

WHEN RULES OVERWHELM THEIR PURPOSE

Regulatory Fatigue: The Hidden Failure Mode of Modern Governance and How to Reverse It Without Sacrificing Protection

José M. López· R+D+i Director

Interdisciplinary Research & Policy Analysis

The problem is not the existence of obligations. It is the point at which their cumulative design undermines the capacity to meet their protective purpose.

Abstract

The expansion of regulatory systems is not, in itself, evidence of regulatory failure. In food safety, medicine, environmental protection, finance and artificial intelligence, sophisticated rules can correct market failures, protect rights and create the trust on which investment depends. Yet the cumulative interaction of obligations, amendments, interpretations, audit demands and fragmented enforcement can produce a distinct failure mode we call *regulatory fatigue*. This article proposes a precise, operational definition of regulatory fatigue as the progressive deterioration of an actor's capacity for meaningful compliance when total regulatory demands outstrip its sustainable capacity to understand, prioritize, implement and verify them. It distinguishes the phenomenon from lawful regulatory complexity, administrative burden, conventional red tape, legal uncertainty and individual burnout. Drawing on public-administration scholarship, regulatory-economics research and primary European Union and OECD materials available as of 12 October 2026, it develops a multilevel causal model, identifies observable warning signals, proposes a pilot diagnostic instrument and sets out a risk-preserving simplification agenda for legislators, regulators, companies and assurance bodies. A food-and-biotechnology case illustrates how layered obligations can undermine innovation without necessarily increasing real-world safety. The central proposition is deliberately conditional: where the marginal administrative and cognitive demands of a regulatory system displace the resources needed for substantive risk control, additional compliance activity may reduce -not increase- the system's effective protective capacity. This is a testable hypothesis, not an established universal law. The appropriate response is therefore not indiscriminate deregulation, but better evidence, greater coherence, proportionality, legal certainty, digital interoperability and continuous evaluation of regulatory outcomes.

Keywords: *regulatory fatigue; compliance fatigue; administrative burden; regulatory accumulation; legal certainty; proportionality; better regulation; risk-based supervision; SMEs; food safety; biotechnology; artificial intelligence; regulatory governance.*

1. A paradox at the centre of contemporary regulation

Imagine two laboratories. The first operates under a demanding but stable, intelligible legal framework, with unambiguous responsibilities, interoperable evidence requirements and inspectors focused on actual risk. The second is subject to a similar level of public protection, but must simultaneously interpret overlapping obligations, maintain several near-identical document repositories, reconcile divergent audit questionnaires and respond to frequent procedural changes. The second laboratory may spend more on compliance while developing less actual capability to prevent harm.

That distinction -between *protective regulation* and the *transaction costs of demonstrating compliance*- is the starting point for understanding regulatory fatigue. More rules do not necessarily mean more protection as it's certain that fewer rules do not necessarily mean less. What matters is the relationship between the risks being controlled, the quality of the rules, the cost and reliability of their implementation, and the outcomes delivered.

The challenge has become especially visible in Europe. The OECD's *Regulatory Policy Outlook 2025* stresses the need to evaluate existing regulatory stocks, not merely assess proposed rules [1]. Its 2026 report *Smart Regulations, Strong Business* highlights cumulative burdens, duplication and the inadequacy of measurement techniques that ignore interactions between requirements [2]. The European Commission's Communication of 28 April 2026, *A Simpler, Clearer and Better Enforced EU Rulebook*, likewise places coherence, implementation and enforceability at the centre of simplification [3]. These documents do not establish the present article's theory of fatigue, but they document the regulatory conditions from which it may emerge.

Regulatory fatigue should not be used as a rhetorical synonym for disliking regulation. It is a question of institutional performance: *can the regulated actor still identify what matters, act on it correctly and demonstrate that protective outcomes are being achieved?*

2. What regulatory fatigue is and what it is not

Proposed definition. Regulatory fatigue *is a progressive, observable degradation in the ability of an organization, professional, public authority or citizen to interpret, prioritize, implement and substantively verify applicable requirements when cumulative regulatory demands exceed its sustainable adaptive capacity.*

This definition contains six necessary ideas. First, fatigue is **dynamic**: it concerns deterioration over time, not merely the size of a rulebook. Second, it is **relational**: identical requirements can be manageable for one institution and overwhelming for another. Third, it is **multidimensional**: the burden includes learning, implementation, coordination, evidence production and uncertainty. Fourth, it is **behavioural and operational**: a compliance department can be fully occupied while important risks are poorly controlled. Fifth, it is **measurable in principle**: delays, duplicative effort, recurrent reinterpretation and control failures can be observed. Sixth, it is **normatively neutral about the legitimacy of the underlying protective objective**: the objective may be indispensable even when its delivery architecture is defective.

Related concepts must remain distinct:

- **Administrative burden** is the time, money and effort associated with satisfying administrative requirements. The Standard Cost Model provides an established approach to estimating a subset of these costs [4].
- **Red tape** refers more narrowly to rules or procedures that impose costs without adequately serving legitimate organizational purposes. Bozeman's seminal 1993 analysis distinguishes dysfunctional rules from formalization that may be necessary [5].
- **Legal uncertainty** concerns insufficient predictability about applicable requirements or their interpretation and enforcement.
- **Compliance fatigue** is often used informally for exhaustion in compliance teams or recurring conformity tasks; it is not a single universally standardized scientific diagnosis.
- **Administrative burden on citizens** includes learning, compliance and psychological costs, as developed by Herd and Moynihan [6]. These categories inform the present framework, although conclusions about welfare access cannot simply be transferred to industrial regulation.
- **Regulatory capture** concerns the distortion of regulation in favour of particular interests. It can coexist with fatigue but is a different phenomenon.
- **Regulatory fatigue**, as proposed here, describes the emergent failure of the *whole adaptation system*: rules, interpretations, institutions, people and evidentiary processes.

The word *fatigue* is an *analytical metaphoric* concept and not a medical diagnosis, and this article does not assert that a validated, universally accepted Regulatory Fatigue Index already exists. The construct and measurement architecture below are original proposals requiring empirical testing.

3. Why the problem is cumulative rather than merely quantitative

Regulatory load has at least seven interacting sources.

Volume. The number of obligations matters, but counting statutes or pages tells us little about applicability, duplication or actual cost. One precise and consequential rule may be more demanding than many routine provisions.

Velocity. Amendments, implementing measures, guidance changes, evolving interpretations and new technical standards create continuing training and adaptation costs. Stability and adequate transition periods are therefore valuable regulatory assets.

Ambiguity. Open-textured legal definitions are sometimes unavoidable, but prolonged uncertainty about borderline classifications can consume disproportionate resources. Firms may delay innovation, demand repeated external opinions or adopt overly defensive practices.

Overlap. Several legitimate institutions may seek similar information for different purposes. If they require incompatible formats, frequencies or definitions, cumulative burden can be much greater than the sum of individual obligations.

Fragmentation. A business may face EU rules, national implementing and enforcement regimes, regional or municipal requirements, sectoral codes and contractual obligations. These have different legal statuses and should not be conflated.

Verification intensity. Evidence is essential to trustworthy compliance, but repetitive audits, certificates, questionnaires and document retention can become detached from incremental risk reduction.

Capacity asymmetry. Large enterprises can spread specialist costs across multiple product lines. A small supplier may need many of the same competences, but without comparable scale. Legal exemptions do not always eliminate indirect burdens imposed by larger customers along supply chains.

These factors interact. A stable, lengthy rulebook can be manageable while a shorter but continually changing, contradictory regime may not be. Thus, a meaningful analysis should examine the *rate of change, interactions among duties, interpretive certainty, evidence duplication* and *available capacity*, rather than simply proposing a numerical cap on laws.

4. The multilevel model: system, organization and professional

At the **system level**, legislators and public authorities create the normative stock, issue interpretations, operate digital portals and conduct enforcement. Fragmented mandates, weak ex post review and limited inter-agency coordination can unintentionally multiply demands.

At the **organizational level**, boards and managers translate obligations into operating procedures, responsibilities, suppliers' declarations, validation studies, training, data systems and audit evidence. Poor internal design can generate bureaucracy even where the underlying law is clear. A company may therefore suffer from ***self-inflicted compliance complexity***.

At the **individual level**, scientists, legal advisers, quality managers, technicians and inspectors decide what a rule requires and which issue deserves attention. The relevant limiting resources include specialist time, concentration, institutional memory, professional judgement and the opportunity to challenge questionable assumptions.

A fourth layer deserves explicit attention: **private assurance**. ISO management systems, retailer standards, customer qualification packs and certification schemes may add commercial discipline but are not automatically legal requirements. Failure to distinguish statutory duties from contract-based and voluntary standards can inflate the apparent legal burden and create unnecessary duplication. ISO 37301:2021 provides an established compliance-management framework; applying it intelligently should reduce confusion rather than merely produce more documents [7].

Feedback loop, proposed: *increased demands → less available interpretive capacity → more defensive documentation and procedural duplication → longer cycle times → more perceived*

uncertainty → further demands for evidence. This loop is a hypothesis about organizational behaviour, not a causal finding established by the sources cited here. Its value is that it identifies intervention points that can be investigated.

5. The regulatory-fatigue paradox: when compliance activity competes with protection

Consider a food manufacturer with a fixed number of qualified technical hours. Those hours can be devoted to hazard analysis, process validation, environmental monitoring, corrective actions, supplier assessment, documentation or customer questionnaires. All may have legitimate purposes. However, when repetitive paperwork expands without supplying genuinely new information, it competes with activities that directly reduce food-safety risk.

The proposed **regulatory-fatigue paradox** is therefore:

Beyond a context-dependent threshold, additional procedural obligations may lower effective protection if their marginal consumption of scarce resources exceeds the additional reduction in risk that they generate.

This is *not* a claim that regulatory obligations generally undermine safety. A requirement for preventive controls, traceability or independent verification may yield benefits far exceeding its cost. A poorly controlled pathogen hazard is not made acceptable by invoking burden reduction. The paradox applies only where marginal procedural costs crowd out material risk control and requires empirical demonstration in each setting.

The most dangerous failure is **false compliance confidence** where impressive records and high audit scores coexist with weak controls over the actual hazard. Conversely, an efficient system may maintain extensive evidence while automating routine collection and concentrating specialist review on deviations. The criterion is not fewer documents, but a demonstrable connection between each information request and a protective decision.

6. The evidence: what current institutions actually report

The 2025 OECD *Regulatory Policy Outlook* identifies ex post evaluation of existing rules as an area requiring improvement, and reports that systematic requirements for such evaluations exist in fewer than one-third of OECD members, under the definition used in that report [1]. This indicates a governance imbalance: policymakers often examine proposed obligations more carefully than the accumulated effects of older ones.

The OECD's 2026 *Smart Regulations, Strong Business* -based in part on its 2025 Simplifying for Success survey across 34 jurisdictions- explicitly identifies overlapping obligations, repeated requests for information and measurement gaps concerning cumulative burden [2]. Its analysis cautions that the Standard Cost Model predominantly captures administrative costs rather than the full substantive, dynamic and indirect effects of regulation [2,4].

In its December 2025 economic analysis, the OECD also associates regulatory stock and features of regulation with weaker productivity and business dynamism, while recognizing the substantial social and economic benefits of sound regulation [8]. These findings should not be converted into a universal claim that every additional legal rule reduces GDP: causality, institutional quality and the composition of regulation matter.

On 28 April 2026 the European Commission adopted Communication COM(2026) 380 final and an associated Regulatory Deep Cleaning Action Plan [3]. The Commission has also stated targets to lower

administrative burdens by at least 25% overall and 35% for SMEs [9]. These are **policy targets**, not evidence that all promised reductions have already reached individual companies. Similarly, projected savings from simplification proposals are not interchangeable with independently verified realized savings.

The empirical record therefore supports the premise that cumulative burden and poor coordination are important policy problems. It does not yet validate a unique, cross-sector physiological or organizational syndrome called regulatory fatigue. That distinction between **documented background conditions** and a **proposed explanatory model** is essential to scientific integrity.

7. European regulatory design: legitimate protection versus avoidable friction

The European Union is an especially valuable case study because common substantive objectives coexist with multi-level implementation. Harmonization can reduce divergence across national markets, but divergent enforcement practice, additional national requirements where legally permitted, and inconsistent administrative systems can reintroduce complexity.

Under Article 6 of Regulation (EC) No 178/2002, EU food law is generally based on risk analysis, with scientific evidence central to risk assessment [10]. Regulation (EU) 2017/625 requires official controls to be undertaken regularly, on a risk basis and at appropriate frequency, considering specified factors including an operator's compliance record and the reliability of its own controls [11]. These provisions demonstrate that proportionate and risk-informed application is not an argument against food protection: it is embedded in the legal design.

Proportionality does not allow operators to waive mandatory duties. It guides the design and exercise of regulatory powers within their proper legal limits. The distinction is decisive. Simplifying an audit questionnaire, accepting interoperable evidence where authorized or improving guidance may be lawful and beneficial; disregarding a mandatory authorisation or substituting an unvalidated safety claim is neither simplification nor sound compliance.

The same distinction extends to the EU's developing digital acquis. Regulation (EU) 2024/1689 establishes a risk-oriented legal framework for artificial intelligence, but the exact application of obligations depends on the system, role, use case, relevant dates and subsequent amendments [12]. Firms should determine applicability from the current consolidated legal text and official guidance rather than rely on fixed checklists based on the original 2024 text alone.

A modern legal system should be judged by the *quality of its protective decisions*: clarity of obligations, access to competent interpretations, proportional control intensity, consistency of enforcement, procedural fairness and effective remedies.

8. Food science, biotechnology and innovation: a particularly concrete stress test

Food and biotechnology reveal why the issue is deeper than general business paperwork. Introducing a new ingredient, a modified processing step or an emerging antimicrobial technology may require several parallel assessments: hazard identification, toxicology where relevant, technological function, product classification, hygiene implications, exposure, process validation, labelling and claims, jurisdictional applicability and supplier traceability.

The relevant instruments depend on the facts. The General Food Law framework [10], hygiene rules such as Regulation (EC) No 853/2004, official-controls rules [11], and, **where the specific statutory criteria are met**, the Novel Foods Regulation (EU) 2015/2283 [13] can arise in different combinations. An ingredient is not automatically a novel food because its process is innovative; nor is a technology

exempt from scrutiny merely because its ingredients are familiar. A legal classification requires an evidence-based examination of the actual product, function, process, exposure and applicable definitions.

Now add commercial layers as retailer questionnaires, food-safety certification, bespoke specifications, supplier portals and repeated client requests for information already validated elsewhere. These private requirements may be rational from each buyer's perspective while creating expensive duplication at supply-chain level.

A more intelligent approach begins with a **regulatory applicability tree**, not an indiscriminate collection of certificates. First identify the marketed product and intended use; then determine the applicable legal category and territorial market; next establish the scientific proof and authorizations genuinely required; only then map relevant voluntary and contractual standards. A single controlled evidence dossier can serve multiple legitimate inquiries, provided confidentiality and legal obligations are respected.

Hypothetical example, not industrial trial data. Suppose a small ingredients company spends 400 technical hours per quarter responding to seven substantially overlapping customer packs. After mapping the requirements, it creates a traceable, version-controlled dossier and reduces duplicated work by 120 hours per quarter while maintaining every required control. The recovered 120 hours could be reassigned to validation, supplier oversight or method development. This is an illustrative mechanism, **not** a measured gain for any named company, and its safety value would have to be tested by appropriate outcome indicators.

The lesson is not to create special legal immunity for biotechnology. It is to differentiate carefully between *scientific evidence necessary to establish safety or efficacy* and *administrative duplication that adds no meaningful assurance*.

9. SMEs, supply chains and the invisible transmission of obligations

Small enterprises can face an asymmetric burden even when legislation formally distinguishes them from larger entities. Major customers may extend their own documentation preferences down the supply chain, require compliance with additional codes or set evidence thresholds that are commercially unavoidable for suppliers.

This creates a critical methodological problem when statutory obligations, binding contractual duties and voluntary assurances have different sources and different routes to correction. A regulation can be improved by the legislature, a customer's duplicative qualification requirements should be addressed through contractual standardization and mutual recognition and a company's own redundant controls can be eliminated internally.

Failure to separate these sources produces policy errors. A government could repeal a reporting duty but businesses might continue filing nearly identical information to multiple private purchasers. Conversely, a firm could blame regulation for documentation that it imposed on itself years earlier and never reviewed.

A responsible burden inventory must therefore record **who requires the action, its legal basis, the associated risk, the original evidence owner, the recipient, the frequency, the time cost and whether the information is reused elsewhere**. This simple discipline transforms complaints into a verifiable diagnosis.

10. The role of uncertainty and defensive compliance

Uncertainty has a distinct economic cost if it can delay technical investment, freeze product decisions and increase dependence on external advice. Some uncertainty is inherent in novel technologies and must be managed through evidence and precaution. Other uncertainty arises from poorly coordinated interpretations or the absence of accessible procedural guidance.

In such settings, defensive compliance can become self-reinforcing. Teams maintain old forms because deleting them feels risky. Audit findings create new checks without removing obsolete checks. Internal manuals accumulate multiple versions of similar procedures. Nobody is explicitly responsible for retiring documents. Eventually, the existence of documentation becomes confused with the existence of control.

The corrective principle is **documentary proportionality** that rely on simple ideas as to retain evidence that is legally required, scientifically material, contractually necessary or demonstrably useful, to identify records that can be consolidated and remove only what a competent, documented review finds redundant. The removal decision itself should be traceable.

11. Regulatory fatigue can also affect regulators

The regulatory authority is not immune to cumulative overload. Inspectors, legislative drafters, laboratory assessors and licensing personnel may face expanding mandates without equivalent analytical resources. Consequences can include longer response times, inconsistent interpretations, procedural conservatism and a shift from risk assessment to readily countable administrative activity.

Importantly, delays or inconsistent decisions do not prove individual fatigue. Staffing shortages, statutory complexity, contested evidence, inadequate IT systems and procedural guarantees may explain them. The purpose of the framework is to identify potentially remediable institutional constraints, not to diagnose or blame public servants.

Regulators need integrated rule inventories, shared interpretation protocols, protected technical training time, quality-controlled digital systems, clear escalation channels and performance indicators focused on outcomes rather than inspection volume alone.

12. Measuring the phenomenon: a proposed Regulatory Fatigue Diagnostic (RFD)

The central research challenge is turning an attractive concept into a falsifiable instrument. I propose a **Regulatory Fatigue Diagnostic**, initially implemented as a dashboard rather than an authoritative numerical ranking.

Its seven exposure dimensions are: **D1-change velocity** (material legal and interpretive updates per quarter); **D2-interpretive friction** (elapsed specialist time needed to settle applicability questions); **D3-evidence duplication** (repeated information requests with equivalent content); **D4-coordination burden** (handoffs and rework among authorities, departments or customers); **D5-verification load** (hours spent demonstrating rather than performing control); **D6-capacity saturation** (mandatory regulatory work relative to qualified capacity); and **D7-uncertainty persistence** (unresolved material compliance questions weighted by age and impact).

Each dimension should be expressed as an observable indicator with a clear denominator and stratification by firm size, sector, legal risk and complexity. Raw numbers of laws or audits should not be used in isolation. Indicators must be benchmarked against comparable entities and their own longitudinal baselines.

A simple operational measure is the **Regulatory Load-to-Capacity Ratio (RLCR)**:

RLCR = qualified hours required to execute, interpret and verify applicable obligations during a period / qualified hours sustainably available for those functions during the same period.

A ratio above 1.0 indicates an infeasible workload *under the assumptions and period measured*; a ratio below 1.0 does not prove that risk is adequately controlled or that working conditions are sustainable. Measurement must account for task criticality, skill differences, peak workload, outsourcing and non-substitutable expert functions.

For research pilots, a composite 0–100 score could be constructed by rating each of the seven dimensions from 0 to 5 using pre-registered, evidence-based thresholds, assigning justified weights w_i , and calculating **RFD = $20 \times \sum(w_i \times D_i)$** , where $\sum w_i = 1$. This is a **proposed formula only**. No sensitivity, reliability, predictive validity, normative thresholds or cross-sector comparability have been demonstrated. Before public use, the measure would need expert content validation, inter-rater reliability testing, longitudinal calibration, factor analysis where appropriate and correlation with independent risk outcomes.

Crucially, outcome measures should remain outside the initial fatigue score to avoid circular validation. They might include material non-conformities, recurrence of critical findings, corrective-action effectiveness, validated process deviations, unplanned market delays, staff turnover and inspection outcomes. Privacy, confidentiality and incentive effects must be addressed.

13. A balanced economic model: net protective value, not paperwork volume

Conventional burden reviews may emphasize hours and direct costs because these are readily measured. A defensible evaluation must also estimate the benefits of regulation: prevented injuries, reduced contamination, avoidance of fraud, environmental gains, consumer confidence, fair competition and improved legal predictability. The World Bank's B-READY methodology likewise distinguishes the quality of formal regulation, the availability of relevant public services and operational efficiency in practice—three different objects of assessment [15].

One conceptual decision rule is:

Net regulatory value = expected protective and market-enabling benefits – implementation, verification, uncertainty and opportunity costs.

This equation is an analytical identity, not a ready-made welfare measurement method. Prevented harm may be rare, benefits may be diffuse, ethical constraints may forbid certain trade-offs, and precise monetization may be inappropriate. Where high-consequence risks or fundamental rights are at stake, the decision must respect binding legal standards and precautionary obligations.

Policy appraisal should compare real alternatives: retain the requirement as designed, clarify its interpretation, integrate it with another data flow, substitute an equivalent and legally permissible assurance mechanism, adjust review frequency to risk, change the transitional timetable or repeal a demonstrably obsolete duty through the proper legislative process.

The correct target is **maximum justified protection per unit of avoidable administrative friction**, not minimum regulation in the abstract.

14. Ten solutions for public policymakers and regulators

- 1. Manage the regulatory stock, not only the regulatory flow.** Maintain inventories of active obligations, implementing instruments and interpretive guidance. Identify provisions whose protective rationale has become obsolete or whose practical interaction is defective.
- 2. Make ex post evaluation routine and consequential.** Establish review dates, measurable protective objectives, stakeholder evidence and public responses explaining what will be retained, amended or removed. Ex post review must assess benefits as well as burdens [1,2].
- 3. Design for clarity from the outset.** Publish plain-language applicability tools alongside authoritative legal texts; make terminology consistent; state scope, responsible actor, evidence standard and enforcement consequences. Explanatory materials should distinguish legally binding rules from guidance.
- 4. Build risk-based enforcement.** Prioritize inspection and verification according to material hazards, control reliability and relevant compliance history, within the limits of governing legislation. In EU food law, this orientation already has a clear statutory basis [11].
- 5. Reduce duplicate reporting through interoperable data.** Use common vocabularies, secure interfaces and the once-only principle where legally permitted. A single submission should satisfy multiple authorized purposes only when the data are adequate for each purpose and privacy safeguards are respected.
- 6. Require cumulative-impact assessment.** Before adopting a new duty, examine how it interacts with existing obligations and with the practical capacities of its intended addressees. A low-cost measure considered alone may be costly in combination with twenty others.
- 7. Adopt intelligent transition periods.** Provide realistic implementation windows, early technical guidance and phased compliance where consistent with urgent protective needs. Avoid simultaneous major changes without a credible transition plan.
- 8. Limit unnecessary national or local overlays.** Where EU law leaves discretion, authorities should identify and publicly justify additional procedural requirements; where national measures are lawful and necessary, their burden should be measured rather than presumed illegitimate.
- 9. Allow challenge and correction.** Create accessible mechanisms to resolve conflicting interpretations, correct administrative error and obtain reasoned decisions. Predictable, accountable enforcement is a form of regulatory quality.
- 10. Evaluate simplification by outcomes.** Count eliminated hours, but also examine injury, contamination, fraud, compliance and innovation indicators. A reduced form count is not success if protection worsens.

Administrative 'one-in, one-out' rules can sharpen attention to costs, but they are not substitutes for risk-benefit analysis; essential new protection must not be blocked merely to meet an arithmetic quota. The OECD explicitly treats offset mechanisms as one tool within a broader stock-management architecture [8].

15. Ten solutions for companies, laboratories and industrial groups

- 1. Build a legally classified obligation register.** For each requirement record: source, jurisdiction, article or clause, legal status, applicable product/process, owner, due date and evidence of compliance.

2. Separate three layers of requirements. Distinguish binding law; contractual and certification conditions; and internal management preferences. This prevents accidental elevation of voluntary routines into supposed statutory duties.

3. Map obligations to real risks. Every significant control should identify the hazard, legal or contractual purpose, preventive action, acceptance criterion and evidence owner.

4. Establish a controlled regulatory-change process. Assign one accountable coordinator; monitor authoritative sources; classify changes by impact; document applicability decisions; and verify that training and procedures are updated.

5. Create a reusable evidence architecture. Maintain version-controlled core dossiers for formulations, safety assessments, specifications, supplier qualifications and validation reports, with permissions appropriate to confidentiality and trade secrets.

6. Consolidate audits where permissible. Map overlapping customer and certification requests against existing assurance evidence; offer structured cross-references rather than generating essentially identical documents each time.

7. Protect critical technical capacity. Reserve qualified personnel time for hazard assessment, monitoring, corrective actions and development. Do not measure compliance performance solely by the volume of reports completed.

8. Use technology as decision support. Regulatory-information systems and AI-assisted retrieval can identify changes, extract candidate obligations and compare evidence packages, but material applicability decisions require expert verification against authoritative, current texts. Maintain citations, dates, access controls and audit trails.

9. Adopt disciplined document retirement. Review obsolete SOPs, redundant checklists and duplicative forms through authorized change control. Never remove a legally required record or scientifically necessary control for convenience.

10. Escalate structural overload to governance. Management boards should receive both workload indicators and risk-outcome indicators. When capacity is demonstrably inadequate, the answer may be staffing, revised product scope, better systems or process redesign and not exhortations to 'work harder'.

ISO 37301 may provide a helpful management-system reference [7]; for AI support, NIST's voluntary AI Risk Management Framework offers complementary functions -Govern, Map, Measure and Manage- without replacing applicable EU legal duties [14].

16. Artificial intelligence: a cure that can become another cause

AI could substantially reduce regulatory search, document comparison, obligation mapping and version tracking. A well-designed system could link each internal requirement to a current legal provision, identify amendments and show where identical evidence is requested by multiple parties.

But an AI layer can itself generate regulatory fatigue when it produces uncontrolled alerts, plausible but inaccurate legal conclusions, undocumented interpretations or excessive governance documentation. Automated extraction from an outdated version of a regulation creates the appearance of precision without legal reliability.

An appropriate architecture has five safeguards: **authoritative source retrieval, temporal version control, explicit uncertainty and abstention, qualified human approval for material decisions and**

reproducible audit logs. Sensitive supplier data and intellectual property must not be exposed through unapproved services.

The objective is not to remove human expertise. It is to preserve expert attention for questions that actually require legal, scientific or ethical judgement. Technological automation should be evaluated by time saved, error rates, false alerts, corrected interpretations and substantive compliance outcomes and not by the number of generated summaries.

17. A practical 90-day implementation pathway

Days 1–30: discover and classify. Select a defined scope, as for example, one plant, one product category or one regulatory function. List applicable requirements, identify sources, distinguish law from private assurance and measure actual work hours. Interview the specialists who perform the work. Establish a baseline for critical risk controls and incident outcomes.

Days 31–60: remove friction without removing protection. Map recurring requests, duplicate files, unnecessary handoffs and unresolved interpretation questions. Consolidate evidence repositories, update ownership, request authoritative clarification where needed and redesign two or three high-friction workflows. Approve every removal through change control.

Days 61–90: validate and institutionalize. Repeat workload measurement and review errors, control effectiveness and critical findings, independently verify that no mandatory obligation has been lost, and report the balance of recovered capacity and maintained or improved protection. Establish quarterly review of regulatory changes and documentary redundancies.

A plausible *pilot success criterion* would be demonstrably lower avoidable administrative time **with no deterioration in pre-specified critical protection indicators**. This is a suggested evaluation criterion, not a guarantee or a substitute for sufficient statistical power. Some hazards require long observation windows; a 90-day pilot can validate processes, not prove all long-term safety outcomes.

18. Research agenda: how to test the theory rather than merely repeat it

The concept becomes scientifically useful only if it yields predictions that can fail. Four studies are proposed.

Study A: construct validity. Compare the seven diagnostic dimensions across matched firms and regulatory settings. Test whether fatigue is distinguishable from total compliance cost, organizational inefficiency and legal uncertainty.

Study B: longitudinal change. Track firms before and after a major regulatory amendment. Measure interpretation time, duplicated evidence production, staffing and objective compliance outcomes. Include comparison groups where feasible and adjust for demand, firm size and sectoral risk.

Study C: process-level randomized or quasi-experimental simplification. Where lawful, introduce interoperable reporting or consolidated evidence requirements in selected workflows while preserving substantive rules. Examine administrative time, errors and safety metrics relative to comparable workflows.

Study D: regulator–industry dyads. Collect data from both sides of the compliance interaction to identify where burdens are generated, transferred or amplified by particular information requests.

Three testable hypotheses follow:

H1. Holding substantive risk and operational complexity approximately constant, greater interpretive ambiguity and evidence duplication predict higher qualified-hours consumption.

H2. Where staffing capacity is constrained, increases in non-value-adding verification work predict deterioration in time available for direct risk controls; whether protective outcomes worsen is an empirical question.

H3. Consolidation of equivalent evidence flows, under maintained legal and scientific controls, reduces administrative burden without reducing independently measured protective performance.

Strong confounding risks remain. Poorly managed firms may simultaneously accumulate documentation and suffer control failures, high-risk enterprises may legitimately require more evidence and inspectors may increase scrutiny in response to detected problems rather than cause them. Any credible evaluation must account for these alternatives.

19. The objection: is regulatory fatigue merely an excuse for deregulation?

This objection must be taken seriously. The language of **burden reduction** can be used to weaken environmental safeguards, reduce labour protections, diminish transparency or transfer risk to consumers. Businesses can also understate the value of rules that constrain profitable but harmful conduct. The distribution of costs and benefits is not neutral: savings to one firm may become risks imposed on workers, competitors, communities or future generations.

A rigorous regulatory-fatigue framework therefore needs **non-negotiable protective constraints**, transparent evidence and scrutiny of who gains and who bears the risk. It cannot presuppose that the cheapest system is the best one. Nor should it confuse the cost of complying with a justified legal standard with the avoidable cost of proving the same fact repeatedly.

The mirror-image mistake is to assume that every procedure must be protective simply because it belongs to a regulatory system. A redundant certificate is not inherently safer than an authenticated, reusable dataset. A second checklist is not automatically more rigorous than a demonstrated control. Serious protection demands evidence of effectiveness, not reverence for bureaucracy.

This perspective is compatible with the European Commission's 2026 insistence on simplification alongside better enforcement [3], and with the OECD's emphasis on safeguarding societal protections while improving regulatory efficiency [2].

20. Conclusion: from accumulating obligations to governing outcomes

The mature regulatory state should not be measured by how many rules it produces, how many inspections it performs or how many documents it collects. It should be measured by whether legitimate objectives are actually achieved, rights remain protected, risks are controlled and regulated actors can determine -with reasonable effort and legal certainty- what they must do.

Regulatory fatigue names a possible failure of that system in special if the progressive diversion of scarce interpretive and operational capacity falls into accumulating, fragmented and sometimes duplicative compliance demands. The distinction is fundamental. An excess of requirements may be costly as fatigue begins when the architecture of those requirements impairs the ability to comply meaningfully.

The solution is neither reflexive deregulation nor blind accumulation. It is **regulatory intelligence** defined as continuous review of the rulebook, disciplined risk assessment, differentiated obligations, integrated evidence, well-designed implementation, professional accountability, reliable digital tools and measurement of real outcomes.

A regulator should be able to explain what hazard a requirement addresses. A company should be able to show how its control reduces that hazard. An auditor should be able to identify what additional assurance an information request actually creates. When none can answer, the proper response is not necessarily to remove the underlying rule, but to redesign the path from legal intention to operational protection.

The future of good regulation will depend less on the accumulation of obligations than on the demonstrable intelligence with which they are designed, applied and reviewed.

Methodological and editorial note

This article is an **original conceptual and policy-analysis paper**, not a systematic literature review, commissioned legal opinion, validated psychometric study or report of completed laboratory or industrial experiments. The proposed definition, feedback loop, paradox, diagnostic instrument and research hypotheses are authorial propositions. All quantitative examples are explicitly hypothetical. Sources cited below are used for their respective legal, policy or scholarly propositions, not as purported empirical validation of the novel framework. The legal discussion describes selected EU instruments and is not a substitute for product and jurisdiction-specific advice.

Source review cut-off: October 2026.

References

- [1] OECD (2025). *OECD Regulatory Policy Outlook 2025*. OECD Publishing. [Open source](#)
- [2] OECD (2026). *Smart Regulations, Strong Business: Enabling Growth and Societal Protection in a Fast-Changing World*. OECD Publishing, 3 June 2026. [Open source](#)
- [3] European Commission (2026). *A Simpler, Clearer and Better Enforced EU Rulebook*. Communication COM(2026) 380 final, 28 April 2026. [Open source](#)
- [4] OECD (2026). *Smart Regulations, Strong Business*, sections on regulatory burden measurement and the Standard Cost Model. [Open source](#)
- [5] Bozeman, B. (1993). A Theory of Government “Red Tape.” *Journal of Public Administration Research and Theory*, 3(3), 273–304. [Open source](#)
- [6] Herd, P., & Moynihan, D. P. (2018). *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation. [Open source](#)
- [7] ISO (2021, current edition confirmed in 2026). *ISO 37301:2021—Compliance management systems: Requirements with guidance for use*. [Open source](#)
- [8] OECD (2025). *OECD Economic Outlook, Volume 2025 Issue 2*, chapter “Time for a Regulatory Reset?” [Open source](#)
- [9] European Commission. *Better Regulation and Simplification* policy pages; targets and programme status as consulted 12 October 2026. [Open source](#) ; [Open source](#)
- [10] European Parliament and Council (2002). Regulation (EC) No 178/2002 laying down the general principles and requirements of food law, especially Articles 6–7. [Open source](#)
- [11] European Parliament and Council (2017). Regulation (EU) 2017/625 on official controls, especially Article 9(1). [Open source](#)
- [12] European Parliament and Council (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence, as subsequently amended. [Open source](#)
- [13] European Parliament and Council (2015). Regulation (EU) 2015/2283 on novel foods. [Open source](#)
- [14] Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. [Open source](#)
- [15] World Bank. *Business Ready (B-READY) Methodology*, framework differentiating regulatory quality, public services and operational efficiency. [Open source](#)