The equilibrium gap satisfies the first-order condition

$$\underbrace{\lambda(\mu_0)\, K_X(X_H^{**}, \Delta_1^{**}) + K_\Delta(X_H^{**}, \Delta_1^{**})}_{\text{static margin}} + \beta \underbrace{\left.\frac{\partial \Phi}{\partial \Delta_1}\right|_{\Delta_1^{**}}}_{\text{reputational margin}>0} = 0. \tag{2}$$

(iii) Δ_1^{**} is strictly decreasing in β: a more patient government enforces more strictly.

(iv) Δ_1^{**} is strictly decreasing in $-C_2'(\mu_0)$: a steeper future value function magnifies discipline.

Proposition 4 gives the cleanest statement of *reputational discipline*. Part (i) shows that once future legitimacy matters, the strategic type chooses a smaller enforcement gap than in the static hidden-enforcement benchmark. Intuitively, the government now sacrifices some current cost savings in order to protect its reputation. Part (ii) makes the mechanism transparent: the optimal gap balances the same static tradeoff as before, but now an additional positive reputational term enters the first-order condition. The government is no longer choosing only how much contemporaneous violation to tolerate; it is also choosing how much legitimacy to preserve.

Parts (iii) and (iv) provide the comparative statics that clarify what strengthens this discipline. Patience matters because future legitimacy is valuable only to the extent that the government cares about future costs. Likewise, the slope of C_2 matters because reputation is more worth preserving when a loss of legitimacy has sharper future consequences. Together, these results show that reputational incentives are strongest when the government is both forward-looking and highly exposed to future compliance costs.

Remark 1 (Pooling Boundary)**.** The two-period model admits the possibility that $\Delta_1^{**} = 0$, i.e., the strategic type fully mimics the committed type in period 1. This *complete pooling* equilibrium arises when the marginal reputational benefit of choosing $\Delta_1 = 0$ relative to Δ_H^* exceeds the foregone static cost savings. Formally, complete pooling occurs if and only if

$$\mathcal{C}_1(0; \mu_0) \leq \mathcal{C}_1(\Delta_H^*(\mu_0); \mu_0).$$

In that case, the strategic type pays the full short-run cost of strict enforcement in order to preserve legitimacy for the final period.

Remark 1 identifies the extreme case of reputational discipline. If the continuation value of legitimacy is large enough, the strategic type may fully imitate the committed type in period 1. This is the sharpest possible contrast with the static benchmark. In the static model, the government exploits opacity to relax enforcement; in the two-period model, the same opacity may instead induce the government to overinvest in strictness in order to keep citizens uncertain.

Corollary 1 (Comparative Statics)**.** *At an interior two-period equilibrium:*

(i) Higher initial legitimacy μ_0 generates a smaller equilibrium gap Δ_1^{**}, provided the slope of C_2 becomes steeper in absolute value as μ decreases.

(ii) A mean-preserving spread in the signal noise ϵ (holding μ_0 fixed) weakens reputational incentives and increases Δ_1^{**}.

The corollary highlights two further intuitions. First, legitimacy itself creates a *legitimacy premium*: a government with more reputation to lose is disciplined more strongly by the prospect of erosion. Second, informativeness matters because reputation can constrain behavior only if policy choices are at least partially visible through signals. Noisier signals blur the mapping from enforcement to posterior beliefs, thereby weakening the benefit of strict enforcement. Put differently, reputational discipline is strongest when legitimacy is valuable and policy is observable enough to affect it.

The two-period model therefore teaches a precise lesson. Opacity still improves deterrence on the citizen side, but once beliefs persist, laxity is no longer merely a static saving. It becomes an intertemporal liability. Yet this discipline is still limited: the government protects legitimacy only for one future period. The infinite-horizon model is what turns this one-step reputational force into a genuinely dynamic process of norm formation.

6 Infinite Horizon

We now return to the full infinite-horizon environment. This extension is not merely technical. It is what allows legitimacy to become a true *state variable* and enforcement to become a dynamic investment decision. In the static model, legitimacy affects current behavior but cannot be built or depleted over time. In the two-period model, legitimacy matters tomorrow, but only once. In the infinite-horizon model, by contrast, current enforcement choices shape a persistent stock of public belief, which in turn shapes future compliance, future enforcement costs, and future incentives. This is where the model delivers a theory of norm formation rather than merely a theory of one-shot deterrence.

Let μ_{t-1} denote the public belief at the start of period t. The strategic type's value function satisfies

$$V(\mu) = \min_{\Delta \in V} \left\{ K\big(X_H(\mu, \Delta),\, \Delta\big) + \beta \int_{-\infty}^{\infty} V\big(\mu'(s; \mu, \Delta)\big)\, g(s - \Delta)\, ds \right\}. \tag{3}$$

Let T denote the Bellman operator on the right-hand side, acting on the space of bounded continuous functions on $[0, 1]$.

6.1 Equilibrium

Proposition 5 (Existence of MPBE)**.** *Under Assumptions 1–3, there exists a unique bounded continuous solution $V^* : [0, 1] \to \mathbb{R}_+$ to the Bellman equation* (3)*. Moreover:*

(i) $V^*(\mu)$ is strictly decreasing in μ;

(ii) $V^*(\mu)$ *is convex in* μ*;*

(iii) $V^*(1) = K(0,0) = 0$*;*

(iv) $V^*(0) = \min_{\Delta \in V} K(\lambda(0) v_H^*(0,\Delta), \Delta)/(1-\beta)$.

Proposition 5 establishes the dynamic foundation of the model. The value function is decreasing because higher legitimacy lowers future governance costs period by period. Its convexity is equally important: losses of legitimacy become especially painful when legitimacy is already low. This means that the marginal value of preserving belief is state-dependent, which is what ultimately generates endogenous variation in equilibrium enforcement across legitimacy levels. The boundary conditions also have natural interpretations. At $\mu = 1$, citizens fully trust the government's commitment, voluntary compliance is maximal, and the government can achieve the first best. At $\mu = 0$, legitimacy has fully collapsed, and the government faces the worst sustainable governance environment forever.

Lemma 6 (Monotone Discipline)**.** *Under Assumptions 1–3, in the MPBE:*

(i) The equilibrium enforcement gap $\sigma^*(\mu)$ *is non-increasing in* μ*: more legitimate governments enforce more strictly.*

(ii) $\sigma^*(\mu) \leq \Delta_H^*(\mu)$*: the infinite-horizon gap is weakly below the static optimum at every belief.*

(iii) $\sigma^*(\mu;\beta)$ *is strictly decreasing in* β *for any fixed* $\mu \in (0,1)$ *where* $\sigma^* > 0$.

Lemma 6 extends the two-period reputational logic into a stationary dynamic policy rule. The first result shows that legitimacy disciplines the government state by state: when the government is more trusted, it chooses a smaller enforcement gap in equilibrium. This is not because high-legitimacy governments have lower current costs of laxity; rather, it is because they have more legitimacy to lose. Strict enforcement becomes a way of protecting a valuable stock of belief. The second result says that dynamic discipline always pushes policy weakly below the static hidden-enforcement benchmark, since the strategic type now internalizes future belief erosion. The third result reinforces the interpretation: patience strengthens discipline exactly because future legitimacy matters more.

Taken together, the results imply that enforcement policy is endogenously history-dependent. The same government type may behave differently at different points in time not because its preferences change, but because its legitimacy state changes.

6.2 Long-run Outcomes

Proposition 6 (Legitimacy Erosion)**.** *Under Assumptions 1–3, in any MPBE in which* $\sigma^*(\mu) > 0$ *for all* $\mu \in (0,1)$*:*

(i) If the government is the strategic type ($\theta = S$)*, then* $\mu_t \xrightarrow{a.s.} 0$ *as* $t \to \infty$.

(ii) If the government is the committed type ($\theta = C$)*, then* $\mu_t \xrightarrow{a.s.} 1$ *as* $t \to \infty$.

Proposition 6 gives the model's stark long-run prediction under interior separation. If the government is strategic and continues to allow positive slack, its type is gradually revealed, and legitimacy erodes toward zero. If instead the government is truly committed, repeated strict enforcement eventually persuades citizens, and legitimacy converges to one. In this sense, legitimacy is not cheap talk and not a purely exogenous cultural endowment. It is the long-run informational consequence of an enforcement history.

This proposition is where the paper moves beyond a repeated deterrence story. Citizens do not merely react to current policy; they learn from policy over time, and those beliefs feed back into future compliance. The result is a dynamic process in which enforcement choices shape the very normative environment within which future enforcement operates.

Remark 2 (Equilibrium Selection)**.** Theorem 6 applies to any MPBE with interior enforcement. The existence of pooling equilibria (in which $\sigma^*(\mu) = 0$ for all or some μ) requires the reputational benefit of full enforcement to dominate the current cost of enforcement strictness. As shown in Remark 1, such equilibria exist when β is sufficiently large or when the cost function is sufficiently flat in Δ.

Remark 2 clarifies that long-run erosion is the natural outcome when the strategic type maintains some persistent slack. Pooling remains possible, but only when reputational considerations are powerful enough to induce full imitation. Thus the model accommodates both partial-discipline and full-discipline regimes; the distinction depends on whether the future value of legitimacy is strong enough to overcome current enforcement costs.

Proposition 7 (Low-Legitimacy Trap)**.** *Define the* low-legitimacy trap *as the self-reinforcing dynamic that occurs as $\mu_t \to 0$ under type S. Formally:*

(i) As $\mu_t \to 0$: the equilibrium enforcement gap $\sigma^*(\mu_t) \to \sigma^*(0) \equiv \bar{\Delta}_0 > 0$.

(ii) Aggregate violations converge: $X_H(\mu_t, \sigma^*(\mu_t)) \to \lambda(0)\, v_H^*(0, \bar{\Delta}_0) > 0$.

(iii) The per-period cost converges to $K(\lambda(0)\, v_H^*(0, \bar{\Delta}_0), \bar{\Delta}_0) > K(0, 0) = 0$: the strategic government bears strictly positive costs forever.

(iv) The trap is self-reinforcing*: once μ_t is near 0, citizens have weak voluntary compliance incentives, opportunist violations are high, and enforcement costs are high—all of which persist.*

Proposition 7 delivers the model's substantive punchline. Once legitimacy falls sufficiently low, the government becomes trapped in a bad equilibrium environment. Citizens no longer comply voluntarily, opportunistic violations remain high, and enforcement becomes persistently costly. The government would like a more compliant society, but precisely because it has already lost credibility, it faces a harder and more expensive governing problem. The trap is therefore not just a low-belief state; it is a *self-sustaining governance regime*.

This result helps clarify what the infinite horizon buys relative to the shorter-horizon benchmarks. The static model shows the immediate tension between deterrence and laxity. The two-period model shows that reputation can discipline current enforcement. The infinite-horizon model

shows that these forces cumulate into long-run paths. Enforcement policy today does not just affect today's violations or tomorrow's beliefs; it shapes whether the polity converges toward a high-legitimacy, low-cost order or toward a low-legitimacy, high-cost trap.

6.3 Comparison

Across the three horizons, the model tells a unified story. In the static benchmark, hidden enforcement improves deterrence but tempts the government to exploit opacity. In the two-period model, reputational concerns begin to restrain that temptation. In the infinite horizon, legitimacy becomes an endogenous state, so enforcement is simultaneously a coercive instrument and an investment in future compliance. The central implication is that norm formation is not separate from enforcement policy. It is produced by it. Governments shape compliance not only by punishing violations directly, but by building or dissipating the credibility that makes citizens comply voluntarily.

7 Discussion

7.1 Normative Foundations of Legality: Fuller and Hart

The model's logic that enforcement is both an instrument and a message has deep roots in jurisprudence. The classic debate between Lon Fuller and H.L.A. Hart provides the normative foundation for understanding why congruence between law and enforcement is essential to legitimacy.

Hart, the leading legal positivist, held that the validity of law depends on its social sources rather than its moral content. Law, for him, is a matter of institutional fact: it is valid when it meets the criteria defined by a legal system's rule of recognition (Hart & Green, 2012, pp. 94–110). Morality may influence lawmaking, but it is not necessary for a rule to count as law. Fuller rejected this separation. He argued that legality itself possesses a moral structure that a set of procedural standards without which law ceases to be law. His eight principles of legality—generality, publicity, prospectivity, clarity, consistency, possibility of compliance, constancy, and congruence—form what he called the Internal Morality of Law (Fuller, 1964, p. 39)[1]. Of the eight principles, congruence between declared rule and official action is the central focus of this paper. A system that regularly issues secret, retroactive, or self-contradictory directives cannot claim fidelity to law at all.

For Fuller, legality is not merely a technique of governance but a moral relationship grounded in reciprocity between ruler and ruled. The state promises to govern through publicly knowable rules; citizens, in turn, accept those rules as reasons for action. As he writes,

> [T]here is a kind of reciprocity between government and the citizen with respect to the observance of rules. Government says to the citizen in effect, "these are the rules we

[1] "There must be rules, they must be promulgated, they must be retroactive as seldom as possible, they must be clear, they must not be contradictory, they must not require the impossible, they should not be changed too frequently, and there must be congruence between declared rule and official action."

> expect you to follow. If you follow them, you have our assurance that they are the rules that will be applied to your conduct." When this bond of reciprocity is finally and completely ruptured by government, nothing is left on which to ground the citizen's duty to observe the rules (Fuller, 1964, p.39-40).

This principle of congruence requires officials to act in accordance with the rules they announce, even when deviation might advance short-term goals. Citizens, observing this fidelity, gain assurance that their own sacrifices will be met with institutional restraint. In this way, law functions not as command backed by threat but as moral communication that coordinates expectations.

Fuller's framework clarifies what is at stake in the model. The enforcement gap Δ formalizes the ideal of *incongruence*, the distance between de jure commitment and de facto behavior. The public belief μ_t ope rationalizes *fidelity to law*: citizens' confidence that officials will honor the rules they declare. When governments consistently align action with rule (Δ small), they build trust ($\mu_t \uparrow$); when they deviate, trust erodes ($\mu_t \downarrow$). The resulting dynamics reproduce Fuller's moral intuition within a formal political-economic environment.

This connection also illuminates the divergence between Fuller and Hart. In a Hart's world, law's authority stems from social acceptance of official rules, which is a static condition. In the Fuller's world captured by the model, authority is dynamic: it must be sustained through consistent practice that keeps μ_t within the trust-supporting region. The model thus translates Fuller's moral theory into a dynamic process of legitimacy formation and decay.

Congruence also distinguishes legality from other forms of governance, such as democratic majorities, traditional authority, or charismatic leadership. As Waldron (1994) notes, law has a unique capacity to structure public expectations through impersonal and repeatable norms. These qualities allow legal subjects to treat the law as a source of reasons—as something they can adopt as a guide to action, even when they disagree with particular outcomes. In this reciprocal framework, legal rules do not function as mere commands backed by threat; they are moral signals that invite individuals to participate in a shared normative order.

Compliance is not secured through coercion alone; it depends on citizens' ability to interpret official behavior as consistent with the public meaning of law. When enforcement is regular, fair, and congruent with declared rules, it affirms the law's credibility as a reason-guiding institution. But when enforcement becomes erratic, opaque, or discriminatory, it undermines this moral ecology. Citizens cease to see the law as a trustworthy guide and begin to view it as arbitrary force. In Fuller's terms, this breakdown of congruence represents a collapse of legality itself. He states that when legality is ignored, "the bond of reciprocity is finally and completely ruptured by government" (Fuller, 1964, p.39-40) Without this bond, citizens no longer feel obligated to follow the rules.

The model captures this collapse precisely: when Δ_t grows large and beliefs fall below μ^C, voluntary compliance vanishes and the regime slides into the coercion-dependent state. Fuller's "bond of reciprocity" becomes a mathematical inequality: when congruence is persistently broken, the citizen's best response is defection. Thus, the model provides a behavioral foundation for

Fuller's jurisprudential claim: legality sustains itself only when the state's fidelity to its own rules remains credible in the eyes of those governed.

7.2 Empirical Illustrations and Policy Implication

The theoretical framework explains a range of real-world settings where gaps between written law and actual enforcement erode institutional trust. These cases reveal how credibility and compliance co-evolve and how restoring congruence can stabilize legality.

Traffic regulation. Consider speeding on U.S. highways, is a persistent issue driven by time-saving incentives, weak enforcement perceptions, and prevailing social norms. Route 130 in New Jersey exemplifies the challenge: "Drivers often speed past you at rates of 80 MPH or more in areas that have posted speed limits of 45-55 MPH," [2] creating significant public safety risks. In response to recurring traffic fatalities, lawmakers enacted "Antwan's Law," permanently lowering speed limits and tripling fines for violations.[3] Yet compliance remained weak. The explanation is not legal ignorance or moral opposition but lack of enforcement. The state's failure to match its expressed rules with visible action signals to drivers that speed limits are not truly binding. Without sufficient deterrence, opportunistic drivers continue to speed and crowd-followers conform to prevailing norms rather than statutory commands.

Policing and constitutional rights. This disconnect is not unique to traffic law. In policing, constitutional standards like "reasonable suspicion" were long undermined by discretionary practices. The New York City stop-and-frisk litigation in *Floyd v. City of New York*[4] exposed a systemic divergence between doctrine and practice. Though the state maintained formal legality, its operational behavior became arbitrary and discriminatory. The court intervened only after years of public protest and empirical evidence. Structural reforms, like mandatory body cameras and judicial monitoring, were imposed to reestablish congruence. This case illustrates the cost of sustained incongruence and the institutional mechanisms required to restore legitimacy.

Environmental enforcement. A similar story emerges in environmental regulation. Under the *Clean Air Act*,[5] the EPA sets strict emissions limits, but enforcement has often been discretionary. In the early 2010s, petrochemical plants in Houston routinely violated these limits with minimal consequence. Instead of levying full penalties, the EPA settled cases through negotiated agreements that prioritized industrial continuity. This selective enforcement bred distrust, particularly among affected communities. In response, Congress mandated reforms such as the public-facing ECHO[6]

[2] What Are the Most Dangerous Roads in South Jersey? https://ferraragable.com/blog/what-are-the-most-dangerous-roads-in-south-jersey/.

[3] 2024 New Jersey Revised Statutes Title 39 — Motor Vehicles and Traffic Regulation, Section 39:4-98.11 - Findings and declarations regarding "Antwan's Law", https://law.justia.com/codes/new-jersey/title-39/section-39-4-98-11/.

[4] Floyd v. City of New York, No. 13-3088 (2d Cir. 2014), https://law.justia.com/cases/federal/appellate-courts/ca2/13-3088/13-3088-2014-10-31.html

[5] Summary of the Clean Air Act, https://www.epa.gov/laws-regulations/summary-clean-air-act

[6] Environment and Compliance History Online, https://echo.epa.gov/. See more specific example from "Why did EPA make this drinking water data available?" on the Safe Drinking Water Act (SDWA) Resources and FAQs, https://echo.epa.gov/help/sdwa-faqs.

database, which made enforcement records transparent and signaled a recommitment to regulatory congruence.

Legal legitimacy is not self-sustaining; instead, it must be cultivated through integrity, transparency, and self-correction. Across traffic, policing, and environmental regulation, the pattern is consistent: when people see a persistent gap between what law says and what the state does, they begin to withdraw that allegiance. Rule-followers become disillusioned, and opportunists exploit the gap. These are not abstract risks but features of everyday governance. As Fuller stressed, legitimacy hinges on congruence between declared rules and official conduct. Law draws its authority not only from coercion or procedure, but from its capacity to guide behavior through moral communication. Enforcement is therefore not merely instrumental; it is expressive of whether the system is principled, sincere, and worthy of allegiance. Legitimacy is not self-sustaining, but it must be continually rebuilt through consistent and transparent administration.

From a policy perspective, the model suggests that preserving legitimacy requires treating trust as a depletable asset. Institutions must minimize the enforcement gap even when strictness is costly, because each act of incongruence consumes credibility that is hard to replenish. Routines such as independent audits, publication of enforcement metrics, and regular judicial review can make congruence observable and self-reinforcing. The aim is not maximal punishment but visible integrity: enforcement that signals sincerity and earns voluntary compliance.

7.3 Possible Extensions

Several extensions could deepen the model's behavioral and institutional scope.

Crowd-Followers The baseline model assumes two behavioral types, rule-followers and opportunists, but real societies also include crowd-followers: individuals whose behavior depends on what they observe others doing. These actors respond less to moral conviction than to perceived social regularities. Incorporating them would make aggregate compliance endogenous to social feedback. When the share of visible compliance rises, crowd-followers imitate restraint, amplifying the stabilizing effect of legitimacy; when visible violations spread, they imitate defection, accelerating decay. Formally, the compliance share would depend not only on beliefs about government fidelity but also on observed compliance among peers, which results from government's enforcement effort rather than gap, introducing a coordination externality. This addition would create richer dynamics mirroring the empirical fragility of legitimacy in large, diverse societies.

De jure and De facto Divergence between the *de jure* and *de facto* laws can result not only from intent but also from limited state capacity, bureaucratic discretion, or interpretation. For example, citizens may perceive judicial decisions as partisan, undermining the assumption of neutrality and the interpretive coherence of the legal order. When people losing trust in the legal system, the legal decisions will not be interpreted as a product of impartial deliberation based on the facts and the law for each particular case, but as in favor of or against a particular party or position, and as predetermined by the composition of the decision-making panel (Gibson, 2024). The model could

also incorporate two-way feedback between state and citizens through democratic mechanisms like voting or public deliberation.

External Morality Moral diversity could be built into the model. Unlike the public enforcement of everyday behaviors, where we assume citizens prefer more freedom, moral-related issues often involve citizens having their own ideal levels of restriction. In such cases, external morality, as defined by Fuller, such as on intellectual property, censorship, or reproductive rights, comes into play as citizens evaluate whether a rule is inherently good and designed for the common good. The government, in turn, tends to favor the design of moral laws, as citizens morally assess their appropriateness. However, some laws, though codified, are rarely enforced due to the high cost of enforcement. For example, in developing countries, citizens may have differing preferences regarding restrictions on pirated products—balancing respect for intellectual property against viewing such resources as free tools for education. Despite strict regulations and significant fines for pirated resources like books and movies on the Internet, the government seldom enforces these laws effectively. Some citizens naturally support laws that align with their ideal points, requiring no additional institutional incentives to support the endeavor. However, for those whose preferences diverge from the codified rule, the perception that the government respects them as individuals and maintains consistency between the written rule and the enforced becomes critical for enforcement. If the government abandons this institutional principle of congruency, it risks losing the only foundation on which some citizens are willing to cooperate.

8 Conclusion

Enforcement is both an instrument and a message. This paper shows that how governments apply the law shapes not only immediate compliance but the legitimacy that sustains legal order over time. When enforcement is congruent with declared rules, it signals institutional fidelity, builds trust, and gradually converts obedience into an internalized norm. When it drifts from the rule it purports to uphold, it erodes credibility, normalizes opportunism, and forces the state to rely increasingly on coercion.

The dynamic model formalizes this feedback between enforcement and belief. A government's enforcement gap serves as a public, though noisy, signal of its underlying fidelity to law. Citizens update beliefs from observed enforcement choices, and their compliance behavior responds endogenously. In a high-legitimacy regime, credible enforcement sustains trust, learning shuts down, and compliance becomes self-enforcing. In an intermediate regime, the government strategically alternates between strict enforcement and limited tolerance, using enforcement as reputational investment while trading off short-run cost savings against long-run credibility. In a low-legitimacy regime, signaling incentives are weak, enforcement behaves approximately myopically, and legality is sustained only through persistent coercion.

These dynamics reinterpret legitimacy, and more broadly reputation, as equilibrium objects sustained by credible signaling of fidelity to law. They explain why identical statutes can produce

divergent enforcement norms across jurisdictions: small early differences in practice steer societies onto distinct long-run trajectories of trust and compliance. The same legal code can become self-enforcing in one polity yet brittle and costly in another.

Normatively, legitimacy is not a luxury of governance but its foundation. Effective legal systems maintain congruence between law on the books and law in action, treating credibility as a scarce institutional asset. Strategic deviations may conserve resources in the short run, but they deplete the moral capital that allows law to rule through consent rather than coercion. To sustain legality's authority, institutions must protect transparency, consistency, and visible fidelity to their own rules. Credible enforcement does more than deter: it invites citizens to comply, transforming obedience into affirmation and coercion into consent.

Proof Appendix

A Static Benchmark

Proof of Lemma 1

Proof. The maximizer $\hat{v}(\mu)$ satisfies the first-order condition $b'(\hat{v}) - \mu w p'(\hat{v}) = 0$. By implicit differentiation, $\partial \hat{v}/\partial \mu = wp'(\hat{v})/[b''(\hat{v}) - \mu w p''(\hat{v})] < 0$, since $b'' < 0$, $p'' > 0$, and the denominator is negative by second-order conditions. The inequality $\hat{v}(\mu) \geq v^s$ follows from $\mu < 1$: reducing the effective sanction from wp to μwp increases the unconstrained optimum. □

A.1 Observed Enforcement

Proof of Lemma 2

Proof. For $\Delta \leq \Delta^O$: the maximum attainable "safe" payoff $b(\Delta) \leq b(\Delta^O) = b(v^s) - wp(v^s)$ is weakly less than the sanction-exposed optimum $b(v^s) - wp(v^s)$, so the citizen prefers $v = v^s$. For $\Delta > \Delta^O$: $b(\Delta) > b(v^s) - wp(v^s)$, so the citizen strictly prefers boundary-hugging at $v = \Delta$ to any sanction-exposed violation. Uniqueness of Δ^O follows from strict monotonicity of b. □

Proof of Proposition 1

Proof. (i) For $\Delta \in (0, \Delta^O)$, $v_O^*(\Delta) = v^s$ and $X_O = \lambda(\mu) v^s$, which is independent of Δ. Hence $\partial K/\partial \Delta = K_\Delta < 0$, so increasing Δ strictly reduces cost—no optimum in this region. (ii) For $\Delta \in (\Delta^O, \bar{v})$, $v_O^*(\Delta) = \Delta$ and the cost becomes $K(\lambda(\mu)\Delta, \Delta)$. The FOC follows from differentiating. (iii) At $\Delta = \bar{v}$, $K_\Delta(\cdot, \bar{v}) = 0$ and $K_X(\cdot, \bar{v}) > 0$ with $X > 0$, so $\frac{d}{d\Delta} K\big|_{\Delta=\bar{v}} = \lambda(\mu) K_X > 0$, meaning cost is increasing in Δ at $\bar{v}$—not optimal. □

A.2 Hidden Enforcement

Proof of Lemma 3

Proof. Existence and uniqueness of $\Delta^H(\mu)$. Define $\phi(v) \equiv b(v) - \mu w p(v)$, the payoff from boundary violation $v \leq \Delta$. We have $\phi(0) = 0$, $\phi(\hat{v}(\mu)) = \max_v \phi(v) = b(\hat{v}(\mu)) - \mu w p(\hat{v}(\mu))$. The sanction-exposed payoff at the interior optimum is $b(v^s) - wp(v^s) \equiv \underline{\pi}$. Since $\mu < 1$, $\phi(\hat{v}(\mu)) > \underline{\pi}$, and $\phi(0) = 0 < \underline{\pi}$, continuity implies a unique $\Delta^H(\mu) \in (0, \hat{v}(\mu))$ with $\phi(\Delta^H(\mu)) = \underline{\pi}$.

Best response structure. For $\Delta < \Delta^H(\mu)$: $b(\Delta) - \mu w p(\Delta) < \underline{\pi}$, so the boundary payoff is inferior; the citizen chooses v^s. For $\Delta^H(\mu) \leq \Delta < \hat{v}(\mu)$: boundary payoff is $b(\Delta) - \mu w p(\Delta) \geq \underline{\pi}$ and the unconstrained within-tolerance optimum $\hat{v}(\mu) > \Delta$, so the binding constraint is $v \leq \Delta$; the citizen chooses $v = \Delta$. For $\Delta \geq \hat{v}(\mu)$: the constraint is non-binding; citizen chooses $\hat{v}(\mu)$.

$\Delta^H(\mu) < \Delta^O$. Since $\mu < 1$, $\phi(\Delta^O) = b(\Delta^O) - \mu w p(\Delta^O) > b(\Delta^O) - wp(\Delta^O) = b(v^s) - wp(v^s) = \underline{\pi}$ (the last equality defining Δ^O). By strict monotonicity of ϕ near $\Delta^H(\mu)$, $\Delta^H(\mu) < \Delta^O$.

Monotonicity. Differentiating $\phi(\Delta^H(\mu); \mu) = \underline{\pi}$ implicitly: $\phi_\Delta \cdot \frac{d\Delta^H}{d\mu} + \phi_\mu = 0$. We have $\phi_\mu = -wp(\Delta^H) < 0$ and $\phi_\Delta = b'(\Delta^H) - \mu w p'(\Delta^H) > 0$ (since $\Delta^H < \hat{v}(\mu)$), so $\frac{d\Delta^H}{d\mu} = -\phi_\mu/\phi_\Delta > 0$. □

Proof of Proposition 2

Proof. (i) For $\Delta \in (0, \Delta^H(\mu))$, $v_H^* = v^s$ and X_H is independent of Δ; $K_\Delta < 0$ implies cost is strictly decreasing in Δ, ruling out an optimum. (ii) On $(\Delta^H(\mu), \hat{v}(\mu))$, $v_H^* = \Delta$ and differentiation gives the stated FOC. (iii) On $[\hat{v}(\mu), \bar{v}]$, violations plateau at $\lambda(\mu)\hat{v}(\mu)$, so aggregate-violation costs do not change with Δ. Since $K_\Delta < 0$ for $\Delta < \bar{v}$, the government minimizes cost by setting $\Delta = \bar{v}$. □

Proof of Proposition 3

Proof. (i) For $\Delta < \Delta^H(\mu)$: $v_H^* = v^s = v_O^*$, so violations are equal. For $\Delta \in [\Delta^H(\mu), \Delta^O)$: $v_H^* = \Delta < v^s = v_O^*$, giving strict improvement in deterrence. For $\Delta \geq \Delta^O$: $v_H^* = \min(\Delta, \hat{v}(\mu)) \leq \Delta = v_O^*$, so $X_H \leq X_O$. (ii) Under hidden enforcement, for $\Delta \geq \hat{v}(\mu)$, X_H is constant in Δ, and $K_\Delta < 0$ for $\Delta < \bar{v}$, making $\Delta = \bar{v}$ strictly preferable within this region. Under observed enforcement, $X_O = \lambda(\mu)\Delta$ continues to increase, and Proposition **??**(iii) rules out $\Delta = \bar{v}$. □

B Two-period Model

B.1 Period-2 Subgame

Proof of Lemma 4

Proof. Differentiability follows from the envelope theorem and the regularity of K and the optimal policy $\Delta_2^*(\mu_1)$. For the monotonicity: by the envelope theorem,

$$C_2'(\mu) = K_X \frac{\partial X_H}{\partial \mu} + K_\Delta \frac{\partial \Delta_2^*}{\partial \mu}.$$

Higher μ reduces $\lambda(\mu)$ (fewer opportunists), directly reducing X_H and hence $K_X \partial X_H / \partial \mu < 0$. Higher μ also raises $\Delta^H(\mu)$, shifting the boundary of the feasible interior region upward. The government can enforce at a less costly gap while maintaining deterrence. Both effects reinforce each other, and the net effect is $C_2'(\mu) < 0$.

More explicitly, for an interior optimum:

$$\begin{aligned} C_2'(\mu) &= \underbrace{K_X(\cdot)\,\lambda'(\mu)\,\Delta_2^*}_{<0} + \underbrace{K_X(\cdot)\,\lambda(\mu)\,\frac{\partial \Delta_2^*}{\partial \mu} + K_\Delta(\cdot)\,\frac{\partial \Delta_2^*}{\partial \mu}}_{\text{FOC}=0} \\ &= K_X(\cdot)\,\lambda'(\mu)\,\Delta_2^* < 0, \end{aligned}$$

where the FOC from Proposition 2(ii) zeroes out the terms involving $\partial \Delta_2^* / \partial \mu$. □

B.2 Period-1 Problem

Proof of Lemma 5

Proof. (i) At $\Delta_1 = 0$, the strategic type mimics the committed type: $s_1 \sim g(\cdot) = g(\cdot - 0)$, so the posterior $\mu'(s; \mu_0, 0) = \mu_0$ for all s (the likelihood ratio is identically 1). Thus the expectation equals $C_2(\mu_0)$.

(ii) We show $\partial\Phi/\partial\Delta_1 > 0$ for $\Delta_1 > 0$. Write $\Phi(\Delta_1) = \int C_2(\mu'(s;\mu_0,\Delta_1))g(s-\Delta_1)\,ds$. Differentiating (using dominated convergence):

$$\frac{\partial\Phi}{\partial\Delta_1} = \int C_2'(\mu')\,\frac{\partial\mu'}{\partial\Delta_1}\,g(s-\Delta_1)\,ds - \int C_2(\mu')\,g'(s-\Delta_1)\,ds. \tag{4}$$

By MLRP, $\partial\mu'/\partial\Delta_1 < 0$ at each s when $\Delta_1 > 0$: a larger gap shifts the signal distribution right, lowering posterior at each realization s. Since $C_2' < 0$, the first integral is positive. For the second integral, $-\int C_2 g'(s-\Delta_1)ds = \int C_2(-g'(s-\Delta_1))ds$. An increase in Δ_1 shifts the distribution of s_1 rightward. Since higher s means lower posterior (by MLRP) and C_2 is decreasing in posterior, higher s means higher C_2. Hence a rightward shift in the distribution of s_1 raises $\int C_2 g(s-\Delta_1)ds$. Both effects of increasing Δ_1 raise Φ, confirming $\partial\Phi/\partial\Delta_1 > 0$.

(iii) At $\Delta_1 = 0$, the committed and strategic types are indistinguishable and the posterior is degenerate at μ_0 for all s. However, any positive deviation $\Delta_1 = \varepsilon > 0$ introduces a positive likelihood ratio differential: $g(s-\varepsilon)/g(s) \neq 1$ for a positive-measure set of s (since g has full support and $\varepsilon > 0$). The KL divergence between $g(\cdot)$ and $g(\cdot - \varepsilon)$ is positive, generating strictly positive expected posterior dispersion and hence strictly positive marginal reputational cost. □

B.3 Equilibrium

Proof of Proposition 4

Proof. (i) and (ii): The two-period objective $\mathcal{C}_1(\Delta_1)$ is C^2 in Δ_1 on the interior region. At the static optimum $\Delta_H^*(\mu_0)$, the static FOC holds: $\lambda(\mu_0)K_X + K_\Delta = 0$. Hence $\frac{\partial\mathcal{C}_1}{\partial\Delta_1}\big|_{\Delta_H^*} = 0 + \beta\,\frac{\partial\Phi}{\partial\Delta_1}\big|_{\Delta_H^*} > 0$ by Lemma 5(ii). Thus the two-period objective is strictly increasing at the static optimum, so the two-period minimizer lies strictly to the left: $\Delta_1^{**} < \Delta_H^*(\mu_0)$. The FOC (2) follows by differentiating $\mathcal{C}_1$.

(iii): Differentiating the FOC (2) implicitly with respect to β:

$$\frac{\partial\Delta_1^{**}}{\partial\beta} = -\frac{\partial\Phi/\partial\Delta_1}{\partial^2\mathcal{C}_1/\partial\Delta_1^2} < 0,$$

since the numerator is positive (Lemma 5) and the denominator is positive (second-order conditions at the minimum).

(iv): The marginal reputational cost $\partial\Phi/\partial\Delta_1$ is proportional to $-C_2'(\mu_0)$ through the chain rule: a steeper C_2 amplifies the cost of posterior dispersion. The same implicit differentiation argument yields $\partial\Delta_1^{**}/\partial(-C_2') < 0$. □

C Infinite Horizon

C.1 Equilibrium

Proof of Proposition 5

Proof. Existence and uniqueness: The Bellman operator T is a contraction on the space of bounded continuous functions on $[0,1]$ (with the sup-norm), by the discount factor $\beta \in (0,1)$ and the boundedness of K. Banach's fixed-point theorem gives a unique fixed point V^*.

(i) Monotonicity: We verify that the set of bounded continuous decreasing functions is T-invariant. Suppose V is decreasing. Then $TV(\mu) = \min_\Delta\{K(X_H, \Delta) + \beta \int V(\mu')\, g(s-\Delta)\, ds\}$. Higher μ reduces $\lambda(\mu)$, hence $K(X_H, \Delta)$ weakly for any fixed Δ. Moreover, higher μ stochastically dominates the posterior distribution (in the sense of first-order stochastic dominance in μ' with respect to μ at fixed Δ and fixed signal realization). Since V is decreasing, $\int V(\mu')g(s-\Delta)ds$ is decreasing in μ. Thus $TV(\mu)$ is decreasing in μ, and the set is T-invariant.

(ii) Convexity: Analogous T-invariance argument using the fact that a mixture of decreasing convex functions is convex.

(iii) At $\mu = 1$, all citizens comply voluntarily ($\lambda(1) = 0$), so $X_H = 0$ for any Δ, and $K(0, \cdot)$ is minimized at $\Delta = 0$ with $K(0,0) = 0$.

(iv) At $\mu = 0$, the committed type is known to be absent, so beliefs remain at 0 for all future periods. The problem reduces to a stationary optimization of $K(\lambda(0)v_H^*(0,\Delta), \Delta)$ each period, and the value is the discounted sum. □

Proof of Lemma 6

Proof. (i): By convexity of V^* and the separability of the current cost and future cost, the optimality condition at interior solutions is:

$$\lambda(\mu)K_X(X_H, \sigma^*) + K_\Delta(X_H, \sigma^*) + \beta \frac{\partial}{\partial \Delta} \int V^*(\mu'(s;\mu,\Delta))g(s-\Delta)\, ds \Big|_{\Delta=\sigma^*} = 0. \tag{5}$$

Denote the reputational term $\Psi(\mu, \Delta) \equiv \beta \frac{\partial}{\partial \Delta} \int V^*(\mu')g(s-\Delta)ds$. By analogous arguments to Lemma 5, $\Psi > 0$: a higher gap shifts the signal distribution right, lowering posteriors (by MLRP) and raising future costs (since V^* is decreasing in μ).

For monotonicity of $\sigma^*(\mu)$: differentiating (5) implicitly and using the second-order conditions plus the comparative statics of K and Ψ with respect to μ, one can show that $d\sigma^*/d\mu \leq 0$. The key force is that higher μ reduces $\lambda(\mu)$ (fewer opportunists), which reduces the current marginal cost of tightening enforcement (the $\lambda(\mu)K_X$ term is smaller), making strict enforcement more attractive. Simultaneously, higher μ means more future legitimacy at stake, amplifying the reputational incentive term Ψ.

(ii): The argument parallels Theorem 4(i): at the static optimum $\Delta_H^*(\mu)$, the static FOC holds, but $\Psi > 0$ means the infinite-horizon objective is still increasing. The minimizer lies strictly to the left.

(iii): The reputational term $\beta\Psi$ is proportional to β; higher β amplifies the reputational cost of laxity. The same implicit differentiation argument as in Theorem 4(iii) gives $\partial\sigma^*/\partial\beta < 0$. □

C.2 Long-run Outcomes

Proof of Proposition 6

Proof. The sequence $\{\mu_t\}$ is a bounded martingale with respect to the public filtration $\{\mathcal{H}_t\}$: by the law of iterated expectations and Bayes' rule, $\mathbb{E}[\mu_t \mid \mathcal{H}_{t-1}] = \mu_{t-1}$. By the Martingale Convergence Theorem, $\mu_t \to \mu_\infty$ a.s. for some $[0,1]$-valued random variable μ_∞.

$\mu_\infty \in \{0,1\}$ *a.s.*: Suppose $\mu_\infty = \mu^* \in (0,1)$ with positive probability. At a limit point μ^*, the optimal gap is $\sigma^*(\mu^*) > 0$ by assumption. The signal distributions under C and S are $g(\cdot)$ and $g(\cdot - \sigma^*(\mu^*))$ respectively, which are distinct (positive KL divergence under MLRP with $\sigma^* > 0$). By the law of large numbers for martingales (see, e.g., **?**), continued distinct signal distributions imply that beliefs do not stabilize at $\mu^* \in (0,1)$—a contradiction. Hence $\mu_\infty \in \{0,1\}$ a.s.

Identification of the limit: Under type C, signals $s_t \sim g(\cdot)$ each period, and the likelihood ratio $L_t = \prod_{k=1}^{t} g(s_k - \sigma^*(\mu_{k-1}))/g(s_k) \to 0$ a.s. by the strong law, since $\mathbb{E}[\log(g(s - \sigma^*)/g(s))] < 0$ when $\sigma^* > 0$ (negative cross-entropy). This forces $\mu_t \to 1$ a.s. under C. Symmetrically, under type S, $L_t \to \infty$ a.s., forcing $\mu_t \to 0$ a.s. □

Proof of Proposition 7

Proof. (i) Continuity of σ^* at 0, established in Theorem 5, and the Assumption **??**(v) boundary condition ruling out full tolerance imply $\sigma^*(0) = \bar{\Delta}_0 > 0$.

(ii) Follows from $\lambda(0) = 1$ (all citizens are opportunists when $\mu = 0$) and continuity of v_H^*.

(iii) $\lambda(0) > 0$ and $v_H^*(0, \bar{\Delta}_0) > 0$ imply $X > 0$, and $K_X > 0$ gives positive cost.

(iv) The reinforcing mechanism: low $\mu_t \Rightarrow$ high $\lambda(\mu_t) \Rightarrow$ more opportunists $\Rightarrow$ higher violations $\Rightarrow$ higher enforcement cost $\Rightarrow$ government continues to choose $\sigma^* > 0 \Rightarrow$ continued belief erosion $\Rightarrow$ lower μ_{t+1}, perpetuating the cycle. □

References

Acemoglu, D. & Jackson, M. O. (2017). Social norms and the enforcement of laws. *Journal of the European Economic Association*, 15(2), 245–295.

Bassetto, M. & Phelan, C. (2008). Tax riots. *The Review of Economic Studies*, 75(3), 649–669.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of political economy*, 76(2), 169–217.

Becker, G. S. & Stigler, G. J. (1974). Law enforcement, malfeasance, and compensation of enforcers. *The Journal of Legal Studies*, 3(1), 1–18.

Benabou, R. & Tirole, J. (2006). Belief in a just world and redistributive politics. *The Quarterly journal of economics*, 121(2), 699–746.

Benabou, R. & Tirole, J. (2011). *Laws and norms*. Technical report, National Bureau of Economic Research.

Besley, T. & Dray, S. (2024). Trust and state effectiveness: The political economy of compliance. *The Economic Journal*, 134(662), 2225–2251.

Chen, Y., Du, K., Stocken, P. C., & Wang, Z. (2025). Peer learning, enforcement, and reputation. *Enforcement, and Reputation (May 17, 2025)*.

Dickson, E. S., Gordon, S. C., & Huber, G. A. (2022). Identifying legitimacy: Experimental evidence on compliance with authority. *Science Advances*, 8(7), eabj7377.

Dilmé, F. & Garrett, D. F. (2019). Residual deterrence. *Journal of the European Economic Association*, 17(5), 1654–1686.

Fenske, J., Haseeb, M., & Kala, N. (2023). *How rules and compliance impact organizational outcomes: Evidence from delegation in environmental regulation*. Technical report, National Bureau of Economic Research.

Fuller, L. L. (1964). The morality of law.

Gibson, J. L. (2024). Losing legitimacy: The challenges of the dobbs ruling to conventional legitimacy theory. *American Journal of Political Science*.

Gilbert, M. D. (2015). Insincere rules. *Va. L. Rev.*, 101, 2185.

Hart, H. L. A. & Green, L. (2012). *The concept of law*. oxford university press.

Holland, A. C. (2016). Forbearance. *American political science review*, 110(2), 232–246.

Huber, J. D. & McCarty, N. (2004). Bureaucratic capacity, delegation, and political reform. *American Political Science Review*, 98(3), 481–494.

Kartik, N., Lipnowski, E., & Pei, H. (2025). Replacement and reputation. *arXiv preprint arXiv:2512.13351*.

Kreps, D. M. & Wilson, R. (1982). Reputation and imperfect information. *Journal of economic theory*, 27(2), 253–279.

Ma, J., Schnakenberg, K. E., & Turner, I. R. (2025). *Democratic Accountability with Citizen Coproduction*. Technical report.

Milgrom, P. & Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of economic theory*, 27(2), 280–312.

Milgrom, P. R., North, D. C., & Weingast*, B. R. (1990). The role of institutions in the revival of trade: The law merchant, private judges, and the champagne fairs. *Economics & Politics*, 2(1), 1–23.

Patty, J. W. & Penn, E. M. (2011). A social choice theory of legitimacy. *Social Choice and Welfare*, 36(3), 365–382.

Penn, E. M. & Patty, J. W. (2023). Algorithms, incentives, and democracy. *arXiv preprint arXiv:2307.02319*.

Pettit, P. (2023). *The State*. Princeton: Princeton University Press.

Shavell, S. (2004). *Foundations of economic analysis of law*. Harvard University Press.

Stokes, L. C. & Hunnicutt, P. (2018). Evaluating water conservation policy in california.

Stone, R. (2016). Legal design for the good man. *Va. L. Rev.*, 102, 1767.

Tyler, T. R. (2001). Trust and law abidingness: A proactive model of social regulation. *BUL Rev.*, 81, 361.

Tyler, T. R. (2006). *Why people obey the law*. Princeton university press.

Tyler, T. R. (2025). Legitimacy-based policing. *Criminology & Public Policy*, 24(2), 165–187.

Tyler, T. R. & Mitchell, G. (1993). Legitimacy and the empowerment of discretionary legal authority: The united states supreme court and abortion rights. *Duke LJ*, 43, 703.

Waldron, J. (1994). Why law-efficacy, freedom, or fidelity? *Law & Phil.*, 13, 259.