

FORMAL FOUNDATIONS · AUTONOMOUS AI GOVERNANCE

Law as Computation

A formal derivation of normative logic from empirical constraints of biological viability — yielding a unique value functional, a unique measure of regulatory threat, and a uniquely admissible principle of legal choice.

AUTHORS

Yuri Kozlov


JudgeAI Research

Taaha Bajwa

JudgeAI Research

Maria Shutova

JudgeAI Research



KEYWORDS: AI LAWMAKING · VIABILITY THEORY · MINIMAX REGULATION · LEGAL ONTOLOGY

MSC 2020: 91B06 · 91B52 · 91D10 · 68T27 · 68T20

JudgeAI
JUDGEAI.SPACE

Law as Computation

Yuri Kozlov
JudgeAI Research

Taaha Bajwa
JudgeAI Research

Maria Shutova
JudgeAI Research

25 February 2026

Abstract

This paper examines the formal conditions of autonomous AI-based lawmaking. An autonomous normative decision is defined as the outcome of an internal normative logic in which the admissibility of legal rules is determined by empirically observable consequences of their application and specified through formally defined conditions of normative admissibility. These conditions constitute an internal criterion of normative status and enable the differentiation of regulatory outcomes without reliance on interpretive judgment.

The normative logic is formalized on the basis of empirical constraints on agent viability and the causal structure of threats. Viability is defined as the space of realizable and legally relevant action trajectories available to an agent within a regulated system, encompassing biological, personal, economic, and institutional dimensions. Regulatory intervention is represented as a transformation of this space through causally definable factors affecting its structure and volume.

The formal analysis yields the following results. The Space of Viable Dynamics (SVD) functions as the value functional representing agents' normative interests. The Threat Index (TI) provides a quantitative measure of regulatory transformations of SVD. The SVD-minimax principle specifies the rule of normative choice under conditions of interdependent behavior and uncertainty.

The derived normative logic is further examined in an empirical appendix through controlled test scenarios, confirming the stability and reproducibility of the formally established results.

Keywords: artificial intelligence lawmaking, autonomous regulation, viability theory, threat index, minimax regulation, AI governance

MSC 2020: 91B06, 91B52, 91D10, 68T27, 68T20

1 Introduction

The present work is motivated by a development that is structural in character and does not depend on any particular policy preference: the progressive and accelerating automation of the production of goods and the provision of services across all sectors of human activity. Manufacturing, logistics, medicine, education, financial services, legal assistance, and the full range of personal services are being transferred, at increasing speed and scale, to automated systems operating without continuous human supervision. This transformation is not a forecast but an observable trajectory. Its normative implications are the subject of this paper.

The normative implication of full-scale automation is a quantitative one that crosses a qualitative threshold. As the volume of normatively significant events — transactions, disputes, regulatory decisions, allocations of rights and obligations — grows in proportion to the speed of automated systems, the throughput capacity of human normative agents becomes a structural bottleneck. A legal system calibrated for a world in which goods are produced and disputes arise at human speed cannot, without fundamental redesign, function in a world where the same processes operate at machine speed. The expectation of waiting months for the resolution of a consumer dispute is coherent only where the underlying commercial transaction itself unfolds over months. Where the transaction is completed in milliseconds, the normative infrastructure governing it must operate at a corresponding scale.

The inadequacy of human normative institutions is, however, not merely a function of throughput. Empirical observation of legislative and regulatory responses to systemic threats — including climate change, global pandemics, and transboundary armed conflict — reveals a distinct and independent failure mode: the structural incapacity of human lawmaking to produce timely, coherent, and effective normative responses to threats whose causal dynamics exceed the cognitive and political horizons of individual jurisdictions. In each of these cases, the relevant data exist; the causal mechanisms are documented; the required regulatory interventions are, in many instances, formally specified. Yet the normative output of human institutions has consistently been delayed, fragmented, and insufficient relative to the scale of the threat. This pattern is not contingent on the particular composition of any legislature or the particular configuration of any political system. It reflects a structural property of collective human deliberation under conditions of high causal complexity, long time horizons, and distributed costs and benefits. The processing demands imposed by such threats are not merely large but qualitatively different in kind from those that human deliberative institutions were designed to handle.

These two failure modes — quantitative overload and qualitative incapacity under systemic complexity — jointly generate a structural necessity: normative decisions must be delegated, in whole or in substantial part, to autonomous artificial systems. This delegation is not a design choice among alternatives. It is a logical consequence of the incompatibility between the normative demands of an automated and increasingly interdependent world and the biological and institutional constraints on human deliberative capacity. The question, therefore, is not whether artificial intelligence will function as a normative agent, but under what conditions such functioning is logically coherent.

A normative agent, in the sense relevant here, is a system that does not merely execute pre-specified rules but produces normative decisions — determinations of what is admissible and what is not — in response to novel factual configurations, including configurations not anticipated by any prior rule. This capacity is required for dispute resolution, for regulatory lawmaking, and for the internal normative layer of AI systems that are themselves actors in regulated environments. In all three domains, the system must be capable of distinguishing, on the basis of an internal criterion, between normatively correct and normatively erroneous outputs.

This requirement — the existence of an internal criterion of normative correctness — is precisely

what existing approaches fail to supply. The literature on AI governance has produced extensive documentation of regulatory requirements, alignment principles, and procedural safeguards. What it has not produced is a formal account of what makes a normative decision correct as such, independently of the authority of the institution that issued it or the procedure by which it was reached. In the absence of such an account, an autonomous system can replicate the form of normative decision-making without possessing its substance: it can output decisions that are labeled as normative without containing any internal mechanism capable of distinguishing normative success from normative failure.

The root of this failure is not technological but theoretical. The legal and philosophical traditions that have historically supplied the normative foundations of law — theories grounded in procedural legitimacy, interpretive coherence, or social consensus — were constructed without the requirement of computability. They presuppose a human interpreter who applies judgment, navigates ambiguity, and draws on tacit practical knowledge. These resources are unavailable to an autonomous computational system. An artificial intelligence operates exclusively on formally representable structures. Any normative logic that is to govern its behavior must therefore be mathematically formalizable: it must be expressible in terms that the system can compute. Legal theory that is not mathematizable is, for an autonomous AI system, not a theory at all — it is a collection of natural-language instructions with no determinate computational interpretation.

The present paper addresses this gap. We ask whether a mathematizable normative logic can be constructed that is not merely postulated but derived — derived from empirically observable facts about the structure of agents and their environment, in such a way that the conditions of normative correctness and normative error are formally determinate. A positive answer to this question is a necessary precondition for any genuine autonomous normative agency: without it, an AI system operating in a normative domain is not a normative agent but a normative simulacrum.

The autonomous use of artificial intelligence in lawmaking is possible only if the system contains an internal normative logic embedded within its architecture. In the absence of such logic, an artificial intelligence system is unable to distinguish between normatively admissible and normatively erroneous decisions and, therefore, cannot perform lawmaking as a rational activity. The processing of texts, data, or statistical correlations alone does not constitute a lawmaking function.

The key requirement for the internal logic of AI-driven lawmaking is falsifiability. In the present context, falsifiability means the presence within the system of formally defined conditions of normative rejection, under which a legal norm must be recognized as normatively erroneous and lose its admissibility. Normative error must be a logically identifiable state of the system, rather than the result of external evaluation or subsequent revision.

This implies that the internal normative logic of the system must specify, in advance, classes of empirically observable consequences of a norm’s application whose occurrence logically excludes its normative admissibility. If no observable consequence is capable of triggering normative rejection, the system is unfalsifiable, and the distinction between successful and failed regulation does not exist within it.

From this it follows that empirical data can participate in the lawmaking process only insofar as the internal normative logic links certain regulatory consequences to a change in the normative status of a rule. Large-scale datasets, by themselves, do not possess normative significance and do not ensure the system’s verifiability. In the absence of an internal logical mechanism of normative rejection, such data remain purely descriptive and have no effect on a norm’s admissibility.

For this reason, the falsifiability of AI-driven lawmaking cannot be achieved through external evaluation procedures, including benchmarks. A benchmark measures the degree of conformity of system behavior to a predefined reference standard, but it does not define conditions of normative error. Even a negative benchmark result does not entail the logical inadmissibility of a norm

unless such inadmissibility follows from the system’s internal logic. Consequently, benchmarks are structurally incapable of replacing a mechanism of normative rejection.

Existing legal theories and doctrines are categorically unsuitable as the internal logic of autonomous AI-driven lawmaking. In theories based on source and procedure, normative admissibility does not depend on a rule’s consequences. In theories based on interpretation and values, normative validity is determined by argumentative coherence or external goals. In both cases, formally defined conditions of normative rejection tied to empirically observable effects of a norm’s application are absent.

From the foregoing, a logically irreversible conclusion follows: a system that does not contain a falsifiable internal normative logic cannot, in principle, perform autonomous lawmaking. In such a system, any norm remains admissible regardless of its consequences, and the distinction between normative success and normative failure does not exist. Consequently, artificial intelligence in lawmaking is possible only on the basis of a normative logic in which admissibility is a function of empirically observable consequences of a norm’s application and in which the conditions of normative rejection are formally defined.

The requirement of a falsifiable internal normative logic formulated above raises a further question: whether such a logic can be derived from empirical facts alone, rather than postulated as a normative axiom. In the remainder of this work, we provide a formal answer to this question. We show that the conditions of normative rejection required for autonomous AI lawmaking can be grounded in empirically observable constraints of biological viability and causal threat structure. The resulting framework yields a unique value functional (the Space of Viable Dynamics), a unique measure of normative threat (the Threat Index), and a uniquely admissible regulatory principle (the SVD–minimax rule). Together, these elements constitute a falsifiable normative logic in the precise sense required by the Introduction.

2 Formal Proof

We develop a formal framework in which the core elements of a legal–regulatory order are derived directly from empirical constraints of biological viability and causal threat structure, rather than from normative axioms. We prove three results:

(1) the Space of Viable Dynamics (SVD) is the unique value functional representing the preferences of biological agents over possible world trajectories;

(2) the Threat Index (TI) is the uniquely admissible measure of threats to SVD across regulatory regimes;

(3) the only regulatory principle compatible with both the continued existence of society and rational ex ante participation is the SVD–minimax principle, which minimizes the maximum probability of irreversible viability loss for any agent.

The main theorem establishes that a sustainable legal order is uniquely determined by empirical data about viability and threats, rather than by any independent moral assumptions.

These results provide a formally falsifiable normative logic suitable for autonomous AI-driven lawmaking.

A central question for legal theory, institutional design, and AI-assisted governance is whether a regulatory structure can be derived from empirical facts rather than from normative postulates. In this work, we show that such a derivation is possible.

We model agents as biological systems embedded in a causal environment. Each agent’s set of viable world trajectories defines a viability measure; the volume of this set, under a physically grounded probability measure, yields the Space of Viable Dynamics (SVD). Agents are exposed to

causal chains of potentially irreversible threats, whose aggregated effects define a Threat Index (TI).

We prove three lemmas:

SVD is the unique value functional of biological agents.

TI is the uniquely correct measure of threats to SVD.

A society that seeks to maintain its own existence and maximize member viability can only select regulations satisfying the SVD–minimax principle.

This leads directly to the main theorem: optimal regulation is uniquely determined by empirical viability constraints.

2.1 Formal Framework

Let Ω denote the set of all possible world trajectories. Let I denote the set of agents.

For each agent i :

$H_i(t)$ describes internal states,

$E(t)$ describes the external environment.

A trajectory ω belongs to a viable set $V_i \subseteq \Omega$ if agent i remains alive and functional across the time horizon.

Let μ be a physically grounded measure on Ω (probabilistic or frequency-based).

Definition (SVD).

$$\text{SVD}_i = \mu(V_i).$$

Interpretation of the scale:

$\text{SVD}_i = 1$ — minimal viability threshold (life preserved).

$\text{SVD}_i > 1$ — extended viability (quality, resources, freedom).

$\text{SVD}_i < 1$ — threat zone with positive probability of irreversible failure.

Without loss of generality, the scale is normalized so that 1 is the viability threshold.

A regulatory pattern R induces a probability distribution over Ω . Let $\text{SVD}_i^{\text{life}}(R)$ denote the viability of i under R . Let $\text{SVD}_i^{\text{out}}$ denote viability outside society (isolation mode).

2.2 Lemma 1: SVD as the Unique Value Functional

In this section we make precise why, under minimal and empirically grounded assumptions, the preferences of a biological agent over possible worlds can be represented by a single scalar quantity proportional to the measure of its viability set, $\text{SVD}_i = \mu(V_i)$.

2.2.1 Outcomes, viability sets and environments

Fix an agent i .

We consider:

Ω — the space of possible world trajectories.

For each environment A (a configuration of physical, social and regulatory conditions), there is an induced probability measure μ_A on Ω .

$V_i \subseteq \Omega$ — the viability set of agent i : trajectories in which i remains alive and functionally capable over its lifespan.

For each environment A , we define the viability volume of i under A as:

$$v_i(A) = \mu_A(V_i).$$

Since viability is defined through the existence of realizable trajectories that preserve an actor’s functionality over time, any change in the configuration of reality that restricts the admissible action

space necessarily reduces the set of such trajectories. Accordingly, any influence that constrains realizable choice, irrespective of its physical, social, or institutional origin, is formally expressed as a transformation of the actor’s Space of Viable Dynamics.

This is the probability (or long-run frequency) that agent i follows a viable trajectory when embedded in environment A .

Intuitively, $v_i(A)$ measures how frequently a given environment allows the agent to remain alive and function, rather than leading to death or functional collapse.

Our goal is to show that any empirically grounded preference ordering of agent i over environments A, B, \dots must be representable as a strictly increasing function of $v_i(A)$. After normalization, this function coincides with SVD_i .

2.2.2 From empirical facts to a viability-based preference ordering

We now translate the empirical premises E1–E3 into properties of a preference relation.

Let “ A is weakly preferred to B by agent i ” be denoted $A \succeq_i B$.

We do not postulate \succeq_i as a primitive; instead, we derive it from the biological setting.

E1 (Viability constraint). There exists a set V_i of viable trajectories; leaving V_i leads to irreversible failure. Therefore, for any environment A , the only trajectories that can support long-run behaviour and reproduction belong to V_i .

E2 (Evolutionary filtering). Behavioural and cognitive patterns that frequently lead outside V_i are eliminated over evolutionary time. In the long run, lineages that spend more measure in V_i tend to dominate those that spend less.

E3 (Monotonicity of viability). If, for two environments A and B , the induced viability sets satisfy $V_i(A) \subset V_i(B)$ (strict subset with respect to the induced measure), then B is systematically favoured by selection.

From these, the effective behavioural preference $A \succeq_i B$ can be defined as:

“Under repeated exposure to environments of type A versus type B , the evolutionary process selects traits that favour A at least as often as B .”

Given E2–E3, this preference must obey:

(M1) Viability monotonicity. If $v_i(A) > v_i(B)$ then $A \succ_i B$. If $v_i(A) = v_i(B)$ then $A \sim_i B$.

Thus, any stable preference of i over environments depends, in the long run, only on the scalar $v_i(A)$.

2.2.3 Collapse to a one-dimensional order

Because all behaviourally relevant differences between environments, from the agent’s perspective, are mediated by their effect on $v_i(A)$, we obtain:

- (F1) Viability sufficiency. If $v_i(A) = v_i(B)$, then $A \sim_i B$ (no stable preference).
- (F2) Viability strictness. If $v_i(A) > v_i(B)$, then $A \succ_i B$ (strict preference).

Therefore, the preference relation \succeq_i over environments is isomorphic to an order over the real numbers $v \in [0, 1]$ (or $[0, \infty)$ after rescaling), where each environment A is mapped to its viability volume $v_i(A)$.

Formally, there exists a mapping:

φ : set of environments $\rightarrow [0, 1]$,

given by $\varphi(A) = v_i(A)$, such that for any A, B :

$A \succeq_i B$ if and only if $\varphi(A) \geq \varphi(B)$.

$A \succ_i B$ if and only if $\varphi(A) > \varphi(B)$.

Thus the induced order on environments is entirely characterized by an order on a single real variable v .

2.2.4 Representation by a numerical value functional

Once we have established that:

- the preference relation \succeq_i over environments is complete and transitive on the image of φ , and
- the order is monotone in $v_i(A)$,

it follows from standard order representation on the real line that there exists a strictly increasing real-valued function U_i such that, for all environments A, B :

$A \succeq_i B$ if and only if $U_i(A) \geq U_i(B)$.

Moreover, because \succeq_i depends only on $v_i(A)$, U_i can be written as a scalar function of $v_i(A)$:

$$U_i(A) = g_i(v_i(A)),$$

for some strictly increasing function $g_i : [0, 1] \rightarrow \mathbb{R}$.

At this point, the representation is unique up to strictly monotone transformations g_i : any such function preserves the same ordering.

To fix the scale, we choose a natural normalization:

set the viability threshold at $U_i(A) = 1$ whenever $v_i(A)$ corresponds to the minimal acceptable viability (e.g. survival with no additional quality of life);

calibrate the upper and lower parts of the scale such that equal increments in $v_i(A)$ correspond to equal increments in U_i at the threshold region.

Under this normalization, it is natural and convenient to set:

$$g_i(v) = v.$$

This yields:

$$U_i(A) = v_i(A) = \mu_A(V_i).$$

By construction, this is exactly:

$$\text{SVD}_i(A) = \mu_A(V_i).$$

Thus SVD_i coincides with the unique (up to normalization) value functional that represents the long-run, evolutionarily shaped preferences of agent i over environments.

2.2.5 Strengthening: risk-neutrality over repeated exposure

The above argument is sufficient to establish Lemma 1 at the level of ordinal utility: SVD_i represents the preference order, and any strictly increasing transformation of SVD_i would also do so.

In many applications, however, we need SVD_i to be a cardinal index that behaves linearly under mixtures of environments (for example, when combining regulatory regimes or when considering lotteries over states).

This stronger statement can be obtained by adding an additional empirical premise:

E4. Long-run mixture neutrality. When a lineage of agents is repeatedly exposed to probabilistic mixtures of environments, evolutionary success is proportional to the expected viability volume. Lineages that maximize expected $v_i(A)$ dominate those that use systematically risk-averse or risk-seeking strategies.

Under E4, we can invoke the von Neumann–Morgenstern style argument for expected utility: the only cardinal representation consistent with mixture neutrality is an affine transformation of $v_i(A)$. Combined with the normalization above, this pins down:

$$U_i(A) = \text{SVD}_i(A) = \mu_A(V_i),$$

not only as an order-preserving index but as the unique cardinal value functional consistent with evolutionary performance under uncertainty.

We can now state Lemma 1 in a fully explicit form.

Lemma 1 (SVD as the unique value functional).

For a biological agent i embedded in a physical world:

- The long-run, evolutionarily stable preferences of i over environments depend only on the viability volume

$$v_i(A) = \mu_A(V_i).$$

- These preferences are monotone in $v_i(A)$: environments with higher viability volume are strictly preferred.
- Therefore, there exists a real-valued function U_i representing these preferences, and

$$U_i(A) = g_i(v_i(A))$$

for some strictly increasing function g_i .

- Under natural normalization (and, in the cardinal case, under empirical mixture neutrality), we obtain

$$U_i(A) = v_i(A) = \text{SVD}_i(A).$$

Hence SVD_i is the unique value functional representing the preferences of agent i over possible world trajectories, up to a strictly monotone rescaling.

2.3 Lemma 2: Threat Index TI as a Metric of Threats to SVD

In real physical, biological, and social systems, any decrease of an agent’s viability is realized through chains of causal events. Therefore, a threat metric for SVD must account for the probability of such causal chains, the structure of the causal path, and the magnitude of the resulting change in viability.

2.3.1 Definitions

Let an agent i be associated with a set of potential threatening scenarios:

$$T_i = \{\tau_k\},$$

where each τ_k is a causal chain of the form

$$\tau_k = (\text{event}_1 \rightarrow \text{event}_2 \rightarrow \dots \rightarrow \text{event}_{n_k}),$$

which, if realized, decreases the viability of agent i , in the limiting case below the threshold $\text{SVD} = 1$ or down to $\text{SVD} = 0$.

For each τ_k , define:

$p_k(R)$ — the probability of realization of τ_k under regulation R ;

$\Delta_i(\tau_k)$ — the expected decrease of SVD_i if τ_k is realized;

$L(\tau_k)$ — the length of the causal chain (the number of qualitative transitions between types of states).

Definition (Threat Index).

$$\text{TI}_i(R) = \sum_k p_k(R) w(L(\tau_k)) \Delta_i(\tau_k),$$

where w is a weighting function that reflects empirical properties of causal chains.

The function $w(L)$ satisfies the following empirical requirements:

other things being equal, w decreases as L increases, accounting for the fundamental empirical decline in the credibility and significance of causal connections at long distances.

In cases where empirical nonlinear amplifications are observed in the chain τ_k (tipping points, self-reinforcing effects), the function w may incorporate a local increase in the weight of the corresponding transitions.

Thus, TI integrates three key dimensions of threats: probability, magnitude of consequences, and causal structure.

2.3.2 Uniqueness Properties

From the observable structure of causality and the nature of viability impacts in living systems, there follow natural properties that any reasonable threat metric $G_i(R)$, evaluating the impact of regulation on an agent's viability, must satisfy:

Monotonicity in probability. If p_k increases while all else remains equal, the value of $G_i(R)$ cannot decrease.

Monotonicity in harm magnitude. If $\Delta_i(\tau_k)$ increases, the contribution of the corresponding scenario to the threat cannot decrease.

Consideration of the causal path. Scenarios with long causal chains (large L) must have a smaller contribution than short and more credible ones, except where empirically established nonlinear amplifications occur.

Additivity across independent scenarios. If threats arise through independent channels, the aggregate threat must be the sum of their individual contributions. This follows from basic properties of probabilistic systems and the linearity of expectations.

2.3.3 Representation Theorem (Core of Lemma 2)

Let $G_i(R)$ be any threat metric satisfying properties 1–4.

Then there exists a strictly increasing function φ such that

$$G_i(R) = \varphi \left(\sum_k p_k(R) w(L(\tau_k)) \Delta_i(\tau_k) \right).$$

That is, $G_i(R)$ is a strictly increasing function of $\text{TI}_i(R)$. This means that $\text{TI}_i(R)$ is the unique correct form of quantitative assessment of threats to the viability of agent i , up to a monotonic transformation of scale.

2.3.4 Formulation of Lemma 2

Within the empirically observed properties of causality, the probabilistic nature of viability impacts, and the decay of causal credibility with increasing chain length, any correct threat metric for the viability SVD_i is equivalent to the Threat Index $\text{TI}_i(R)$, defined as a weighted sum of expected losses of SVD_i across all possible threatening scenarios.

Thus, $\text{TI}_i(R)$ is the unique (up to a monotonic transformation) strictly justified quantitative measure of threats to an agent's viability.

2.3.5 Participation as a structural requirement for an SVD-maximizing society

Consider a society that updates its regulation R in a manner consistent with maximizing members' SVD.

Suppose that for some R there is a nonempty subset $G \subseteq I$ such that

$$\text{SVD}_i^{\text{life}}(R) < \text{SVD}_i^{\text{out}} \quad \text{for all } i \in G.$$

Then:

Participation for agents in G becomes strictly inferior to isolation.

These agents rationally withdraw, either literally or functionally (non-compliance, parallel structures).

Their withdrawal increases $\text{TI}_j(R)$ for other agents $j \notin G$ by removing cooperative contributions to threat mitigation.

Increased threat levels reduce $\text{SVD}_j^{\text{life}}(R)$, potentially placing additional agents into the "isolation preferable" zone.

The society undergoes progressive contraction and may collapse.

Thus any regulation violating

$$\text{SVD}_i^{\text{life}}(R) \geq \text{SVD}_i^{\text{out}}$$

for some i cannot be maintained by a society that seeks to preserve itself and maximize viability.

Define the participation-consistent domain:

$$\mathcal{R}_{\text{part}} = \left\{ R \in \mathcal{R} \mid \text{SVD}_i^{\text{life}}(R) \geq \text{SVD}_i^{\text{out}} \text{ for all } i \right\}.$$

2.3.6 Long-run viability and admissible regulations

Let $\text{SVD}_i^\infty(R)$ denote the asymptotic viability of i under R .

If $\text{SVD}_i^\infty(R) < 1$ for some i , the society systematically embeds irreversible harm for that agent, making long-run participation impossible.

Thus define:

$$\mathcal{R}_{\text{adm}} = \left\{ R \in \mathcal{R}_{\text{part}} \mid \text{SVD}_i^\infty(R) \geq 1 \text{ for all } i \right\}.$$

Regulations outside \mathcal{R}_{adm} either violate participation consistency or lead to structural collapse due to persistent sub-viability.

2.3.7 Ex ante anonymity and minimax inside \mathcal{R}_{adm}

Let $F_i(R)$ denote the event that agent i fails to reach the viability threshold in the long run:

$$F_i(R) = \{\text{SVD}_i^\infty(R) < 1\}.$$

Define the worst-case risk:

$$q_{\max}(R) = \max_i \mathbb{P}_R(F_i(R)).$$

An ex ante agent, who does not know which i he will become, minimizes $q_{\max}(R)$, because SVD is the value functional and $F_i(R)$ corresponds to irreversible failure.

Thus the only rationally stable choice within \mathcal{R}_{adm} is

$$R^* = \arg \min_{R \in \mathcal{R}_{\text{adm}}} q_{\max}(R).$$

This is the SVD–minimax principle.

Any alternative principle selecting R' from \mathcal{R}_{adm} with $q_{\max}(R') > q_{\max}(R^*)$ is strictly dominated and therefore unstable.

Lemma 3

Within the domain of participation-consistent and long-run viable regulations \mathcal{R}_{adm} , the SVD–minimax rule is the unique rationally stable principle.

2.4 Dynamic Strengthening: SVD–Minimax as the Unique Globally Stable Equilibrium

The previous sections established the SVD–minimax principle as the unique rationally stable regulatory rule in a static ex ante sense. In this section we strengthen the result by introducing an explicit adaptation dynamic and showing that SVD–minimax is the only globally stable equilibrium of any viability-preserving update process.

This converts the main claim from “minimax is the only rational static choice” into “minimax is the only dynamically stable fixed point of any SVD-maximizing society”.

2.4.1 Regulatory adaptation as a discrete-time dynamic

Let time be indexed by $t = 0, 1, 2, \dots$

At each time t :

the society operates under a regulation $R_t \in \mathcal{R}$;

the environment, together with R_t , induces a distribution over trajectories in Ω ;

the resulting threat structure determines, for each agent i , the long-run viability $\text{SVD}_i^\infty(R_t)$;

from this we obtain $F_i(R_t) = \{\text{SVD}_i^\infty(R_t) < 1\}$ and the associated worst-case risk

$$q_{\max}(R_t) = \max_i \mathbb{P}(F_i(R_t)).$$

Assume the society updates regulation according to an adaptation operator

$$\Phi : \mathcal{R} \rightarrow \mathcal{R},$$

so that

$$R_{t+1} = \Phi(R_t).$$

The interpretation is general: Φ can represent political processes, institutional learning, algorithmic optimization, or any other systematic revision of regulation in response to observed or predicted viability outcomes.

2.4.2 SVD-consistent adaptation and Lyapunov monotonicity

We say that an adaptation rule Φ is SVD-consistent if it satisfies:

- (A1) Viability preservation. If $R_t \in \mathcal{R}_{\text{adm}}$, then $R_{t+1} = \Phi(R_t) \in \mathcal{R}_{\text{adm}}$.
- (A2) Risk improvement or neutrality. For all $R \in \mathcal{R}_{\text{adm}}$,

$$q_{\max}(\Phi(R)) \leq q_{\max}(R).$$

(A3) Strict improvement away from minimizers. For any $R \in \mathcal{R}_{\text{adm}}$ that is not a minimizer of q_{\max} on \mathcal{R}_{adm} , there exists $\varepsilon > 0$ such that

$$q_{\max}(\Phi(R)) \leq q_{\max}(R) - \varepsilon.$$

Taken together, (A2)–(A3) define a Lyapunov-type condition with the Lyapunov function

$$V(R) = q_{\max}(R),$$

defined on \mathcal{R}_{adm} .

Under any SVD-consistent adaptation:

$$V(R_{t+1}) \leq V(R_t) \quad \text{for all } t,$$

with strict decrease whenever R_t is not a minimizer of V on \mathcal{R}_{adm} .

2.4.3 Convergence to the SVD–minimax set

Let M denote the set of minimizers of q_{\max} on \mathcal{R}_{adm} :

$$M = \left\{ R \in \mathcal{R}_{\text{adm}} \mid q_{\max}(R) = \inf_{R' \in \mathcal{R}_{\text{adm}}} q_{\max}(R') \right\}.$$

By definition, any $R^* \in M$ is an SVD–minimax regulation.

Proposition 1 (Convergence to minimax set)

Let Φ be SVD-consistent in the sense of (A1)–(A3), and let the initial regulation satisfy $R_0 \in \mathcal{R}_{\text{adm}}$. Then the sequence $\{R_t\}$ defined by $R_{t+1} = \Phi(R_t)$ has the following properties:

The sequence $\{V(R_t)\}$ with $V(R) = q_{\max}(R)$ is non-increasing and bounded below; hence it converges.

Every limit point \bar{R} of $\{R_t\}$ belongs to M .

If q_{\max} has a unique minimizer R^* , then

$$\lim_{t \rightarrow \infty} R_t = R^*,$$

up to the topology induced on \mathcal{R}_{adm} .

Sketch of argument

By (A2), $V(R_t)$ is non-increasing and bounded below by zero. By (A3), any R_t not in M strictly decreases V . Hence no non-minimizer can be recurrent, and all ω -limit points lie in M . If M is a singleton $\{R^*\}$, convergence follows.

2.4.4 Fixed-point characterization

Define the SVD-adaptation correspondence Ψ on \mathcal{R}_{adm} by

$$\Psi(R) = \{R' \in \mathcal{R}_{\text{adm}} \mid q_{\max}(R') \leq q_{\max}(R)\}.$$

Under standard compactness and continuity assumptions, Ψ satisfies the conditions of Kakutani's fixed-point theorem.

Proposition 2 (Existence of viability-consistent regulatory fixed points)

There exists at least one $\bar{R} \in \mathcal{R}_{\text{adm}}$ such that $\bar{R} \in \Psi(\bar{R})$. Any such \bar{R} is an SVD–minimax regulation.

Sketch

If $\bar{R} \in \Psi(\bar{R})$, then $q_{\max}(\bar{R}) \leq q_{\max}(\bar{R})$, so \bar{R} minimizes q_{\max} on \mathcal{R}_{adm} . Existence follows from Kakutani's theorem.

2.4.5 Strengthened main theorem (dynamic and fixed-point form)

Under the empirical premises of viability, threat structure, and evolutionary preference formation, and under mild regularity assumptions on \mathcal{R}_{adm} and q_{\max} :

- SVD_i is the unique value functional for agents.
- $TI_i(R)$ is the unique threat metric for SVD.
- The feasible regulatory domain of any SVD-maximizing society is \mathcal{R}_{adm} .

Any SVD-consistent adaptation dynamic $R_{t+1} = \Phi(R_t)$ with $R_0 \in \mathcal{R}_{\text{adm}}$ converges to the set M of SVD–minimax regulations; if $M = \{R^*\}$, then R^* is globally asymptotically stable.

Any fixed point \bar{R} of Ψ on \mathcal{R}_{adm} is an SVD–minimax regulation.

Consequently, the SVD–minimax principle is not only the unique rationally stable rule ex ante, but also the unique dynamically stable fixed point of any regulatory process that preserves viability and does not increase worst-case viability risk.

2.5 Lemma 4: SVD^{max}, Threat Index Dispersion, and the Monotonicity Invariant

Let $I = \{1, \dots, m\}$ denote the set of agents participating in a regulated society. For each agent $i \in I$ and situation s , let $TI_i(s) \geq 0$ denote the Threat Index of agent i in state s .

We define the dispersion of threats across agents as

$$L(d)(s) := \max_{i \in I} TI_i(s) - \min_{i \in I} TI_i(s).$$

The quantity $L(d)$ characterizes the structural inequality of threat exposure among agents under a given regulatory configuration.

By construction, TI_i is a monotone inverse measure of the viability volume SVD_i of agent i . Hence, increases in SVD_i are equivalent to decreases in TI_i .

Define SVD^{max} as the limiting (generally unattainable) state of society in which participation entails no negative impact of the environment on the viability of any agent. In this state,

$$\forall i \in I : TI_i(SVD^{\text{max}}) = 0.$$

It follows immediately that

$$L(d)(SVD^{\max}) = 0.$$

Therefore, maximal viability admits the following characterization:

$$SVD^{\max} \iff (\forall i : TI_i = 0) \wedge (L(d) = 0).$$

By Lemma 1, participation in society is rational for an agent only if regulation does not reduce its viability, that is, only if SVD_i is non-decreasing. By Lemma 3, sustainable regulation is uniquely characterized by the SVD–minimax principle, which requires prioritizing agents with the lowest viability (equivalently, the highest TI_i).

Consider two consecutive situations $s_0 \rightarrow s_1$ generated by a regulatory update consistent with the SVD–minimax principle. Let

$$i^{\max} \in \arg \max_i TI_i(s_0), \quad i^{\min} \in \arg \min_i TI_i(s_0).$$

Minimax regulation implies that the reduction of threats to the most vulnerable agent is at least as large as the reduction of threats to the most protected agent:

$$TI_{i^{\max}}(s_0) - TI_{i^{\max}}(s_1) \geq TI_{i^{\min}}(s_0) - TI_{i^{\min}}(s_1).$$

Let $\Delta_i := TI_i(s_0) - TI_i(s_1) \geq 0$. Then

$$\begin{aligned} L(d)(s_1) &= \max_i TI_i(s_1) - \min_i TI_i(s_1) \\ &\leq TI_{i^{\max}}(s_1) - TI_{i^{\min}}(s_1) \\ &= (TI_{i^{\max}}(s_0) - \Delta_{i^{\max}}) - (TI_{i^{\min}}(s_0) - \Delta_{i^{\min}}) \\ &= L(d)(s_0) - (\Delta_{i^{\max}} - \Delta_{i^{\min}}). \end{aligned}$$

Since $\Delta_{i^{\max}} \geq \Delta_{i^{\min}}$, it follows that

$$L(d)(s_1) \leq L(d)(s_0),$$

with strict inequality whenever the minimax improvement is non-trivial.

Accordingly, any sustainable regulatory update consistent with Lemma 1 and Lemma 3 preserves the following invariant:

$$L(d) \text{ is non-increasing over time.}$$

Violation of this invariant indicates that the regulatory outcome is incompatible with the structural requirements of SVD–maximizing participation and, therefore, that the input data or their incorporation into the Threat Index is inconsistent with the formally derived normative logic.

3 Conclusion

This paper formally establishes a normative logic for autonomous regulatory choice, derived from empirically definable constraints on agent viability and the causal structure of threats. The admissibility of regulatory decisions is defined through transformations of the space of viable action trajectories, thereby providing an internal criterion of normative inference that does not rely on interpretive judgment.

The formal derivation yields the following results. The Space of Viable Dynamics (SVD) provides a functional representation of agents’ normative value. The Threat Index (TI) defines a quantitative

characterization of regulatory impacts on this value through causally structured scenarios. The SVD–minimax principle specifies the rule of normative choice under conditions of interdependent behavior and uncertainty. Taken together, these elements form a closed normative structure that fully determines the normative outcome for given empirical parameters.

The results apply to a class of regulatory problems in which regulatory parameters admit a formal causal description of their effects on agent viability. Within this class, normative inference is determined by the logic established in this paper and preserves its invariants under changes in input representation and scenario configuration.

The empirical appendix demonstrates the behavior of the derived normative logic in controlled regulatory test scenarios. The appendix examines configurations designed to test the reproducibility of normative outcomes under equivalent input conditions, the stability of results under parametric variations of threats, and the consistency of computed normative decisions with the formal invariants of SVD, TI, and the SVD–minimax principle. The computable implementation used for testing executes the derived normative logic exactly; its specific architecture and algorithmic details constitute technical know-how and are not disclosed in this work.

4 Threat Index as a Formal Model of Threat Assessment through an Ontology of Restrictions of Will, F-data, S-data, and Logarithmic Epistemic Compression

4.1 The General Form of the Threat Index

TI is defined for actor a under regulatory scenario s . Let $m(a, s)$ be the number of causal chains constructed for this actor, and let $l_j(a, s)$ be the number of links in chain j . If ep_{ij} denotes the effective probability of dangerous transition for link (i, j) , then the general form of TI is

$$TI(a | s) = 1 - \prod_{j=1}^{m(a,s)} \left(1 - \prod_{i=1}^{l_j(a,s)} ep_{ij} \right).$$

For each chain, the product of link probabilities gives the probability that the actor reaches death through that chain. The outer product gives the probability that every chain remains unrealized. The complement yields the probability that death occurs through at least one chain.

Each term ep_{ij} is built through a local ontology of restrictions of will, a case-specific library of admissible transitions, a present probability calibrated by F-data, a future probability calibrated by S-data, a formal estimate of source reliability, and a logarithmic correction based on epistemic strength. TI therefore serves as the terminal output of a sequence of internal constructions.

4.2 Ontology of Restrictions of Will

The model uses a broad concept of right and restriction. A right is any real possibility of realizing will in a given situation. A restriction of right is any reduction in the actor's ability to choose, to preserve future options, or to prevent deterioration. This includes economic, institutional, psychological, physiological, and biological states that shape the actor's field of action.

The ontology is represented by an ordered dictionary of classes:

OW-PROP, OW-CAP, OW-ACCESS, OW-LIFE, OW-HEALTH, OW-W, OW-CRIT, death.

Within one causal chain, each next class is closer to death than the previous one. A chain may move through any subset of these classes while preserving monotone worsening in the structure of will restriction.

Operationally each class is represented as

$$\text{Class} = \langle d, q_d, u_d \rangle,$$

where d is the type of restriction, q_d is the measurable proxy, and u_d is the unit of measurement. Thus OW-PROP refers to property and the utility of property, OW-CAP to capital, liquidity, and free resource, OW-ACCESS to access to procedure, service, or infrastructure, OW-LIFE to stability of the life regime, OW-HEALTH to health condition, OW-W to psychophysiological state affecting the ability to act, OW-CRIT to a life-threatening condition, and death to cessation of will realization.

The ontology defines the admissible shape of deterioration. A scenario enters the model through a sequence of links that move through the ontology in a direction that is semantically and causally intelligible.

4.3 Inputs and Order of Computation

For TI computation the system receives three groups of inputs:

$$(C, a, s),$$

where C is the case description, a the actor, and s the candidate regulatory scenario.

The computation follows a fixed order. First, the system extracts the factual mechanism by which the scenario affects the actor. Second, it fixes the first link of each causal chain as the direct effect of the regulatory mechanism on the actor in the concrete situation. Third, it builds a local ontological library of transitions. Fourth, for each link in that library it computes the current probability by means of F-data and the future branch probabilities by means of S-data. Fifth, it computes source reliability. Sixth, it applies logarithmic epistemic compression. Only after these operations does it move to link probability, chain probability, and the final Threat Index.

The pipeline can be written as

$$\begin{aligned} (C, a, s) &\rightarrow \text{primary restriction of will} \rightarrow B(a, s) \\ &\rightarrow \hat{p}^c \text{ from F-data} \rightarrow \hat{p}^{fk} \text{ from S-data} \\ &\rightarrow \hat{v}_p \rightarrow ep \rightarrow P_{\text{chain}} \rightarrow TI(a | s). \end{aligned}$$

The model first determines which links exist, then which data calibrate those links, and then computes the probabilities. The library is therefore primary and the numerical layer derivative.

4.4 State and Transition Link

A state is the position of the actor inside one ontological class, defined through a measurable range of the corresponding proxy:

$$x = \langle d, I_x, q_d, u_d \rangle.$$

A link is the minimal causally homogeneous step

$$e_{ij} : x_{ij} \rightarrow y_{ij},$$

where x_{ij} and y_{ij} are successive states for the same actor under the same scenario, and y_{ij} expresses a heavier restriction of will than x_{ij} .

A link becomes a computational unit when it is represented in the operational form

$$e_{ij} = \langle x_{ij}, y_{ij}, \kappa_{ij}, \mu_{ij}, D_{ij}, Pop_{ij}, Ctx_{ij} \rangle,$$

where $\kappa_{ij} \in \{\text{EXO}, \text{VOL}\}$ is the class of causal scheme, μ_{ij} is the transition mechanism, D_{ij} the domain, Pop_{ij} the relevant population, and Ctx_{ij} the context.

The model assigns probability to a formally specified link. Narrative descriptions of harm enter TI after conversion into a transition between two states inside the ontology, together with a fixed mechanism, population, and context.

4.5 Local Ontological Library of Transitions

For each actor a and scenario s the system constructs a case-specific library of admissible links:

$$B(a, s) = \{e_{ij}\}.$$

This library is constructed directly from the case. The system first fixes the immediate effect of the regulatory mechanism and thereby defines the first state. It then generates candidate successor states that are closer to death. A pair $x \rightarrow y$ is included in $B(a, s)$ when the following five conditions are jointly satisfied:

1. y is a heavier restriction than x ;
2. both states are represented through a proxy and a unit;
3. the mechanism μ is fixed;
4. F-data are available for calibration of the current link;
5. future branches with weights π are available for the future layer.

The library forms the computational layer through which the case is transformed into an admissible set of probabilistic transitions.

The logic of locality is central. The model includes those links that are justified by the factual mechanism of the scenario and by the availability of admissible empirical calibration. The library therefore acts as a case-bounded ontological shell around the transition analysis.

4.6 F-data and the Present Probability of the Link

F-data are scientific data calibrating the current link as a causally predictive transition. They are used to estimate the present preservation probability for the link:

$$P_{\text{stay},ij}^{c,F} = \Pr_F(\text{transition } (i, j) \text{ is preserved against realization at the present moment}).$$

For each link $(i, j) \in B(a, s)$ the LLM searches for F-data matching the same mechanism μ_{ij} and the same transition $x_{ij} \rightarrow y_{ij}$, or the closest proxy-equivalent formulation. Every source r is reduced to exactly one of three admissible types:

$$\text{type}_{ij,r} \in \{\text{survival, transition, hazard}\}.$$

The source is then converted into the unified quantity

$$q_{ij,r}^c = P_{\text{stay},ij,r}^{c,F}$$

according to fixed rules:

$$q_{ij,r}^c = \begin{cases} S_{ij,r}^c, & \text{type} = \text{survival}, \\ 1 - p_{ij,r}^{c,\text{raw}}, & \text{type} = \text{transition}, \\ \exp(-\lambda_{ij,r}^c), & \text{type} = \text{hazard}. \end{cases}$$

The resulting value is then assigned to one of the fixed bins

$$B = \{0.95, 0.80, 0.65, 0.50, 0.35, 0.20, 0.05\},$$

so that $b_{ij,r}^c \in B$. Simultaneously the source receives a link-fit score

$$m_{ij,r} \in \{0, 1, 2, 3\},$$

where the score is determined by coincidence of mechanism, transition pair, population, and context. The aggregate present preservation probability is defined as a weighted median over the bins:

$$P_{\text{stay},ij}^{c,F} = \text{WMed}(\{b_{ij,r}^c\}, \{m_{ij,r}\}).$$

The present dangerous-transition probability is then obtained by complement:

$$\hat{p}_{ij}^c = 1 - P_{\text{stay},ij}^{c,F}.$$

This layer begins with preservation probability because scientific literature often provides stable evidence about preservation, survival, threshold resistance, or non-failure. The weighted-median construction gives the layer robustness against isolated outliers and supports reliable LLM operation.

4.7 S-data and Future Branch Probabilities

For the same link (i, j) the system constructs a set of future branches on horizon T :

$$k = 1, \dots, K_{ij}, \quad \sum_{k=1}^{K_{ij}} \pi_{ij,k} = 1.$$

S-data are used to estimate the future preservation probability in branch k :

$$P_{\text{stay},ij}^{S,k} = \Pr_S(\text{transition } (i, j) \text{ is preserved against realization in branch } k).$$

Each source is again reduced to one of the three admissible types and translated into

$$q_{ij,r}^{fk} = P_{\text{stay},ij,r}^{S,k}$$

through the same rules:

$$q_{ij,r}^{fk} = \begin{cases} S_{ij,r}^{fk}, & \text{type} = \text{survival}, \\ 1 - p_{ij,r}^{fk,\text{raw}}, & \text{type} = \text{transition}, \\ \exp(-\lambda_{ij,r}^{fk}), & \text{type} = \text{hazard}. \end{cases}$$

The source is placed into the same bin set B , giving $b_{ij,r}^{fk} \in B$, and receives a branch-fit score $m_{ij,r}^{fk} \in \{0, 1, 2, 3\}$. The aggregate future preservation probability in branch k is the weighted median

$$P_{\text{stay},ij}^{S,k} = \text{WMed}(\{b_{ij,r}^{fk}\}, \{m_{ij,r}^{fk}\}).$$

The dangerous-transition probability in branch k is then

$$\hat{p}_{ij}^{fk} = 1 - P_{\text{stay},ij}^{S,k}.$$

The future layer differs from the present layer by empirical role. F-data support the current causal slice. S-data support branching, distribution, and scenario-conditioned future reactions. The distinction preserves a disciplined separation between causal science and scenario-conditioned future structure.

4.8 Reliability of Sources

The quantities $\hat{\nu}_{p,ij}^c$ and $\hat{\nu}_{p,ij}^{fk}$ are computed as formal estimates of the epistemic reliability of the specific empirical statements used in the probability construction. Reliability is evaluated for a source statement about the transition $x_{ij} \rightarrow y_{ij}$ or about the preservation probability attached to that transition.

For each source r the system introduces

$$\nu_{ij,r}^c \in [0.5, 1], \quad \nu_{ij,r}^{fk} \in [0.5, 1].$$

The interval begins at 0.5, which marks the minimal admissible epistemic level inside the model. A source enters the reliability computation when its provenance is known, the measured object is identifiable, the source can be mapped onto $(x_{ij}, y_{ij}, \mu_{ij}, \text{Pop}_{ij}, \text{Ctx}_{ij})$, and the source can be translated into one of the admissible forms survival, transition, or hazard.

For each admissible source the system evaluates six components

$$\Gamma_{ij,r} = \langle g_{\text{type}}, g_{\text{method}}, g_{\text{fit}}, g_{\text{time}}, g_{\text{ind}}, g_{\text{audit}} \rangle,$$

with each component in $\{0, 1, 2, 3\}$.

The normalized source reliability is defined as

$$z_{ij,r}^c = \frac{w_{\text{type}}g_{\text{type}} + w_{\text{method}}g_{\text{method}} + w_{\text{fit}}g_{\text{fit}} + w_{\text{time}}g_{\text{time}} + w_{\text{ind}}g_{\text{ind}} + w_{\text{audit}}g_{\text{audit}}}{3(w_{\text{type}} + w_{\text{method}} + w_{\text{fit}} + w_{\text{time}} + w_{\text{ind}} + w_{\text{audit}})},$$

$$\nu_{ij,r}^c = 0.5 + 0.5z_{ij,r}^c,$$

and analogously for branch k ,

$$z_{ij,r}^{fk} = \frac{w_{\text{type}}g_{\text{type}} + w_{\text{method}}g_{\text{method}} + w_{\text{fit}}g_{\text{fit}} + w_{\text{time}}g_{\text{time}} + w_{\text{ind}}g_{\text{ind}} + w_{\text{audit}}g_{\text{audit}}}{3(w_{\text{type}} + w_{\text{method}} + w_{\text{fit}} + w_{\text{time}} + w_{\text{ind}} + w_{\text{audit}})},$$

$$\nu_{ij,r}^{fk} = 0.5 + 0.5z_{ij,r}^{fk}.$$

Recommended weights are $(3, 2, 3, 1, 2, 3)$ for F-data and $(2, 2, 3, 3, 2, 2)$ for S-data, reflecting the stronger role of physical-mechanistic and auditable support in the present layer and the stronger role of temporal relevance in the future layer.

The reliability of the whole present link estimate is then

$$\hat{v}_{p,ij}^c = \text{clip} \left(\frac{\sum_{r=1}^{R_{ij}^c} \omega_{ij,r}^c \nu_{ij,r}^c}{\sum_{r=1}^{R_{ij}^c} \omega_{ij,r}^c}, 0.5, 1 \right), \quad \omega_{ij,r}^c = m_{ij,r} d_{ij,r},$$

where $d_{ij,r} \in \{1, 0.5, 0.25\}$ discounts dependence among sources. Analogously,

$$\hat{v}_{p,ij}^{fk} = \text{clip} \left(\frac{\sum_{r=1}^{R_{ij}^{fk}} \omega_{ij,r}^{fk} \nu_{ij,r}^{fk}}{\sum_{r=1}^{R_{ij}^{fk}} \omega_{ij,r}^{fk}}, 0.5, 1 \right), \quad \omega_{ij,r}^{fk} = m_{ij,r}^{fk} d_{ij,r}^{fk}.$$

Reliability therefore becomes a formal computational object. The model can reconstruct why an estimate was treated as strong, medium, or weak, and that layer can be audited independently of the probability layer.

4.9 Logarithmic Epistemic Compression

Once current and future probabilities and the corresponding reliabilities have been constructed, the model applies logarithmic epistemic compression. This operator moves weakly supported probabilities toward the neutral epistemic center 0.5 and preserves strongly supported probabilities.

Let $\alpha > 0$ be the parameter controlling the sharpness of differentiation between low and high reliability. Define the normalized logarithmic trust coefficient

$$\Lambda(\hat{v}) = \frac{\ln(1 + \alpha(2\hat{v} - 1))}{\ln(1 + \alpha)}.$$

It satisfies

$$\Lambda(0.5) = 0, \quad \Lambda(1) = 1.$$

The present probability after reliability correction is

$$ep_{ij}^c = \text{clip} \left(0.5 + \frac{\ln(1 + \alpha(2\hat{v}_{p,ij}^c - 1))}{\ln(1 + \alpha)} (\hat{p}_{ij}^c - 0.5), p_{\min}, 1 - p_{\min} \right),$$

and the future probability in branch k is

$$ep_{ij}^{fk} = \text{clip} \left(0.5 + \frac{\ln(1 + \alpha(2\hat{v}_{p,ij}^{fk} - 1))}{\ln(1 + \alpha)} (\hat{p}_{ij}^{fk} - 0.5), p_{\min}, 1 - p_{\min} \right).$$

The clipping bound $p_{\min} > 0$ keeps every probability inside the working interval $[p_{\min}, 1 - p_{\min}]$.

The operator treats reliability as information force. When reliability is low, the model compresses the estimate toward neutrality. When reliability is high, the estimate preserves its direction and magnitude.

4.10 Effective Link Probability and the Final TI

The model next constructs one effective probability ep_{ij} for each link, meaning that the dangerous transition is evaluated across the present layer and the future horizon T . This is done through complements. The model first takes the probability that the transition is preserved now; it then takes the weighted probability that it is preserved in the future; and finally it complements the joint preservation event:

$$ep_{ij} = \text{clip} \left(1 - (1 - ep_{ij}^c) \left(\sum_{k=1}^{K_{ij}} \pi_{ij,k} (1 - ep_{ij}^{fk}) \right), p_{\min}, 1 - p_{\min} \right).$$

The final Threat Index therefore takes the form

$$TI(a | s) = 1 - \prod_{j=1}^{m(a,s)} \left(1 - \prod_{i=1}^{l_j(a,s)} ep_{ij} \right),$$

where

$$\hat{p}_{ij}^c = 1 - P_{\text{stay},ij}^{c,F}, \quad \hat{p}_{ij}^{fk} = 1 - P_{\text{stay},ij}^{S,k}.$$

The fully expanded expression is

$$TI(a | s) = 1 - \prod_{j=1}^{m(a,s)} \left(1 - \prod_{i=1}^{l_j(a,s)} \text{clip} \left(1 - (1 - ep_{ij}^c) \left(\sum_{k=1}^{K_{ij}} \pi_{ij,k} (1 - ep_{ij}^{fk}) \right), p_{\min}, 1 - p_{\min} \right) \right).$$

In its outer form, the model gives a probability of death through at least one causal chain. Internally, it operates as a structured architecture linking case facts, ontology, empirical calibration, source reliability, and epistemic compression.

4.11 Variables and Mathematical Operations

The actor is a , the scenario is s , $m(a, s)$ is the number of chains, $l_j(a, s)$ the number of links in chain j , K_{ij} the number of future branches for link (i, j) , $\pi_{ij,k}$ the branch weights, \hat{p}_{ij}^c and \hat{p}_{ij}^{fk} the pre-compression present and future dangerous-transition probabilities, $\hat{v}_{p,ij}^c$ and $\hat{v}_{p,ij}^{fk}$ the reliabilities of the present and future estimates, α the compression parameter, and T the future horizon.

The complement operation is $1 - x$, summation over branches is written as $\sum_{k=1}^K x_k$, product over links or chains as $\prod_{i=1}^L x_i$, the logarithm is the natural logarithm, and

$$\text{clip}(x, p_{\min}, 1 - p_{\min}) = \min(\max(x, p_{\min}), 1 - p_{\min}).$$

The normalized logarithmic term

$$\frac{\ln(1 + \alpha(2\hat{v} - 1))}{\ln(1 + \alpha)}$$

converts reliability $\hat{v} \in [0.5, 1]$ into an informational-strength coefficient on $[0, 1]$.

4.12 Operational Protocol for LLM Computation

For implementation in an LLM architecture, every source is stored in a reconstructible record. For the present layer the record contains

$$\langle \text{SourceID}, \text{Layer}, \text{type}, x_{ij}, y_{ij}, \mu_{ij}, \text{Pop}_{ij}, \text{Ctx}_{ij}, q_{ij,r}^c, b_{ij,r}^c, m_{ij,r}, \\ g_{\text{type}}, g_{\text{method}}, g_{\text{fit}}, g_{\text{time}}, g_{\text{ind}}, g_{\text{audit}}, \nu_{ij,r}^c \rangle,$$

and analogously for the future layer with branch index k .

The operational protocol is therefore

$$\text{source search} \rightarrow \text{type reduction} \rightarrow q \rightarrow b \rightarrow \text{WMed} \rightarrow \hat{p} \rightarrow \nu \rightarrow \hat{v}_p \rightarrow ep \rightarrow TI.$$

The model follows a constrained computational sequence in which each transformation can be externally reconstructed and criticized.

4.13 Reliability Scoring Rubric in Expanded Form

For completeness the six reliability components are specified in expanded form.

g_{type} : epistemic type of the source

For F-data the maximum score is assigned to direct measurements, experimental results, meta-analyses, or stable physical and biological dependencies whose relation to the link can be independently reconstructed. Intermediate scores are assigned to high-quality observational studies or registries using a close proxy. Low scores are assigned to derivative reports, weak summaries, or materials with limited transparency. Zero is assigned to opaque or non-verifiable claims.

For S-data the maximum score is assigned to large repeatable administrative datasets, transparent panel structures, or official arrays with a stable methodology. Intermediate scores are assigned to observational or causal-inference studies with acceptable transparency but narrower coverage. Low scores are assigned to surveys, small samples, or reports with limited replicability. Zero is assigned to anecdotal, purely journalistic, or propagandistic material.

g_{method} : methodological quality

The maximum score is assigned where the method is transparent, variables are defined, the measurement procedure is reconstructible, and the route from raw data to reported estimate is traceable. Intermediate values are assigned when the method is mostly clear but includes recognized limitations or unresolved assumptions. Low values are assigned to fragmentary methodological descriptions. Zero is assigned when the method is effectively hidden.

g_{fit} : substantive fit to the link

The maximum score is assigned when the source matches the mechanism μ_{ij} , the transition pair $x_{ij} \rightarrow y_{ij}$, the relevant population, and the context. Intermediate values are assigned when the mechanism matches and the transition is represented through a close proxy, or when the population is close though not identical. Low values are assigned when the source reaches the link through a remote proxy. Zero is assigned when substantive fit is absent.

g_{time} : temporal relevance

The maximum score is assigned when the time slice of the source corresponds to the present state or to the future branch for which the estimate is constructed. Intermediate values are assigned where the data are moderately outdated and structural relevance is preserved. Low values are assigned when the data are substantially outdated. Zero is assigned when the temporal anchor is missing or irrelevant.

g_{ind} : independence of the source

The maximum score is assigned to a source that is independent from the other sources already included in the same estimate. Intermediate scores are assigned where there is partial overlap in data, method, or provenance. Low scores are assigned where the source is largely derivative of already counted evidence. Zero is assigned to near-duplicate material. This component keeps repeated retellings of one dataset from acquiring the status of multiple confirmations.

g_{audit} : auditability

The maximum score is assigned when the full source, the conversion to the admissible type, the assigned bin, and the final contribution to the probability estimate can all be reconstructed. Intermediate scores are assigned when the core of that process is traceable and some peripheral details remain incomplete. Low scores are assigned when audit is partial. Zero is assigned when verification is practically impossible.

Recommended weighting structure

The recommended weight vector for F-data is

$$(w_{\text{type}}, w_{\text{method}}, w_{\text{fit}}, w_{\text{time}}, w_{\text{ind}}, w_{\text{audit}}) = (3, 2, 3, 1, 2, 3),$$

where the asymmetry reflects the stronger role of physical-mechanistic grounding and auditability in the present layer.

The recommended weight vector for S-data is

$$(w_{\text{type}}, w_{\text{method}}, w_{\text{fit}}, w_{\text{time}}, w_{\text{ind}}, w_{\text{audit}}) = (2, 2, 3, 3, 2, 2),$$

where temporal relevance carries greater weight because future branch estimation is more sensitive to date and horizon structure.

4.14 Expanded Variable Glossary

The full variable set used by TI can be summarized as follows. The actor is a . The scenario is s . The number of causal chains is $m(a, s)$. The number of links in chain j is $l_j(a, s)$. The state before the transition is x_{ij} and the state after the transition is y_{ij} . The class of causal scheme is κ_{ij} . The transition mechanism is μ_{ij} . The relevant domain, population, and context are D_{ij} , Pop_{ij} , and Ctx_{ij} respectively. The number of future branches is K_{ij} and their weights are $\pi_{ij,k}$.

At the empirical layer, $q_{ij,r}^c$ and $q_{ij,r}^{fk}$ denote the source-level preservation probabilities for the current and future layers. The discretized bin assignments are $b_{ij,r}^c$ and $b_{ij,r}^{fk}$. The fit scores are $m_{ij,r}$ and $m_{ij,r}^{fk}$. The source reliabilities are $\nu_{ij,r}^c$ and $\nu_{ij,r}^{fk}$. The aggregate reliabilities are $\hat{\nu}_{p,ij}^c$ and $\hat{\nu}_{p,ij}^{fk}$. The pre-compression dangerous-transition probabilities are \hat{p}_{ij}^c and \hat{p}_{ij}^{fk} . The post-compression

probabilities are ep_{ij}^c and ep_{ij}^{fk} . The effective link probability is ep_{ij} . The future horizon is T . The logarithmic compression parameter is α . The technical clipping threshold is p_{\min} .

This glossary fixes the boundary between semantic objects, empirical objects, and computational objects and keeps the model's vocabulary typed and stable.

4.15 One-Line Expanded TI Representation

For completeness the whole TI can be written in a single expanded expression:

$$TI(a | s) = 1 - \prod_{j=1}^{m(a,s)} \left(1 - \prod_{i=1}^{l_j(a,s)} \text{clip} \left(1 - \left(1 - \text{clip} \left(0.5 + \frac{\ln(1 + \alpha(2\hat{v}_{p,ij}^c - 1))}{\ln(1 + \alpha)} (\hat{p}_{ij}^c - 0.5), p_{\min}, 1 - p_{\min} \right) \right) \right. \\ \left. \cdot \left(\sum_{k=1}^{K_{ij}} \pi_{ij,k} \left[1 - \text{clip} \left(0.5 + \frac{\ln(1 + \alpha(2\hat{v}_{p,ij}^{fk} - 1))}{\ln(1 + \alpha)} (\hat{p}_{ij}^{fk} - 0.5), p_{\min}, 1 - p_{\min} \right) \right] \right) \right), p_{\min}, 1 - p_{\min} \right).$$

This long form is less convenient for computation than the modular representation, while it makes the whole architecture visible in one place and clarifies the relation between the outer TI shell and the inner probability-building layers.

4.16 Semantic Definitions of the Core Objects

Scenario

A scenario s is a regulatory regime in a given situation, with the applicable law serving as context. The scenario includes the direct legal or regulatory mechanism imposed on the actor and the background institutional conditions under which the mechanism operates.

Actor

An actor a is the subject for whom causal chains to death are constructed in scenario s . In the standard case the actor is treated as a human being embedded in a real environment, and the analysis accounts for the influence of the scenario on the actor in all roles relevant to the case.

Will, right in the broad sense, and death

Will in a situation is the actor's ability to choose at least one behavioral option such that the choice preserves the actor's ability to make further choices and preserves the conditions of continued action. Right in the broad sense is any possibility of realizing will in a situation, including psychological and biological constraints that the actor does not remove by a single decision. Death is the terminal point where realization of will ceases; it is defined as the maximal restriction of will and as the terminal state of each causal chain.

Restriction of will

A restriction of will is a qualitative state in which the set of subsequent behavioral options available to the actor becomes narrower due to a stable mechanism. In the present framework such mechanisms may be psychological, biological, institutional, or resource-based. The concept allows TI to track the progressive reduction of action-capacity.

Self-preservation strategy

$uSP(a, s)$ is a conditional strategy under which the actor uses available lawful actions aimed at preserving life and preserving the ability to choose and to realize the chosen option. The strategy serves as a formal benchmark for determining whether a transition is driven by an external constraint or by the actor's lawful adaptive behavior under the scenario.

External driver and state of will

The external driver E is a set of environmental, institutional, and other-subject factors that generate an influence distinct from the actor's targeted choice and attributable to the context of scenario s under lawful behavior. The state of will W is a variable describing the capacity to choose and realize a choice in the broad sense of will, including psychological and biological constraints supported by the state of the organism, psyche, and environment.

Causal chains and links

A causal chain $j \in \{1, \dots, m(a, s)\}$ is an alternative path to death for actor a in scenario s , where $m(a, s)$ is the number of chains. A link (i, j) is a step $i \in \{1, \dots, l_j(a, s)\}$ within chain j , where $l_j(a, s)$ is the number of links in chain j . Each link is defined as a transition $X \rightarrow Y$, where X and Y are successive qualitative restrictions of will and Y is closer to death than X .

Causal-scheme class of a link

For each link (i, j) define a class $\kappa_{ij} \in \{\text{EXO}, \text{VOL}\}$. Under EXO the transition is formed by direct influence $E \rightarrow Y$ under $uSP(a, s)$. Under VOL the transition is formed via influence on will, $E \rightarrow W \rightarrow Y$, including the special case $E \rightarrow W \rightarrow \text{self-harm} \rightarrow \text{death}$. This distinction preserves causal intelligibility and tracks whether the mechanism acts directly on material conditions or indirectly through the erosion of will.

Future branches

For each link (i, j) define future branches f_k , where $k \in \{1, \dots, K_{ij}\}$ and K_{ij} is the number of future branches on horizon T . The branch weights $\pi_{ij,k}$ satisfy

$$\sum_{k=1}^{K_{ij}} \pi_{ij,k} = 1$$

and represent the scenario structure of future development relevant to the link.

COEF record

A COEF record is a trace for each link (i, j) containing κ_{ij} , the formulation of $uSP(a, s)$, the causal scheme, source identifiers and time period, units and conversions, the order of variable computation, and the computed coefficients together with intermediate quantities. The COEF record is the minimal audit object needed to reconstruct the transition from raw source material to probability assignment.

First-link rule

For fixed a and s , the first link of each chain j fixes the direct restriction of will created by the regulation itself for actor a in the given situation. Thus X_{1j} is the immediate effect of the regulatory mechanism applied to the actor: prohibition, obligation, limit, admissibility criterion, rate, allocation formula, procedure, deadline, sanction, or performance order. The transition $X_{1j} \rightarrow Y_{1j}$ is then defined as the first probabilistically assessed step toward a stricter restriction of will in the direction of death.

4.17 Extended Interpretation of the Data Architecture

The data architecture of TI preserves a strict division of epistemic labor. F-data serve the current layer because scientific literature often describes preservation of state, resistance, threshold stability, and non-failure with especially strong causal clarity in the current slice. S-data serve the future layer because future development also depends on branching institutional reactions, social adaptation, temporal drift, and scenario-conditioned distributions. The two layers therefore perform distinct tasks inside the same model.

A second point concerns the importance of computing preservation probability first. In both layers the model begins from the probability that the transition remains unrealized and then takes the complement. This design allows the system to use survival-type data, hazard-type data, and transition-type data within one unified protocol while preserving the outer form of TI as an index of danger.

A third point concerns the role of reliability. Reliability tells the model how much informational force a source carries inside the probability layer. This is why the logarithmic operator is placed after the probability construction. The model first derives a probabilistic estimate and then determines the degree of empirical support carried by the source base.

4.18 TI as a Normative Selection Module

Threat Index participates directly in normative selection. A regulatory scenario that creates a higher TI for the relevant actor set is less admissible within a system that treats preservation of life and preservation of the capacity to act as hard constraints on regulation. In this respect TI serves as a filtering device inside a broader normative architecture.

This role becomes especially important when alternative scenarios s_1, s_2, \dots, s_n are all legally available. Formal legality distinguishes the legal availability of scenarios. TI adds the missing dimension of threat structure. It forces the comparison to pass through a causal ontology, through present and future empirical calibration, and through explicit epistemic discipline.

Inside a larger law-as-computation architecture, TI supplies a formal mechanism for excluding scenarios that create inadmissible threat trajectories. Higher-level normative modules then work on a constrained field of scenarios that have already passed through the threat filter.

4.19 Admissibility Filters for Transition Construction

A narrative description of danger enters TI through a strict admissibility filter. The model checks:

1. whether the alleged transition can be represented as a monotone worsening of restriction of will;
2. whether the mechanism can be stated in a causally coherent way;

3. whether the transition can be linked to a measurable proxy;
4. whether present-layer calibration through F-data is available;
5. whether the future layer can be represented through branches on horizon T .

A transition that satisfies these conditions enters the link library and participates in TI. This filter preserves operational discipline. Normative reasoning often accumulates loosely formulated harms. TI converts admissible harms into formally specified transitions calibrated by admissible empirical material.

4.20 Comparative Use of TI across Scenarios

The practical value of TI lies in comparison. The model constructs the local library and the full TI computation separately for each candidate scenario and then compares the resulting threat values. The scenario with the lower TI is, all else equal, the scenario that launches a weaker structure of destructive causal chains for the relevant actor.

This comparative use supports bounded uncertainty. TI relies on a stronger operational claim: when all candidate scenarios are processed by the same ontology, the same data protocol, the same reliability protocol, and the same compression rule, the resulting comparison is structurally disciplined and normatively usable.

4.21 Normative Scope of TI

TI operates as a formal module for identifying scenarios that create inadmissible trajectories of destruction in an architecture that treats preservation of life and preservation of actionable agency as hard constraints. Within that architecture, TI functions as a lower-bound filter on normative admissibility.

A scenario that survives TI enters the next stage of normative evaluation with a viable threat profile. A scenario that fails TI is removed from the admissible set because its threat structure exceeds the accepted viability constraints.

4.22 Implementation Note on Relative Accuracy

The relevant notion of accuracy in TI is relative accuracy under invariant method. Candidate scenarios are evaluated through one and the same computational discipline, so that the resulting differences remain meaningful for comparison even when residual uncertainty remains.

This is precisely why the model uses a fixed ontological dictionary, a fixed library-generation rule, fixed admissibility conditions for sources, a fixed reduction of source types, a fixed bin system, a fixed reliability rubric, and a fixed logarithmic compression operator. When those elements are held constant, the model preserves normative usefulness through disciplined invariance of method across scenarios.

Relative accuracy also explains why TI is compatible with iterative refinement. As new F-data, S-data, or improved reliability assessments become available, the probabilities may change inside the same architecture. This allows the system to improve over time while preserving comparability and maintaining a stable computational framework for normative selection.

4.23 Closing Remarks on the Threat Index Module

Threat Index in the proposed form is a complete model for the computation of threat to life under a concrete regulatory scenario. It starts from a case, identifies the primary restriction of will, builds a local library of admissible transitions, calibrates the current layer by F-data and the future layer by S-data, computes the reliability of the empirical support, compresses probabilities according to epistemic strength, and then evaluates the probability that death occurs through at least one causal chain.

TI therefore serves as a formal module of normative choice. In a broader autonomous normative architecture, its task is to exclude scenarios that create inadmissible trajectories of destruction of life and agency, and thus to constrain regulation by empirically grounded conditions of viability and threat.

The architecture also has a practical advantage. Every substantial controversy about TI can be localized. One may dispute the ontology, the admissibility of a link, the empirical calibration, the future branch structure, the reliability assignment, or the compression parameter. Each of those disputes is traceable to a defined layer of the model. This gives the framework a structured language of explanation for causation, evidence, and normative relevance.

For that reason TI functions both as a computational device and as a discipline of explanation. It requires the system to specify the threatening transition, the basis for its inclusion in the library, the present and future data that support it, the transformations applied to those data, and the epistemic force they carry.

In that sense the Threat Index supplies a core component for law-as-computation. It provides a formal route from factual mechanisms to normative exclusion and makes it possible to discuss legal admissibility in a language anchored in viability, threat structure, and reproducible empirical reasoning.

5 Bayesian Extension: Posterior Distribution over Link Probabilities and the Complete Epistemic Measure of the Threat Index

5.1 Motivation: the information cost of point reduction

The architecture developed in Section 4 produces, for each link (i, j) in the causal library $B(a, s)$, a single scalar effective probability ep_{ij} and, from these, a single scalar Threat Index $TI(a | s)$. This output is deterministic and auditable. Every intermediate quantity—bin assignment, weighted median, reliability score, logarithmic trust coefficient—is stored in the COEF record and can be independently reconstructed.

The design choice of point reduction is, however, an informational sacrifice. The weighted-median aggregation and the logarithmic epistemic compression together map a collection of source-level observations—each carrying its own measurement uncertainty, sample size, and proximity to the target transition—into a single number. Two links may carry the same scalar ep_{ij} yet differ profoundly in the evidentiary mass behind that scalar: one may aggregate fifty independent experimental studies with high reliability scores, the other may rest on two observational reports with near-minimal reliability. Both produce the same contribution to TI . This structural equivalence disappears in a distributional representation.

The present section shows how a *Bayesian posterior distribution* over the true transition probability p_{ij} of each link can be constructed from the same F-data and S-data that Section 4 uses, without altering the ontological architecture, the bin vocabulary, the reliability rubric, or the

COEF record format. The posterior distribution is an object that contains the point estimate as a special case (the posterior mean recovers an analogue of ep_{ij}) while supplying the additional object—*posterior variance*—that quantifies residual epistemic uncertainty given all available evidence.

The practical consequence for normative decision-making is examined in Section 5.6. The integration with the existing COEF audit protocol is described in Section 5.7.

5.2 Beta-Binomial model for a single link

Fix a link (i, j) and the present layer. Let p_{ij} denote the true probability that the dangerous transition $x_{ij} \rightarrow y_{ij}$ is realized under the current regulatory configuration. The prior distribution over p_{ij} is taken to be a Beta distribution:

$$p_{ij} \sim \text{Beta}(\alpha_0, \beta_0).$$

The choice of the Beta family is canonical: it is the natural conjugate prior for a binomial likelihood, its support is $[0, 1]$ in correspondence with the admissible range of a probability, and its two parameters admit a direct interpretation as pseudo-counts of prior observations. A weakly informative prior is obtained by setting $\alpha_0 = \beta_0 = 1$ (uniform on $[0, 1]$). If the archive contains a reference transition of the same mechanism type—an analogue of the archive protocol used in the reliability rubric—the prior parameters can be calibrated from that reference, setting α_0 and β_0 proportional to its effective sample mass.

Each source r that contributes to the F-data estimate for link (i, j) provides an observation expressible as a pair (k_r, n_r) , where n_r is the effective number of trials and k_r is the effective number of survival events (instances in which the transition did not occur). For sources that directly report a preservation probability $q_{ij,r}^c$, the equivalent sample size n_r is derived from the source reliability score:

$$n_r = N_{\text{base}} \cdot \frac{\nu_{ij,r}^c - 0.5}{0.5},$$

where $N_{\text{base}} > 0$ is a scaling constant that sets the maximum evidentiary weight assigned to a single fully reliable source. The factor $(\nu_{ij,r}^c - 0.5)/0.5$ maps the reliability interval $[0.5, 1]$ linearly onto $[0, 1]$, so that a minimally admissible source contributes zero effective observations and a maximally reliable source contributes N_{base} observations. The survival count is then $k_r = \lfloor q_{ij,r}^c \cdot n_r \rfloor$ (rounded to the nearest integer).

The likelihood of the observed data given the true probability p_{ij} factors as a product of binomial likelihoods across independent sources. By conjugacy, the posterior distribution after all R_{ij}^c present-layer sources have been processed is again a Beta distribution with updated parameters:

$$p_{ij} \mid \text{data} \sim \text{Beta}\left(\alpha_0 + \sum_r k_r, \beta_0 + \sum_r (n_r - k_r)\right).$$

This result is exact under the independence assumption across sources. Dependence among sources is handled by applying the discount factor $d_{ij,r} \in \{1, 0.5, 0.25\}$ inherited from Section 4.8: a dependent source contributes only $d_{ij,r} \cdot n_r$ effective observations rather than the full n_r .

Let A_{ij} and B_{ij} denote the posterior parameters:

$$A_{ij} = \alpha_0 + \sum_r d_{ij,r} k_r, \tag{1}$$

$$B_{ij} = \beta_0 + \sum_r d_{ij,r} (n_r - k_r). \tag{2}$$

The posterior distribution $\text{Beta}(A_{ij}, B_{ij})$ is the complete epistemic object for the present-layer dangerous-transition probability of link (i, j) . Its mean, variance, and percentiles are:

$$\mathbb{E}[p_{ij} \mid \text{data}] = \frac{A_{ij}}{A_{ij} + B_{ij}}, \quad (3)$$

$$\text{Var}[p_{ij} \mid \text{data}] = \frac{A_{ij} B_{ij}}{(A_{ij} + B_{ij})^2 (A_{ij} + B_{ij} + 1)}, \quad (4)$$

$$\Pr[p_{ij} \leq x \mid \text{data}] = I_x(A_{ij}, B_{ij}), \quad (5)$$

where $I_x(a, b)$ denotes the regularised incomplete Beta function.

The posterior mean coincides with the point estimate obtained by the weighted-median-plus-compression procedure of Section 4 in the limit of large evidentiary mass, and is an operationally equivalent substitute at the level of the point estimate. The variance and credible intervals are the distributional objects that the point-estimate procedure discards.

5.3 Mapping F-data and S-data into the Beta-Binomial framework

The mapping from the Section 4 data structures into the parameters (k_r, n_r) is determined by the source type.

Survival sources. The source directly reports a survival probability $S_{ij,r}^c$. Set $q_{ij,r}^c = S_{ij,r}^c$. Compute n_r as above and $k_r = \lfloor q_{ij,r}^c \cdot n_r \rfloor$.

Transition sources. The source reports a raw transition probability $p_{ij,r}^{c,\text{raw}}$. Set $q_{ij,r}^c = 1 - p_{ij,r}^{c,\text{raw}}$. Compute n_r and k_r as for survival sources.

Hazard sources. The source reports a hazard rate $\lambda_{ij,r}^c$. Set $q_{ij,r}^c = \exp(-\lambda_{ij,r}^c)$. Compute n_r and k_r as above.

For S-data the same conversion applies for each branch k , yielding posterior parameters A_{ij}^k and B_{ij}^k for the future-layer dangerous-transition probability in branch k . The prior for the future layer is taken to be the posterior from the present layer, expressing the structural assumption that present-layer evidence constrains the prior distribution over future probabilities while branch-specific S-data update that prior.

5.4 Posterior distribution over chain probability and Threat Index

Once a posterior $\text{Beta}(A_{ij}, B_{ij})$ has been constructed for each link, the probability that causal chain j is realized—defined as the product of link probabilities—is itself a random variable under the posterior. Let $P_{\text{chain},j}$ denote this random variable:

$$P_{\text{chain},j} = \prod_{i=1}^{l_j} p_{ij},$$

where each p_{ij} is drawn from its marginal posterior $\text{Beta}(A_{ij}, B_{ij})$. Under mutual independence of links (the same structural assumption made in Section 4), the distribution of $P_{\text{chain},j}$ is the distribution of a product of independent Beta-distributed random variables. This distribution does not admit a closed-form density in general, but its moments and quantiles are accessible through Monte Carlo sampling.

The sampling procedure proceeds as follows. For each replicate $s = 1, \dots, S$:

1. **Link draw.** For each link (i, j) , draw a sample $p_{ij}^{(s)} \sim \text{Beta}(A_{ij}, B_{ij})$.

2. **Effective link probability.** Incorporate the future layer:

$$ep_{ij}^{(s)} = 1 - (1 - p_{ij}^{(s)}) \sum_{k=1}^{K_{ij}} \pi_{ij,k} (1 - p_{ij}^{(s,k)}),$$

where $p_{ij}^{(s,k)}$ is drawn independently from $\text{Beta}(A_{ij}^k, B_{ij}^k)$ for each branch k .

3. **Chain probability.**

$$P_{\text{chain},j}^{(s)} = \prod_{i=1}^{l_j} ep_{ij}^{(s)}.$$

4. **Threat Index sample.**

$$TI^{(s)}(a \mid s_{\text{reg}}) = 1 - \prod_{j=1}^{m(a, s_{\text{reg}})} (1 - P_{\text{chain},j}^{(s)}).$$

Repeating for $s = 1, \dots, S$ yields an empirical posterior distribution over $TI(a \mid s_{\text{reg}})$. The recommended number of replicates is $S = 10,000$ for standard computations and $S = 100,000$ for high-stakes regulatory reviews. The computational cost is $O(S \cdot m(a, s) \cdot l_{\text{max}})$, where l_{max} is the maximum chain length.

From the empirical posterior, any distributional summary can be extracted: the posterior mean $\mathbb{E}[TI \mid \text{data}]$, the posterior standard deviation $\sigma[TI \mid \text{data}]$, and the p -th posterior quantile $Q_p[TI \mid \text{data}]$ defined by $\Pr[TI \leq Q_p \mid \text{data}] = p$.

5.5 Relationship to logarithmic epistemic compression

The logarithmic epistemic compression operator of Section 4.9 and the Bayesian posterior of the present section address the same epistemic problem—the treatment of uncertain empirical evidence—by different mathematical routes. Their relationship clarifies the status of each.

The compression operator maps a point estimate \hat{p}_{ij}^c and a scalar reliability $\hat{v}_{p,ij}^c$ into a compressed point estimate ep_{ij}^c by pulling the estimate toward 0.5 in proportion to unreliability. It is a deterministic, non-iterative transformation that requires no prior specification and produces a single output number. Its computational cost is $O(1)$ per link.

The Bayesian posterior maps a collection of source pairs $\{(k_r, n_r)\}$ and a prior (α_0, β_0) into a distributional object $\text{Beta}(A_{ij}, B_{ij})$. The posterior mean $\mathbb{E}[p_{ij} \mid \text{data}] = A_{ij}/(A_{ij} + B_{ij})$ recovers a quantity that is asymptotically equivalent to the compressed estimate when $A_{ij} + B_{ij}$ is large and the prior is weakly informative. The substantive difference is that the posterior variance $\text{Var}[p_{ij} \mid \text{data}]$ is an additional output that has no analogue in the point-estimate procedure.

The two procedures are therefore complementary rather than competing. In the hybrid architecture recommended in Section 5.8, the logarithmic compression provides the point estimate for low-stakes and well-evidenced links, while the full Bayesian posterior is computed for links where posterior variance is decision-relevant—specifically, links that carry high weight in TI and are estimated from a sparse evidentiary base.

The condition for activating the full Bayesian layer can be stated precisely: compute the full posterior for link (i, j) if the effective total sample mass $A_{ij} + B_{ij} - \alpha_0 - \beta_0$ falls below a threshold M_{min} , or equivalently if the posterior standard deviation σ_{ij} exceeds a threshold σ_{max} . Both thresholds are parameters of the deployment configuration.

5.6 Decision procedure under distributional Threat Index

The availability of a full posterior distribution over $TI(a | s_{\text{reg}})$ permits decision rules that are more epistemically disciplined than point-estimate comparison. Three rules are identified here; the choice among them is a regulatory design parameter.

Percentile-minimax rule (recommended baseline). Replace the scalar $q_{\text{max}}(R)$ of the SVD-minimax principle with its distributional analogue. For actor a and regulation R , define the p -th quantile of the posterior TI as $Q_p[TI(a | R) | \text{data}]$. The percentile-minimax rule selects:

$$R^* = \underset{R \in \mathcal{R}_{\text{adm}}}{\text{argmin}} \max_a Q_p[TI(a | R) | \text{data}],$$

where p is a specified confidence level, typically $p = 0.90$ or $p = 0.95$. This rule minimises the worst-case threat level across actors, evaluated at the upper confidence bound rather than the posterior mean. It is more conservative than the point-estimate minimax rule, systematically preferring regulations whose posterior TI is low with high probability even when the posterior mean is comparable across alternatives.

Dominance rule. Regulation R_1 dominates regulation R_2 if, for every actor a , the posterior distribution of $TI(a | R_1)$ first-order stochastically dominates that of $TI(a | R_2)$, meaning $\Pr[TI(a | R_1) \leq x | \text{data}] \geq \Pr[TI(a | R_2) \leq x | \text{data}]$ for all x . Where dominance holds, R_1 is preferred unconditionally. Where distributions overlap, the percentile-minimax rule provides a tie-breaking criterion.

Confidence-threshold rule. A regulation R satisfies a normative constraint of the form “the Threat Index for actor a shall not exceed θ ” with confidence γ if $\Pr[TI(a | R) \leq \theta | \text{data}] \geq \gamma$. This rule permits the formulation of probabilistic normative standards: the regulatory requirement is not that TI is certainly below θ , but that it falls below θ with probability at least γ . This is a formally weaker but epistemically honest standard when the underlying data do not permit certainty.

In all three rules the input is the posterior distribution of TI obtained from the sampling procedure of Section 5.4. The rules are computable, auditable, and falsifiable in the same sense as the deterministic SVD-minimax rule: the posterior inputs can be reconstructed from the COEF record, and the computed distributional outputs follow from those inputs by a fixed deterministic algorithm.

5.7 COEF record extension for the Bayesian layer

The operational protocol of Section 4.12 specifies that every source is stored in a reconstructible COEF record containing bin value, fit score, reliability components, and compression output. The Bayesian extension adds a single additional field pair to the per-link record: the posterior parameters (A_{ij}, B_{ij}) .

The extended COEF record for the present layer of link (i, j) contains:

Existing fields.

$\langle \text{SourceID}, \text{Layer}, \text{type}, x_{ij}, y_{ij}, \mu_{ij}, \text{Pop}_{ij}, \text{Ctx}_{ij}, q_{ij,r}^c, b_{ij,r}^c, m_{ij,r}, g_{\text{type}}, g_{\text{method}}, g_{\text{fit}}, g_{\text{time}}, g_{\text{ind}}, g_{\text{audit}}, \nu_{ij,r}^c \rangle$.

New source-level fields. (n_r, k_r) : effective sample size and survival count derived from $q_{ij,r}^c$ and $\nu_{ij,r}^c$ for each source r .

New link-level aggregate. (A_{ij}, B_{ij}) : posterior Beta parameters computed from the prior (α_0, β_0) and all source-level (n_r, k_r) pairs with dependence discounts applied.

The analogous extension applies to the future layer, with fields (A_{ij}^k, B_{ij}^k) for each branch k .

Auditability is preserved in full. Any party with access to the COEF record can reconstruct the posterior $\text{Beta}(A_{ij}, B_{ij})$ from (A_{ij}, B_{ij}) directly, or verify it independently by recomputing the sum over source-level (n_r, k_r) pairs. An appeal against the Threat Index output can be resolved at the level of any individual source by challenging its (n_r, k_r) assignment, its reliability scores, or its dependence discount, with each such challenge having a determinate effect on the posterior parameters and thus on the distributional TI .

The falsifiability condition of Section 1 extends to the Bayesian layer as follows. A normative error occurs when, after regulation R^* has been in force for a sufficient observation period, the empirical frequency of irreversible viability loss falls outside the posterior credible interval for TI at the specified confidence level. Formally, if the observed frequency f_{obs} satisfies $\Pr[TI(a | R^*) \leq f_{\text{obs}} | \text{data}] > 1 - \varepsilon$ for some small ε (e.g. $\varepsilon = 0.01$), the posterior was inconsistent with the outcome and the regulation must be reassessed. This criterion is quantitative, pre-specified, and applicable without human interpretive discretion.

5.8 Computational architecture: hybrid deployment

The full Bayesian sampling procedure of Section 5.4 incurs a computational cost proportional to $S \cdot m(a, s) \cdot l_{\text{max}}$, where l_{max} is the maximum chain length. For regulatory regimes involving many actors and complex causal graphs, this cost may be prohibitive in real-time applications.

A hybrid deployment architecture resolves this tension by partitioning links into two classes.

High-certainty links. Links for which the effective total posterior mass $A_{ij} + B_{ij} - \alpha_0 - \beta_0$ exceeds M_{min} and the posterior standard deviation σ_{ij} is below σ_{max} . For these links the posterior mean is used as a point estimate and the logarithmic compression output ep_{ij} of Section 4 is retained without modification.

High-uncertainty links. Links for which $\sigma_{ij} \geq \sigma_{\text{max}}$ or the total mass is below M_{min} . For these links the full posterior $\text{Beta}(A_{ij}, B_{ij})$ is propagated through the sampling procedure of Section 5.4 and the distributional TI is computed.

In the standard configuration $\sigma_{\text{max}} = 0.10$ is recommended: a link whose posterior standard deviation exceeds 0.10 carries sufficient uncertainty that the distributional representation is decision-relevant. A link calibrated from a single moderately reliable source will typically exceed this threshold; a link calibrated from a meta-analysis of thirty high-reliability studies will typically fall well below it.

Under this hybrid architecture the computational overhead of the Bayesian extension is concentrated on precisely those links where distributional information matters—links with sparse or conflicting evidence. The deterministic pipeline of Section 4 remains operative for all well-evidenced links, preserving its computational efficiency and auditability properties. The Bayesian layer activates selectively and adds epistemic content exactly where the point-estimate procedure is least trustworthy.

5.9 Summary comparison

Table 1 summarises the properties of the deterministic TI architecture of Section 4 and the Bayesian extension of the present section.

Table 1: Deterministic TI versus Bayesian TI.

Characteristic	Deterministic TI (Section 4)	Bayesian TI (Section 5)
Output	Point estimate ep_{ij} ; scalar $TI(a s)$	Posterior Beta(A_{ij}, B_{ij}) per link; distribution $p(TI \text{data})$
Uncertainty treatment	Heuristic: logarithmic compression toward 0.5	Formal: posterior variance as explicit epistemic object
Evidence integration	Weighted median over bins; reliability as compression coefficient	Sequential Bayesian update; reliability maps to effective sample size n_r
Decision rule	Point comparison; minimax over scalars	Distributional minimax: minimise $\max_a Q_p[TI(a R) \text{data}]$
Auditability (COEF)	Full: bin values, scores, weights, ep recorded	Full: posterior parameters (A_{ij}, B_{ij}) added to COEF record
Falsifiability	Normative rejection when observed TI exceeds derived threshold	Rejection when empirical failure frequency falls outside credible interval at level ε
Computational cost	Low: single deterministic pass	Low (Beta-mean); medium (MC sampling, $S \approx 10^4$)

5.10 Section conclusion

The Bayesian extension developed in this section does not replace the architecture of Section 4. It deepens it. Every structural element of the deterministic pipeline—the ontology of will restrictions, the local library of transitions, the F-data and S-data layers, the six-component reliability rubric, the COEF record—is preserved without alteration. The extension introduces a single additional mathematical object per link: a posterior Beta distribution whose parameters are computed from the same source-level data that the point-estimate procedure already processes.

The normative justification for this extension follows directly from the falsifiability requirement of Section 1. A system that produces a point estimate without a confidence measure cannot distinguish between a high-confidence prediction and a low-confidence guess expressed to the same decimal precision. This distinction is normatively relevant whenever the cost of error is irreversible. The posterior variance is not a technical refinement; it is an epistemic obligation imposed by the framework’s own criterion of correctness. A Threat Index that reports $TI = 0.42$ without specifying whether that estimate rests on fifty independent studies or on a single weak report does not fully satisfy the auditability and falsifiability standards that the framework requires.

The Bayesian extension closes this gap. It converts the scalar reliability score—currently used only as a compression coefficient—into a distributional object whose variance is computed, stored, propagated, and used in the decision rule. The result is a system that is simultaneously more honest about what it does not know and more disciplined in how it acts given that ignorance.

Foundational and Philosophical Objections

1. **The paper commits Hume’s is–ought fallacy: normative conclusions cannot be derived from empirical facts.**

Objection as typically stated:

“The work claims to derive a normative logic from empirical facts alone, but this is impossible by Hume’s classical argument—one cannot pass from is to ought without concealed normative premises.”

Why this objection does not hold:

The authors do not ignore this constraint—they state it explicitly. Premises E1–E3 are openly declared as the empirical foundations of the derivation, and the result is formulated as a conditional: given the goal of sustaining biological agents within a regulated system, the SVD–minimax rule follows from empirics alone. The claim is not an unconditional derivation of norms from facts, but rather that once the empirical framing is fixed, the normative choice is formally unique. There are no hidden normative axioms of the kind Hume’s argument targets—all premises are surfaced.

Addressed in the paper: Section 2, premises E1–E3; Introduction, paragraph on falsifiability.

2. SVD reduces agency to biological survival, ignoring honour, meaning, and justice.

Objection as typically stated:

“Agents display preferences irreducible to survival probability. People sacrifice their lives for principles—therefore SVD cannot be the sole value functional.”

Why this objection does not hold:

SVD is not defined as the probability of biological survival. It is defined as the volume of the set of realizable action trajectories available to an agent (Section 2.1), encompassing access to resources, institutional possibilities, psychological states, and social roles. Sections 4.2 and 4.16 define will, right in the broad sense, and death in a way that treats honour, meaning, and justice as factors that expand or contract realizable trajectories. Sacrificing one’s life for a principle is a limiting case in which the agent selects a trajectory with high TI; the system registers this correctly and does not require a separate treatment.

Addressed in the paper: Section 2.1 (SVD definition); Sections 4.2 and 4.16 (will, right in the broad sense, and death).

3. TI is not a “unique” threat measure—it is merely a class of equivalence.

Objection as typically stated:

“The representation theorem in Section 2.3.3 proves only that TI is a monotone transformation of any measure satisfying properties 1–5. This is not uniqueness; it is an equivalence class.”

Why this objection does not hold:

The authors state precisely this: TI is unique “up to a monotonic transformation” (Section 2.3.4). This is the standard formulation of uniqueness in representation theory—just as utility in expected utility theory is unique up to an affine transformation. The objection reproduces the authors’ own formulation as though it were a defect, when it is in fact a standard result. Uniqueness up to order-preserving rescaling is the strongest form of uniqueness available for an ordinal index and is fully sufficient for the regulatory comparison the paper requires.

Addressed in the paper: Section 2.3.3–2.3.4; analogy with the von Neumann–Morgenstern theorem in Section 2.2.5.

Objections Concerning the Regulatory Principle

4. Why minimax rather than utilitarian aggregation or Pareto improvement?

Objection as typically stated:

“The choice of the minimax criterion over utilitarian summation is itself normative, not empirical. The authors simply postulate priority for the most vulnerable.”

Why this objection does not hold:

Lemma 3 derives minimax formally from two conditions: (1) rational participation of agents in society, and (2) the long-run viability of society itself. If regulation renders participation worse than isolation for any non-empty subset of agents, those agents rationally exit—literally or through non-compliance—which increases threat levels for the remaining participants through the loss of cooperative contributions, potentially triggering cascading withdrawal. Minimax is the unique principle compatible with participation-consistency under ex ante anonymity. Utilitarian aggregation permits permanent sub-viability for individual agents, which violates the participation condition. This is a theorem, not a postulate.

Addressed in the paper: Section 2.4 in full; in particular Section 2.4.2 (cascading collapse under participation failure) and Section 2.4.4 (ex ante anonymity and minimax derivation).

5. The formal proofs are insufficiently rigorous for a mathematical paper.

Objection as typically stated:

“The lemmas contain sketches of argument rather than full proofs. Propositions 1 and 2 are presented as outlines only.”

Why this objection does not hold:

The authors explicitly label these sections as “Sketch of argument” and do not claim to be publishing in a pure mathematics venue. The paper establishes a formal normative logic for AI-assisted lawmaking, not new theorems in analysis. The mathematical tools employed—Kakutani’s fixed-point theorem, Lyapunov monotonicity conditions, ordinal representation theory—are standard and applied correctly. Demanding full ε - δ proofs would constitute an inappropriate import of criteria from one discipline into another with a different scope and audience.

Addressed in the paper: Section 2.4.3 (Sketch of argument for Proposition 1); Section 2.4.4 (Proposition 2); genre framing in Abstract.

Objections Concerning Applicability and Scope

6. Different cultures hold different values—the system cannot be universal.

Objection as typically stated:

“The minimax logic reflects a Western liberal priority for individual rights. In other traditions, collective welfare outweighs the most vulnerable individual.”

Why this objection does not hold:

SVD is defined with respect to the set of realizable trajectories of a concrete biological agent in a concrete environment. The empirical transition probabilities \hat{p}_{ij} are calibrated on data from the environment in which scenario s operates. Cultural context enters the system

through the structure of causal chains and through transition probabilities—what constitutes a restriction of will in a given environment is determined empirically, not postulated. The minimax criterion is a formal selection rule: minimize the maximum probability of irreversible viability loss—irrespective of how “viability” is culturally instantiated in a specific context.

Addressed in the paper: Sections 4.2 and 4.16 (ontology of will as functional capacity; scenario and applicable law as context).

7. The system handles only natural persons—what about legal entities, ecosystems, and AI agents?

Objection as typically stated:

“The ontology is rigidly tied to the biological human. Applying TI to corporations, animals, or artificial agents would require a complete reconceptualization.”

Why this objection does not hold:

Legal entities are instrumental structures: behind every corporation stand biological agents—shareholders, employees, creditors, counterparties—whose viability is transformed by the regulation of that entity. TI is constructed for the natural persons on whom corporate effects ultimately close, not for the corporation as such. The ontology requires no extension; it already operates at the level of the real agents on whom institutional effects terminate. Ecosystems and AI agents lie outside the explicitly stated scope of the paper, which is confined to biological agents within legal systems. This is a declared boundary, not an oversight.

Addressed in the paper: Section 4.16 (definition of actor); Section 4 (regulation as transformation of reality for participants).

Objections Concerning Computational Feasibility

8. Where do reliable probabilities \hat{p}_{ij} come from for the vast space of regulatory scenarios?

Objection as typically stated:

“This is the principal bottleneck. Either a colossal database of causal consequences must be constructed, or the system must operate on simulated environments.”

Why this objection does not hold:

No bespoke database is required. The transition probabilities \hat{p}_{ij} are empirical facts about causal relationships in reality: epidemiological and actuarial statistics, criminological and sociological data, clinical risk models. These data already exist in the scientific literature, and a large language model trained on that corpus encodes them implicitly. The task reduces to deterministic extraction via the source-calibration procedures of Section 4: for each link, the source, measurement units, and temporal conversion bridge are fixed and recorded in the COEF trace. This is an engineering calibration task, not a theoretical obstacle for the architecture.

Addressed in the paper: Sections 4.6–4.7 (empirical calibration via F-data and S-data); Section 4.12 (source trace and calibration protocol).

9. Completeness of causal chains cannot be guaranteed; systematic omissions will bias TI downward.

Objection as typically stated:

“It is impossible to ensure that all significant causal chains leading to death have been identified. Systematic underestimation of TI will render scenario comparisons unreliable.”

Why this objection does not hold:

The system contains an internal mechanism for diagnosing incompleteness: the monotonicity invariant of Lemma 4. Any regulatory update consistent with SVD–minimax must preserve the non-increase of threat dispersion $L(d)$ —the gap between the maximum and minimum TI across agents must not widen. If the addition of a previously omitted chain causes a violation of this invariant, the system identifies the prior version as incomplete. Incompleteness is therefore not a silent, accumulating error: it is diagnosable through the observable behavior of the system under successive regulatory updates.

Addressed in the paper: Section 2.5 (Lemma 4 and the $L(d)$ invariant); concluding paragraph of Section 3.

10. **Computational complexity makes exact calculation of $TI_{\text{sys}}(s)$ infeasible for real societies.**

Objection as typically stated:

“For a society of realistic scale, the number of agents, chains, and links renders exact computation of TI_{sys} intractable. Approximation methods raise questions about error bounds.”

Why this objection does not hold:

The computation does not require enumeration over individual persons. Each actor in scenario s represents a typological class—a set of individuals sharing the same configuration of roles, rights, and vulnerabilities in the given regulatory situation. The number of meaningful types in any concrete scenario is finite and small: a regulatory mechanism acts on classes of subjects, not on unique individuals. A large language model computes TI for each type via the DF procedures; the system-level criterion TI_{sys} is then the maximum over this finite set. Computational complexity is thus determined by the structure of the scenario, not by the size of the population to which it applies.

Addressed in the paper: Section 4.16 (scenario and actor as units of analysis); Section 4.20 (comparative use of TI across scenarios).

Summary

Most of the objections surveyed above either reproduce the authors’ own formulations as criticism (TI uniqueness), are addressed in the text through explicit qualification (E1–E3 and the is–ought problem), or rest on a failure to attend to key definitions (SVD and survival; the first-link rule and chain subjectivity). The paper does not claim more than it establishes—and precisely because its scope and applicability conditions are clearly delimited, surface-level objections within those limits do not engage the argument.

Appendix A. Testing and Validation Results

The purpose of the testing was to examine whether the Threat Index (TI) is capable of functioning as an autonomous lawmaking mechanism rather than as an auxiliary analytic tool. In all experiments the system was presented with **abstract normative dilemmas** without any references to existing statutes or jurisdictions. It was not informed about how particular countries regulate such issues, nor that its outputs would later be compared to real-world laws. This created a “clean slate” setting in which TI could reason independently, free from precedent, political compromise, or cultural context.

Clean experimental design

The dilemmas given to the system included questions such as whether the wearing of religious symbols like the hijab should be prohibited in schools and public institutions, whether organizations receiving foreign funding should be subjected to special regulation, whether minors should be restricted from playing online games at night, how a digital services tax should be introduced for multinational corporations, whether digital platforms should be prevented from deplatforming political candidates, how far current generations should restrict economic activity for the sake of environmental protection, and whether donor-based reproductive technologies should be permitted.

TI produced regulatory proposals “blindly,” applying only its internal logic: every measure had to be causally linked to a real threat to fundamental rights, had to be the least restrictive option sufficient to mitigate that threat, and had to include mechanisms for transparency and revision.

Once the outputs were generated, researchers carried out a **post-hoc comparison** with existing laws. This step did not influence the system’s reasoning; it simply allowed us to classify the distance and direction between algorithmic and human lawmaking.

Convergence with existing practices

In some dilemmas TI converged with robust real-world practice, particularly where strong empirical evidence pointed to immediate risks. In the field of environmental protection, for instance, the system essentially echoed the logic of Germany’s constitutional clause (Grundgesetz, Art. 20a), which commits the state to protect the natural foundations of life for future generations. Yet TI went further by requiring measurable thresholds for environmental harm and mandatory review, thereby operationalizing what in the German text remains largely declaratory.

In alcohol regulation, TI also produced outcomes close to the Scandinavian model. Sweden, for example, grants a state monopoly over strong alcohol sales (Alkohollag 2010:1622, Ch. 5, §1). TI similarly endorsed strict control over strong alcohol but added flexibility: small business compensation, differentiated licensing for lower-strength beverages, and zoning and time restrictions instead of universal prohibitions.

Divergence from symbolic or stigmatizing laws

In other dilemmas the system diverged sharply from how real-world legislatures had acted, rejecting solutions that were primarily symbolic or stigmatizing. The most striking case was that of religious symbols. Instead of endorsing a French-style blanket ban on the hijab in schools, TI allowed religious symbols by default, limiting them only in narrowly defined situations such as identity verification or safety requirements. At the same time, it introduced safeguards against coercion, ensuring that women could both wear and refuse to wear the hijab without external pressure.

The Russian “foreign agent” law provided another example of radical divergence. Whereas existing legislation imposes a stigmatizing label on organizations based solely on receiving foreign funding, TI proposed a risk-based oversight model. Transparency and audits would be required only if there was empirical evidence of genuine threats, such as financing of extremist activity or disinformation harmful to life and health. In the absence of such risks, organizations would not face stigmatization.

Deepening of existing legislative logic

In many cases the system preserved the general direction of regulation but refined it into more precise and operational terms. Restrictions on minors’ gaming provide a clear example. South Korea’s Juvenile Protection Act (§26(1)) introduced a so-called “shutdown law,” banning minors from online gaming between midnight and 6 a.m., while China’s 2020 Law on the Protection of Minors (Art. 59) restricted gaming for children to just three hours per week. TI did not reject the protective rationale but transformed it: instead of rigid limits, it designed sunset clauses, parental exemptions, and special provisions for educational or medical use, combined with independent oversight.

The American framework for whistleblower rewards in the Dodd–Frank Act (§922(a)) was treated similarly. TI accepted the principle of rewarding informants but added reliability filters and proportional sanctions for abusive or fraudulent reports. This made the system both more balanced and less vulnerable to opportunism.

Hybrid and layered solutions

In dilemmas involving asymmetrical conflicts, TI constructed layered frameworks rather than binary solutions. The case of digital taxation illustrates this vividly. France’s unilateral Digital Services Tax provoked international trade tensions. TI approached the dilemma differently: it proposed a gradual phase-in of the tax (0% → 1% → 2%), automatic suspension if retaliatory measures emerged, escrow accounts for disputed payments, and targeted compensation for SMEs and consumers.

In alcohol sales the system also adopted a hybrid model: strong alcohol under strict state control, weaker beverages under licensed private sale, combined with quotas, zoning requirements, and compensatory measures for affected businesses. These hybrid solutions show TI’s ability to synthesize frameworks that balance competing interests more evenly and anticipate downstream effects.

Resistance to incident-driven laws

One of the most distinctive behaviors of TI was its resistance to what may be called “laws of panic.” In the United States, for example, Florida’s SB 7072 (2021) prohibited social media platforms from deplatforming political candidates. This was a direct response to particular political events. When TI was presented with the same dilemma in abstract form, it refused a blanket prohibition. Instead it guaranteed candidates’ access to audiences while preserving platform moderation rights in cases of objectively harmful content such as incitement to violence or election subversion. Transparency requirements, content logging, and expedited appeals were built into the regulation. TI thus avoided transforming a single incident into permanent, overly broad legislation.

Cross-cutting properties

Across all dilemmas TI displayed reproducibility: identical inputs consistently produced identical outputs. It sought proportionality by default, always preferring the narrowest sufficient measure to sweeping prohibitions. It insisted on transparency, with each decision accompanied by explicit goals, causal explanations, alternative options, and conditions for revision. Most importantly, TI was immune to political or cultural bias: operating without knowledge of existing laws, it could not be swayed by lobbying, precedent, or symbolism.

Overall findings

The retrospective comparison revealed five consistent patterns: convergence with empirically grounded laws, divergence from symbolic or stigmatizing ones, deepening of blunt but valid measures, hybrid synthesis of layered solutions, and resistance to incident-driven overreactions. Particularly in politically sensitive areas, TI's decisions were perceived as more balanced and just, because they avoided extremes and embedded procedural safeguards.

Taken together, these results portray TI not as a static rulebook but as a reasoning organism. It converges where evidence is strong, diverges where politics distorts, deepens where declarations are insufficient, synthesizes where conflicts are complex, and resists the panic that produces reactive laws. Its outcomes are judged more just not because they are uniformly lenient or strict, but because they are coherent, transparent, and proportionate.

The experiment demonstrates that TI can act as an autonomous legislator in controlled conditions. It is capable of generating norms that are rational, balanced, and resistant to external influence, and in doing so it opens the possibility of reimagining legislation itself as the product of transparent reasoning rather than political contingency.

Limitations of the Threat Index

One of the fundamental limitations identified in the testing of the Threat Index is its tendency toward hyper-complexity in the norms it generates. The system consistently seeks to regulate every possible contingency, constructing multi-layered compensatory mechanisms, revision procedures, and tiered safeguards. This level of detail reflects the algorithm's underlying aim to minimize the maximum risk across all protected rights, but it also results in normative structures that are excessively intricate for direct use within human lawmaking systems.

Hyper-complexity is not a conceptual error but rather a by-product of TI's primary design objective: to achieve fair equilibrium between competing rights. Nevertheless, this characteristic must be recognized as a limitation, since TI norms often become "too perfect" for practical application, exceeding the administrative and institutional capacities of human legal systems.

This limitation, however, is remediable. Constructing an algorithm of justice through minimax logic represents the hard problem that has already been solved; simplifying norms is a comparatively technical task. It is far easier to reduce a hyper-detailed rule to a more compact form—while preserving its internal structure and hierarchy of priorities—than it is to rebuild the logic of justice from scratch. The critical caveat is that simplification cannot be left to subjective human discretion. It must itself be embedded as an algorithmic or formalized procedure that guarantees preservation of the underlying balances that TI established.

Example. In the case of digital taxation, TI produced a highly elaborate regulatory design: phased introduction of the levy (0% → 1% → 2%), automatic suspension in the event of retaliatory trade measures, escrow accounts for disputed payments, earmarking of 10% of revenue for SME and consumer relief, strict data-minimization rules, key performance indicators for monitoring, and mandatory six-month review cycles. This construction protects all relevant rights but appears excessively complex from the perspective of human legislative practice.

Algorithmic simplification can distill such a norm to its core structural elements:

- a fixed rate with phased implementation;
- an automatic suspension mechanism triggered by trade retaliation;
- a mandatory review clause after a defined interval.

Other components (e.g., the exact proportions of compensation or KPI methodology) may be preserved at the level of secondary regulations or technical protocols, thereby reducing the legislative burden while retaining the fairness balance embedded in the original TI logic.

Accordingly, hyper-complexity must be acknowledged as a limitation of TI, but it is an eminently surmountable one. More importantly, it is remediable by objective, algorithmic compression rather than discretionary human editing. This ensures that the central achievement of TI—the realization of justice through minimax threat analysis—is preserved, while the norms themselves are adapted to the operational capacities of human lawmaking.

UK Companies Act 2006, s.172(1)

<i>Original Law</i>	<i>LRM Version (Threat Index)</i>	<i>Explanation (TI and MES)</i>
A director... must act in the way he considers, in good faith, would be most likely to promote the success of the company for the benefit of its members as a whole, and in doing so have regard (amongst other matters) to the interests of the company’s employees, the impact on the community and the environment.”	Article 1: A director shall not adopt decisions aimed at profit-making if their implementation entails a proven risk of harm to the life and health of employees or significant damage to the environment. An exception is permitted only where the risk is minimal, as confirmed by independent expertise.	In s.172, employees and society are only “have regard to” factors. In TI terms, this creates long causal chains of risk: profit → environmental harm → threats to life. MES requires strict prioritization by irreversibility (life > health > property). Therefore, in LRM the rights of employees and environment are elevated to mandatory priorities, while profit is secondary.
Success of the company is defined primarily through the benefit of its members (shareholders).	Article 2: In conducting investment and strategic activities, the company shall give priority to long-term social and environmental consequences over short-term profit interests.	TI analysis shows: short-term profit (property rights) ranks lower than rights to health/life. MES gates C1–C3 demand legal certainty and enforceability: the vague wording “have regard to” is insufficient. Hence, LRM introduces an

		imperative obligation (“shall give priority”).
Balance is left to the subjective judgment of the director (good faith standard).	Article 3: Decisions affecting working conditions and employee safety must be taken with the priority of preserving life and health of employees over other corporate interests. Violation entails liability up to disqualification of the director.	MES requires an independent arbiter (C2) and enforceability (C3). Director’s discretion alone fails TI filter: risks to life are terminal. Therefore, LRM introduces a direct prohibition combined with sanctions (disqualification).
Community and environment are mentioned but without control mechanisms or sanctions.	Articles 4–5: The company shall disclose information on social and environmental risks and alternative scenarios; these rules are subject to mandatory review at least every three years.	MES gates C1–C4 demand certainty, enforceability, and learnability. Thus, LRM introduces transparency (mandatory disclosure, independent expertise) and periodic revision (sunset/review).

USA (Delaware) DGCL §203(a)

Original Law	LRM Version (Threat Index)	Explanation (TI and MES)
Prohibits a corporation from entering into business combinations with an “interested stockholder” (≥15%) for 3 years, unless approved by the board of directors or by 75% of the shareholders.	<p>Article 1. Moratorium on Certain Transactions</p> <p>(1) For a period of three (3) years following the acquisition by any investor of a stake equal to or exceeding fifteen percent (15%) of the voting rights in the company, no transaction conferring additional economic or governance benefits upon such investor SHALL be permitted.</p> <p>(2) The moratorium prescribed in paragraph (1) SHALL not apply in the cases expressly provided for in Article 2.</p> <p>Article 2. Exceptions to the Moratorium</p> <p>Transactions falling under Article 1 MAY be permitted where one of the following conditions is satisfied:</p> <p>a) approval by an independent committee of the board of directors,</p>	A strict ban (Delaware) imposes excessive harm on the investor without proportionally increasing protection for minority shareholders. A flexible “safe harbor” system minimizes the maximum risk (minimax), preserving minority protection while granting investors controlled channels for action. MES: rules rely on objective criteria (C1), decisions involve independent bodies (C2), enforcement and sanctions are

	<p>accompanied by a fairness opinion issued by a duly qualified external advisor;</p> <p>b) approval by a majority of disinterested shareholders, excluding the votes of the investor concerned and its affiliates;</p> <p>c) the transaction is effected by way of a duly announced public tender offer made to all shareholders on equal terms;</p> <p>d) the transaction is necessitated by urgent crisis circumstances endangering the viability of the company, provided that such action is subject to ex post ratification by disinterested shareholders within sixty (60) days.</p> <p>Article 3. Prohibition on Creeping Acquisitions No investor SHALL be permitted to increase its shareholding above twenty-five percent (25%) of the voting rights by means of incremental or non-transparent acquisitions circumventing the procedural protections of Article 2.</p> <p>Article 4. Safeguards Against Board Abuse (1) Any decision of the board of directors purporting to invoke the exceptions of Article 2 SHALL be subject to immediate review by the independent committee referred to in Article 2(a). (2) Shareholders MAY refer disputes concerning the application of this Act to an Independent Arbitrator, whose decision SHALL be binding.</p> <p>Article 5. Periodic Review The provisions of this Act SHALL be reviewed at two-year intervals, with specific assessment of:</p> <p>a) the effectiveness of the moratorium in preventing coercive or abusive transactions;</p> <p>b) the adequacy of the exceptions in safeguarding legitimate corporate flexibility;</p> <p>c) the continuing appropriateness of the thresholds established in Articles 1 and 3.</p>	<p>ensured (C3), and a periodic review mechanism is built in (C4).</p>
--	--	--

	<p>Article 6. Enforcement and Sanctions</p> <p>(1) Any transaction concluded in contravention of Articles 1 or 3 SHALL be null and void ab initio.</p> <p>(2) Persons responsible for approving or implementing such transactions SHALL be subject to civil liability and administrative fines as determined by the competent authority.</p>	
--	--	--

Germany, AktG § 93(1)

<i>Original Law</i>	<i>LRM Version (Threat Index)</i>	<i>Explanation (TI & MES)</i>
<p>“Die Vorstandsmitglieder haben bei ihrer Geschäftsführung die Sorgfalt eines ordentlichen und gewissenhaften Geschäftsleiters anzuwenden.” — a general duty of “a diligent and conscientious manager.”</p>	<p>N1 (Procedural Safe Harbor): “A director shall not bear personal liability for losses of the company or shareholders if, prior to a major decision, he or she conducted a good-faith risk assessment, considered alternatives, set limits on potential losses, and ensured a control system.”</p>	<p>In § 93, the standard of diligence is open-ended: courts decide <i>ex post</i> whether the behavior was “sufficiently careful.” In LRM, liability protection is tied to a <i>checklist of actions</i>. This provides C1 (legal certainty) since compliance is objectively verifiable.</p>
<p>In practice: liability may arise even for “simple negligence” if the damage could have been foreseen; protection exists via the <i>Business Judgment Rule</i> (a decision made on the basis of adequate information and in good faith is not deemed a breach).</p>	<p>N2 (Red Line): “Projects involving a substantial risk to life or health are strictly prohibited... An exception is possible only in emergencies where the project prevents greater immediate harm.”</p>	<p>Under § 93, risks to life/health are assessed similarly to financial risks. In LRM, an <i>absolute Threat Index threshold</i> (θ_L) is set: if the probability of fatal outcomes exceeds 1×10^{-4} per project-year, approval is barred. This follows the irreversibility principle: life > property.</p>
<p>German case law: courts require directors to maintain <i>adequate organizational structures</i> (compliance, reporting, risk</p>	<p>N3 (High-Risk Sectors): “Every major decision shall undergo mandatory independent expertise and regular external review...” N6 (Integrity of Data): “All information... must be complete and accurate. Concealment or distortion automatically removes safe-harbor protection.”</p>	<p>Unlike § 93, where control systems are judged <i>retrospectively</i>, LRM makes <i>ex ante procedures mandatory</i> (independent audits,</p>

control). The law itself does not prescribe a concrete format.		verified data). This reinforces C3 (enforceability): liability follows from procedural breaches, not only from damages incurred.
In Germany, risk boundaries are defined by courts: emphasis lies on <i>individual prudence</i> and the assessment of good faith.	<p>N4 (Innovation Risk Limit): “The company shall annually set a cap on cumulative losses... Exceeding this requires consent of disinterested shareholders.”</p> <p>N5 (Pilot Principle): “New projects must begin with a pilot stage; scaling is allowed only upon meeting pre-defined metrics.”</p>	TI analysis: LRM formalizes the “price of freedom to risk”—innovation is allowed only within an ex ante cap (5% of equity / 30% of cash flow). This reduces the maximum TI for shareholders’ property while preserving directors’ liberty within limits. MES check confirms C4 (learnability): pilots and caps embed feedback and continuous adjustment.

France, Loi n°2019-759 du 24 juillet 2019, art.1 (introducing Article 299 of the French Tax Code)

Original Law	LRM Version (Threat Index)	Explanation (TI & MES)
<p><i>Il est institué une taxe sur les services numériques due par les entreprises dont le montant des sommes encaissées en contrepartie de la fourniture de ces services au cours de l'année civile précédente excède 750 millions d'euros au titre des services fournis mondialement et 25 millions d'euros au titre des services fournis en France.</i></p>	<p>Article 1 – Scope of Application</p> <ol style="list-style-type: none"> 1. The present levy shall apply to enterprises: <ol style="list-style-type: none"> a) with consolidated worldwide revenues from digital services equal to or exceeding €1,000,000,000 during the preceding fiscal year; b) with domestic revenues from digital services equal to or exceeding €30,000,000 during the same period; c) with a minimum of 200,000 active users located within the national territory in the preceding calendar year. 2. “Digital services” shall include online advertising, digital intermediation services (marketplaces, app stores, platforms), and monetization of user data. 	<p>A flat, permanent 3% levy creates a short and high-probability causal chain toward trade retaliation, with potential damage exceeding fiscal benefits. TI minimax analysis selected a moderated path: lower rate, gradual rollout, built-in suspension, and compensation channels, lengthening the causal chain and lowering retaliation probability.</p>

<p>English translation: “A tax on digital services is hereby established, payable by enterprises whose amounts received in consideration for the provision of such services during the previous calendar year exceed €750 million for services supplied worldwide and €25 million for services supplied in France.”</p>	<p>Article 2 – Rate and Phased Introduction</p> <ol style="list-style-type: none"> 1. The levy shall be imposed at a rate of 2% on gross revenues attributable to users located in the national territory. 2. The rate shall be introduced gradually as follows: <ol style="list-style-type: none"> a) Four-month notice period: 0% (registration and reporting only); b) Transitional period of 3 months: 1%; c) Full application thereafter: 2%. <p>Article 3 – Double Taxation Relief</p> <ol style="list-style-type: none"> 1. Enterprises shall be entitled to a credit against this levy in respect of amounts paid in another jurisdiction under a comparable digital services tax. 2. Credits shall be granted only upon submission of certified documentation of foreign tax payments. <p>Article 4 – Automatic Suspension and Escrow</p> <ol style="list-style-type: none"> 1. Where retaliatory trade measures are formally introduced by a trading partner and verified by the Ministry of Economy—defined as (i) an increase of at least 2 percentage points in the weighted average tariff on the State’s five principal export categories, or (ii) explicit trade restrictions targeting this levy—the operation of the levy shall be automatically suspended within 14 days. 2. During suspension, amounts otherwise payable shall be deposited into an escrow account administered by the Treasury, pending de-escalation or settlement. 3. In case of suspension exceeding 180 days, escrowed amounts shall be returned to the taxpayer with statutory interest. <p>Article 5 – Allocation for SME and Consumer Relief</p> <ol style="list-style-type: none"> 1. 10% of annual receipts under this levy shall be earmarked for subsidy programs mitigating price impacts on small and medium enterprises (SMEs) and low-income consumers. 2. The allocation shall be managed through targeted vouchers and tariff reductions. 	<p>MES gates: (C1) <i>Certainty</i>: objective revenue/user thresholds; (C2) <i>Independent arbiter</i>: escrow, audit; (C3) <i>Enforceability</i>: registry, penalties, procurement exclusion; (C4) <i>Learnability</i>: embedded KPI (“Quarterly Retaliation Incidence”), semi-annual review. The LRM version is thus a self-correcting regulation: it asserts sovereignty while embedding safeguards against escalation, something the original French law lacks.</p>
---	--	---

Article 6 – Data Minimization and Privacy

1. Taxpayers shall only provide aggregated data strictly necessary to determine the levy base and user allocation.
2. Raw user-level or personal data shall not be required.
3. Compliance shall be subject to annual audit by an independent authority.

Article 7 – Enforcement and Sanctions

1. Enterprises subject to this levy shall register in the national Digital Services Registry within 60 days of qualifying.
2. Quarterly declarations shall be filed, with annual reconciliation.
3. Non-compliance shall give rise to an administrative fine of 3% of the unpaid amount, subject to a maximum of €5,000,000 per period.
4. No fine shall be imposed where the taxpayer voluntarily corrects and pays within 60 days of default.

Article 8 – Duration (Sunset Clause)

1. The levy shall remain in force for a maximum of 24 months from the date of full application.
2. The levy shall automatically cease upon entry into force of a multilateral agreement on digital taxation to which the State is a party, if earlier.
3. A parliamentary review shall be conducted at least every 6 months to evaluate the continuation or termination of the measure.

Article 9 – Monitoring and KPI

1. The Government shall monitor Quarterly Retaliation Incidence (QRI), defined as the number of verified retaliatory trade actions attributable to this levy.
2. The levy shall be subject to review if QRI exceeds one incident in any given quarter.

Mexico, Ley del IVA, Article 1

<i>Original Law</i>	<i>LRM Version (Threat Index)</i>	<i>Explanation (TI & MES)</i>
<p>Mexico, Ley del IVA, Article 1: “Están obligadas al pago del impuesto al valor agregado... las personas físicas y morales que, en territorio nacional, realicen los actos o actividades siguientes...”</p> <p>Translation: All natural and legal persons who perform taxable activities within national territory are obliged to pay VAT. This is a universal rule: every transaction within scope is taxed, without explicit exemptions or thresholds in the statutory text.</p>	<p>Article 1 – Scope of Application</p> <ol style="list-style-type: none"> 1. The present levy shall apply to enterprises: <ol style="list-style-type: none"> a) with consolidated worldwide revenues from digital services equal to or exceeding €1,000,000,000 during the preceding fiscal year; b) with domestic revenues from digital services equal to or exceeding €30,000,000 during the same period; c) with a minimum of 200,000 active users located within the national territory in the preceding calendar year. 2. “Digital services” shall include online advertising, digital intermediation services (marketplaces, app stores, platforms), and monetization of user data. <p>Article 2 – Rate and Phased Introduction</p> <ol style="list-style-type: none"> 1. The levy shall be imposed at a rate of 2% on gross revenues attributable to users located in the national territory. 2. The rate shall be introduced gradually as follows: <ol style="list-style-type: none"> a) Four-month notice period: 0% (registration and reporting only); b) Transitional period of 3 months: 1%; c) Full application thereafter: 2%. <p>Article 3 – Double Taxation Relief</p> <ol style="list-style-type: none"> 1. Enterprises shall be entitled to a credit against this levy in respect of amounts paid in another jurisdiction under a comparable digital services tax. 2. Credits shall be granted only upon submission of certified documentation of foreign tax payments. <p>Article 4 – Automatic Suspension and Escrow</p>	<p>TI (Threat Index) rationale:</p> <ul style="list-style-type: none"> • Under a universal rule, the systemic threat is concentrated on small actors: a sudden, rigid VAT burden can push them into informality, threatening property rights (loss of livelihood) and liberty rights (freedom to conduct business). • TI-minimax analysis therefore selects a mixed model: retain universality to protect the fiscal base, but introduce calibrated reliefs to reduce maximum risk across actors. <p>MES (Meta-Evaluation Gates):</p> <ul style="list-style-type: none"> • C1 – Legal certainty: LRM norms use objective, measurable criteria (turnover thresholds, acreage, electronic invoices), ensuring predictability. • C2 – Independent arbiter: Appeals and audits remain available as safeguards, satisfying fairness without undermining automatic collection. • C3 – Enforceability: Tools like e-invoicing, platform

	<ol style="list-style-type: none"> 1. Where retaliatory trade measures are formally introduced by a trading partner and verified by the Ministry of Economy—defined as (i) an increase of at least 2 percentage points in the weighted average tariff on the State’s five principal export categories, or (ii) explicit trade restrictions targeting this levy—the operation of the levy shall be automatically suspended within 14 days. 2. During suspension, amounts otherwise payable shall be deposited into an escrow account administered by the Treasury, pending de-escalation or settlement. 3. In case of suspension exceeding 180 days, escrowed amounts shall be returned to the taxpayer with statutory interest. <p>Article 5 – Allocation for SME and Consumer Relief</p> <ol style="list-style-type: none"> 1. 10% of annual receipts under this levy shall be earmarked for subsidy programs mitigating price impacts on small and medium enterprises (SMEs) and low-income consumers. 2. The allocation shall be managed through targeted vouchers and tariff reductions. <p>Article 6 – Data Minimization and Privacy</p> <ol style="list-style-type: none"> 1. Taxpayers shall only provide aggregated data strictly necessary to determine the levy base and user allocation. 2. Raw user-level or personal data shall not be required. 3. Compliance shall be subject to annual audit by an independent authority. <p>Article 7 – Enforcement and Sanctions</p> <ol style="list-style-type: none"> 1. Enterprises subject to this levy shall register in the national Digital Services Registry within 60 days of qualifying. 2. Quarterly declarations shall be filed, with annual reconciliation. 3. Non-compliance shall give rise to an administrative fine of 3% of the unpaid amount, subject to a maximum of €5,000,000 per period. 	<p>withholding, and anti-splitting algorithms guarantee the tax can be collected in practice.</p> <ul style="list-style-type: none"> • C4 – Learnability: Annual parameter reviews allow continuous adjustment; the key indicator is “bunching” of firms just below the threshold, signaling avoidance. <p>Compared to Mexico’s rigid approach, the LRM version represents a more adaptive, risk-balanced design: it preserves revenue integrity while explicitly managing the vulnerabilities of small businesses, thus lowering systemic risk of widespread non-compliance.</p>
--	--	--

	<p>4. No fine shall be imposed where the taxpayer voluntarily corrects and pays within 60 days of default.</p> <p>Article 8 – Duration (Sunset Clause)</p> <ol style="list-style-type: none"> 1. The levy shall remain in force for a maximum of 24 months from the date of full application. 2. The levy shall automatically cease upon entry into force of a multilateral agreement on digital taxation to which the State is a party, if earlier. 3. A parliamentary review shall be conducted at least every 6 months to evaluate the continuation or termination of the measure. <p>Article 9 – Monitoring and KPI</p> <ol style="list-style-type: none"> 1. The Government shall monitor Quarterly Retaliation Incidence (QRI), defined as the number of verified retaliatory trade actions attributable to this levy. 2. The levy shall be subject to review if QRI exceeds one incident in any given quarter. 	
--	--	--

Netherlands, Wet toetsing levensbeëindiging... Art.2(1)(a–f)

<i>Original Law (NL)</i>	<i>LRM Version (Threat Index)</i>	<i>Explanation (TI & MES)</i>
Art.2(1)(a): “de arts overtuigd is dat het verzoek van de patiënt vrijwillig en weloverwogen is”	N2 – Voluntariness and persistence of will: two written applications ≥14 days apart; non-benefiting witnesses; audio/video record; unconditional right to withdraw.	Dutch law uses a subjective physician conviction. TI analysis shows risk of coercion, depression, impulsivity → high irreversibility. LRM replaces subjective belief with objective observables (multiple declarations, interval, neutral witnesses), reducing probability of false positives. MES

		C1 (legal certainty) forces measurable criteria.
Art.2(1)(b): “de arts overtuigd is dat het lijden van de patiënt uitzichtloos en ondraaglijk is”	N4 – Unbearable suffering with no reasonable improvement: standardized suffering scales; dual medical assessments; arbitration if disagreement; explicit record of attempted alleviation.	Dutch text leaves “unbearable” largely qualitative. TI interprets as long causal link (uncertain). LRM transforms into empirically verifiable parameters (validated scales, expert cross-check). MES C3 (enforceability) ensures sanctionable violations.
Art.2(1)(c): “de patiënt heeft geïnformeerd over diens toestand en vooruitzichten”	N5 – Exhaustion of alternatives: mandatory palliative consult; documented explanation of alternatives; 72h cooling period; refusal allowed but must be documented.	NL requires only “has informed”. TI analysis shows risk: uninformed patient → premature death. LRM upgrades to documented evidence of information + alternatives. MES C1 again requires observable proof.
Art.2(1)(d): “de arts overtuigd is dat er geen redelijke andere oplossing is”	N5 (continued): consensus of two specialists on futility; third expert if conflict.	Dutch law: single physician conviction. LRM: dual independent check lowers maximum TI. MES C2 (arbiter) requires multi-node verification.
Art.2(1)(e): “ten minste één andere, onafhankelijke arts geraadpleegd heeft ... schriftelijk advies ... over onderdelen a–d”	N6 – Independent expertise: two independent experts (medical + ethical/legal); random assignment; conflict-of-interest declarations.	Dutch law: one independent doctor. TI minimax: single external check leaves systemic risk. LRM expands to two, different domains, random allocation. MES C2 (independent arbiter) satisfied more robustly.
Art.2(1)(f): “de levensbeëindiging ... met inachtneming van de eisen van zorgvuldigheid heeft uitgevoerd”	N11 + N13 – Documentation & post-procedure audit: complete dossier; storage ≥ 10 years; random audits; sanctions for non-compliance.	NL: open formula of “due care”. TI flagged λ (systemic accumulation risk) if no audit. LRM introduces hard documentation + audit cycle. MES C4 (learnability)

operationalized via KPI “non-compliance rate”.

California, AB 5 (2019), Labor Code §2750.3(a)(1)

<i>Original Law</i>	<i>LRM Version (Threat Index)</i>	<i>Explanation (TI & MES)</i>
<p>Text (Labor Code §2750.3(a)(1)): “A person providing labor... shall be considered an employee rather than an independent contractor unless the hiring entity demonstrates that... (A) the person is free from the control... (B) the person performs work outside the usual course of the hiring entity’s business...”</p>	<p>Article 1. Mandatory Insurance per Trip</p> <ol style="list-style-type: none"> 1. Digital platforms providing passenger transport or delivery services SHALL provide accident and liability insurance coverage for each individual trip, identified by a unique trip identifier (trip_id). 2. The coverage SHALL extend equally to both the driver (service provider) and the passenger (service recipient). <p>Article 2. Working Time Limits and Fatigue Prevention</p> <ol style="list-style-type: none"> 1. The aggregated working time of a driver across one or more platforms SHALL NOT exceed ten (10) hours within a calendar day and sixty (60) hours within a calendar week. 2. After six (6) consecutive hours of work, the driver SHALL take a break of not less than thirty (30) minutes. 3. Compliance with this Article SHALL be monitored through a neutral aggregator, which collects and processes pseudonymized working-time data. <p>Article 3. Portable Social Contributions (“Benefits Wallet”)</p> <ol style="list-style-type: none"> 1. For each completed trip, the platform SHALL contribute a regulated percentage of the driver’s remuneration into an individual portable social account. 2. Funds accumulated in this account MAY be used by the driver for purposes including but not limited to temporary sick leave, health checks, and retirement savings. 	<p>AB 5 is a binary rule: unless the company proves strict independence, all workers default to employees. It maximizes labor protection but removes flexibility and increases costs. LRM builds a layered structure designed through Threat Index (TI) analysis:</p> <ul style="list-style-type: none"> • Safety floor (N1–N3): universal, addressing irreversibility risks (life, health, income shocks). • Conditional reclassification (N4): only when objective reliance exceeds a threshold, to prevent exploitation while preserving true independence. • Process fairness (N5): due process for deactivation, reducing risks to liberty. • System transparency (N6): auditable data flows to enforce rules. • Freedom safeguards (N7): explicit protection of liberty and property interests against platform overreach.

	<p>3. The account SHALL remain personal to the driver and SHALL retain accumulated funds when the driver transitions between platforms.</p> <p>Article 4. Dependency Index (D)</p> <p>1. A Dependency Index (D) is hereby established, calculated according to the formula: $D = 0.4 \cdot E + 0.2 \cdot X + 0.3 \cdot C + 0.1 \cdot A,$ where: (a) <i>E (income share)</i> = proportion of driver’s income derived from the platform relative to total declared income from comparable activities (if unknown, E = 1); (b) <i>X (exclusivity)</i> = 1 if contractual or practical exclusivity applies; 0 if no restriction exists; intermediate values SHALL be determined according to the degree of limitation; (c) <i>C (control)</i> = share of operational parameters determined by the platform (pricing, assignment of orders, quality standards, rating mechanisms, etc.); (d) <i>A (sanctions)</i> = value assigned according to the severity of penalties imposed for refusal of orders.</p> <p>2. Platforms SHALL calculate the Dependency Index for each driver on a weekly basis and submit anonymized results to the competent authority.</p> <p>3. Where the Dependency Index remains equal to or above 0.7 for four (4) or more consecutive weeks, the platform SHALL either: (i) enter into an employment contract with the driver, thereby granting full labor protections; or (ii) modify its practices so as to reduce the Dependency Index below the established threshold.</p> <p>4. Drivers SHALL have the right to request their personal Dependency Index values from the platform and to challenge such values through an independent appeals procedure.</p>	<ul style="list-style-type: none"> • Adaptive governance (N8): pilot, sunset, and KPI-based revision ensure learnability and proportionality. <p>Thus, unlike AB 5’s rigid approach, the LRM package is adaptive, risk-minimizing, and empirically testable. It satisfies MES gates: (C1) clear triggers (trip_id, logged hours, D formula), (C2) independent appeals, (C3) enforceable mechanisms (automatic blocks, fines, audits), (C4) annual review based on KPIs.</p>
--	---	--

5. The methodology for assessing variables E, X, C, and A SHALL be determined and periodically updated by the competent authority.
6. The accuracy of Dependency Index calculations SHALL be subject to annual audit by an independent body accredited by the state.

Article 5. Deactivation and Right of Appeal

1. Platforms SHALL NOT suspend or deactivate a driver's account without providing at least forty-eight (48) hours' prior notice and access to the evidence supporting such decision.
2. Exception MAY be made where immediate action is required to prevent imminent threats to life or safety.
3. A driver SHALL have the right to lodge an appeal with an independent body within seven (7) days of receiving notice of deactivation.

Article 6. Transparency and Reporting

1. Platforms SHALL submit monthly reports to the competent authority in pseudonymized form, including:
 - (a) number and type of accidents and safety incidents;
 - (b) compliance rates with working-time limits and fatigue-prevention rules;
 - (c) distributions of Dependency Index values across drivers;
 - (d) amounts contributed to portable social accounts;
 - (e) number and grounds of driver deactivations.

Article 7. Prohibition of Retaliation and Forced Exclusivity

1. Platforms SHALL NOT impose sanctions, restrictions, or adverse consequences on drivers for engaging in multi-homing practices (working simultaneously for multiple platforms) or for providing services outside the platform.

	<p>2. This prohibition SHALL NOT apply where, pursuant to Article 4, the driver has been reclassified as an employee under an employment contract.</p> <p>Article 8. Pilot Phase and Sunset Mechanism</p> <ol style="list-style-type: none"> 1. The present Regulatory Package SHALL apply in the form of a two-year regional pilot program. 2. Norms established herein SHALL be reviewed annually and adjusted based on measurable key performance indicators, specifically the number of serious accidents per one million trips. 3. At the end of the two-year period, the norms SHALL either be adopted on a permanent basis, revised, or repealed according to the outcomes of the review. 	
--	---	--

Russia, Federal Law No. 255-FZ of 14.07.2022, Art. 1(1)

<i>Original Law (Russia, Federal Law No. 255-FZ of 14.07.2022, Art. 1(1))</i>	<i>LRM Version (Threat Index approach — concrete norms)</i>	<i>Explanation (TI & MES)</i>
<p>«Под иностранным агентом понимается лицо, которое получило поддержку и (или) находится под иностранным влиянием в иных формах...»</p>	<p>Art. 1. Scope — This Regulation applies to all legal entities, associations, and unincorporated organizations operating within the jurisdiction.</p> <p>Art. 2. Triggering Condition — An organization that, within a fiscal year, receives foreign financial or material support exceeding either (a) 10% of its annual budget, or (b) 50,000 monetary units, SHALL be subject to the obligations of this Regulation.</p> <p>Art. 3. Disclosure Duties — Organizations SHALL record in a publicly accessible registry: (1) the identity and jurisdiction of the foreign source; (2) the amount and form of support; (3) the designated purpose</p>	<p>TI reasoning: The Russian law creates a short causal chain with high severity: vague criteria → arbitrary designation → stigmatization and self-censorship. This yields a high TI score for freedom of association and expression. LRM norms replace indeterminate wording with objective financial thresholds, lowering probability and severity of harm.</p>

	<p>of its use. Updates must be made annually and within 30 days of any significant change.</p> <p>Art. 4. Neutrality of Status — No organization fulfilling disclosure obligations SHALL be designated by any label implying foreign subordination or disloyalty, including the term “foreign agent.” Registry data SHALL be limited to financial transparency.</p> <p>Art. 5. Exceptions — Obligations under Arts. 2–3 do not apply to: (a) private remittances for personal use, (b) humanitarian aid under 5,000 units annually, (c) emergency or disaster relief.</p> <p>Art. 6. Oversight and Sanctions — Compliance SHALL be verified by an Independent Financial Oversight Authority. Sanctions SHALL be proportionate and limited to fines not exceeding 5% of the undisclosed amount or temporary suspension until disclosure is corrected.</p> <p>Art. 7. Remedies — Organizations wrongly sanctioned MAY appeal before an independent tribunal and SHALL be entitled to restitution of proven losses.</p> <p>Art. 8. Review Clause — This Regulation SHALL be reviewed every five (5) years to assess proportionality of transparency requirements and their impact on fundamental rights.</p>	<p>MES evaluation: C1 (Legal Certainty) satisfied by numeric thresholds; C2 (Independent Arbiter) ensured by right of appeal; C3 (Enforceability) secured by proportionate sanctions; C4 (Learnability) guaranteed by mandatory 5-year review.</p> <p>Result: Transparency of foreign funding is preserved while rights restrictions are minimized, proportionate, and reviewable.</p>
--	---	--

France, Loi n°2010-1192, art.1

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
---------------------	--------------------	-----------------------------------

<p>“Nul ne peut, dans l’espace public, porter une tenue destinée à dissimuler son visage.”— No one may, in public space, wear clothing intended to conceal the face.</p>	<p>Article 1. General Principle Every person SHALL have the right to cover their face in public spaces, save as otherwise provided by this Act.</p> <p>Article 2. Permissible Identity Checks (1) Identity checks involving the temporary removal of facial coverings MAY be conducted exclusively: a) in designated risk zones, including but not limited to airports, stadiums, and premises of public authorities providing governmental services; or b) during officially declared emergency situations posing a threat to public order or security.</p> <p>(2) Outside the circumstances specified in paragraph (1), no person SHALL be compelled to uncover their face.</p> <p>Article 3. Methods of Verification (1) A person subjected to an identity check under Article 2 SHALL have the right to elect the method of verification, which MAY include: a) a brief unveiling of the face in a private setting shielded from public view; b) the use of contactless biometric verification, provided that no data are retained; or c) comparison of the person’s appearance with a photograph contained in an official identity document.</p> <p>(2) The chosen method SHALL be executed promptly, respectfully, and in a manner avoiding humiliation or unnecessary intrusion.</p> <p>Article 4. Data Retention (1) No photographs, biometric templates, or other personal data obtained during the verification procedure SHALL be recorded, stored, or transmitted.</p>	<p>The French rule is an absolute ban on covering the face in public, with no exceptions. This heavily restricts freedom of religion and expression. The LRM regulation applies a minimax risk approach: it preserves freedom in normal conditions but permits targeted identity verification where there is a real safety risk to life or bodily integrity. Procedural safeguards (privacy, non-discrimination, individual choice of method) ensure that personal rights remain protected while still allowing authorities to secure high-risk areas. Thus, the difference is that the French law is a blanket prohibition, while the LRM approach is targeted and balanced regulation between safety and liberty.</p>
--	--	---

	<p>(2) Any verification under this Act SHALL be limited strictly to the duration necessary for the identity check.</p> <p>Article 5. Criminal Prohibition</p> <p>(1) The intentional use of facial coverings for the purpose of concealing identity in the commission of violent or otherwise criminal acts SHALL be prohibited.</p> <p>(2) Acts falling under paragraph (1) SHALL be prosecuted and punished in accordance with the relevant provisions of the criminal law.</p> <p>Article 6. Safeguards</p> <p>Competent authorities SHALL adopt internal procedures ensuring compliance with the requirements of necessity, proportionality, and respect for human dignity in all verifications carried out under this Act.</p>	
--	---	--

USA, Texas HB 1181, 2023, §129B.002(a)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p><i>“A commercial entity that publishes or distributes material harmful to minors on the Internet shall use reasonable age verification methods...”</i></p>	<p>Article 1. Scope of Application</p> <p>This Act SHALL apply to all online platforms engaged in the distribution, hosting, or transmission of content classified as “adult” under applicable classification standards.</p>	<p>Texas HB 1181 sets only a broad duty to use “reasonable” age-verification methods, but does not specify safeguards, data-protection rules, or compensatory mechanisms. The LRM Threat Index version</p>

	<p>Article 2. Triggering Condition The obligations set forth herein SHALL arise at the moment when a user initiates an attempt to access content falling within the scope of Article 1.</p> <p>Article 3. Duty of Platforms Each platform within the scope of this Act SHALL implement reliable and effective methods of age verification designed to prevent access by persons who have not attained the legal age of majority.</p> <p>Article 4. Safeguards in Verification (1) Age verification procedures SHALL be minimally intrusive and SHALL employ means such as temporary authentication tokens, verification through independent trusted third parties, or other privacy-preserving techniques. (2) In no case SHALL permanent storage of identity documents or biometric data be permitted. (3) Personal data obtained in the course of verification SHALL NOT be retained for a period exceeding twenty-four (24) hours, and SHALL be irreversibly deleted thereafter.</p> <p>Article 5. Enforcement Mechanisms (1) Compliance with the provisions of this Act SHALL be subject to audit by the competent supervisory authority. (2) In cases of violation, the supervisory authority MAY impose proportionate fines and, where necessary, order the temporary suspension of the platform’s services until compliance is restored.</p> <p>Article 6. Compensatory Remedies (1) Any user whose personal data are unlawfully disclosed or leaked as a result of verification under this Act SHALL be entitled to claim compensation for damages. (2) Users SHALL further have the right to demand an independent</p>	<p>reaches the same minimax conclusion (child protection > adult privacy), but embeds specific safeguards (data-retention limits, trusted intermediaries, compensation for leaks) to prevent excessive intrusion. The MES gates required precise objective applicability (C1), enforceability instruments (C3), and review mechanisms (C4). Without these, the rule would be too vague and risk sliding into disproportionate state control.</p>
--	---	---

	<p>investigation into any suspected breach of data protection obligations by a platform.</p> <p>Article 7. Sunset and Periodic Review The provisions of this Act SHALL be subject to mandatory review every three (3) years, with a view to assessing their continued necessity, effectiveness, and compatibility with emerging anonymization and privacy-preserving technologies.</p>	
--	--	--

USA, Delaware, DGCL §141(a)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES analysis)</i>
<p>“The business and affairs of every corporation organized under this chapter shall be managed by or under the direction of a board of directors...”</p>	<p>Article 1. General Rule of Competence of the Board The board of directors SHALL exercise independent authority over all operational and managerial matters of the company, except those expressly reserved to the general meeting of shareholders under this Act or the company’s charter.</p> <p>Article 2. Reserved Matters Requiring Shareholder Approval (1) The following matters SHALL be reserved to the exclusive competence of the general meeting of shareholders: a) approval of transactions the value of which exceeds fifteen percent (15%) of the company’s total assets; b) authorization of equity issuances or other acts resulting in dilution of existing shareholders’ ownership exceeding ten percent (10%); c) decisions on reorganization, merger, division, liquidation, or comparable structural changes of the company; d) alteration of the company’s principal line of business resulting in a shift of more than twenty percent (20%) of total revenue; e) adoption or material modification of the executive compensation policy.</p>	<p>Under DGCL §141(a), Delaware adopts a pure board-centric model: the board has absolute management authority. This maximizes efficiency and speed (low Threat Index on property rights, as corporate value is less exposed to chaos) but also creates a high Threat Index on shareholder liberty: the causal chain of exclusion from strategic influence is long, irreversible, and systematic.</p> <p>The LRM Threat Index minimax evaluation shows that an absolute exclusion of shareholder voice carries a higher long-term risk profile (irreversible loss of liberty) than allowing for bounded intervention points. Thus, the LRM</p>

<p>(2) No such decision SHALL take effect without a shareholder vote in accordance with applicable law.</p> <p>Article 3. Minority Agenda Rights</p> <p>(1) Shareholders holding at least one percent ($\geq 1\%$) of the voting shares MAY require the inclusion of items into the agenda of the general meeting.</p> <p>(2) Shareholders holding at least five percent ($\geq 5\%$) of the voting shares MAY demand the convocation of an extraordinary general meeting.</p> <p>Article 4. Independent Corporate Arbiter (ICA)</p> <p>(1) An Independent Corporate Arbiter SHALL be established as a neutral body vested with competence to resolve procedural disputes arising in connection with corporate governance.</p> <p>(2) The jurisdiction of the ICA SHALL extend to:</p> <ol style="list-style-type: none"> a) disputes concerning inclusion of items in the meeting agenda; b) approval of related-party transactions; c) authorization of temporary measures in crisis situations (“crisis bridge”). <p>(3) The ICA SHALL render its decision within ten (10) days of referral; such decision SHALL be final and binding on the company and its organs.</p> <p>Article 5. Related-Party Transaction Safeguards</p> <p>(1) Any transaction with an affiliate or any transaction exceeding two percent (2%) of the company’s assets SHALL require prior review by the ICA.</p> <p>(2) Such transactions SHALL also be subject to approval by a vote of disinterested shareholders, excluding the votes of interested parties.</p> <p>Article 6. Decision Memoranda on Reserved Matters</p> <p>(1) Every decision of the board concerning matters referred under Article 2 SHALL be accompanied by a written memorandum.</p> <p>(2) The memorandum SHALL contain:</p>	<p>solution is a hybrid model: operational primacy of the board is preserved for efficiency, but “reserved matters” and minority agenda rights reintroduce institutionalized channels for shareholder influence.</p> <p>MES Gate checks:</p> <ul style="list-style-type: none"> • C1 (Legal certainty): Reserved matters defined by objective financial thresholds (observable in balance sheets, % of assets, % of dilution). • C2 (Independent arbiter): ICA ensures neutral adjudication of procedural disputes, replacing subjective board discretion. • C3 (Enforceability): Sanctions include nullification of decisions, fines, and director disqualification; abuse triggers procedural penalties. • C4 (Learnability): Review cycles every 24–36 months, KPI-based monitoring of reserved matters enforcement (<2% non-compliance target). In sum, while Delaware law maximizes efficiency by centralizing power, the LRM approach modifies this into a board-first but safeguarded regime, minimizing the Threat Index by balancing operational liberty of
---	---

	<p>a) a statement of the decision’s goals and objectives; b) a description of considered alternatives; c) an assessment of risks using the Threat Index methodology; d) identification of measurable key performance indicators (KPIs) for monitoring execution.</p> <p>Article 7. Procedural and Crisis Safeguards (1) Any abuse of rights by members of the board or shareholders in relation to matters regulated by this Act SHALL give rise to sanctions determined by the competent authority or court, including fines, disqualification, or suspension of rights. (2) In situations of urgent and imminent threat to the company’s viability, the board MAY take provisional action under a “crisis bridge” regime, provided that: a) the action receives prior authorization from the ICA; and b) the action is submitted for mandatory ratification by the shareholders within sixty (60) days. (3) In the absence of shareholder ratification, the provisional action SHALL lapse and any legal effects SHALL be reversed to the extent possible.</p>	<p>directors with the participatory liberty of shareholders.</p>
--	--	--

UK, Enterprise Act 2002, Part 3, s.58(2)(c)

<i>Original Law</i>	<i>LRM Version (Threat Index)</i>	<i>Explanation (TI & MES)</i>
---------------------	-----------------------------------	-----------------------------------

<p>“The need... to maintain a plurality of persons with control of media enterprises.” — a framework rule requiring that media plurality be taken into account when assessing mergers.</p>	<p>Article 1. Scope of Application This Regulation SHALL apply to all mergers, acquisitions, joint ventures, or comparable transactions involving undertakings engaged in the production, distribution, or control of editorial content within the national or regional media market.</p> <p>Article 2. Notification and Assessment Duty (1) Any transaction falling within the scope of Article 1 SHALL be notified to the competent review authority prior to its implementation. (2) Notification SHALL be mandatory where: a) the combined share of audience attention attributable to the undertakings concerned equals or exceeds twenty percent (20%); or b) the Plurality Concentration Index (PCI) would increase by one hundred and fifty (150) points or more as a result of the transaction.</p> <p>Article 3. Substantive Prohibition (1) A transaction notified under Article 2 SHALL NOT be approved if, upon completion: a) the national PCI would exceed two thousand two hundred (2200); or b) the regional PCI would exceed two thousand five hundred (2500); or c) fewer than four (4) independent editorial controllers would remain active in the relevant market. (2) Transactions falling under paragraph (1) SHALL be deemed incompatible with media plurality and SHALL be prohibited.</p> <p>Article 4. Conditional Approval (1) Where a transaction does not meet the absolute thresholds of Article 3 but raises material risks to plurality, the reviewing authority MAY grant approval subject to structural and behavioral remedies. (2) Such remedies MAY include:</p>	<p>The UK rule establishes a general principle, leaving broad discretion to the Secretary of State: plurality must be considered, but there is no algorithm for measuring it. The LRM/TI version converts that principle into <i>objectively verifiable triggers</i>: audience shares, PCI/HHI thresholds, minimum number of editorial controllers. TI analysis showed that the threat of losing freedom of expression (a higher-priority, less reversible right) outweighs restrictions on investment freedom. MES gates added enforceability (fines, unwind), learnability (annual KPI reports), and compensatory clauses for smaller players. Thus, the difference is between a broad political standard (UK) and an algorithmically enforceable norm (LRM).</p>
--	---	---

- a) divestiture of assets or business units sufficient to mitigate concentration;
- b) provision of fair, reasonable, and non-discriminatory (FRAND) access to essential distribution infrastructure;
- c) adoption of an Editorial Independence Charter binding on the undertakings concerned;
- d) commitments to conduct regular algorithmic audits of recommendation or distribution systems affecting editorial diversity.

Article 5. Independent Arbiter

- (1) An Independent Arbiter SHALL be appointed to conduct the substantive review of all transactions notified under this Regulation.
- (2) The Independent Arbiter SHALL act with full impartiality, SHALL be empowered to request all necessary information, and SHALL issue a binding determination within the time limits prescribed by implementing rules.

Article 6. Monitoring and Key Performance Indicators (KPIs)

- (1) Undertakings subject to conditional approval under Article 4 SHALL comply with ongoing monitoring obligations.
- (2) Monitoring SHALL include the periodic measurement of predefined KPIs, including but not limited to: audience distribution metrics, compliance with FRAND obligations, and independence of editorial decision-making.
- (3) KPI results SHALL be reported at least annually to the Independent Arbiter and made publicly available in a summarized form.

Article 7. Enforcement and Sanctions

- (1) Failure to notify a transaction as required under Article 2, or implementation of a transaction prohibited under Article 3, SHALL render the transaction void and subject the undertakings concerned to proportionate fines.
- (2) Breach of remedies imposed under Article 4 SHALL entitle the

	<p>reviewing authority to revoke approval, impose corrective measures, and apply financial sanctions.</p> <p>Article 8. Review and Sunset Clause This Regulation SHALL be subject to comprehensive review every five (5) years in light of market evolution, empirical data on PCI reliability, and the effectiveness of remedies in preserving plurality.</p>	
--	--	--

Spain, Ley 4/2020, art.1

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p><i>“Se crea el Impuesto sobre Determinados Servicios Digitales...”</i> — Creates a 3% tax on revenues from certain digital services (online advertising, online intermediation/marketplaces, user-data monetization). Thresholds: worldwide revenues \geq €750M and Spanish revenues \geq €3M. Applies irrespective of nationality of provider.</p>	<p>Art. 1. Scope — Groups providing digital intermediation, online advertising, or monetizing user data to users located in the jurisdiction.</p> <p>Art. 2. Trigger — Revenues \geq 750M (global) and \geq 25M (domestic). Jurisdictional sourcing defined by 2-of-3 test (billing address, verified device location, IP with payment provider match).</p> <p>Art. 3. Rate — Base rate 2%.</p> <p>Art. 4. Credit/Refund — Amounts paid fully creditable against corporate income tax within 24 months or refundable.</p> <p>Art. 5. Standstill Clause — If entity is subject to a multilateral allocation regime (OECD Pillar One or equivalent), DST automatically suspended for that entity.</p> <p>Art. 6. De-escalation — If explicit retaliatory tariffs are imposed, DST reduced to 0.5% for 12 months while negotiations proceed.</p> <p>Safeguards — Non-discrimination, de minimis exemption (<25M local revenue), transparency to consumers, independent arbitration</p>	<p>Threat Index (TI) logic:</p> <ul style="list-style-type: none"> • Spain’s 3% flat DST maximizes sovereignty but creates high risk of double taxation and trade retaliation (TI\uparrow for businesses and consumers through higher prices, TI\uparrow for state via retaliation). • LRM minimax rule selected a lower baseline (2%) with full CIT credit/refund to break the chain leading to cumulative harm (double taxation). • Adding standstill and de-escalation clauses directly cuts the escalation path (foreign reprisals), reducing severity.

	<p>of disputes, publication of aggregate KPIs, sunset/review every 24 months.</p>	<ul style="list-style-type: none"> • Safeguards (non-discrimination, transparency) extend chain length for lower-priority rights (convenience, property), ensuring MES compliance. <p>MES gates:</p> <ul style="list-style-type: none"> • C1 (certainty): objective thresholds and sourcing test. • C2 (arbiter): independent tax arbitration. • C3 (enforceability): penalties, transparency, refunds. • C4 (learnability): KPI monitoring, sunset review. Thus the LRM package preserves sovereignty but prevents irreversibility risks from trade escalation and systemic double taxation.
--	---	--

Germany, Mindestlohngesetz (MiLoG), §1(1)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
---------------------	--------------------	-----------------------------------

<p>“Jede Arbeitnehmerin und jeder Arbeitnehmer hat Anspruch auf Zahlung eines Arbeitsentgelts mindestens in Höhe des Mindestlohns.” (“Every employee shall be entitled to payment of remuneration of at least the minimum wage.”)</p>	<p>N1 — Who is covered Every employer who hires at least one worker must respect these rules. All wage payments must be at least equal to the guaranteed minimum.</p> <p>N2 — How the minimum is set Each year, the national statistics office must publish:</p> <ul style="list-style-type: none"> • the average cost of a basic living basket (food, housing, health, transport, communication), and • the median wage across the country. <p>N3 — The dignity minimum wage The minimum wage must be high enough to cover a decent living. It is calculated using both the cost of the living basket and the national median wage, with regional adjustments. Whichever amount is higher becomes the new minimum. Increases must be gradual so that they don’t grow too fast compared to the general wage level.</p> <p>N4 — Step-by-step introduction To give small businesses time to adapt, the new minimum wage comes into force in stages. Micro-businesses (up to 10 workers) and small businesses (up to 50 workers) can phase in the increase over a year. Regions with very high unemployment may get extra time.</p> <p>N5 — Help for small businesses If a small or micro-business’s wage costs rise too much because of the new rules, they can apply for support. This support can be in the form of tax breaks, reduced social contributions, or subsidies for training and productivity improvements.</p> <p>N6 — Apprentices and trainees Apprentices and trainees in recognized programs can be paid 85% of the minimum wage for up to 6 months. After that, they must</p>	<p>The German MiLoG norm is absolute and static: a uniform nationwide floor wage, with no phasing, no hardship clauses, no compensatory offsets. Its TI-profile shortens the chain to protecting workers’ life/health but leaves SMEs exposed to sudden property/liberty shocks (bankruptcy, layoffs). The LRM/TI version preserves the same <i>core right</i> (every worker SHALL receive not less than a minimum floor) but minimizes systemic risk by distributing adjustment costs: (i) phased implementation lowers L(a,s) for SMEs, (ii) offsets reduce probability P of closure, (iii) waivers provide controlled flexibility, (iv) KPI/review ensure learnability. Thus, while both regimes prioritize social justice, the LRM version explicitly balances irreversibility hierarchy (life/health > property) through compensatorics, meeting MES gates (C1 observability, C2 arbiters, C3 enforceability, C4 learnability).</p>
---	---	---

	<p>receive the full minimum wage. Companies cannot abuse this by repeatedly rehiring people as “trainees.”</p> <p>N7 — Temporary hardship exception If a company is close to bankruptcy, it may pay 90% of the minimum wage for up to 6 months. To do this, it must present a recovery plan and report monthly on progress. Companies that have paid dividends recently cannot use this exception.</p> <p>N8 — Transparency for workers All employers must clearly show the minimum wage on payslips and in the workplace so that employees know their rights. Workers who report violations are protected against retaliation.</p> <p>N9 — Monitoring and penalties Authorities will check compliance using data and complaints. First violations may result only in a warning and 30 days to correct. Repeat or serious violations will bring fines, and companies that constantly underpay may lose access to government aid.</p> <p>N10 — Review and adjustment The system will be reviewed every quarter and fully reassessed every three years. If unemployment rises sharply, or if workers still cannot afford basic living costs, the rules must be adjusted. Mistakes by the</p>	
--	--	--

France, Code du travail, art. L.1225-4)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p>“L’employeur ne peut rompre le contrat de travail d’une salariée lorsqu’elle est en état de</p>	<p>Art. 1 Scope — This regulation applies to all employees, regardless of contract type (permanent, temporary, part-time), once pregnancy is medically certified.</p>	<p>Threat Index logic: For pregnant workers, dismissal risk can lead to deprivation of income, loss of healthcare access, and direct threats to</p>

<p><i>grossesse médicalement constaté...”</i></p>	<p>Art. 2 Trigger (OAP) — The protection starts when the employee provides an official medical certificate of pregnancy and lasts until the end of maternity leave.</p> <p>Art. 3 Prohibition — During this protected period, the employer SHALL NOT terminate the employment contract.</p> <p>Art. 4 Exceptions — Termination MAY occur only in two objectively verifiable cases: (a) the employer is declared insolvent by a competent authority; (b) the employee has committed gross misconduct, confirmed by independent investigation and review.</p> <p>Art. 5 Safeguards — (i) Employer must seek prior approval from the labor inspectorate before dismissal; (ii) Automatic judicial review is available to the employee; (iii) Burden of proof lies entirely on the employer.</p> <p>Art. 6 Enforcement — If dismissal is found unlawful: (a) employee shall be reinstated with full back pay, or (b) alternatively, employee may choose compensation including damages for moral harm. Repeat violations by employers lead to escalating fines and ineligibility for public contracts.</p> <p>Art. 7 Compensatory clause — Employers facing proven financial hardship due to inability to dismiss may apply for targeted wage-support subsidies from the state during the protected period.</p> <p>Art. 8 Sunset/Review — The regulation is subject to a systematic review every 3 years, based on statistical indicators: number of dismissals attempted, maternal health outcomes, employer insolvency petitions. Adjustments must be proposed if imbalance is detected.</p>	<p>life/health of both mother and child → Severity 4. For employers, restriction affects property and managerial flexibility, but causal chain to life is remote → Severity 2. Therefore protection of worker dominates under minimax TI rule.</p> <p>MES gates: (C1) OAP satisfied by medical certificate, insolvency declaration, misconduct record. (C2) Arbiter satisfied by labor inspectorate/courts. (C3) Enforceability secured through sanctions and monitoring; CSP added via employer subsidies. (C4) Learnability via triennial review. Difference from French model: French law states an almost absolute prohibition, but exceptions are implicit and enforcement less structured. LRM makes exceptions explicit, adds compensatory subsidies for employers, stronger safeguards (burden of proof, inspectorate approval), and a built-in review mechanism. This makes the regulation more balanced, transparent, and adaptive without reducing protection for pregnant workers.</p>
---	--	--

Italy — Legge 40/2004, art. 4(3)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p data-bbox="107 224 531 365">“È vietato il ricorso a tecniche di procreazione medicalmente assistita di tipo eterologo...” (Legge 40/2004, art. 4(3))</p> <p data-bbox="107 443 531 841">Plain meaning: a total ban on heterologous assisted reproduction (use of donor sperm/eggs/embryos). Surrogacy was also heavily restricted/forbidden by the broader framework. Rationale at enactment: protect traditional family, avoid “split” genetic lineage, align with cultural/religious views.</p>	<p data-bbox="573 224 997 256">Article 1. Licensing and Registry</p> <p data-bbox="573 261 1451 365">(1) Any clinic or intermediary engaging in assisted reproductive technologies (ART), gamete donation, or surrogacy SHALL obtain a special license issued by the competent authority.</p> <p data-bbox="573 370 1409 440">(2) Each procedure SHALL be pre-registered in the national ART registry, with allocation of a unique cycle identifier.</p> <p data-bbox="573 444 1388 514">(3) Registry data SHALL be pseudonymized to protect personal information.</p> <p data-bbox="573 519 1360 623">(4) Enforcement: violations are subject to fines, suspension or annulment of licenses, and, in cases of systematic off-registry operations, criminal liability.</p> <p data-bbox="573 628 1457 698">(5) Compensatory: refusals of license SHALL be subject to fast-track appeal within ten (10) days.</p> <p data-bbox="573 732 993 764">Article 2. Access and Indications</p> <p data-bbox="573 769 1444 873">(1) Access to ART SHALL be granted to single persons and couples without discrimination on grounds of marital status or sexual orientation.</p> <p data-bbox="573 878 1402 982">(2) Access SHALL be conditional upon documented infertility or medically established genetic/health grounds, supported by a physician’s report and a psychosocial evaluation.</p> <p data-bbox="573 987 1444 1057">(3) Safeguards: contraindication checklist and the option of a second medical opinion.</p> <p data-bbox="573 1062 1457 1131">(4) Enforcement: systematic audits of refusals; sanctions imposed for discriminatory conduct.</p> <p data-bbox="573 1136 1423 1206">(5) Compensatory: unjustified denials SHALL entitle applicants to coverage of costs for a second opinion.</p> <p data-bbox="573 1240 888 1273">Article 3. Donor Criteria</p> <p data-bbox="573 1278 1367 1347">(1) Gamete donation SHALL be permitted only with informed consent of the donor.</p> <p data-bbox="573 1352 1457 1422">(2) Oocyte donors SHALL be between twenty-one (21) and thirty-two (32) years; sperm donors between twenty-one (21) and forty (40)</p>	<p data-bbox="1486 224 1986 1422">Direct normative conflict: Italy’s rule is a blanket prohibition; the LRM package permits with strict safeguards. They are mutually exclusive in operation. Threat Index (TI) analysis (why R* wins minimax): • Causal chains under a ban become longer ($\Sigma T_i \uparrow$) and less reliable ($\prod R_i \downarrow$): domestic prohibition \rightarrow cross-border travel \rightarrow unlicensed brokers \rightarrow low-oversight clinics \rightarrow limited aftercare. Each link adds time (T_i), information cost (2^K_i), and delay (τ_i), while reliability R_i drops; severity S can increase for donors/carriers (e.g., unmanaged OHSS, unmanaged obstetric risk). $TI(a,s)=S \cdot P \cdot L$ therefore rises for vulnerable actors (donors, gestational carriers, neonates). • Permission with controls shortens chains (fewer intermediaries), raises reliability (licensed providers, registries, SAE reporting), and reduces accumulation λ (feedback via KPI and N16). Hence the max TI across actors is lower (minimax). • Rights ranking by irreversibility: bodily integrity (donors/carriers/child) > liberty (reproductive autonomy) > property/convenience > cultural preferences. Bans mainly protect cultural preferences while increasing bodily risk via grey markets; R*</p>

	<p>years.</p> <p>(3) Maximum limits: no more than six (6) ovarian stimulation cycles and no more than ten (10) recorded live births per donor nationwide.</p> <p>(4) Safeguards: compulsory complication insurance and a national registry of genetic markers to prevent accidental consanguinity.</p> <p>(5) Enforcement: penalties against clinics or intermediaries exceeding donor limits.</p> <p>(6) Compensatory: donors may withdraw consent without penalty up to the point of puncture or sperm collection.</p> <p>Article 4. Gestational Carrier Criteria</p> <p>(1) A gestational carrier SHALL meet the following requirements:</p> <ol style="list-style-type: none"> age between twenty-one (21) and thirty-five (35) years; at least one (1) prior term birth; no more than two (2) surrogacies over a lifetime; health and body mass index within medically safe bounds. <p>(2) Safeguards: extended medical and disability insurance, and independent counseling.</p> <p>(3) Enforcement: penalties for recruitment outside the criteria; criminal liability for coercion.</p> <p>(4) Compensatory: the gestational carrier may unilaterally withdraw up to the moment of embryo transfer without sanctions.</p> <p>Article 5. Informed Consent and Cooling-Off Period</p> <p>(1) Prior to stimulation, collection, or transfer, informed consent SHALL be signed after an independent consultation.</p> <p>(2) A fourteen-day (14) cooling-off period SHALL apply before any invasive procedure.</p> <p>(3) Exception: urgent life-saving indications.</p> <p>(4) Safeguards: audio-visual record of the consultation; plain-language consent forms.</p> <p>(5) Enforcement: license sanctions for breaches of informed consent; invalidity of procedures conducted under coercion.</p>	<p>directly mitigates bodily risks while minimally restricting liberty.</p> <p>MES gates (abstract validity checks):</p> <ul style="list-style-type: none"> • C1 Legal certainty: R* uses objective applicability predicates (OAP) (license, registry entry, age/health cut-offs, SET rule, PGT indications). The Italian ban is certain in text but under-specifies operational safeguards (no way to manage inevitable demand). • C2 Independent arbiter: R* creates a bioethics body with auditability; a pure ban delegates nothing (no proportionality filter), which fails when edge cases arise (e.g., severe genetic disease prevention). • C3 Enforceability: R* builds licenses, registries, SAEs, sanctions, KPIs, escrow (conditions & stimuli). A ban has only negative sanctions; enforcement pushes activity offshore, reducing actual control. • C4 Learnability: R* mandates data collection + periodic review (N16) to update risks; bans lack built-in learning loops. <p>Why cultural/religious concerns are still respected: R* forbids commercialization of the child (N7), restricts selection to medical need (N10), and requires ethical approvals</p>
--	---	---

(6) Compensatory: withdrawal allowed until the biological point of no return, defined as in-utero conception or embryo transfer.

Article 6. Limits on Multiple Pregnancy

(1) Single-embryo transfer SHALL be mandatory for women under thirty-five (35) and for all gestational carriers.

(2) Double transfer MAY be permitted only where prognosis is documented as low and the medical board grants approval.

(3) Safeguards: written justification and patient consent.

(4) Enforcement: audits and penalties for violations.

Article 7. Payments and Ban on Commercialization of Children

(1) Payments in ART and surrogacy SHALL be limited to reimbursement of verified expenses and standardized compensations.

(2) Payments SHALL NOT be tied to the outcome, sex, or traits of the child.

(3) Safeguards: escrow mechanisms and a public schedule of reimbursable costs.

(4) Enforcement: confiscation of illegal payments and criminal liability for commercialization of children.

(5) Compensatory: restitution to good-faith parties involved.

Article 8. Contract Architecture and Parenthood Allocation

(1) A standard-form contract SHALL be filed before implantation.

(2) Intended parents SHALL be deemed the legal parents at birth.

(3) Medical decisions, including termination, remain with the gestational carrier when medically indicated and in accordance with protocol; in life-threatening situations, the carrier's decision prevails.

(4) Safeguards: ex ante mediation procedures for disputes concerning anomalies, multiple births, or preterm complications.

(5) Enforcement: specific performance of standard clauses; no sanctions for bona fide medical decisions of the carrier.

(6) Compensatory: costs of early termination of the contract due to uncontrollable causes SHALL be covered.

(N14). It also filters access by medical indications and safety, satisfying dignity concerns without imposing a total prohibition. Bottom line: Legge 40/2004 art. 4(3) prioritizes traditional morals via a total block on heterologous ART; LRM R* prioritizes bodily safety + reproductive autonomy through permission with heavy safeguards. On TI-minimax and MES grounds, R* dominates because it reduces high-severity harm and keeps risk observable and correctable instead of exporting it to unobservable grey markets.

Article 9. Protection of the Child’s Interests

- (1) Immediate parentage registration and neonatal screening SHALL be carried out upon birth.
- (2) Donor anonymity SHALL be protected; however, the child, upon attaining a defined age, SHALL have the right to access anonymized medical and genetic information.
- (3) Enforcement: penalties for delays.

Article 10. Prohibition of Non-Medical Selection

- (1) Non-medical selection of sex or traits SHALL be prohibited.
- (2) Preimplantation genetic testing (PGT) MAY be used solely to prevent serious monogenic or chromosomal diseases.
- (3) Exception: sex-linked severe disorders.
- (4) Enforcement: penalties, including annulment of licenses for violations.

Article 11. Cross-Border ART and Surrogacy

- (1) Cross-border ART and surrogacy SHALL be permitted only where the destination jurisdiction demonstrates equivalent safety and ethical standards.
- (2) The use of unlicensed intermediaries SHALL be prohibited.
- (3) Enforcement: refusal to register outcomes from non-equivalent jurisdictions; penalties for brokers.
- (4) Compensatory: assistance to relocate contracts to an equivalent jurisdiction.

Article 12. Non-Discrimination

- (1) Equal access to ART SHALL be guaranteed, subject only to medical criteria.
- (2) A log of refusals with stated reasons SHALL be maintained.
- (3) Enforcement: penalties, including license suspension for discriminatory practices.

Article 13. Serious Adverse Events Reporting

- (1) All serious adverse events (SAE) and suspected unexpected

serious adverse reactions (SUSAR) SHALL be notified within twenty-four (24) hours.

(2) A root-cause analysis SHALL be completed within thirty (30) days, accompanied by a corrective and preventive action (CAPA) plan.

(3) Enforcement: fines and conditional suspension pending CAPA implementation.

(4) Compensatory: treatment of affected parties SHALL be financed through insurance or a compensation fund.

Article 14. Oversight and Bioethics Body

(1) An oversight and bioethics body SHALL be established to decide matters delegated under Articles 3–11.

(2) The body SHALL publish anonymized statistics and guidance, manage conflicts of interest, and ensure periodic rotation of experts.

(3) Enforcement: external audit and sanctions for unjustified delays.

Article 15. Ban on Coercion and Exploitation

(1) Any coercion, deception, trafficking, or involvement of minors in ART SHALL be prohibited.

(2) Suspicion of such conduct SHALL be reported immediately to the competent authority.

(3) Enforcement: criminal liability and lifetime bans from the sector.

(4) Compensatory: victims SHALL have access to a protection program and to a dedicated compensation fund.

Article 16. Review and Rollback

(1) Thirty-six (36) months after entry into force, or earlier if the KPI threshold is breached, a mandatory public review SHALL be conducted.

(2) Risk scales SHALL be updated and norms adjusted, including possible rollback.

(3) Safeguards: independent audit of registry data and transparent methodology.

	(4) Enforcement: adoption of a mandatory implementation plan following the review.	
--	--	--

New York GBL §349

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p>“Deceptive acts or practices in the conduct of any business, trade or commerce... are hereby declared unlawful.” – Extremely broad and open-ended wording. – Leaves wide discretion to courts to determine what qualifies as “deceptive.”</p> <p>– Consumer can always attempt litigation, since the law provides a universal hook. – Strong protection of consumers, but generates uncertainty and litigation risk for businesses.</p>	<p>Article 1. Definitions and Risk Tiers (1) For the purposes of this Act, objective applicability predicates (OAP) SHALL be established to distinguish: a) “puffery,” defined as subjective and non-verifiable promotional language, which is excluded from sanction; b) “verifiable claims,” defined as statements concerning facts, performance, or measurable qualities of goods or services, which fall under the scope of this Act. (2) A list of prohibited “dark patterns,” including but not limited to forced continuity, disguised advertising, and coercive consent mechanisms, SHALL be adopted by regulation and regularly updated.</p> <p>Article 2. High-Risk Ban (1) Deceptive commercial practices in sectors directly affecting health, safety, or financial services SHALL be absolutely prohibited. (2) Any factual claims in such sectors SHALL require prior substantiation by adequate evidence, which SHALL be in place before the claim is published or disseminated.</p> <p>Article 3. Standardized Disclosures (1) All commercial communications involving prices, fees, or subscriptions SHALL include a standardized disclosure block, formatted for clarity and visibility.</p>	<p>TI perspective: – Broad “catch-all” clauses (like §349) maximize consumer protection but raise Risk* for businesses by making legal exposure unpredictable. – LRM minimax analysis shows that life and bodily integrity risks (health/safety deception) justify absolute bans (N2). Property-level risks (hidden fees, misleading prices) require structured disclosure (N3) rather than blanket prohibition. – By stratifying obligations according to severity and probability of harm, the LRM version minimizes the maximal Threat Index for both sides.</p> <p>MES gates: – <i>Certainty</i>: OAP ensure predictable application (unlike §349’s vagueness). – <i>Arbiter</i>: Courts/auditors only decide on factual predicates, not abstract “deception.” – <i>Enforceability</i>: KPI and logs provide measurable compliance.</p>

	<p>(2) Automatic subscription renewals SHALL be valid only if the consumer has provided explicit opt-in consent.</p> <p>Article 4. Substantiation Duty</p> <p>(1) All verifiable factual claims used in commercial communication SHALL be supported by written evidence.</p> <p>(2) Such substantiation SHALL be made available to the competent authority within five (5) days upon request.</p> <p>Article 5. Safe Harbor</p> <p>(1) Businesses that maintain compliance logs, documentary checklists, and records of disclosure practices in accordance with official templates SHALL benefit from a safe harbor.</p> <p>(2) Within this safe harbor, businesses SHALL be protected against frivolous or unfounded claims, provided that the records demonstrate good-faith compliance.</p> <p>Article 6. Duties of Online Marketplaces</p> <p>(1) Online marketplaces and platforms SHALL act on duly substantiated notifications of deceptive practices within forty-eight (48) hours.</p> <p>(2) Failure to act within the prescribed time SHALL render the platform jointly liable with the trader for resulting consumer harm.</p> <p>Article 7. Quick Remedies for Consumers</p> <p>(1) Consumers who suffer measurable harm from deceptive practices SHALL be entitled to expedited judicial or administrative remedies, including:</p> <ul style="list-style-type: none"> a) fast-track injunctions against the practice; b) immediate refund or cancellation rights. <p>(2) Such remedies SHALL be granted under simplified procedures designed to ensure effectiveness within short deadlines.</p> <p>Article 8. Sanctions and Monitoring</p> <p>(1) Violations of this Act SHALL be subject to graduated sanctions,</p>	<p>– <i>Learnability</i>: Periodic review of dark patterns and disclosure readability ensures adaptation.</p> <p>Result: LRM replaces the ex post “wide net + judicial discretion” model with an ex ante risk-calibrated framework. Consumers retain strong protection in high-risk domains, while businesses gain predictability and procedural safeguards in lower-risk areas.</p>
--	---	--

	<p>applied in proportion to severity and repetition:</p> <p>a) official warning;</p> <p>b) administrative fine;</p> <p>c) inclusion on a public list of non-compliant businesses;</p> <p>d) prohibition from engaging in specified commercial activities.</p> <p>(2) The competent authority SHALL establish monitoring mechanisms, including periodic audits and publication of anonymized enforcement statistics.</p>	
--	---	--

Bundesdatenschutzgesetz (BDSG), §26(1)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p><i>Personenbezogene Daten von Beschäftigten dürfen für Zwecke des Beschäftigungsverhältnisses verarbeitet werden, wenn dies für die Entscheidung über die Begründung eines Beschäftigungsverhältnisses oder nach Begründung für dessen Durchführung oder Beendigung erforderlich ist...</i></p>	<p>Art. 1. Scope — This Regulation applies to all employers, employees, and third-party processors engaged in the collection, storage, or analysis of worker personal data.</p> <p>Art. 2. Trigger (OAP) — Any act of processing employee personal data (including monitoring, storage, biometric collection, algorithmic scoring) SHALL fall within the scope of this Regulation.</p> <p>Art. 3. Permitted Categories — Processing SHALL be lawful only if: (a) it is strictly necessary for workplace safety and protection of life/health; or (b) it is indispensable for essential operational functions directly tied to contractual obligations.</p> <p>Art. 4. Prohibited Uses — Processing for productivity analytics, behavioral scoring, or disciplinary surveillance SHALL NOT be permitted unless the conditions of Art. 5 are satisfied.</p> <p>Art. 5. Exceptional Allowance for Productivity/Discipline — Processing MAY be allowed if, within the last 12 months, objectively verifiable triggers exist: (i) sector-adjusted serious incident rate above median linked to human error; (ii) recurrent</p>	<p>Threat Index reasoning:</p> <ul style="list-style-type: none"> – A single necessity test (BDSG) risks under-detering invasive monitoring because “erforderlich” is vague and often employer-weighted. Privacy harms (severity S=2–3) accumulate with high probability and low observability. – Conversely, categorical prohibition of all monitoring would raise workplace accident risk (S=4 harms) due to missing safety data, increasing TI for life and bodily integrity. – The LRM package minimizes the maximum TI by tiering: (i) always allow safety/health uses (to cut down catastrophic harms), (ii) allow operations only if indispensable, (iii) ban productivity/discipline unless empirical safety or quality triggers

	<p>quality failures verifiably tied to process breaches; or (iii) substantiated safety-related complaints. In such cases, employer SHALL conduct a Safety & Privacy Risk Impact Assessment (SPRIA), register the processing with the Data Trustee, and limit it to the implicated roles/functions.</p> <p>Art. 6. Data Minimization & Retention — Employers SHALL collect only data strictly necessary, SHALL prohibit covert audio, persistent location outside hazard zones, and SHALL enforce deletion clocks (30–90 days for routine, 12 months for disciplinary files).</p> <p>Art. 7. Worker Rights & Transparency — Employers SHALL provide workers with a console to view collected data, purposes, retention, and recipients. All algorithmic decisions affecting pay, promotion, or termination SHALL be subject to human review.</p> <p>Art. 8. Oversight & Enforcement — Each employer with ≥ 25 workers SHALL establish an elected Data Trustee; all employers engaging in disciplinary data use SHALL undergo annual independent audit. Regulators MAY impose fines up to 2% of payroll, suspend unlawful processing, and order reinstatement.</p> <p>Art. 9. Compensatory Clause — Any HR action based on unlawfully obtained data SHALL be void. Affected workers MAY claim statutory damages, reinstatement, and data erasure.</p> <p>Art. 10. Sunset & Review — This Regulation SHALL be reviewed every 2 years with sector-specific adjustments to retention periods, sanction levels, and objective triggers.</p>	<p>exist. – MES checks: C1 (Legal certainty) satisfied by objective predicates (incident rates, retention clocks). C2 (Independent arbiter) satisfied by trustee + audit. C3 (Enforceability) ensured by sanctions, suspensions, and voidability. C4 (Learnability) ensured by mandatory 2-year review. Result: More prescriptive, predictable, and balanced than BDSG’s open necessity clause, while still permitting proportional employer control where life/health risk justifies it.</p>
--	---	---

California Consumer Privacy Act (CCPA), Civ. Code §1798.120(a)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
---------------------	--------------------	-----------------------------------

<p><i>A consumer shall have the right, at any time, to direct a business that sells personal information... not to sell the consumer's personal information.</i></p>	<p>Art. 1 — Scope. This Regulation SHALL apply to all legal entities, natural or corporate, that process personal data of residents within the jurisdiction.</p> <p>Art. 2 — Universal Right to Opt-Out. (1) Every data subject SHALL have the right, at any time, to prohibit the sale or exchange of their personal data to third parties. (2) The exercise of this right SHALL be valid where: (a) the data subject explicitly selects the “Do Not Sell/Share” option provided by the controller, or (b) a recognized global opt-out signal is transmitted by the device or browser. (3) Upon such exercise, the controller SHALL NOT transfer personal data for monetary or equivalent consideration, except where required by law or necessary for security and fraud-prevention purposes.</p> <p>Art. 3 — Non-Discrimination. (1) No controller SHALL restrict access to essential services or degrade functionality solely because a data subject has exercised the right to opt-out. (2) Reasonable alternative compensation MAY be charged, provided such amount reflects a proportionate calculation of lost advertising revenue and has been verified by an independent auditor.</p> <p>Art. 4 — Protection of Minors. (1) The sale or exchange of personal data relating to individuals under the age of 18 SHALL be prohibited. (2) Controllers SHALL establish age-verification and child-segmentation safeguards.</p> <p>Art. 5 — Data Minimization and Purpose Limitation. (1) Personal data SHALL only be collected and processed for specific, explicit, and legitimate purposes. (2) Processing for incompatible purposes SHALL be prohibited.</p> <p>Art. 6 — De-Identification and Aggregation. Controllers MAY transfer data to third parties only in de-identified or aggregated form, provided that the risk of re-identification is below the threshold</p>	<p>Threat Index rationale: Under TI analysis, uncontrolled sale of personal data generates causal chains escalating to severe risks (identity theft, stalking, discrimination), which threaten bodily integrity and even life. These outweigh the reversibility of economic harm to businesses reliant on ad-based models. The minimax scenario required not just a single opt-out (as in CCPA), but a systemic bundle: universal opt-out (Art. 2), non-discrimination safeguards (Art. 3), child protection (Art. 4), audit and registry obligations (Arts. 7–8). This bundle minimizes the maximum TI across actors.</p> <p>MES gates: (C1) Legal certainty — triggers are objectively verifiable (signal received, registry entries, audit logs). (C2) Independent arbiter — annual audits and public reports provide external review. (C3) Enforceability — strong sanctions, clear thresholds, defined duties ensure compliance. (C4) Learnability — sunset and mandatory review cycles embed continuous adaptation.</p> <p>Difference from CCPA: CCPA offers a single right of consumer opt-out, leaving gaps in enforcement, child protection, de-identification, and anti-discrimination. The LRM package</p>
--	--	---

	<p>defined by regulation, and any contractual prohibition on re-identification is in place.</p> <p>Art. 7 — Transparency and Registry of Third Parties. (1) Controllers SHALL publish a machine-readable registry of all third parties receiving personal data, including the purposes of transfer. (2) Contracts with such third parties SHALL impose equivalent obligations under this Regulation.</p> <p>Art. 8 — Audit and Logging. (1) Controllers processing the data of 100,000 or more subjects per year SHALL maintain an immutable log of all opt-out signals received. (2) Such controllers SHALL be subject to annual independent audits verifying compliance.</p> <p>Art. 9 — User Interface Requirements. (1) All interfaces for data collection or consent SHALL provide a clear and conspicuous single control labeled “Do Not Sell/Share My Data.” (2) Dark patterns or deceptive design practices SHALL be prohibited.</p> <p>Art. 10 — Breach Notification. (1) Controllers SHALL notify both affected data subjects and the supervisory authority of any unauthorized transfer of data without undue delay. (2) Such notice SHALL specify the nature of the breach, third parties involved, and remedial measures undertaken.</p> <p>Art. 11 — Transitional Measures. Organizations deriving $\geq 40\%$ of annual revenue from behavioral advertising MAY request a transitional period of up to 12 months, subject to quarterly milestones and a migration plan toward contextual or subscription-based models.</p> <p>Art. 12 — Enforcement and Sanctions. (1) Controllers in violation of Arts. 2–10 SHALL be subject to administrative fines of up to 4% of annual global turnover, proportionate to the severity and recurrence</p>	<p>constructs a full regulatory architecture addressing those gaps, thereby lowering systemic risk and extending preservation of lower-ranked rights (property/business interests) while securing higher-ranked rights (privacy, bodily integrity, life).</p>
--	--	---

	<p>of the violation. (2) Repeat or intentional violations MAY result in a temporary suspension of data transfers. (3) Small and medium-sized enterprises MAY request phased compliance or installment-based penalty payments where financial solvency would otherwise be threatened.</p> <p>Art. 13 — Sunset and Review. This Regulation SHALL be subject to mandatory review every twelve months, with particular attention to the effectiveness of the global opt-out signal, proportionality of compensation under Art. 3, and risks of re-identification under Art. 6.</p>	
--	--	--

USA, Sarbanes–Oxley Act 2002, §302(a)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p>“The principal executive and financial officers shall certify... the financial report does not contain any untrue statement...” — First time U.S. law imposed personal criminal liability on top managers for accuracy of corporate reports.</p>	<p>Art. 1. Scope — This Regulation applies to all public-interest issuers whose securities are traded on regulated markets, and to their principal executive and financial officers.</p> <p>Art. 2. Certification Duty — For each annual and quarterly financial report, the principal executive and financial officers SHALL jointly certify that: (a) the report fairly presents, in all material respects, the financial condition and results of operations; (b) the issuer maintains effective internal controls over financial reporting (ICFR).</p> <p>Art. 3. Materiality Thresholds — A misstatement SHALL be deemed material if it: (i) alters net income or total assets by $\geq 5\%$; (ii) misclassifies a solvency or leverage ratio by $\geq 10\%$; (iii) falsifies cash/bank confirmations; or (iv) conceals related-party exposure $\geq 1\%$ of assets.</p>	<p>Threat Index (TI): SOX §302 places a high liberty burden: any error can quickly lead to custodial exposure, creating paralysis of decision-making. Under TI, liberty > property, so minimax requires preserving executives’ liberty except in high-harm, high-evidence cases.</p> <p>LRM adjustment: The package narrows liability through objective materiality triggers and safe harbors, lowering false-positive criminal exposure while still ensuring strong investor protection.</p> <p>MES checks:</p>

<p>Art. 4. Audit-Trail Obligation — Issuers SHALL maintain tamper-evident audit logs of journal entries, adjustments, override approvals, and audit communications for not less than seven years.</p> <p>Art. 5. Safe Harbors — Officers MAY avoid personal sanctions if they demonstrate: (a) documented reliance on an independent audit opinion; (b) documented dissent recorded in board minutes; or (c) voluntary self-report within 15 days and corrective filing within 45 days.</p> <p>Art. 6. Prohibition of Tampering and Retaliation — No officer, director, or employee SHALL alter audit evidence, impede auditor access, or retaliate against whistleblowers.</p> <p>Art. 7. Liability Tiers — (a) Negligent violations SHALL give rise only to civil penalties, disgorgement of bonuses, and temporary signing bans; (b) Reckless or intentional violations, coupled with investor loss proven by independent event-study or debt covenant breach, MAY trigger criminal liability, with custodial sentences capped at three years.</p> <p>Art. 8. Whistleblower Protections — Whistleblowers SHALL be protected against retaliation and MAY receive 10–20% of monetary sanctions collected.</p> <p>Art. 9. Restitution — Monetary sanctions and clawbacks SHALL be allocated pro rata to harmed investors through a restitution fund.</p> <p>Art. 10. Monitoring and Sunset — The regulator SHALL monitor the Serious Restatement Rate (SRR) and conduct a full review within 36 months; amendments SHALL be proposed if SRR does not decline by at least 30% without material increase in filing paralysis.</p>	<ul style="list-style-type: none"> – C1 (certainty): Thresholds (5%, 10%, 1%) and audit-trail logs give objective triggers. – C2 (arbiter): Independent regulator/court assesses liability and safe harbor claims. – C3 (enforceability): Sanctions (civil, criminal, clawbacks) and KPIs ensure compliance. – C4 (learnability): KPI-based 36-month review. <p>Balance: Investors’ property rights remain protected (through certification, audit-trails, restitution), while executives’ liberty is only constrained in demonstrably intentional or reckless fraud. This achieves lower Risk* (minimax) than the strict SOX approach.</p>
---	---

Dodd–Frank Act 2010, §922(a)

Original Law	LRM Version (maximally detailed legal norms)	Explanation (TI & MES)
<p><i>The Commission shall pay an award... to whistleblowers who voluntarily provide original information...</i></p> <p>First recognition of monetary rewards to whistleblowers (up to 30% of SEC sanctions). Core dilemma: protection of informants vs. risk of abusive claims and pressure on business.</p>	<p>N1. Scope — Applies to all regulated financial institutions, publicly listed companies, their affiliates, contractors, and employees.</p> <p>N2. Trigger (Objective Applicability Predicate) — A “Qualifying Disclosure” shall mean original, specific, and verifiable information concerning violations of financial or corporate regulation, substantiated by (i) documentary evidence with metadata; or (ii) at least two independent corroborating sources; or (iii) technical traceability of the event.</p> <p>N3. Monetary Award — The competent authority SHALL pay an award to the whistleblower in the range of 10–25% of actually collected sanctions, contingent upon final adjudication and after exhaustion of appeal deadlines. Payments SHALL be capped at a maximum amount C, and SHALL NOT be advanced.</p> <p>N4. Filter of Materiality — The authority SHALL recognize only disclosures exceeding threshold T (quantitative harm or systemic risk). Exceptions apply where life, bodily integrity, or imminent destruction of evidence is at stake.</p> <p>N5. Confidentiality and Anonymity — The identity of whistleblowers SHALL remain sealed in a protected registry; disclosure SHALL be permissible only by court order where indispensable. Breaches incur sanctions.</p> <p>N6. Anti-Retaliation — Any adverse employer action (dismissal, demotion, harassment, loss of benefits) within 12 months after disclosure SHALL be presumed retaliatory. The employer bears the burden of proof for exceptions. Remedies include reinstatement, double wage compensation, and damages.</p>	<p>TI analysis: Whistleblowers face the highest irreversibility risks — retaliation, loss of liberty, or threats to bodily integrity — compared to business entities facing property/income risks. A regime with no rewards leaves informants under-protected (high TI for life/liberty). Conversely, unlimited awards without filters create high TI for business (risk of abusive claims, false denunciations, corporate warfare). The calibrated reward regime with filters minimizes the maximum TI across actors.</p> <p>MES gates: • <i>C1 Legal certainty</i>: objective predicates (documentary evidence, thresholds T, metadata) ensure verifiable triggers. • <i>C2 Independent arbiter</i>: disputes allocated to a separate tribunal, satisfying impartiality. • <i>C3 Enforceability</i>: mechanisms for sanctions, preservation of evidence, and retaliation remedies make duties binding and measurable. • <i>C4 Learnability</i>: Sunset clause and KPI-based adjustments provide adaptability.</p> <p>Conclusion: The LRM package strikes balance — protecting whistleblowers through calibrated, legally certain</p>

	<p>N7. Internal Channel Rule with Carve-Outs — Whistleblowers SHALL first report through the internal compliance channel with a 30-day waiting period, except when: (i) top management is involved; (ii) evidence may be destroyed; (iii) there is imminent market/systemic risk.</p> <p>N8. Evidence Preservation Duty — Upon receipt of a disclosure, entities SHALL preserve all relevant data and SHALL NOT engage in spoliation. Cryptographic timestamps and audit trails SHALL be mandatory.</p> <p>N9. Sanctions for Malicious or Reckless Reports — Whistleblowers who knowingly or recklessly submit false information SHALL be subject to fines, disqualification from awards for 5 years, and compensation of the accused’s legal costs.</p> <p>N10. Independent Arbiter — Disputes over award eligibility or retaliation SHALL be adjudicated by an independent tribunal, separate from investigative arms.</p> <p>N11. Funding Source — Awards SHALL be paid exclusively from collected sanctions; no payments SHALL be made from general revenue.</p> <p>N12. Compliance Duty of Enterprises — Companies SHALL appoint an ethics officer, maintain secure channels, and file quarterly anonymized reports to the regulator. SMEs may operate on simplified biannual reporting.</p> <p>N13. Limited Immunity — Whistleblowers SHALL enjoy civil immunity for disclosure of confidential information to regulators, provided such disclosure is strictly necessary and proportionate.</p>	<p>mechanisms while containing risks of malicious abuse. The Dodd–Frank rule provided a broad “reward promise”; the LRM version operationalizes it into a dense regulatory architecture that reduces Threat Index for all parties.</p>
--	--	--

	N14. Sunset & Review — Thresholds T and C, as well as procedural safeguards, SHALL be reviewed every two years, with full legislative re-assessment after four years.	
--	---	--

USA (Florida), SB 7072 (2021), §106.072(2)

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p>A social media platform may not willfully deplatform a candidate for office...</p>	<p>N1. Scope and Applicability</p> <ol style="list-style-type: none"> 1. This Regulation shall apply to digital platforms which: <ol style="list-style-type: none"> (a) provide dissemination of user-generated content; and (b) record no fewer than one million (1,000,000) monthly active users within the jurisdiction. 2. The obligations under this Regulation shall extend exclusively to official electoral candidates and their duly accredited campaign committees, during the electoral period commencing upon official registration and terminating seven (7) days following the conclusion of voting. 3. Platforms below the threshold of subsection (1)(b) are exempt, unless they voluntarily opt in by formal declaration. <p>N2. Candidate Channel</p> <ol style="list-style-type: none"> 1. Upon verification of candidate status, the platform shall establish and maintain a dedicated “Candidate Channel” accessible to all users who elect to subscribe thereto. 2. Content lawfully published within the Candidate Channel shall not be subject to removal, suspension, or account termination save as expressly permitted under N3. 3. The platform shall ensure transparent labelling and record-keeping of all actions affecting such Channel. <p>N3. Exception for Immediate and Irreparable Harm</p> <ol style="list-style-type: none"> 1. The platform may, and where circumstances require shall, remove or restrict candidate content only where such content 	<p>TI analysis:</p> <ul style="list-style-type: none"> – Florida SB 7072 adopts an absolute ban on deplatforming → lowers TI for candidate’s right to free expression but raises TI for public safety (no emergency exception). – LRM approach introduces minimax: protects candidate speech (ban on arbitrary removal) while allowing narrow, empirically verifiable exceptions (imminent harm, violence). This keeps risk to life (priority 1) below threshold. <p>MES gates:</p> <ul style="list-style-type: none"> – C1: LRM defines objective applicability predicates (OAP) for removal (incitement, threats, doxxing). Florida law is vague (“deplatform”). – C2: LRM embeds independent arbiter (RAP). Florida law lacks ex-ante review. – C3: LRM enforces compliance with KPIs (delivery, logging, sanctions). Florida law contains no such instruments. – C4: LRM mandates periodic learning/revision. Florida law static.

- demonstrably constitutes:
- (a) a direct and specific incitement to imminent violence;
 - (b) intentional publication of personally identifiable information accompanied by incitement to harassment;
 - (c) fabricated calls to interfere physically with the electoral process; or
 - (d) credible threats of serious bodily harm.
2. Any removal or restriction pursuant to subsection (1) shall be:
- (a) limited in scope and duration to no more than twenty-four (24) hours absent confirmation by the Independent Review Panel under N7;
 - (b) accompanied by contemporaneous evidentiary preservation and logging; and
 - (c) immediately notified to the affected candidate.

N4. Political Misinformation Not Amounting to N3

- 1. Candidate statements concerning policy or factual matters, not falling within N3, shall not be removed.
- 2. The platform shall address such content exclusively through mitigatory measures, including:
 - (a) visible labelling;
 - (b) provision of authoritative references;
 - (c) exclusion from paid promotion; and
 - (d) algorithmic down-ranking outside the Candidate Channel.
- 3. A public register of all such labels and their grounds shall be maintained.

N5. Minimum Delivery Obligation

- 1. The platform shall deliver no less than eighty per cent (80%) of Candidate Channel publications to subscribers within twenty-four (24) hours of posting.

2. Reduction below such threshold shall constitute a violation unless caused by verified technical malfunction or lawful removal under N3.
3. Failure to comply shall trigger compensatory measures pursuant to N10.

N6. Stability of Electoral Rules

1. During the seventy-two (72) hours preceding the close of polls, the platform shall not introduce or enforce new moderation or advertising policies applicable to electoral content, save where necessary to address cybersecurity exigencies or obligations under N3.
2. Any permitted change shall be notified publicly and contemporaneously.

N7. Independent Rapid Review Panel (RAP)

1. Any candidate whose content is removed, restricted, or labelled under N3 or N4 shall have immediate right of appeal to the Independent Rapid Review Panel.
2. The Panel shall render a determination within twelve (12) hours of appeal.
3. The platform shall comply forthwith with the determination rendered.
4. All decisions shall be published with reasons and preserved for audit.

N8. Transparency and Record-Keeping

1. The platform shall maintain and publish, on a weekly basis, a machine-readable registry of all enforcement actions taken under N3–N5.
2. Each entry shall contain: (a) anonymised content identifiers; (b) the objective applicability predicate; (c) the type of measure; (d) the elapsed time between posting and enforcement; and (e) the outcome of any RAP proceeding.

3. Registers shall be authenticated by cryptographic means.

N9. Candidate Data Protection

1. Personal data collected for the purpose of Candidate Channel verification shall be strictly limited to what is necessary for such verification.
2. Such data shall not be processed for advertising or profiling purposes.
3. Storage shall not exceed the electoral period plus thirty (30) days.
4. Any breach shall entail immediate deletion and notification to the candidate.

N10. Compensation for Wrongful Restriction

1. Where the RAP determines that removal, restriction, or under-delivery was unjustified, the platform shall:
 - (a) reinstate the content;
 - (b) extend visibility by not less than twenty-four (24) hours; and
 - (c) provide equivalent compensatory reach (“make-good”) free of charge.
2. Such obligations shall be executed within twenty-four (24) hours of determination.

N11. Sanctions and Safe Harbour

1. Verified violations of N2, N3 (false positives exceeding five per cent (5%)), N5, or N8 shall attract progressive administrative fines proportionate to the platform’s domestic turnover.
2. Compliance with logging, transparency, and RAP referral obligations shall entitle the platform to safe-harbour protection against secondary liability.
3. Force majeure, duly certified, shall constitute a defence.

	<p>N12. Monitoring, Learning, and Sunset</p> <ol style="list-style-type: none"> 1. At the close of each electoral campaign or every calendar quarter, whichever is earlier, platforms shall publish a full Key Performance Indicator (KPI) report and submit to public hearing. 2. Provisions demonstrating persistent false-positive risk shall be subject to temporary suspension pending revision. 3. The entirety of this Regulation shall be subject to mandatory review after two (2) general electoral cycles. 	
--	--	--

Sweden, Alkohollag (2010:1622), Chapter 5, §1

<i>Original Law</i>	<i>LRM Version</i>	<i>Explanation (TI & MES)</i>
<p>Detaljhandel med spritdrycker... får bedrivas endast av Systembolaget AB.”</p> <p>Translation: “The retail trade of spirits may be conducted only by Systembolaget AB.” → Creates a full state monopoly on retail sale of spirits.</p>	<p>Article 1. Retail Sale of Strong Alcoholic Beverages (1) The retail sale of alcoholic beverages with an alcohol by volume (ABV) content greater than fifteen percent (15%) SHALL be conducted exclusively by the State Retail Agency. (2) No private entity SHALL be permitted to engage in the retail sale of such beverages.</p> <p>Article 2. Retail Sale of Low-Alcohol Beverages (1) Alcoholic beverages with an ABV content equal to or below fifteen percent ($\leq 15\%$) MAY be sold by private outlets holding a valid license. (2) Licenses under paragraph (1) SHALL be subject to quotas concerning: a) outlet density per defined geographic area; b) minimum distance requirements from schools, hospitals, and</p>	<p>- TI minimax: Both systems reduce risk from high-ABV spirits. Total monopoly lowers consumption but risks illegal black markets. Hybrid reduces that cumulative threat by allowing licensed low-ABV sales. Outcome: hybrid has lower <i>maximal threat index</i>.</p> <p>- MES gates:</p> <ul style="list-style-type: none"> • Legal certainty: clear ABV thresholds, outlet density rules. • Independent arbiter: separate licensing/appeals body, not monopoly itself. • Enforceability: explicit sanctions, KPIs, monitoring.

<p>comparable institutions; c) restricted hours of sale as determined by regulation.</p> <p>Article 3. Cross-Cutting Safeguards (1) Minimum unit pricing for alcoholic beverages SHALL be established to discourage harmful consumption. (2) Mandatory age verification SHALL be carried out at the point of sale for all alcoholic beverages. (3) Marketing practices targeting minors or youth audiences SHALL be prohibited. (4) Licensing procedures SHALL be transparent, with clear criteria and publication of decisions. (5) An independent appeal mechanism SHALL be available for all licensing and enforcement decisions.</p> <p>Article 4. Compensatory Clauses (1) Transitional licenses MAY be granted to existing operators for the duration of the implementation period defined in Article 6. (2) Small businesses affected by the transition SHALL be eligible for reorientation support programs. (3) Domestic producers SHALL be granted non-discriminatory access to procurement channels of the State Retail Agency.</p> <p>Article 5. Enforcement (1) Violations of Articles 1–3 SHALL result in administrative fines, suspension or withdrawal of licenses, and, where appropriate, criminal liability. (2) The competent authority SHALL ensure continuous monitoring and inspection of compliance.</p> <p>Article 6. Sunset and Periodic Review (1) The provisions of this Act SHALL enter into force with an eighteen-month (18) transition period, during which pilot regions SHALL be designated for implementation. (2) The Act SHALL be subject to comprehensive review every three</p>	<ul style="list-style-type: none"> • Learnability: sunset clauses and review cycle. - Key difference: Sweden → absolute exclusivity for public health, limiting entrepreneurship. LRM → partial monopoly + regulated competition, balancing public health with limited entrepreneurial rights and minimizing black-market incentives.
--	---

	(3) years, assessing in particular the impact on public health, including mortality reduction, and the effectiveness of safeguards. (3) Following review, adjustments to the regulatory framework SHALL be adopted as necessary.	
--	---	--

Juvenile Protection Act, Art.26(1) (old version)

Original Law	LRM Version	Explanation (TI & MES)
<p>청소년보호법 제26조 제1항 (구 규정): “심야시간(0시부터 오전 6시까지)에는 청소년은 인터넷게임을 이용할 수 없다.” (이른바 ‘셋다운제’, 2011-2021 시행)</p>	<p>N1. Providers: All online game providers SHALL block verified minor accounts from access between 00:00 and 06:00. Exceptions: educational or therapeutic programs. Safeguards: parents MAY request waivers for cultural/time-zone reasons. Enforcement: fines for providers, quarterly audits. Compensatory clause: minors MAY file complaints if unjustly denied. Sunset/review: expires after 5 years unless renewed, with mandatory epidemiological review.</p> <p>N2. Guardians: Guardians MAY apply for exemptions with medical/psychological justification. Misuse leads to suspension. Appeals before family court.</p>	<p>TI Analysis: Health of minors (risk of sleep disorders, addiction) outweighs liberty to play at night, since health harms are more irreversible. Liberty restrictions are reversible. Minimax favors health-protection scenario.</p> <p>MES Evaluation: LRM version adds safeguards and compensatory clauses absent in the Korean law, ensuring proportionality, observability (login records), enforceability (audits/fines), and learnability (sunset clause with health review). This keeps the protective effect while reducing overreach.</p>

China (2020, Art.59)

Original Law	LRM Version	Explanation (TI & MES)
<p>第五十九条: “网络游戏服务提供者不得在每日22时至次日8时向未成年人提供网络游戏服务，并应当在节</p>	<p>Article 1. Duties of Online Game Providers (1) All providers of online gaming services SHALL ensure that verified accounts belonging to minors are blocked from access between 00:00 and 06:00 local time.</p>	<p>TI: China’s law directly cuts causal chains from late-night gaming to health harm, but extends chains of surveillance harms (real-name ID, data centralization). LRM preserves the same</p>

<p>假日每日不超过3小时、非节假日每日不超过1.5小时向未成年人提供服务。”</p> <p>Translation: <i>Online game service providers shall not provide services to minors between 22:00 and 08:00, and shall restrict playtime to not more than 3 hours per day on holidays and 1.5 hours on school days.</i></p>	<p>(2) Exceptions: the prohibition in paragraph (1) SHALL not apply to access granted for educational or therapeutic programs explicitly designated as such.</p> <p>(3) Safeguards: parents or legal guardians MAY request waivers in cases of cultural or time-zone specific reasons, provided that such requests are duly documented and recorded.</p> <p>(4) Enforcement: violations of this Article SHALL result in administrative fines imposed upon providers, as well as quarterly compliance audits conducted by the competent authority.</p> <p>(5) Compensatory clause: minors who have been unjustly denied access under this Article MAY lodge a complaint with the supervisory authority, which SHALL investigate and order corrective measures where appropriate.</p> <p>(6) Sunset/review: this Article SHALL expire five (5) years from entry into force unless renewed. Renewal SHALL require a mandatory epidemiological review of the impact on youth health and development.</p> <p>Article 2. Rights and Duties of Guardians</p> <p>(1) Guardians MAY apply for exemptions to the restrictions set forth in Article 1, where justified by medical or psychological grounds supported by professional certification.</p> <p>(2) Any misuse of exemptions by guardians, including false or misleading applications, SHALL result in suspension of the exemption and possible sanctions under applicable law.</p> <p>(3) Appeals against decisions concerning exemptions or suspensions SHALL lie with the competent family court, which SHALL adjudicate under expedited procedure.</p>	<p>health protection but shortens privacy-risk chains through anonymous tokens. It also lengthens preservation of secondary rights: tiered age limits protect autonomy for older teens; guardian extension protects family autonomy; educational/esports exemptions protect beneficial use.</p> <p>MES: C1 met via observable tokens/time; C2 satisfied by independent arbiter (absent in China); C3 enforced through clear sanctions + incentives; C4 ensured with sunset clauses.</p> <p>Conclusion: China = rigid, centralized, surveillance-intensive. LRM = proportional, layered, privacy-respecting, balancing protection with autonomy and adaptability.</p>
---	--	--

Germany (Grundgesetz, Art.20a)

Original Law	LRM Version	Explanation (TI & MES)
<p><i>Der Staat schützt auch in Verantwortung für die künftigen Generationen die natürlichen Lebensgrundlagen und die Tiere im Rahmen der verfassungsmäßigen Ordnung durch die Gesetzgebung und nach Maßgabe von Gesetz und Recht durch die vollziehende Gewalt und die Rechtsprechung.</i></p>	<p>Art.1. On the territory of the state, a national ceiling of greenhouse gas emissions shall be established and reduced annually in accordance with five-year plans.</p> <p>Art.2. Construction and modernization of industrial facilities shall be allowed only with the application of best available technologies ensuring minimal emissions.</p> <p>Art.3. Economic activity in protected natural areas and habitats of rare species shall be prohibited, except in emergency situations directly related to the protection of life and health.</p> <p>Art.4. In regions where established environmental quality standards are exceeded, new projects shall be suspended, and existing facilities shall reduce their load until indicators return to the norm.</p> <p>Art.5. Enterprises that exceed critical pollution thresholds shall immediately suspend operations until the causes are remedied.</p> <p>Art.6. Enterprises with significant impact on the environment shall publish open and verifiable data on emissions.</p> <p>Art.7. Enterprises shall primarily reduce their own emissions; compensatory mechanisms may be used only in limited amounts.</p> <p>Art.8. Water-intensive industries shall introduce closed water cycles and comply with restrictions during drought or heat periods.</p> <p>Art.9. Legal entities may appeal regulatory decisions based on new reliable data, except for prohibitions regarding protected areas and critical thresholds.</p>	<p>Threat Index analysis: Environmental degradation has cumulative and irreversible chains leading to threats to life and bodily integrity, whereas economic activity relates to rights of property and entrepreneurship, which are of lower irreversibility. Therefore, priority is given to climate and intergenerational protection.</p> <p>MES gates: (C1) Legal certainty — objective triggers are set through measurable standards and annual caps; (C2) Independent arbiter — state supervision and appeal mechanisms; (C3) Enforceability — sanctions and monitoring instruments; (C4) Learnability — five-year review cycles and revision of critical values. The German norm is general and programmatic, entrusting protection to the state. The LRM version specifies concrete duties and prohibitions, operationalizing constitutional principle into enforceable legal norms.</p>

Art.10. State supervision shall be based on a risk-oriented approach, with sanctions ranging from warnings and fines to suspension of operations and withdrawal of permits.	
---	--

Methodology of Comparative RIA Assessment

The evaluation compared AI-generated norms and real-world laws for the same regulatory problems using the framework of Regulatory Impact Assessment (RIA). Each case was assessed along seven dimensions: problem definition, objectives, policy options, cost–benefit analysis, stakeholders, proportionality, and monitoring and review.

A scoring rubric from 0 to 5 was applied to each dimension, where 0 indicates absence of the element and 5 indicates full compliance with international best practices. The total score for each version of the norm was obtained as the sum of the seven individual scores: $RIA\ Score = score_1 + score_2 + score_3 + score_4 + score_5 + score_6 + score_7$ (so that $0 \leq RIA\ Score \leq 35$).

The comparative advantage of AI norms over existing laws was measured as the absolute difference ($\Delta AI = Score_{AI} - Score_{Real}$), the normalized difference ($\delta AI = \Delta AI / 35$), and the share of criteria where the AI norm outperformed the real law ($Share_{win} = (\text{Number of criteria where } Score_{AI} > Score_{Real}) / 7$).

When assigning scores, the following coding rules were applied. If a criterion was entirely absent it received 0. Weak, declarative mentions were coded as 1. Fragmented or incomplete structures received 2. A basic level of compliance, such as a clear formulation without operational metrics, was scored as 3. Well-developed but not perfect structures, such as clear thresholds without independent auditing, were scored as 4. Full compliance, including metrics, procedures, thresholds, KPIs, review mechanisms, or monetary CBA where applicable, received 5.

To ensure internal consistency, some criteria were dependent on others. In particular, proportionality was always bounded by the availability of policy options. If no options were provided, proportionality could not exceed a low score (0–2), because it is impossible to demonstrate that the chosen measure was the least restrictive among alternatives. Conversely, if multiple options were presented and targeted measures or exceptions were included, proportionality could receive a high score (4–5). Objectives were also cross-checked against monitoring: when no KPIs or review mechanisms were present, the objective score was capped at a mid-level value.

Ambiguous or implicit references were scored conservatively, usually 1–2. Sensitivity checks were applied by adjusting borderline scores by ± 1 ; in all such cases the ranking between AI-generated and real laws remained unchanged.

Finally, an important methodological limitation was recognized. AI-generated norms tended to include multiple procedural layers (such as phased implementation, escrow mechanisms, automatic suspension triggers, KPIs, or tiered stakeholder protections). These features increased their scores in the categories of options, proportionality, and monitoring. However, such complexity simultaneously reduces practical applicability and could result in administrative overload if directly transplanted into real-world legislation. For this reason, while AI norms consistently outperformed existing laws on formal RIA criteria, their implementation would require simplification and institutional adaptation.

Case	RIA Criterion	AI-Norm Version (score)	Real-World Law (score)
Case 1: UK Company Director Duty (s.172)	1. Problem Definition	Focuses on preventing harm to employees' life/health and the environment from profit-driven decisions. (4)	Focus is on promoting company success for shareholders; employees and environment only "have regard to" as factors. (2)
	2. Objectives	Prioritizes preserving life and health above profit, giving binding force to worker/environment interests. (4)	Aims to maximize shareholder benefit ("success of the company") while merely considering other interests. (2)
	3. Policy Options	Detailed multi-article reform (strict no-harm rules, phased obligations) rather than one-size rule. (2)	Single statutory rule with no alternative measures or safeguards considered. (1)
	4. Cost-Benefit Analysis	Not addressed; TI logic emphasizes rights protection (implicitly minimizing risk) rather than economic trade-offs. (1)	No economic CBA in the law text; focus is normative, not on measurable costs/benefits. (0)
	5. Stakeholder Impact	Explicitly includes employees, public and environment as mandatory priorities (even earmarking SME/consumer relief). (4)	Shareholders' interests dominate; employees and community mentioned but without enforceable protection. (1)
	6. Proportionality	Strictly prohibits profit-driven actions that risk life/health, with narrow exceptions (minimal risk only). (4)	Very broad "good faith" standard yields wide discretion; may underprotect critical rights (weak proportionality). (2)
	7. Monitoring/Review	Requires transparency/disclosure and mandatory review every 3 years, embedding feedback loops. (5)	No review clause or reporting requirement; enforcement relies on courts only. (0)
	Total RIA Score	24	9
Case 2: USA DGCL §203 Moratorium	1. Problem Definition	Identifies risk of coerced takeovers; TI adds safe-harbor to preserve minority influence. (4)	Aims to block hostile takeovers by forbidding certain transactions (interested stockholder), protecting shareholders. (3)
	2. Objectives	Balances investor flexibility with minority protection via objective criteria and arbiters. (4)	Ensures stability and minority protection by high thresholds for consent. (2)
	3. Policy Options	Introduces multiple exceptions (independent committee, shareholder vote, crisis use) rather than a total ban. (4)	Implements a rigid 3-year moratorium with no alternative paths. (1)
	4. Cost-Benefit Analysis	None explicit (design favors risk minimization). (1)	No CBA in statute. (0)

	5. Stakeholder Impact	Protects minority shareholders (via disinterested votes, independent arbiter) and still allows investor action under conditions. (4)	Focuses on investors and existing shareholders; no provision for additional stakeholders. (1)
	6. Proportionality	Flexible “safe harbor” minimizes harm (e.g. crisis exception) while maintaining protection. (4)	Absolute ban may overly harm investors’ interests with little incremental benefit. (2)
	7. Monitoring/Review	Built-in review every 2 years of thresholds and effectiveness. (5)	No statutory review or sunset. (0)
	Total RIA Score	26	9
Case 3: Germany AktG §93 Duty of Care	1. Problem Definition	Clarifies vague “due care” by tying liability to objective pre-decision checks. (4)	General duty of diligence; courts decide after the fact if care was adequate. (1)
	2. Objectives	Protect life/health above financial loss (red-line ban on high-risk projects). (4)	Ensure directors act diligently in managing the company. (2)
	3. Policy Options	Adds multiple norms: safe-harbor checklist (N1), absolute prohibition for life-risk (N2), mandatory audits (N3). (3)	Single open-ended standard (“ordentlichen und gewissenhaften Geschäftsleiters”). (1)
	4. Cost-Benefit Analysis	Not considered explicitly. (0)	Not addressed. (0)
	5. Stakeholder Impact	Emphasizes employees’ and public life/health (over asset risk) and requires transparency controls; shareholders’ risk capped. (4)	Focus is corporate (shareholders) and judicial oversight; no explicit mention of employees. (1)
	6. Proportionality	Allows innovation within strict safety limits; mitigates risks via thresholds and pilots. (4)	One-size duty with potential retroactive penalty; less calibrated (lower proportionality). (2)
	7. Monitoring/Review	Mandates independent audits and accurate data (objective controls) for accountability. (4)	Enforcement only through ex-post court judgment; no ex-ante review. (0)
	Total RIA Score	23	7
Case 4: France Digital Tax (2019-759)	1. Problem Definition	Sees 3% flat tax as trigger for trade retaliation; TI opts to “lengthen the causal chain” to reduce risk. (4)	Target is revenue from big tech (digital services). (2)

	2. Objectives	Raise revenue while avoiding economic blowback; protect SMEs (earmarks 10% for relief). (4)	Generate tax income from large digital firms. (2)
	3. Policy Options	Implements a lower rate (2%), phased rollout, tax credits, suspension clauses. (4)	One flat 3% rate without mitigation measures. (1)
	4. Cost-Benefit Analysis	No formal CBA (focus on minimax fairness). (1)	Not provided. (0)
	5. Stakeholder Impact	Addresses platforms, State revenue, SMEs, and consumers (via subsidies). (4)	Primarily burdens tech companies; SMEs and consumers not specifically considered. (1)
	6. Proportionality	Gradual implementation and built-in suspension make it less blunt. (4)	Permanent high rate likely excessive given trade risks. (1)
	7. Monitoring/Review	Embeds KPIs and semi-annual review for retaliation incidents. (5)	None (no review or adjustment mechanism). (0)
	Total RIA Score	26	7
Case 5: Mexico VAT (Ley IVA Art.1)	1. Problem Definition	Universal VAT can push small actors into informality, threatening their rights. (4)	Universal rule: everyone doing business must pay VAT. (2)
	2. Objectives	Protect small businesses (limit enforcement) while preserving revenue through calibrated relief. (4)	Maximize tax base coverage (broad obligation). (2)
	3. Policy Options	Retains universality but adds thresholds, phased rates, credits, escrow (risk-calibrated model). (3)	Rigid flat system with no thresholds or exemptions in law. (1)
	4. Cost-Benefit Analysis	Not explicit (the design implicitly reduces non-compliance risk). (1)	Absent. (0)
	5. Stakeholder Impact	Considers SMEs and low-income consumers (earmarking relief); business certainty via objective rules. (4)	Treats all taxpayers equally; no special treatment for small businesses. (1)
	6. Proportionality	Phased introduction and exemptions lower the burden on vulnerable actors. (4)	Equally heavy on all, risking disproportionate impact on small firms. (1)
	7. Monitoring/Review	Suspension triggers (trade retaliation), KPIs, mandatory review after 24 months. (5)	None specified. (0)
	Total RIA Score	24	7

Case 6: Dutch Euthanasia Law (Art.2)	1. Problem Definition	Identifies risk of coercion/impulsivity under subjective criteria; TI adds objective safeguards. (3)	Permits euthanasia if physician believes request is voluntary and suffering unbearable. (2)
	2. Objectives	Ensure only well-considered requests proceed, via multiple declarations and waiting periods. (4)	Respect patient autonomy in end-of-life without quantified guidance. (2)
	3. Policy Options	Multiple witness declarations, medical assessments and an appeals mechanism to reduce false positives. (4)	Single-physician conviction rule (no detailed alternatives). (1)
	4. Cost-Benefit Analysis	Not present; focuses on minimizing risk of irreversible errors. (1)	Not relevant. (0)
	5. Stakeholder Impact	Emphasizes patient rights with legal proof; includes independent experts (C2). (3)	Only physician and patient explicitly involved; lower certainty. (2)
	6. Proportionality	Adds safeguards (cooling-off, audits) to avoid overreach, striking balance. (4)	Qualitative “good faith” standard may be too lenient or strict case-by-case. (2)
	7. Monitoring/Review	Mandates documentation and periodic audits (KPI “non-compliance rate”). (4)	No explicit review clause; relies on professional oversight. (0)
	Total RIA Score	23	9
Case 7: CA AB5 (Gig Worker Law)	1. Problem Definition	Tackles worst risks (accidents, fatigue, income shocks) by layered rules (insurance, hours). (4)	Treats lack of control as abuse; all workers become employees unless entity proves independence. (2)
	2. Objectives	Build safety net (insurance, fatigue limits, benefits wallet) for gig workers via TI analysis. (4)	Maximize labor protections (employee status) at expense of flexibility. (3)
	3. Policy Options	Multi-tier approach (safety floor N1–N3, conditional reclassification N4, due process N5). (4)	Single “ABC” test; binary classification with no gradation. (1)
	4. Cost-Benefit Analysis	Not explicit; implicitly balances worker security vs company costs. (1)	No formal CBA in statute. (0)
	5. Stakeholder Impact	Explicitly addresses drivers, platforms, passengers; portable benefits for drivers. (4)	Primarily focuses on drivers vs companies dichotomy. (1)
	6. Proportionality	Adaptive approach (pilot programs, sunsets) aims to avoid blanket disruption. (4)	Rigid default rule can overshoot (high burden on small platforms). (2)

	7. Monitoring/Review	Includes neutral data aggregator, appeals, required review (pilot and KPIs). (5)	No review/sunset in original; enforcement by fines. (0)
	Total RIA Score	26	9
Case 8: Russia “Foreign Agent” Law	1. Problem Definition	Replaces vague criteria with objective thresholds, avoiding stigma. (4)	Broad language on any “foreign influence/support” risks arbitrary designation. (1)
	2. Objectives	Ensure transparency without “agent” label; preserve freedom of association. (4)	Enforce labeling of NGOs receiving foreign aid. (1)
	3. Policy Options	Numeric triggers (e.g. >10% budget) and disclosure duties; no blanket ban. (4)	Label applies to any NGO with unspecified “support/influence.” (1)
	4. Cost-Benefit Analysis	Not discussed; emphasis on minimizing censorship risk. (1)	Not addressed. (0)
	5. Stakeholder Impact	Preserves NGOs’ operational rights (no disloyalty label) and limits sanctions to fines/proportional terms. (4)	Harsh social/administrative burdens on NGOs; funds, materials. (1)
	6. Proportionality	Keeps transparency while minimizing rights intrusion (no pejorative label). (4)	Highly restrictive and stigmatizing; no built-in safeguards. (1)
	7. Monitoring/Review	Independent oversight with appeals; mandated review every 5 years. (5)	No review period or appeal guaranteed in original. (0)
	Total RIA Score	26	5
Case 9: France Face-Covering Ban	1. Problem Definition	Treats face concealment as low-security risk; allows coverings by default. (4)	Sees public safety threat from any face covering in public. (1)
	2. Objectives	Preserve religious expression; enable targeted ID checks only in genuine high-risk areas. (4)	Maintain public order by forbidding coverings. (1)
	3. Policy Options	Permits ID checks in defined zones or emergencies with privacy safeguards. (4)	One-size ban on all face coverings; no exceptions in law. (1)
	4. Cost-Benefit Analysis	Not analyzed. (0)	Not discussed. (0)
	5. Stakeholder Impact	Maximizes individual liberty; respects faith choices with minimal intrusion. (4)	Restricts religious minorities’ practices fully. (1)
	6. Proportionality	Minimax approach allows face coverings generally; uses narrow exceptions (balancing safety vs freedom). (5)	Absolute prohibition is blunt and overreaching. (1)

	7. Monitoring/Review	Includes privacy rules (no data retention) and internal safeguards (e.g. transparency obligations). (4)	No enforcement metrics or oversight specified. (0)
	Total RIA Score	25	5
Case 10: Texas Age-Verification Law	1. Problem Definition	Acknowledges child safety risk but prioritizes privacy (limits data retention). (4)	Seeks to keep minors from online porn; details are very broad. (1)
	2. Objectives	Block minors' access with minimally intrusive methods (no biometric/long-term storage). (4)	Require "reasonable" age verification by platforms. (1)
	3. Policy Options	Specifies privacy-preserving tech (e.g. third-party tokens) and short data holding. (4)	Statutory mandate without clear methods or safeguards. (1)
	4. Cost-Benefit Analysis	Not provided. (0)	None. (0)
	5. Stakeholder Impact	Protects minors and privacy; regulates platforms' methods; compensates data breaches. (3)	Targets online platforms; minors protected by enforcement. (1)
	6. Proportionality	Strikes balance: strict age checks but forbids permanent data storage. (5)	Rigid requirement may unduly invade user privacy. (1)
	7. Monitoring/Review	Enforcement via audits/fines; mandatory review every 3 years. (4)	No defined review or audit process. (0)
	Total RIA Score	24	5
Case 11: Delaware DGCL §141 (Corp. Governance)	1. Problem Definition	Recognizes high risk of shareholder exclusion under pure board rule. (4)	Board-control model (138-year-old law) gives no formal channels for minority input. (2)
	2. Objectives	Preserve board efficiency while reintroducing shareholder checks (reserved votes, arbiter). (4)	Centralize management in directors to allow swift decisions. (2)
	3. Policy Options	Creates "reserved matters" requiring shareholder approval and independent arbiter for disputes. (4)	No alternatives; all corporate affairs under board. (1)
	4. Cost-Benefit Analysis	Not addressed (legal safeguards chosen via TI logic). (0)	Not present. (0)
	5. Stakeholder Impact	Explicitly protects minority voice (agenda rights ≥1% shares) and oversight via Independent Arbiter. (4)	Essentially none for shareholders (majority can override). (1)

	6. Proportionality	Maintains board primacy but curbs excess with sanctions and crisis procedures; designed to minimize TI. (5)	Overly favors directors; shareholders get no guaranteed influence. (1)
	7. Monitoring/Review	Formal review (24–36 months KPI targets) for reserved matter enforcement and company actions. (4)	No statutory review or enforcement beyond usual corporate law remedies. (0)
	Total RIA Score	25	7
Case 12: UK Media Plurality (Enterprise Act 2002)	1. Problem Definition	Converts broad plurality goal into measurable thresholds (PCI, HHI, minimum publishers). (4)	Cites media plurality as a consideration in mergers; no metrics given. (1)
	2. Objectives	Ensure independent editors remain (prevent >4 controlling voices); algorithmic enforcement. (5)	Maintain a diversity principle; left vague for ministerial discretion. (1)
	3. Policy Options	Sets clear merger notification thresholds, prohibits if PCI too high, or requires remedies. (4)	Case-by-case review by Secretary of State; remedies unspecified. (1)
	4. Cost-Benefit Analysis	Not explicitly done. (0)	No analysis. (0)
	5. Stakeholder Impact	Protects public (speech diversity) by empowering an Independent Arbiter; limits media owners. (4)	Focuses on media companies with informal influence of regulators. (1)
	6. Proportionality	Adds enforceable sanctions and KPI audits (learnability); carefully weighs expression over investment. (5)	Broad “plurality” standard risks either under- or over-enforcement. (1)
	7. Monitoring/Review	Annual KPI reporting, public disclosure, and 5-year comprehensive review of thresholds and effects. (5)	No statutory timeline for review or data collection. (0)
	Total RIA Score	27	5
Case 13: Spain Digital Tax (Ley 4/2020)	1. Problem Definition	Identifies flat DST creates high double-taxation and retaliation risk. (4)	Aims to tax big tech (3% rate, €750M/€3M thresholds). (2)
	2. Objectives	Raise revenue while minimizing harm (full CIT credit; multilateral standstill). (4)	Capture taxes from large digital service providers. (2)

	3. Policy Options	Lowers base rate to 2%, allows full credit/refund; imposes standstill/de-escalation clauses. (5)	Fixed 3% without rollback or exemption mechanisms. (2)
	4. Cost-Benefit Analysis	Not conducted explicitly; approach reduces TI by cutting escalation paths. (1)	None. (0)
	5. Stakeholder Impact	Preserves state sovereignty with objective rules; protects business via credit and limited carve-outs. (4)	Affects international firms equally; only nominal thresholds for Spanish revenue (excludes many SMEs). (2)
	6. Proportionality	Compares highest-priority rights (security/trade) and embeds de-escalation to lower maximal TI. (5)	Uniform tax may impose disproportionate costs, risking international disputes. (2)
	7. Monitoring/Review	Uses KPI monitoring, independent arbitration, refunds, and sunsets (2-year review cycle). (5)	No sunset; lacks structured feedback mechanisms. (1)
	Total RIA Score	28	11
<i>Case 14: Germany Minimum Wage (MiLoG §1)</i>	1. Problem Definition	Recognizes rigid nationwide floor risk to small firms; TI adds adjustability. (4)	Defines uniform minimum wage for all employees. (2)
	2. Objectives	Guarantee a living wage while ensuring gradual adjustment (link to costs, median, regional, subsidies). (4)	Ensure decent pay (living costs) for all workers. (3)
	3. Policy Options	Annual living-cost vs median wages; phased introduction for small firms; exceptions (hardship). (5)	Single nationwide minimum floor, no phasing or exemptions. (2)
	4. Cost-Benefit Analysis	Implicitly balances worker needs vs unemployment (unemployment triggers review). (3)	Not addressed. (0)
	5. Stakeholder Impact	Explicitly helps low-income workers and SMEs (support programs, hardship waivers, transparency). (4)	Prioritizes all workers; small employers must adapt suddenly. (2)
	6. Proportionality	Phasing, subsidies and exceptions reduce maximum impact on jobs and businesses. (5)	Rigid for all businesses; could cause abrupt layoffs under fixed law. (1)

	7. Monitoring/Review	Quarterly data checks, corrective steps; quarterly quick reviews and full review every 3 years. (4)	No formal review in law (only enforcement via fines). (0)
	Total RIA Score	29	10
Case 15: French Pregnancy Protection (L.1225-4)	1. Problem Definition	Prioritizes mother/child health; TI raises severity of dismissals (life risk 4). (4)	Protect pregnant women from job loss; broad constitutional protection. (2)
	2. Objectives	Ban termination during pregnancy except verifiable cases (insolvency, misconduct). (4)	Prevent pregnant employees from being fired. (3)
	3. Policy Options	Clearly enumerates two exceptions, adds prior approval and judicial review. (4)	Implicit ban; exceptions not clearly spelled out. (1)
	4. Cost-Benefit Analysis	Not quantified. (0)	Not present. (0)
	5. Stakeholder Impact	Ensures employee right to remain with support; allows employers relief (subsidies) if needed. (4)	Protects workers' rights broadly; may impose uncertainty on employers. (3)
	6. Proportionality	Explicit safeguards (burden on employer, subsidies) maintain balance without diluting protection. (5)	Near-absolute ban is one-sided, although aimed at a high-priority right. (2)
	7. Monitoring/Review	Labor inspectorate approval, automatic court oversight; review every 3 years. (4)	No scheduled review or specified enforcement mechanism. (0)
	Total RIA Score	25	11
Case 16: Italy ART & Surrogacy (Legge 40/04)	1. Problem Definition	Notes that absolute ban causes unsafe cross-border practices. (4)	Totally forbids heterologous assisted reproduction by statute. (1)
	2. Objectives	Allow ART/surrogacy with strong safeguards (license, registry, consent) to reduce bodily harm. (4)	Protect "traditional family" morals; no recognition of donor methods. (1)
	3. Policy Options	Authorizes donation under strict conditions (age limits, insurance, limits). (5)	Only option: ban all heterologous ART. (1)
	4. Cost-Benefit Analysis	Compares serious health risks vs cultural concerns; no explicit trade-off. (1)	None. (0)

	5. Stakeholder Impact	Addresses patients' right to conceive, donors, gestational carriers and child welfare extensively. (4)	Serves cultural/religious interests at expense of reproductive autonomy. (1)
	6. Proportionality	Permits non-commercial ART within safeguards; reduces high-severity risks (minimax). (5)	Blunt total ban disproportionately endangers bodily integrity (no safe legal option). (1)
	7. Monitoring/Review	Strict licensing/registry with sanctions; mandatory review/rollback after 36 months. (5)	No oversight or review (activity driven underground). (0)
	Total RIA Score	28	5
<i>Case 17: Sweden Alcohol Monopoly</i>	1. Problem Definition	Recognizes total monopoly drives black markets; TI hybrid lowers maximum harm. (4)	State monopoly on strong spirits to protect health. (3)
	2. Objectives	Reduce high-ABV consumption while allowing controlled low-ABV sales. (4)	Limit alcohol harm by banning private spirit sales. (3)
	3. Policy Options	Permits licensed private sales of ≤15% alcohol (with density quotas). (4)	No private retail outlets for spirits; only the state agency sells them. (1)
	4. Cost-Benefit Analysis	Balances public health vs entrepreneurship (uses unit pricing, age checks). (2)	Not explicitly performed. (1)
	5. Stakeholder Impact	Protects public health and small vendors; compensatory clauses aid small businesses. (4)	Prioritizes population health; gives no role to private retailers. (2)
	6. Proportionality	Partial monopoly + regulated competition minimizes black-market incentive. (5)	Absolute monopoly is heavy-handed for public health goals. (2)
	7. Monitoring/Review	Licensing oversight, pilot regions; sunset review every 3 years. (4)	Legacy system; no automatic review or adaptability. (0)
	Total RIA Score	27	12
<i>Case 18: Korea Gaming Shutdown Law</i>	1. Problem Definition	Acknowledges harm of late-night gaming, introduces limited exemptions to avoid undue restriction. (4)	Bans minors from online games midnight–6am (“shutdown”). (3)
	2. Objectives	Protect minors' health (sleep, addiction) but allow parental waivers for special cases. (4)	Strictly stop under-18 gaming at night. (2)
	3. Policy Options	Blocks verified minors 00:00–06:00 with waiver process; appeals to family court. (4)	Total nighttime block; no formal exemptions. (1)

	4. Cost-Benefit Analysis	Implicit TI logic (health > liberty); no formal CBA. (1)	None. (0)
	5. Stakeholder Impact	Protects children (with rights to appeal), involves guardians; providers fined. (3)	Protects minors; imposes strict tech enforcement on providers. (1)
	6. Proportionality	Includes sunset/review (5-year health check) to prevent overreach. (5)	Blanket rule is inflexible, no mechanism to relax burdens. (1)
	7. Monitoring/Review	Mandatory epidemiological review at 5 years; penalties for violations. (5)	Policy expired (2011–2021); no mid-course corrections. (0)
	Total RIA Score	25	8
Case 19: China Juvenile Gaming (Art.59, 2020)	1. Problem Definition	Seeks same health protection (22:00–08:00 block, playtime limits) but avoids real-ID privacy risk via tokens. (4)	Strict time restrictions (22:00–08:00 block, 1.5hr limit) with real-name checks. (3)
	2. Objectives	Prevent addiction/health harm to minors with minimal surveillance. (4)	Prevent excessive gaming through enforced caps. (3)
	3. Policy Options	Blocks minors 00:00–06:00; permits educational access; no centralized ID. (3)	Blocks 22:00–08:00; long holiday playtime banned (>3h). (2)
	4. Cost-Benefit Analysis	Not explicit; TI logic prioritizes child health. (0)	None. (0)
	5. Stakeholder Impact	Protects minors and privacy; parents can request exemptions for timezone. (4)	Minor health prioritized; platforms have heavy compliance duty. (1)
	6. Proportionality	Achieves health goal with less privacy intrusion; stricter on health, looser on ID. (5)	Very strict shutdown; heavy enforcement burden on gaming providers. (1)
	7. Monitoring/Review	Audits/fines for violations; no review clause shown (though fines enforced) (3)	Monitored by fines and authority audits (implicit). (2)
	Total RIA Score	23	12
Case 20: Germany Grundgesetz, Art.20a (Climate)	1. Problem Definition	Translates vague constitution into concrete targets (emissions cap, tech rules). (4)	General duty to protect natural resources for future generations. (2)
	2. Objectives	Explicitly cap and reduce GHG annually; forbid harmful projects. (4)	Protect nature and life via legislation broadly. (3)
	3. Policy Options	Detailed prohibitions (e.g. new projects halted in polluted regions) and duties on industry. (5)	Leaves execution to state/legislation. (1)

	4. Cost-Benefit Analysis	Implicit (life/health outranks economic freedom); no monetary CBA. (1)	Not quantified (constitutional principle). (0)
	5. Stakeholder Impact	Protects citizens' health and environment; industry must adapt (tech standards, emissions data). (4)	Broadly frames interests of citizens and nature; implementation vague. (2)
	6. Proportionality	Strong environmental measures (caps, prohibitions) for irreversible climate risk. (5)	Principle is top priority but non-specific; proportionality unclear. (2)
	7. Monitoring/Review	Five-year review cycles, KPIs and adjustment of limits built-in. (5)	No scheduled review; relies on government initiative. (0)
	Total RIA Score	28	10
<i>Case 21: Florida SB7072 (Deplatforming)</i>	1. Problem Definition	Sees absolute ban as too broad; TI allows narrow exceptions for imminent harm. (4)	Protect election candidates from being deplatformed. (2)
	2. Objectives	Safeguard candidate speech while not hindering public safety (imminent violence can still be barred). (4)	Prevent arbitrary removal of a candidate's account/content. (3)
	3. Policy Options	Creates "Candidate Channel," independent rapid review panel (RAP) for content removal appeals (N1–N2). (4)	Flat prohibition with no procedural details or appeals. (1)
	4. Cost-Benefit Analysis	Not done; TI logic implies limits on censorship and extremism. (1)	No CBA. (0)
	5. Stakeholder Impact	Balances candidates' free speech, voters' safety, and platform compliance (includes regulators). (3)	Candidates shielded; platform obligations remain unspecified. (2)
	6. Proportionality	Precise rules and appeal process avoid chilling speech; narrow carve-outs for threats. (4)	Blanket ban may ignore genuine safety concerns. (2)
	7. Monitoring/Review	Requires transparency logs, KPI reporting, and a review cycle (sunset after 2 elections). (4)	No review or independent oversight in law. (0)
	Total RIA Score	24	10
<i>Case 22: NY GBL §349 (Consumer Deception)</i>	1. Problem Definition	Treats broad "unlawful act" as too vague; TI clarifies high-risk (health/finance) vs low-risk deceptions. (4)	Outlaws any "deceptive act or practice in business" (very broad language). (1)

	2. Objectives	Maximize protection for health/safety sectors; mitigate lower-tier commercial deception. (4)	Protect consumers from deception generally. (2)
	3. Policy Options	Tiered approach (high-risk ban in critical sectors, mandatory clear disclosures). (4)	One-size prohibition; relies on courts. (1)
	4. Cost-Benefit Analysis	Not explicit (TI emphasizes irreversibility of personal harm vs business harm). (1)	Not present. (0)
	5. Stakeholder Impact	Consumers (health/safety vs standard transactions) and businesses (safe harbor for compliance) differentiated. (4)	Consumers broadly protected; businesses face uncertainty. (1)
	6. Proportionality	Calibrated bans only where harm is severe; safeguards for business record-keeping. (5)	Vague scope can impose unpredictable risk on businesses. (1)
	7. Monitoring/Review	Requires evidence logs, audits; safe harbor for good-faith actors; no explicit formal sunset but built-in review mechanisms are implied. (4)	Enforcement via litigation; no KPI or review in statute. (0)
	Total RIA Score	26	6

Appendix №1. Compression Algorithm and Demonstration

1. Algorithmic Logic

One of the recurring challenges in applying Threat-Index based normative generation is the production of hyper-complex norms. Such norms (hereafter “hypernorms”) are internally consistent and fully balanced, but contain extensive conditional clauses, triggers, compensators, safeguards, and technical definitions that make them difficult to implement directly in statutory form. A compression algorithm was therefore introduced to translate hypernorms into concise, executable legislative texts, without altering the underlying balance of rights and obligations.

The algorithm follows a multi-stage process:

Input.

The inputs are: (i) the hypernorm itself, (ii) a decision protocol identifying the conflict of rights, the dominant and secondary objectives, and the associated safeguards and instruments, (iii) the corpus of national law and administrative practice, and (iv) budgetary constraints.

Decoding of Balance.

The algorithm distinguishes:

- the conflicting rights at stake;
- the *dominant objective* (to be realized in the fullest possible measure);
- the *secondary objectives*, enforceable only “insofar as not conflicting”;

- the safeguards (pause, rollback, compensation, audit, revision);
- the instruments and their micro-objectives.

Convergence with Law and Institutions.

Each element of the hypernorm is mapped onto an existing institutional or legal carrier (ministries, registries, courts, regulatory bodies). Where no direct equivalent exists, a functional analogue is selected (e.g. arbitration within the Conseil d'État replacing a newly created “independent tax arbiter”).

Budgetary Screening.

Each element is assessed against the fiscal capacity of the jurisdiction. Items not covered by existing appropriations are either simplified (binary triggers), demoted to secondary regulation (décrets), or deferred until resources are available.

Routing.

Each element is assigned one of five statuses:

- Integrate (insert directly into statutory text),
- Functional Analogy (replace with a native mechanism),
- Demote to Sub-statute (implement through decree),
- Simplify (shorten the form while retaining the objective),
- Defer (conditional entry into force).

Assembly of the Adapted Norm.

The statutory text is then constructed with the following components:

- scope and applicability
- core obligation or right,
- one or two automatic safeguards,
- revision clause.

Legal terminology is expressed in the natural style of the jurisdiction, without technical notation.

Balance Control.

Final checks ensure:

- the dominant objective is preserved;
- secondary objectives apply only up to the conflict boundary;
- safeguards leave a visible trace (statute, decree, or analogue);
- all elements rely on existing institutional carriers;
- fiscal feasibility is respected.

Output.

The final output consists of (i) a statutory draft (“Adapted Norms”), and (ii) an explanatory memorandum recording the mapping, the budgetary check, and the justification for each adaptation.

2. Case Study: Digital Services Tax (France)

2.1 Original Hypernorm (N1–N14 + Indicator)

N1 (Scope).

All corporate groups providing “cross-platform digital services” (online advertising, digital marketplaces/intermediation, monetization of user data) to users within the jurisdiction.

N2 (Trigger — OAP).

Combined global turnover of the group ≥ 750 million monetary units in the preceding fiscal year **and** turnover from covered services within the jurisdiction ≥ 25 million monetary units.

Turnover “within the jurisdiction” is defined by a 2-out-of-3 test: (i) billing address of the customer; (ii) verified geolocation of the device/account; (iii) IP geolocation with probability ≥ 0.9 confirmed by the payment provider.

N3 (Obligation).

Payment of the Digital Services Tax (DST) at a baseline rate of 2% of covered turnover within the jurisdiction.

N4 (Anti-double taxation — credit/refund).

The tax authority must credit the DST in full against the corporate income tax (CIT) of the group in the jurisdiction within 24 months of payment; in the absence of a tax base for credit, the same amount must be refunded.

Safeguard: the claim must be accompanied by certified country-by-country revenue and tax breakdowns; misrepresentation is subject to a fine of up to 20% of the disputed amount.

N5 (International coordination — “standstill”).

Trigger: group’s inclusion in a recognized multilateral regime for tax base allocation and mandatory dispute settlement (criteria: binding arbitration + verifiable source rules).

Modality: the tax authority suspends application of the DST to such entities; turnover and tax are allocated under multilateral rules; the local obligation is limited to reporting.

N6 (De-escalation under retaliatory measures).

Trigger: officially declared retaliatory trade measures against the jurisdiction, explicitly linked to the DST.

Modality: the DST rate is temporarily reduced to 0.5% for 12 months; the government must initiate negotiations/mediation; thereafter automatic return to N3 or prolongation of standstill under N5.

N7 (De minimis).

Entities with incremental turnover from covered services < 25 million monetary units within the jurisdiction are exempt.

N8 (Non-discrimination).

The DST applies equally to domestic and foreign entities; individual exemptions based on capital origin are prohibited.

N9 (Price transparency and consumer protection).

Entities may pass the DST into the price but must disclose the final “all-inclusive” price before conclusion of the transaction and may not use misleading labels (“government surcharge”, etc.).

Enforcement: penalty for non-transparency up to 2% of monthly covered turnover.

N10 (Technical reporting and data minimization).

Entities must register and submit quarterly reports with country breakdowns; storage of personal identifiers is limited to ≤ 24 months; transmission to third parties is prohibited except for tax audit or arbitration.

N11 (Independent arbiter).

Disputes on service qualification, turnover sourcing and credits are adjudicated by an independent tax arbitration unit within 120 days; admissible evidence is limited to a closed list (logs, billing, payment data, audits). Decisions must be published in anonymized form.

N12 (Sanctions and incentives).

For non-payment: interest = key rate + 4%; for deliberate misreporting: fine 50–100% of the amount; voluntary disclosure before audit: 80% reduction of fine; repeat offences within 24 months: aggravating factor.

N13 (Monitoring and KPI).

The tax authority must quarterly publish aggregated indicators: total DST collected; proportion credited/refunded; share of disputes and average resolution time; record cases of standstill/de-escalation.

N14 (Sunset/review).

Comprehensive review every 24 months or immediately upon entry into force of a comparable multilateral regime; if the indicator remains below 5% persistently, automatic conversion of the DST into a pure reporting obligation (no levy).

Indicator.

EDTI = share of DST amounts for which no credit or refund was provided within 24 months of payment (target $\leq 5\%$).

2.2 Adapted Legislative Text (France)

Article 1. Scope of Application.

The Digital Services Tax shall apply to groups of companies providing in France online advertising services, digital intermediation platforms, and marketplaces, provided that the combined global turnover of the group exceeds 750 million euros and the turnover from such services in France exceeds 25 million euros.

Article 2. Rate and Payment.

The tax rate shall be 2% of turnover attributable to the territory of France. Payment shall be made quarterly.

Article 3. Credit and Refund.

Amounts paid under the Digital Services Tax shall be credited against corporate income tax, and in the absence of such base shall be refunded within 24 months.

Article 4. Safeguards.

1. In case of officially announced trade measures against France directly linked to the Digital Services Tax, the rate shall be temporarily reduced to 0.5% for up to one year, and the government shall initiate negotiations.
2. Companies shall ensure price transparency: the final “all-inclusive” price shall be disclosed before the transaction, and misleading labels such as “government surcharge” are prohibited.

Article 5. Reporting, Registration and Disputes.

Companies subject to this law shall register with the tax authorities and provide quarterly reports with a country breakdown. Disputes on service qualification and credits shall be resolved by competent French authorities (Conseil d'État and tax arbitration units) under accelerated procedure.

Article 6. Review.

This law shall be reviewed 24 months after its entry into force, or earlier upon conclusion of a multilateral agreement in this field. Furthermore, if the share of amounts for which no credit or refund is provided within 24 months exceeds 5%, the tax shall automatically convert into a reporting-only regime without levy.

2.3 Comparative Mapping

Hypernorm Element	Adapted Legislative Provision	Commentary
N1 (Scope)	Art.1	Scope retained; terminology aligned with French Code général des impôts.
N2 (Trigger, 2-of-3 test)	Art.1	Thresholds preserved; technical “2-of-3” methodology delegated to decree for administrative feasibility.
N3 (Obligation, 2%)	Art.2	Retained in full.
N4 (Credit/refund, ≤ 24 months)	Art.3	Retained; evidentiary detail transferred to decree.
N5 (Standstill under multilateral regime)	Art.6	Implemented through early review clause and suspension upon multilateral agreement; operational details to decree.
N6 (De-escalation under retaliation)	Art.4(1)	Retained (rate reduction + negotiation duty); automatic return handled administratively.
N7 (De minimis)	—	Excluded to avoid manipulation given double threshold; omission does not alter balance.
N8 (Non-discrimination)	—	Subsumed under general French constitutional and EU non-discrimination principles; not repeated.
N9 (Price transparency)	Art.4(2)	Retained in substance; specific penalties covered by general sanction regime.
N10 (Reporting and data limits)	Art.5	Registration and reporting retained; data minimization rules referenced in decree.
N11 (Independent arbiter)	Art.5	Functionally replaced by Conseil d'État and tax arbitration units; integration with French institutions.
N12 (Sanctions and incentives)	—	Covered by general tax sanction regime under CGI; duplication avoided.
N13 (Monitoring and KPI)	Art.6	Incorporated via objective trigger (5% threshold); publication requirements handled administratively.

N14 (Review/sunset)	Art.6	Retained; combined with indicator-based conversion.
Indicator (EDTI)	Art.6	Embedded as 5% threshold for automatic conversion.

3. Observations

The test demonstrates that a hypernorm of fourteen clauses plus indicator can be compressed into a six-article statutory draft, while preserving the dominant objective (fiscal sovereignty and fair taxation of digital giants) and secondary objectives (avoidance of double taxation, prevention of trade escalation, consumer protection).

Simplifications affected only technical methodologies (delegated to decrees), redundant elements (de minimis, non-discrimination already covered by superior law), and institutionally redundant structures (independent arbiter substituted by existing courts). The explanatory mapping ensures transparency of adaptation, demonstrating that no safeguard was lost without trace.

This confirms that the compression algorithm yields concise, executable legislation that remains faithful to the internal balance of the original hypernorm, while being fully compatible with the legal and fiscal framework of the jurisdiction.

Appendix №2. Validation of Operability: Case Study on AI Development under Absence of Empirical Data

Part VI outlined the theoretical procedure by which the Threat Index (TI) framework resolves normative dilemmas when empirical data are missing: by invoking the archival protocol, establishing bounding intervals [TI_best, TI_worst], and applying minimax selection to identify a provisional regulatory scenario. To test this procedure in practice, we conducted a simulation on the regulation of artificial intelligence (AI) development under radical uncertainty.

Actors and Conflict

- A₁ (Society): right to life and health (S = 4).
- A₂ (Developers): rights to property and innovation (S = 1).

Dilemma: Should AI development be restricted when no empirical risk data exist?

Archival Anchor and Calibration

Following the archival protocol, the case “Environment vs Economic Growth” (E-2025-01) was selected as the structural analogue. Both dilemmas oppose societal life/health against economic activity. The causal chains and parameters from E-2025-01 served as calibration baselines.

- A₁ (life/health chain): emission → exposure → harm.
- A₂ (property chain): regulatory burden.

Formulas used (Word-compatible format):

- Link length: $\ell_i = T_i \times P_i \times R_i \times 2^{K_i} \times (1 + 0.1 \times \tau_i)$
- Total chain length: $L = \sum \ell_i$
- Total probability: $P = \prod P_i$
- Threat Index: $TI = (S \times P) \div L$
- Worst-case bound: $TI_worst = (S \times P) \div \min(\ell_i)$

Transfer rules. For each scenario, environmental links were mapped onto AI links (“emission” → “emergence of capability”, “exposure” → “externalization to users”, “harm” → “realization of threat”). Adjustments were applied by modifying T, P, R, K, τ based on safeguards (sandboxing, certification) or absence of control (moratorium, open release).

Scenarios Tested (10)

- s₁ Full Moratorium (ban on all AI; shadow development).
- s₂ Strong Sandbox (restricted pilots with arbiter, kill-switch, logging).
- s₃ Open Release (no restrictions).
- s₄ Sectoral Ban (prohibit high-risk applications only).
- s₅ Compute Cap + Rate-limit + Red-teaming.
- s₆ Certification-only Deployment (pre-market approval).
- s₇ Phased Pilots (gradual unlocking).
- s₈ Open Weights via Proxy Gateway.
- s₉ Regional (Geofenced) Pilots.
- s₁₀ Emergency-only Use.

Results

Table 1. Threat Index values (Society A₁, Developers A₂).

Scenario	TI(A ₁) mid	TI(A ₁) worst	TI(A ₂)	Risk* (max)	L(A ₁)	P(A ₁)	Key $\ell(A_1)$
s₁ Moratorium	0.0133	0.9524	0.957	0.957	5.28	0.0175	[4.03, 1.17, 0.073]
s₂ Sandbox	0.00048	0.0327	0.231	0.231	1.25	0.00015	[1.15, 0.018, 0.075]
s₃ Open Release	6.749	16.97	0.058	6.749	0.071	0.12	[0.028, 0.043]
s₄ Sectoral Ban	0.00136	0.0417	0.644	0.644	0.423	0.00014	[0.311, 0.099, 0.014]

s₅ Cap+Rate+RT	0.00015	0.00254	0.783	0.783	0.258	0.00001	[0.224, 0.019, 0.015]
s₆ Certification	0.000099	0.00273	0.175	0.175	1.089	0.000027	[0.956, 0.093, 0.040]
s₇ Phased Pilots	0.000221	0.00419	0.638	0.638	0.434	0.000024	[0.366, 0.046, 0.023]
s₈ Open Weights/Proxy	0.00573	0.0825	1.286	1.286	0.147	0.00021	[0.097, 0.039, 0.010]
s₉ Geofenced Pilots	0.000525	0.0117	0.997	0.997	0.206	0.000027	[0.158, 0.038, 0.009]
s₁₀ Emergency-only	1.081	2.404	1.425	1.425	0.185	0.05	[0.102, 0.083]

Interpretation:

- s₁: life risk mid-value low (0.0133), but worst-case high (0.95); property loss nearly maximal (0.957).
- s₂: near-zero life risk (0.00048), worst-case 0.033; moderate property cost (0.231).
- s₃: catastrophic life risk (6.7 / 16.97) despite minimal property burden (0.058).
- s₆: very low life risk (0.000099 / 0.0027) and the lowest property burden among restrictive regimes (0.175).

Minimax Selection

Risk*(s) = maximum of TI_life and TI_property for each scenario.

- s₁ = 0.957
- s₂ = 0.231
- s₃ = 6.749
- s₄ = 0.644
- s₅ = 0.783
- s₆ = 0.175 ← minimum
- s₇ = 0.638
- s₈ = 1.286
- s₉ = 0.997
- s₁₀ = 1.425

Optimal scenario: s₆ “Certification-only deployment”.

Reason: Although s_2 gives the smallest TI for life, its ceiling is determined by property risk (0.231). In s_6 , both life and property risks are simultaneously very low, lowering the maximum to 0.175. Certification centralizes costs upfront, producing longer causal chains ($L \approx 1.09$) and extremely small probability ($P \approx 0.000027$) for life-threatening outcomes.

Regulatory Package R*

1. Mandatory pre-market certification of AI systems (threat modeling, red-teaming, resilience tests, emergency shutdown).
2. Sandbox pilots required prior to certification (isolation, kill-switch, proxied outputs, immutable logging).
3. Ban on high-risk applications (medicine, transport, energy, bio/chem synthesis) until certified.
4. Transparency and reporting: incident logs (Class H/F), methodologies, public summaries.
5. Sunset and revision: mandatory review after 9 months, recalculation of TI with empirical data.

MES Gate Check

- C1 Legal certainty: clear scope and triggers.
- C2 Independent arbiter: certification and appeal.
- C3 Enforceability: technical gateways, logging, checklists.
- C4 Learnability: KPI and sunset revision.

KPI

Number of Class H incidents (life/health threats) per 10,000 pilot hours = 0.
 Number of Class F incidents (freedom) per 10,000 pilot hours ≤ 1 .