

REGULATORY FATIGUE AS CAPACITY EROSION

A Multilevel Theory of Compliance Congestion, Protective Displacement and Risk-Preserving Simplification

José M. López · R+D+I Director INTABIOTECH

Interdisciplinary Research and Regulatory Governance

ORCID: 0009-0006-9238-0080

Article type: Conceptual/theory-development paper with sectoral applications and a prospective empirical research protocol.

Abstract

Regulatory systems may become more elaborate while the practical capacity of regulated organisations to achieve regulatory objectives deteriorates. Existing research examines red tape, administrative burden, regulatory accumulation, complexity, audit proliferation and organisational decoupling, but these literatures do not by themselves specify the dynamic transition from manageable requirements to degraded substantive compliance. This paper advances *regulatory fatigue* as a distinct, conditional, organisation-level phenomenon: progressive erosion of an actor's adaptive compliance capacity under cumulative, interacting and insufficiently absorbable regulatory demands. Through a critical integrative synthesis of established scholarship and European regulatory sources, it proposes a multilevel mechanism connecting obligation load, interpretive volatility, evidentiary duplication, attention allocation and the displacement of risk-control work. It draws on emerging work on *compliance congestion*, distinguishes the proposed mechanism of *protective displacement*, and specifies boundary conditions, rival explanations, falsifiable propositions and a provisional diagnostic architecture. Food ingredients, biotechnology and industrial supply chains provide particularly revealing settings because binding public-law duties coexist with contractual assurance and voluntary certification. Practical proposals include a traceable evidence passport, a protective-floor test for simplification and an enterprise regulatory-capacity review. The paper does not report primary empirical data, assert a validated fatigue scale or establish a universal causal relationship between regulation and harm. Its contribution is a testable theory of when apparently successful procedural compliance can mask declining capacity for prevention, innovation and credible assurance. The central policy implication is not deregulation: it is the redesign of verification and implementation so that scarce expertise remains available for the substantive risks regulation is intended to control.

Keywords: *regulatory fatigue; compliance congestion; protective displacement; administrative burden; regulatory governance; red tape; food regulation; biotechnology; audit duplication; risk-based enforcement; organisational capacity; regulatory technology.*

1. Introduction: an institutional paradox requiring explanation

The debate on regulation is often framed as a dispute between public protection and business freedom. The frame is incomplete. Two regimes may pursue identical levels of consumer, worker or environmental protection while differing sharply in how much attention, technical effort and interpretive work they demand from the same enterprise. A control-oriented regime may require one validated hazard assessment and permit its evidence to be reused while another may repeatedly require the same facts in different forms, at different intervals and under inconsistent interpretations. The distinction is not simply administrative inconvenience. It raises a substantive question about the production of safety and legal compliance.

This article asks: **under what conditions does the cumulative process of demonstrating regulatory compliance begin to undermine an organisation's capacity to achieve the underlying protective objectives?** Related questions follow. *How can the transition be distinguished from ordinary complexity or management failure? Which mechanisms generate it? How can public authorities and enterprises detect it without treating compliance costs as evidence against legitimate standards? Which forms of simplification preserve or strengthen protection rather than silently transferring risk to consumers, workers or communities?*

The policy setting is unusually relevant. The OECD's *Regulatory Policy Outlook 2025* argues for strengthening the effectiveness of rulemaking and reviewing existing frameworks [1]. In June 2026, *Smart Regulations, Strong Business* examined cumulative burden, fragmentation, business perceptions and the structural weaknesses of simplification programmes [2]. The European Commission's April 2026 Communication, *A Simpler, Clearer and Better Enforced EU Rulebook*, proposed a regulatory deep-cleaning agenda and simplicity by design while explicitly maintaining legitimate objectives [3]. These sources establish an important policy problem; they do **not** demonstrate the causal construct advanced here.

The scientific starting point is therefore modest but significant. We should not rename bureaucracy and announce a discovery. Decades of scholarship already address organisational red tape, audit systems, perceived fairness, administrative burden and the decoupling of formal procedure from operational activity [4–9]. The proposed advance is more precise: to theorise **capacity erosion over time**, specify the threshold separating costly-but-functional obligations from deteriorating substantive performance, and identify observations that could disconfirm the theory.

Three contributions structure the paper. **First**, it distinguishes the **stock of obligations** from **compliance congestion**, understood here as competing interpretive and evidentiary tasks at identifiable bottlenecks, while the term and closely related congestion models also appear in contemporary financial-compliance scholarship [25]. **Second**, it defines **protective displacement** as the real locative mechanism through which low-marginal-value verification consumes the scarce personnel, attention or process windows needed for direct hazard control. **Third**, it designs a falsifiable diagnostic and evaluation architecture that requires the joint observation of workload, organisational capacity and protective outcomes. Neither congestion nor displacement should be inferred merely from managers' complaints or from high expenditure.

The discussion is centred on European food and biotechnology enterprises because the relationship among product safety, pre-market authorisation, supplier assurance, traceability, plant controls and innovation is unusually visible there. Nevertheless, the mechanism potentially applies to medical technologies, chemicals, finance, environmental permitting and digital governance, subject to different legal and empirical boundary conditions.

2. Materials, analytical strategy and scope of inference

This is a **critical integrative conceptual analysis**, not a systematic review or an empirical effects study. The materials comprise three groups: (i) established research in public administration, organisational sociology, audit studies and behavioural compliance; (ii) primary EU legislation and institutional policy material available up to 12 October 2026; and (iii) publicly available standards and assurance frameworks relevant to industrial compliance. Sources were selected for conceptual relevance, legal authority and usefulness in testing the proposed mechanism, rather than through a preregistered exhaustive literature search.

The analytical sequence is explicit. Prior constructs are distinguished: a hypothesised process is assembled and its causal dependencies and competing explanations are specified, observable variables and possible instruments are proposed and industrial scenarios are employed as **analytical illustrations**, not case-study findings. No unpublished company data, internal industrial validation, confidential formulation or field observations are represented as scientific evidence. Mathematical expressions are conceptual accounting devices, not estimated structural coefficients. Numerical illustrations, where used, are labelled as assumptions.

This discipline matters because conceptual originality and empirical confirmation are different scholarly achievements. A proposed diagnostic index does not become validated by having a memorable acronym, and the fact that a public authority has embraced simplification does not establish the new theory. Submitting the paper as a theoretical contribution rather than an invented empirical research article protects both scientific integrity and the prospective journal submission.

3. The intellectual antecedents: where regulatory fatigue fits - and where it differs

Red tape and rule density. Bozeman's work distinguishes purposeful formalisation from pathological administrative constraint and explicitly relates rule density to organisational resources [4]. Any account of fatigue must acknowledge this precursor. What is added here is a time-sensitive deterioration mechanism: an initially legitimate obligation can become part of a congested system when interpretations, overlaps and verification demand interact with changing capacity. Fatigue may occur even when every individual rule serves a legitimate purpose; conversely, red tape can exist without measurable fatigue if organisations still have ample capacity.

Administrative burden. Herd and Moynihan conceptualise learning, compliance and psychological costs in citizens' encounters with the state [5]. Their framework helps identify hidden costs, but citizen access to public programmes and industrial assurance are not identical domains. Industrial measures require production-process and safety variables, supply-chain information and externally specified scientific thresholds in addition to perceived effort.

Institutional decoupling. Meyer and Rowan showed how formal structures can become ceremonial, serving legitimacy while becoming separated from operational activity [6]. Regulatory fatigue offers a possible *generating mechanism* for some forms of decoupling: personnel rationally respond to intense, externally visible proof demands by reallocating time toward auditable artefacts. Yet decoupling may also arise from strategic deception, culture or incentives even in the absence of overload. The two concepts must be separately measured.

The audit society. Power's analysis of verification rituals highlights the potential divergence between formal assurance and what can be nowadays known through auditing [7]. The present theory examines when the marginal informational value of one additional assurance request becomes small compared with its effect on scarce verification and hazard-control capacity. That marginal value cannot be presumed zero merely because a document resembles another document.

Compliance with complex regulation. Mendoza, Dekker and Wielhouwer studied 602 Dutch financial intermediaries and found that regulatory knowledge mediated the relationship between perceived fairness and compliance [8]. Their results caution against a simplistic '*complexity necessarily impairs compliance*' proposition. Complexity can be understood and legitimately managed. Subsequent research into firms' responsiveness to regulatory feedback has also emphasised financial capacity and expectations, supporting the need to model organisational resources rather than rules alone [9].

Contemporary congestion modelling. Likas (2026) has already advanced an anti-money-laundering model in which expanding alert volumes under constrained investigative capacity degrade evidence and create exploitable triage patterns [25]. That contribution precedes this paper; accordingly, neither the term **compliance congestion** nor the general possibility of overload-induced enforcement failure is claimed as original here. The proposed incremental contribution is the food- and biotechnology-oriented analysis of regulated organisational capacity, the separation of hazard-protection outcomes from evidentiary queues, a protective-floor intervention test, and a comparative validation agenda across private and public assurance interfaces.

Regulatory accumulation and economic analysis. Research and policy debate on accumulation ask how interacting rules affect economic performance and investment [2,21]. Some published macroeconomic estimates of cumulative costs are highly model-dependent; they should not be imported as established causal evidence of organisational fatigue. Our unit of analysis is the firm–requirement–process interaction, with directly observable tasks, capacity constraints and performance indicators.

These intellectual antecedents expose an important limitation of the word *fatigue*. It risks importing a psychological or medical interpretation. This paper therefore treats fatigue as a **system property and dynamic process**, not a clinical diagnosis, and proposes the more precise descriptive term **regulatory-capacity erosion** for empirical operationalisation. The popular label may aid communication, the scientific claim rests on mechanisms and measurements.

Table 1. Conceptual separation from established constructs

Construct	Primary object	Distinguishing feature	Why it does not alone establish regulatory fatigue
Red tape	Dysfunctional rules and procedures	Weak relationship to legitimate purpose [4]	Dysfunction can persist without capacity deterioration
Administrative burden	Learning, compliance and psychological costs [5]	Transaction costs to the addressee	Costs can remain stable and be absorbed
Regulatory accumulation	Growing stock and interactions of rules [2,21]	System-level change in requirements	More rules can increase protection or be manageable
Institutional decoupling	Gap between formal structures and practice [6]	Symbolic versus operational compliance	Other incentives can cause decoupling without overload
Audit proliferation	Intensification of verification [7]	Demand for documented assurance	Independent checks can have high marginal value
Regulatory-capacity erosion	Organisation's adaptive compliance capability	Dynamic congestion, displacement and deterioration	Requires observation of capacity and attribution, not a claim of burden alone

Analytical classification developed for this paper. The table compares concepts; it is not an empirical ranking.

4. Definition, unit of analysis and boundary conditions

Regulatory fatigue (proposition): a sustained, observable decline in an actor's capacity to interpret, prioritise, implement and substantively verify applicable requirements, attributable in material part to cumulative regulatory demands and their interactions relative to the actor's sustainable adaptive capacity.

Four conditions distinguish the construct. **Temporality:** a high steady workload is insufficient, the key variable is degradation, an inability to recover or increasingly costly maintenance of equivalent performance. **Attribution:** overload must contribute materially to the decline after plausible alternatives

are assessed. **Substance:** diminished protective capacity, timeliness, accuracy or recoverability matters more than subjective annoyance. **Relativity:** the same regulatory environment can affect differently a vertically integrated corporation, a small supplier, an accredited laboratory or a regulator.

The **regulated actor** may be an enterprise, a plant, a laboratory, a professional team or a competent authority. Comparisons must preserve the level of analysis from staff burnout, corporate financial distress and regulatory-capacity erosion are related but nonidentical variables. A stressed legal department does not automatically indicate poorer food safety. Equally, good historical audit scores do not guarantee that the organisation has preserved its future capacity to identify emerging hazards.

Boundary conditions matter. Greater documentation may *improve* safety where the baseline was weak, risks are poorly controlled, or evidence is irreproducible. High compliance costs may reflect intrinsically hazardous operations, not procedural dysfunction. Rapid regulatory updates may be justified by new scientific evidence. Where capacity can scale freely and information systems are interoperable, cumulative demands may not cause fatigue. The theory predicts the greatest vulnerability when skilled capacity is scarce, deadlines collide, the same facts are repeatedly translated between schemes, and managerial incentives favour visible documentation over difficult-to-measure prevention.

An organisation experiencing frequent non-conformities, but no distinctive compliance congestion may simply be poorly governed. A firm receiving additional justified inspections after prior breaches cannot cite those inspections as causal proof of regulatory excess. These counter-examples are essential to the construct's discriminant validity.

5. From obligations to congestion: the proposed multilevel mechanism

An obligation is not itself fatigue. The mechanism runs through six linked stages. **(1) Demand arrives:** a binding provision, guidance update, customer request, audit scheme, product-status question or information request creates work. **(2) Interpretation is required:** qualified personnel determine applicability, version, legal status, evidence, timing and accountability. **(3) Dependencies multiply:** the same data must be reconciled across products, markets, customers, sites and authorities. **(4) Bottlenecks form:** specialist attention becomes the limiting resource, producing queues and context switching. **(5) Substitution occurs:** organisations shift time toward urgent documentary outputs, sometimes away from preventive verification, method development or corrective-action closure. **(6) Adaptive capacity deteriorates:** backlogs, repeated errors, obsolete records, unreviewed deviations, delayed regulatory decisions or avoidable market-access delays accumulate.

The stages are a hypothesis, not a statement that this sequence always occurs. Two possible feedback loops deserve empirical attention. In an **assurance-demand loop**, documentary omissions provoke extra questionnaires or audit scrutiny, producing further requests and larger queues. In a **capacity-loss loop**, increased turnover or delayed training reduces regulatory expertise, making each subsequent obligation more expensive to interpret. A stabilising feedback loop is also possible when better process design or specialist hiring interrupts congestion and restores control.

Seven families of demand are analytically separable: **stock** (active, applicable requirements); **flow** (amendment velocity); **ambiguity** (interpretive uncertainty); **interaction** (cross-regime conflicts and dependencies); **assurance duplication** (multiple proofs of similar facts); **deadlines** (temporal concentration); and **spillover** (downstream private requirements). Each family can be individually benign but jointly consequential. A technical change introduced just before a customer audit, with a novel-food status query and supplier requalification, may consume more specialist capacity than its nominal requirements suggest.

Crucially, legal duties must not be flattened into the same category as private schemes. EU regulations and binding national measures occupy one level since the official guidance and competent-authority interpretations have different legal effects, accredited certification and GFSI-benchmarked programmes belong to the private assurance architecture, customer-specific questionnaires arise from contract and internally invented forms are managerial choices. The *aggregate* burden is real to the firm, but its sources, removal powers and public purposes differ.

This suggests a new analytical object: the **requirement–evidence network**. Requirements are nodes and evidentiary artefacts and technical controls are connecting resources: incompatibilities, repeated translation and approval hand-offs form costly edges. Counting legal instruments alone misses this network's architecture. For example, one validated allergen-control procedure may support several assurance objectives, whereas five separate formats for the same procedure create duplicative transformations. Whether each format adds assurance must be checked, rather than presumed.

6. The protective-displacement hypothesis and a conditional threshold

A firm's finite qualified capacity can be represented, for a specified interval, as:

$$C = D + V + I + U + R$$

where **C** is usable specialist capacity (qualified hours); **D** is direct hazard-control and validation work; **V** is necessary evidentiary verification; **I** is interpretation and regulatory-change management; **U** is avoidable duplication, reformatting and rework; and **R** is protected reserve for anomalies, training, emergencies and innovation. This identity is a classification rule for measured time, not a universal production function. Activities may contribute to multiple objectives, so analysts must allocate mixed tasks consistently and avoid double counting.

Let **P** denote protective performance, assessed using pre-specified indicators appropriate to the process; let **Q** denote information and assurance quality. Extra verification may initially raise both P and Q. The displacement proposition arises only where a constrained increase in U or low-yield V forces a reduction in D or R, and that reduction is consequential. The relevant inequality is **avoidable or low-yield assurance demand > sustainably available discretionary capacity**, subject to skill mix, process criticality and scheduling. There is no universal numerical threshold or optimal 'percentage of compliance hours'.

Proposition scale

P1 - Congestion. Conditional on task complexity and capacity, an increase in interacting reporting and interpretation requirements will be associated with longer decision queues and greater specialist task-switching than an equivalent increase in separable requirements.

P2 - Displacement. In capacity-constrained teams, increased low-marginal-value evidence preparation will be associated with reduced protected time for direct risk controls or corrective-action closure, unless resources or processes are adapted.

P3 - Conditional protection. Displacement will predict worse prospective protective indicators only where displaced work was causally important and alternative controls did not compensate. Zero outcome deterioration is therefore a meaningful possibility, not a falsification avoided by definition.

P4 - Interoperability. Reuse of evidence, with legally adequate provenance and access controls, will reduce administrative time more than merely reducing the number of published rules when proof duplication is the principal bottleneck.

P5 - Heterogeneity. The same demand shock will have larger congestion effects where specialist capacity is thin, product portfolios are diverse, and customer evidence formats are fragmented; proportional obligations and shared services may attenuate them.

P6 - Dynamics. Once queues, departures and version drift reinforce each other, capacity recovery may lag demand reduction. This potential hysteresis should be tested longitudinally, not assumed.

Taken together, these propositions articulate the **regulatory-fatigue paradox**: *additional procedural activity may coincide with weaker substantive control under identifiable scarcity and displacement conditions*. The proposition says neither that laws generally reduce safety nor that businesses should set their own acceptable health risks.

7. Diagnostic architecture: measuring rather than proclaiming fatigue

A composite score can hide hazardous states and falsely suggest scientific validation.

The proposed **Regulatory Fatigue Diagnostic (RFD)** is therefore a provisional **measurement dashboard**, not yet a validated scale or certification standard. It would track exposures, mediators, outcomes and controls separately before any score aggregation is considered.

Exposure domain A - Applicable burden. Measure binding obligation counts only after applicability screening, update frequency, mandatory reporting events, required response time, and the proportion of changes requiring specialist assessment. Preserve the legal source and date for each obligation.

Exposure domain B - Interaction and duplication. Capture overlapping evidence demands, number of transformations of the same validated fact, repeated questionnaires, incompatible definitions, regulator/customer hand-offs and the frequency of unresolved interpretive disputes. Not every overlap is wasteful: a second independent verification may detect error.

Mediator domain C - Capacity congestion. Record qualified-hours utilisation, queues, waiting time for decisions, burstiness of deadlines, task switching, vacancies, available reserve and rework caused by unclear requirements. Time records should be sampled and triangulated with objective logs, not reconstructed solely from management recollection.

Mediator domain D - Compliance quality. Measure version accuracy, traceability from obligation to evidence, reopened corrective actions, overdue revalidations, compliance decisions reversed after legal review, and audit findings attributable to documentation errors versus substantive control failures.

Outcome domain E - Protective performance. Sector-appropriate indicators include critical control excursions, time-to-containment, environmental-monitoring trends, allergen mislabelling events, recall readiness, validated method performance and closure of serious hazards. Rare outcomes require longer windows and appropriate statistical treatment; absence of incidents in one quarter is not proof of safety.

Outcome domain F - Innovation and access. Track time from technically complete dossier to decision, repeated classification queries, justified versus avoidable launch delay and allocation of qualified scientific hours to development. These variables are consequential but not substitutes for safety outcomes.

Context domain G - Confounders and equity. Product hazard level, plant size, market scope, automation, business stress, baseline management quality, previous breaches, customer mix, inspection history and workforce turnover must be recorded before comparative claims.

A cautious research design would initially avoid creating an 'RFD = 72/100' score. Exposure dimensions might form a **formative index** only after their weighting rationale, stability and external validity are tested.

Capacity and protective outcomes should remain separate, independently observed response variables: incorporating failures into the fatigue score and then correlating that score with failures would be circular. Research should assess inter-rater reliability, construct separation, cross-sector measurement invariance and predictive performance against simpler alternatives such as compliance hours, headcount or organisational maturity.

A plausible sentinel dashboard contains eight prospective variables: qualified-hours consumed by reformatting, median obligation-interpretation lead time, repeated evidence requests per validated artefact, overdue high-risk CAPAs, time reserved for hazard verification, evidence-to-source traceability completeness, serious process deviations and days of unresolved classification decisions. These should be interpreted jointly and stratified by risk. A 'good' score cannot excuse a critical food-safety failure.

Table 2. Proposed diagnostic domains and their evidentiary role

Domain	Example observation	Role in the hypothesised mechanism	Essential caveat
A. Applicable demand	Changes requiring qualified assessment per quarter	Exposure	Must be screened for legal applicability
B. Proof interaction	Eligible duplicate requests per evidence object	Exposure	Independent verification may be beneficial
C. Capacity queues	Median high-skill approval waiting time	Mediator	Also affected by staffing and demand
D. Information integrity	Incorrect versions in active evidence sets	Mediator/outcome	Attribute cause rather than assume overload
E. Risk protection	Overdue serious CAPAs; control deviations	Independent outcome	Predefine hazards and observation windows
F. Innovation	Rework before market authorisation	Secondary outcome	Do not conflate justified review with delay
G. Background conditions	Hazard severity, firm size, prior breaches	Confounders	Required for defensible comparisons

These indicators are proposed and not yet psychometrically validated. Exposure and outcomes should not be collapsed into a self-validating composite score.

8. Why food ingredients and biotechnology offer a stringent test

Food businesses operate under a layered structure in which some requirements are scientifically indispensable, others are public-law mechanisms for proving compliance, and still others arise through voluntary private governance or commercial purchasing conditions. Separating these layers is not an anti-regulatory position; it is necessary to understand which institution may legitimately change a requirement and what would be lost by doing so.

The protective floor. Regulation (EC) No 178/2002 establishes the architecture of EU general food law, including business responsibilities, traceability and risk analysis [10]. Regulation (EC) No 852/2004 requires relevant food business operators to implement procedures based on HACCP principles; Article 5 also addresses documentary evidence and the nature and size of the business [11]. Regulation (EU) 2017/625 provides for official controls according to risk and other specified considerations, including prior compliance and the reliability of an operator's own controls [12]. Commission Regulation (EU) 2021/382 strengthened food-safety culture requirements within the hygiene framework, including management commitment, communication, knowledge and adequate resources [13]. These are substantive and procedural safeguards that should not be removed merely to reduce paperwork.

The private assurance layer. Customer qualifications, third-party standards and GFSI-recognised certification programmes may support confidence and market access. GFSI's benchmarking architecture is expressly designed to facilitate recognition across programmes and reduce duplicate audits, but recognition is not itself a law or a stand-alone food-safety certification [16]. The practical outcome varies

with purchasers' acceptance, sectoral scope and contractual demands. A business may hold a respected food-safety certificate yet still be asked to complete separate questionnaires that differ only marginally from information already validated. Such duplication is a plausible exposure variable for our theory; its prevalence and effects must be measured rather than asserted as universal.

The innovation layer. For new ingredients, an early-stage classification error can invalidate substantial development work. Regulation (EU) 2015/2283 provides the EU novel-food regime; the European Commission administers authorisation, and EFSA performs scientific risk assessment where relevant [14]. EFSA's scientific guidance details expected information on identity, manufacture, composition, specifications, proposed uses and exposure [15]. The proper response to a demanding safety dossier is not to waive the safety assessment. It is to identify classification issues early, reuse methodologically valid information where lawful, and separate genuine new safety questions from repeated formatting exercises. The Union list under Implementing Regulation (EU) 2017/2470 remains part of market-status analysis [24].

Market-specific technical categorisation. The same substance or process may be evaluated differently depending on its intended function and legal context - for example, as a food ingredient, food additive, processing aid, feed material or chemical product. The classification cannot be resolved by trade names, supplier statements or a generic 'food grade' label alone. Separate rules may govern authorised food additives under Regulation (EC) No 1333/2008 and the provision of food information to consumers under Regulation (EU) No 1169/2011 [22,23]. Under the applicable legal tests, intended use, technological function, conditions of use, product category and the current market authorisations may matter. Some borderline questions demand expert judgment and competent-authority interaction. Here the fatigue mechanism can appear in repeated interpretive loops and delayed decisions, not necessarily in the volume of legislative provisions.

The scientific assurance layer. A challenge test, contamination-control validation or analytical method may underpin multiple disclosures. Its scientific validity depends on study design, matrices, relevant hazards, sampling uncertainty and appropriate acceptance criteria. It must not be rebranded as evidence for unrelated claims simply because an auditor asks for it. The opportunity for simplification lies in creating a **controlled evidence lineage**: one technical truth, versioned and accompanied by its scope and limitations, which may be referenced repeatedly without falsely broadening its evidential meaning.

Food-safety culture and paradox. This sector offers a useful counterargument to superficial 'less documentation' campaigns. Regulation (EU) 2021/382 explicitly requires management attention and adequate resources [13]. Thus, an enterprise that reduces proof activity while allowing weakened monitoring, poor training or neglected CAPAs would not be improving regulatory performance. Equally, a system that monopolises quality professionals with repetitive customer questionnaires could undermine the organisational attention that safety culture is meant to protect. The empirical question is how to reallocate work without eroding independently assessed controls.

9. Four industry pathways through which fatigue may arise

Scenario A - Multi-market ingredient supplier. Consider an SME exporting a technically unchanged ingredient to several jurisdictions. Its quality, regulatory and commercial teams maintain a master specification, analytical methods, contaminant results, traceability records and certificates. Customers request multiple template versions, partially incompatible supplier declarations and repeated 'confirmation' emails. Public-law obligations and legitimate customer risk assessments coexist, but the same facts pass through many evidence transformations. **Potential signal:** high specialist time spent on document formatting with stable underlying technical requirements. **Candidate intervention:** a machine-readable, version-controlled evidence passport plus customer-specific gap mapping. **Necessary safeguard:** no use outside the validated matrix, grade, intended function or jurisdiction.

Scenario B - Product innovation in a borderline category. A laboratory develops a novel process for an ingredient whose market status is uncertain. Regulatory analysis is deferred until pilot scale, after which repeated questions arise about intended use, history of consumption, processing effects and applicable authorisations. **Potential signal:** serial reversals in classification decisions, post-pilot evidence rework and launch-stage dead time. **Candidate intervention:** a pre-development classification gate, an auditable record of alternative legal pathways, and planned contact with competent authorities when needed. **Necessary safeguard:** scientific safety evidence cannot be treated as a substitute for an authorisation where one is legally required.

Scenario C - Food plant with layered audits. A plant runs HACCP, validated prerequisite programmes, internal checks, official controls and a private certification scheme. Several retail customers impose separate questionnaires and audit visits. **Potential signal:** corrective-action teams spend more time re-narrating already closed observations than testing whether corrective measures remain effective. **Candidate intervention:** audit-evidence mapping, recognition of equivalent independent evidence where allowed, and a jointly agreed annual audit calendar. **Necessary safeguard:** independent inspections, mandatory record retention and genuine additional risk controls remain intact.

Scenario D — Contract processor facing a sudden change. An amended rule, revised guidance or new customer specification affects a labelling or traceability process just as the enterprise is changing an ERP system. **Potential signal:** simultaneous version drift in artwork, specifications and internal work instructions, accompanied by high rework and slowed technical approval. **Candidate intervention:** a regulatory-change impact map linking affected SKUs, labels, SOPs, test plans and contracts, with controlled effective dates. **Necessary safeguard:** all products placed on the market must satisfy the legally applicable provisions at the relevant time.

None of these are reported field studies. They are **structured hypothetical mechanisms**. Their value lies in specifying observable data and interventions. Publishing them as real industrial results, or attaching invented percentages of savings, would compromise scientific reliability.

10. The hidden industrial economics: not all compliance costs are waste

Three economic categories must be separated. **Necessary control costs** are resources intrinsically required to achieve a legitimate protective outcome, such as microbial monitoring, validated allergen controls or hazard-specific testing. **Necessary assurance costs** provide proportionate evidence that these controls are functioning, including justified independent verification. **Avoidable transaction costs** are the increment due to preventable duplicate formatting, repeated evidence requests without incremental assurance, poor interoperability, unclear internal ownership or obsolete procedures. Only the third category is unambiguously a candidate for burden reduction, and the second may be optimised but requires proof-preserving design.

A fourth category is often underestimated: **regulatory option value**. Early classification expertise allows management to decide whether to invest in a product, pivot its intended use or defer an uncertain route before significant capital is committed. The benefit is not only fewer hours of paperwork, more than that, it is better allocation of R&D resources under legal constraints. Equally, a scientifically justified authorisation can create customer trust and protect credible innovators from unsafe competition. Regulation can therefore have positive economic option value as well as costs.

The enterprise-level calculation should be transparent. For a defined process, let annual avoidable document transformation hours be H_u , the fully loaded opportunity cost per qualified hour be c , and annual system maintenance costs be M . A narrow gross administrative saving from a re-engineering intervention is approximately $S_g = (H_{u,before} - H_{u,after}) \times c$, before deducting implementation costs M and independent validation expenditure. This estimates resource consumption or released capacity; it is

not automatically realised cash saving. Unless posts or external spending are genuinely reduced, its immediate value is more defensibly reported as **capacity redeployed**, not profit.

To estimate business benefit, one might additionally document avoided launch rework, faster approved dossier turnaround or fewer duplicate third-party audit fees. These elements should be attributed carefully and not added mechanically where they share causal components. Reduced development delay is not the same as guaranteed additional sales. A cost-of-delay estimate requires a justified opportunity scenario, demand evidence, gross-margin assumptions and uncertainty intervals.

Hypothetical arithmetic illustration (not industry evidence): assume a small technical team spends 14 qualified hours per week translating already validated data into customer-specific formats and can reduce this task to 6 hours through controlled reuse. At 46 operational weeks and an assumed fully loaded cost of €55 per qualified hour, the gross annual released capacity equals $(14-6) \times 46 = 368$ hours, equivalent to **€20,240 in imputed resource cost**. That number does not establish a commercial saving, validate the intervention or justify removing a single required verification. A pilot would first confirm the logged baseline, the actual time saved, implementation cost and unchanged technical adequacy.

An industrial board should therefore view fatigue through a dual dashboard: **resources spent proving compliance** and **protective performance achieved**. A reduction in documentary labour is a win only if mandatory evidence remains complete, scientific controls remain effective and external risks are not transferred to others. Conversely, a costly system may be entirely justified when high-consequence hazards warrant intensive verification.

11. SMEs, supply-chain transmission and the distribution of regulatory burden

Firm size matters through both fixed-cost absorption and asymmetric bargaining power. A central regulatory team at a large corporation may serve hundreds of products, while an SME may carry comparable classification and assurance duties with two specialists. A small upstream supplier may be formally exempt from a reporting scheme but must nevertheless provide information because major customers are covered by it. That is a **contractual transmission of regulatory information demand**, not automatic extension of the legal duty to the supplier. Confusing the two obscures what can be changed through contracts, sectoral standards or law.

This distinction is relevant to market access. Buyers sometimes use high-friction onboarding to screen suppliers for genuine risk, but duplicated forms may also impose disproportionate fixed costs without materially improving due diligence. The result can be concentration among firms able to finance assurance infrastructures. The hypothesis should be evaluated with procurement data, supplier exit patterns and credible controls for quality, scale and product risk; a general anti-competition effect cannot simply be presumed.

The fairness test must be two-sided. A small company should not claim lower testing standards for an equivalent severe hazard merely because resources are limited. But authorities and dominant purchasers should justify why evidence must be regenerated rather than safely referenced. Practical remedies include interoperable supplier declarations, standardised legal-status fields, proportionate reporting frequencies, secure reusable certificates and industry-recognised core datasets. Confidential formulas, trade secrets and personal data require need-to-know access, and the absence of disclosure is not automatically noncompliance where a less intrusive proof is legally sufficient.

A further burden can be shifted onto public authorities. Poor-quality or inconsistent submissions demand correction, generate requests for clarification and consume scarce inspection resources. Effective simplification is therefore not merely an industry efficiency agenda; good upstream evidence design can increase supervisory throughput and make justified enforcement more credible.

12. Regulatory fatigue in the public authority: supervisory congestion

Regulators also face finite analytic and inspection capacity. New duties, fragmented databases, emergency responses, complex product innovation and simultaneous enforcement priorities may result in backlogs or inconsistent interpretation. Yet delay or inconsistent guidance does not prove a regulator is at fault: scientific uncertainty, due process and proportional scrutiny may require time.

A symmetric public-sector version of the mechanism focuses on qualified-hours demand, case complexity, queueing and protective outcomes. Indicators could include time to resolve technically complete cases, variation in equivalent classifications, percentage of inspection effort devoted to the highest validated risks, corrected administrative decisions and timeliness of public guidance. Comparisons would control for staffing, statutory time limits and legal rights of appeal. Authorities should not be penalised for spending more time on intrinsically difficult or high-risk cases.

Regulation (EU) 2017/625 demonstrates that risk-based official controls are already embedded in EU food law [12]. This does not give private operators unilateral authority to substitute their preferred audit regime for mandatory official controls. It does, however, support careful consideration of verified risk information, historical compliance and reliable own-controls evidence within the legal framework. The objective is better targeting, not withdrawal of oversight.

Public reforms should prioritise the **stock-flow interface**: publish authoritative interpretations with version histories; assess the cumulative implications of amendments; resolve contradictory terminology; permit lawful reuse of verified information; provide transition periods appropriate to the underlying risk; and evaluate whether enforcement requests change behaviour or merely generate more files. The Commission's 2026 rulebook Communication gives a contemporary policy setting for such measures [3].

13. Risk-preserving simplification: the Protective-Floor Test

A simplification proposal should pass a written **Protective-Floor Test (PFT)** before implementation. It consists of six questions.

Purpose: Which specific legal, scientific, rights-based or contractual objective does the existing obligation serve? **Coverage:** Which hazards, populations or events would no longer be observed if the procedure were removed? **Equivalence:** Is the replacement evidence demonstrably equivalent in scope, freshness, independence and precision? **Integrity:** Does the new process preserve traceability, anti-fraud controls, record retention and competent-authority access? **Distribution:** Does reduced cost to the applicant transfer risk or burden to workers, consumers, competitors, regulators or vulnerable communities? **Reversibility:** Can the change be monitored, corrected or reversed if assurance deteriorates?

A proposal failing a material legal or protection test is not safe simplification, even if its accounting return is attractive. Some requirements cannot be consolidated because independent sampling or cross-checking provides genuine error detection. Conversely, if several parties request the same verified fact without a defensible incremental purpose, consolidation may preserve both public protection and private capacity.

The test corrects a recurring ambiguity in simplification debates. Regulation has multiple outputs as safety, transparency, equity, rights, environmental quality, legal certainty, innovation, legitimate competition. These are not always commensurable and should not be collapsed into an arbitrary single 'net score'. Where legislation establishes a binding threshold, the analyst's cost-benefit model cannot overrule it. Political institutions, courts and competent authorities retain the roles assigned to them by law.

The PFT can be used by a firm's compliance committee before retiring documentation, by a standards owner reviewing audit duplication, or by legislators considering ex post revision. It is a proposal for governance discipline, not a legally recognised EU instrument.

14. A controlled evidence passport: original solution for industrial assurance

One promising intervention is a **controlled evidence passport (CEP)**, a modular digital dossier that stores validated facts once while displaying exactly what each fact can and cannot prove. The passport is not a blanket 'single certificate for everything', nor a claim to immunity from later inspections. Its value depends on authoritative provenance, version control, reproducible methods, legally appropriate reuse and granular sharing permissions.

For each evidence item, the minimum metadata would be: stable identifier, product or site identity, legal and market scope, source category (binding law, official guidance, recognised standard or contract), responsible owner, scientific method and accreditation where relevant, matrix and sampling basis, analytical uncertainty or acceptance criteria where applicable, issued and expiry dates, supersession history, relation to the hazard or requirement, independent-review status and release permissions. Trade-secret material is separated from the results required to establish conformity.

A customer should be able to request *evidence that allergen-control procedures are current and effective* and receive a controlled, scoped response - not a full confidential formulation or a newly typed general declaration every month. Another customer needing substantively different market authorisation must receive a different legally accurate answer, not an indiscriminate export of the same document. The passport therefore permits **reusability without promiscuous equivalence**.

The proposed evidence architecture also produces useful research data: document reuse count, number of non-equivalent requests, lead time, rejection causes, corrections after source updates and hours of expert review. It would allow a process-level test of proposition P4. Its risks are equally real: common-mode errors, access-control failures, obsolete legal mappings and automation bias. Periodic source reconciliation and independent sampling are therefore integral to the design.

15. Artificial intelligence: opportunity, automation bias and the 2026 European setting

AI-assisted search and structured retrieval could lower the learning cost of regulation, distinguish legal sources from contractual materials, map changes across product portfolios and identify repeated information requests. It can also manufacture apparent certainty from incorrect versions, produce false classifications or overwhelm teams with alerts. The critical question is whether the tool reduces **verified decision work**, not merely whether it generates more regulatory content.

The current EU setting requires particular care. Regulation (EU) 2024/1689 establishes the AI Act [17]; Regulation (EU) 2026/1744, adopted on 8 July 2026 and published on 24 July, amended its implementation architecture through the Digital Omnibus on AI [18]. Any timetable or classification must therefore be checked against the current consolidated legal text and the actual use case, rather than relying on an old guidance summary. The existence of AI regulation does not imply that every compliance chatbot is itself a high-risk AI system; nor does an internal tool's limited scope remove privacy, confidentiality or professional accountability duties.

An enterprise system should be built around **retrieval from authoritative versioned sources, separation of extraction from legal determination, explicit uncertainty and escalation, cross-checking against current legal text, permissioned access to industrial information, and auditable human sign-off for material decisions**. NIST's voluntary AI Risk Management Framework provides an additional conceptual

tool for governance and measurement having account that it does not substitute for EU law [20]. ISO 37301 provides a separately relevant compliance-management reference [19].

Three failure modes warrant focused testing. **Alert multiplication:** systems may flag every amendment without classifying true applicability, increasing interruptions. **Confidence laundering:** a fluent summary may be treated as authoritative despite incomplete retrieval. **Document inflation:** automation may make it cheap to produce ever more questionnaires, generating systemwide costs for suppliers. For this reason, responsible AI deployment must measure verified error rates, false positives, time saved after expert correction and changes in protective performance. Automated output volume is not a benefit indicator.

16. An empirical programme capable of falsifying the theory

Study 1 - Construct development. Interview compliance leads, laboratory managers, safety specialists, procurement teams, auditors and public officials across firms of different sizes. Collect candidate items for exposures, bottlenecks and outcomes. Use cognitive interviewing to ensure that 'duplicate request' means substantively duplicative evidence, not legitimate independent controls. Document negative cases where intensive regulation improves outcomes without overload.

Study 2 - Multi-site observational panel. Follow a defined portfolio of food and biotechnology businesses across at least four repeated measurement waves. Sample obligation changes, staff allocation, turnaround, versions, CAPAs, nonconformities and relevant hazards. Prespecify the primary outcomes and analysis plan. Firm fixed effects may help with unchanging organisational characteristics, but time-varying management changes, demand shocks, customer mix and risk could still confound inference.

Study 3 - Controlled workflow intervention. Where legally possible, randomly allocate comparable evidence requests or supplier onboarding processes to a conventional documentation route or a traceable evidence-passport route, keeping substantive controls and reviewer authority constant. Primary endpoint: qualified minutes to produce **accepted, accurate evidence**. Secondary endpoints: correction rate, regulator/customer rejection, controlled-hazard coverage, security incidents and review time. A faster but less accurate process would fail.

Study 4 - Regulatory amendment natural experiment. Observe exposure to a clearly dated rule change across affected and unaffected product groups, or firms differentially exposed for exogenous reasons. Difference-in-differences or event-study approaches require parallel-trend and anticipation checks as well as enactment timing may coincide with other shocks. Any causal interpretation must remain bounded by identification quality.

Study 5 - Regulator–enterprise dyads. Pair case-handling data on both sides of the information exchange, with confidentiality safeguards, to identify whether delays arise primarily from submission defects, legal ambiguity, supervisory queues or iterative rework. The unit of analysis is the actual decision pathway, not an opinion about regulation generally.

Construct validation should proceed in stages: content validity and expert agreement, reproducibility of time classification, convergent and discriminant relationships with workload and management quality, predictive validity for delays and reserve capacity, and then independent tests of protection. A formative score, if ultimately justified, requires sensitivity analysis on weights and missing data, plus separate reporting of critical failures. Sample-size calculations must use pilot estimates of variation and clustering, inventing a universal n without an outcome variance or effect target would be inappropriate.

Prespecified competing hypotheses: (A) higher obligations reflect greater underlying hazard and therefore justifiably predict better safety; (B) poor management causes both rework and safety failures, making regulation an innocent correlate; (C) customers' assurance requests reveal genuine previous weaknesses and improve supplier performance; (D) better-capitalised companies adopt more auditing

precisely because they have more protective capacity. A credible publication must allow each possibility to explain the data.

A particularly strong falsification test would observe enterprises with high cumulative legal and private requirements but stable or improving direct controls after adjustment for hazard, task mix and resources. If fatigue is present regardless of performance or attribution, the construct has become unfalsifiable and loses scientific value.

17. Limitations, normative safeguards and distributional consequences

This theory is presently unvalidated. Its causal diagram is a proposed explanation, its diagnostic a measurement plan, and its industry pathways illustrative. Any claim that avoidable verification routinely causes safety failures would exceed the available evidence. The terminology may not travel unchanged across jurisdictions with different legal systems, and the capacity categories may overlap in real workflows. Surveys of regulatory annoyance are vulnerable to strategic reporting, recall error and ideological framing. Rare safety events are difficult to use as short-term endpoints. The observed positive effect of intense oversight can reverse a simplistic expectation that fewer reports automatically make enterprises safer.

Distributional analysis is indispensable. A manufacturer may prefer lighter reporting while the information is essential to consumers, enforcement officers or workers. A large buyer's standardised data request may seem costly to a supplier while legitimately addressing fraud risk in an opaque chain. Small-firm proportionality cannot negate the statutory duties necessary to prevent serious harm. The theory therefore incorporates *protective non-regression* as a condition on interventions, with transparent consideration of who gains and whose risks change.

A related policy hazard is opportunistic appropriation. '**Regulatory fatigue**' could become a slogan against environmental limits, food standards, labour protections or evidence obligations that powerful actors dislike. The concept should resist that use. It offers a way to improve **implementation efficiency and protective credibility**, not a normative entitlement to avoid obligations. Effective public rules can strengthen markets and prevent costly harm and any assessment must value their benefits alongside burdens.

The argument is consequently neither 'more regulation is better' nor 'less regulation is better'. It is that the translation of legal objectives into industrial performance is subject to bottlenecks, evidence quality, attention constraints and incentive effects. These must be managed with the same rigour applied to technical hazards.

18. Conclusion: regulation should preserve the capacity to protect

A mature regulatory system cannot be evaluated only by the number of statutes, inspections, forms or certificates it produces. The relevant question is whether actors reliably understand their obligations, document genuine compliance and possess sufficient qualified capacity to detect and control real risks. Complexity may be justified; duplication may sometimes be protective and rigorous verification is indispensable where hazards are consequential. Yet a regime that repeatedly converts scarce scientific and supervisory attention into low-value procedural transactions risks eroding the capacities on which its own legitimacy depends.

The contribution of this paper is a conditional theory: **regulatory fatigue is not regulatory volume but sustained capacity erosion, potentially mediated by compliance congestion and protective displacement**. The argument is testable and can fail. It anticipates contexts where more verification

improves performance, organisations whose primary problem is poor management, and reforms whose apparent savings merely transfer costs or risks elsewhere.

For policymakers, the proposed direction is cumulative-impact review, legal clarity, interoperable evidence and risk-based administration. For industry, it is a classified obligation register, measurable bottleneck management, disciplined evidence reuses, source-aware regulatory technology and non-negotiable preservation of hazard controls. For researchers, it is a programme that observes obligations, capacity and outcomes separately, compares competing explanations and evaluates interventions with legitimate counterfactuals.

The future of effective regulation is not an endless multiplication of proofs; it is the disciplined preservation of the organisational and institutional capacity required to make protection real.

Declarations and scholarly status

Article classification: Original conceptual/theory-development paper, with prospective empirical research design. **Data availability:** No original empirical dataset was generated or analysed for this manuscript. **Quantitative values:** Any numerical illustration is explicitly hypothetical. **Conflicts of interest:** The author is academically affiliated with different educational institutions and professionally with an industrial research and ingredient-technology organisation. This professional standpoint motivates the research question but does not substitute for independent evidence, and journal-specific declarations must be completed by the author before submission. **Funding:** To be confirmed by the author before submission; no funding claim is made here. **Ethics:** No human subjects were recruited, and no confidential industrial or personal records were used in this conceptual article. **Sources review cut-off:** 12 October 2026. **Publication status:** Author-review manuscript.

References

- [1] OECD. (2025). *OECD Regulatory Policy Outlook 2025*. OECD Publishing. <https://doi.org/10.1787/56b60e39-en>
- [2] OECD. (2026). *Smart Regulations, Strong Business: Enabling Growth and Societal Protection in a Fast-Changing World*. OECD Publishing. <https://doi.org/10.1787/93d38770-en>
- [3] European Commission. (2026, 28 April). *A Simpler, Clearer and Better Enforced EU Rulebook*. COM(2026) 380 final. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0380>
- [4] Bozeman, B. (1993). A theory of government “red tape”. *Journal of Public Administration Research and Theory*, 3(3), 273–304. <https://doi.org/10.1093/oxfordjournals.jpart.a037171>
- [5] Herd, P., & Moynihan, D. P. (2018). *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation. <https://doi.org/10.7758/9781610448789>
- [6] Meyer, J. W., & Rowan, B. (1977). Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology*, 83(2), 340–363. <https://doi.org/10.1086/226550>
- [7] Power, M. (1997). *The Audit Society: Rituals of Verification*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198296034.001.0001>
- [8] Mendoza, J. P., Dekker, H. C., & Wielhouwer, J. L. (2016). Firms' compliance with complex regulations. *Law and Human Behavior*, 40(6), 721–733. <https://doi.org/10.1037/lhb0000215>
- [9] van Duin, S. R., Dekker, H. C., Mendoza, J. P., & Wielhouwer, J. L. (2026). Responding to regulatory feedback: Financial capacity, revenue expectations, and firms' responses to the authority's recommendations. *Regulation & Governance*, 20(3), 1194–1208. First published online 1 June 2025. <https://doi.org/10.1111/rego.70034>
- [10] European Parliament and Council. (2002). Regulation (EC) No 178/2002 laying down the general principles and requirements of food law. <https://eur-lex.europa.eu/eli/reg/2002/178/oj>

- [11] European Parliament and Council. (2004). Regulation (EC) No 852/2004 on the hygiene of foodstuffs, Art. 5. <https://eur-lex.europa.eu/eli/reg/2004/852/oj>
- [12] European Parliament and Council. (2017). Regulation (EU) 2017/625 on official controls, Art. 9. <https://eur-lex.europa.eu/eli/reg/2017/625/oj>
- [13] European Commission. (2021). Commission Regulation (EU) 2021/382 amending the Annexes to Regulation (EC) No 852/2004 as regards allergens, food redistribution and food safety culture. <https://eur-lex.europa.eu/eli/reg/2021/382/oj>
- [14] European Parliament and Council. (2015). Regulation (EU) 2015/2283 on novel foods. <https://eur-lex.europa.eu/eli/reg/2015/2283/oj>
- [15] EFSA. (2024). Guidance on the scientific requirements for an application for authorisation of a novel food in the context of Regulation (EU) 2015/2283. *EFSA Journal*, 22(9), e8961. <https://doi.org/10.2903/j.efsa.2024.8961>
- [16] Global Food Safety Initiative. (2024). *Benchmarking Requirements v2024: Implementation Handbook, Parts II–III*. https://mygfsi.com/wp-content/uploads/2024/12/Benchmarking_Requirements_v2024_Implementation_Handbook_Part_II.pdf
- [17] European Parliament and Council. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [18] European Parliament and Council. (2026). Regulation (EU) 2026/1744 of 8 July 2026 (Digital Omnibus on AI). <https://eur-lex.europa.eu/eli/reg/2026/1744/oj>
- [19] International Organization for Standardization. (2021). *ISO 37301:2021 — Compliance management systems: Requirements with guidance for use*. <https://www.iso.org/standard/75080.html>
- [20] Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). <https://doi.org/10.6028/NIST.AI.100-1>
- [21] OECD. (2025). Time for a regulatory reset? In *OECD Economic Outlook, Volume 2025 Issue 2*. https://www.oecd.org/en/publications/oecd-economic-outlook-volume-2025-issue-2_9f653ca1-en/full-report/time-for-a-regulatory-reset_90ca6147.html
- [22] European Parliament and Council. (2008). Regulation (EC) No 1333/2008 on food additives. <https://eur-lex.europa.eu/eli/reg/2008/1333/oj>
- [23] European Parliament and Council. (2011). Regulation (EU) No 1169/2011 on the provision of food information to consumers. <https://eur-lex.europa.eu/eli/reg/2011/1169/oj>
- [24] European Commission. (2017). Commission Implementing Regulation (EU) 2017/2470 establishing the Union list of novel foods. https://eur-lex.europa.eu/eli/reg_impl/2017/2470/oj
- [25] Likas, K. (2026). AML compliance congestion as a factor in trade-based money laundering. *Journal of Financial Regulation and Compliance*. Published online 24 July 2026. <https://doi.org/10.1108/JFRC-04-2026-0206>

FROM REGULATORY OVERLOAD TO PROTECTIVE CAPACITY

Industrial Implementation Companion to “Regulatory Fatigue as Capacity Erosion”

by José M. López

INTABIOTECH Knowledge Center Scientific Director

Technical Management and Compliance Architecture · A practical companion for industrial review

Purpose and limitations

This companion translates a proposed scientific theory into an auditable improvement programme for industrial firms, laboratories, ingredient manufacturers, quality directors and supply-chain teams. It is not a substitute for binding obligations, legal advice, validated process controls or competent-authority decisions. The proposed procedures are management tools, not ISO requirements or recognised regulatory exemptions. The worked numbers are hypothetical. A high-risk operation may legitimately need more, rather than less, documentation.

The central distinction is between **necessary hazard control**, **necessary evidence of that control**, and **avoidable transactions generated while supplying that evidence**. The principal goal is to recover expert time from the third category without damaging the first two.

A. A seven-step implementation architecture

Step 1 - Define the business perimeter. Limit the pilot to a plant, product family, jurisdictional route or supplier onboarding workflow. Name the legal entity, site, covered processes, markets, customer groups and effective date. Do not assemble a meaningless multinational master register that no one can own.

Step 2 - Classify each demand by authority. Use five classes: L = binding law/authoritative administrative duty; O = official nonbinding guidance or published interpretation, with the caveat that its legal effect is jurisdiction-specific; C = contractual customer condition; S = voluntary or privately certified assurance standard; I = internal company requirement. Some controls support several classes; the classification indicates provenance, not priority by itself.

Step 3 - Map the evidence lineage. Identify which result, record, validation report or trained decision establishes each fact. Specify scope, scientific limitations, expiry and confidentiality. Map repeated requests to a common evidence asset only if their scope is equivalent.

Step 4 - Measure qualified time. Sample actual work over four to six weeks. Separate application/legal interpretation, direct risk management, necessary proof, avoidable reformatting and corrective rework. Use time diaries, ticket-system logs and document revision histories. Do not report guessed time as measured savings.

Step 5 - Prioritise bottlenecks. Look for evidence transformations repeated more than once, open status-classification queries, simultaneous audit calendars, approval queues and recurring missing-context questions. Rank candidate reforms by expected recovery *and* by protection risks.

Step 6 - Apply the Protective-Floor Test. Before changing a workflow, show that there is no deterioration in: statutory compliance, critical control coverage, method validity, independent verification where required, record retention, traceability, fraud resistance or access by the competent authority.

Step 7 - Run a controlled pilot. Pre-register primary efficiency and safety endpoints, appoint an independent reviewer, implement on a bounded scope, monitor errors and stop if protective criteria fail. Report released capacity and cash savings separately.

B. Template 1 - Source-classified obligation register

An actual register should hold one record per operationally actionable requirement, not one record per entire regulation. The minimal fields are:

Field	Operational instruction
Unique obligation ID	Permanent identifier, even after an amendment
Legal/contractual class	L, O, C, S or I; name document owner
Exact source	Regulation and provision / customer contract / standard clause
Source version and effective date	Distinguish publication date from applicability date
Scope	Product, process, facility, market, role and exemption analysis
Underlying purpose	Hazard, right, market integrity or verified contractual objective
Owner and approver	Functional accountable person and independent legal/technical reviewer
Control	Physical, scientific, managerial or information action taken
Evidence and quality criteria	Validated methods, record fields, acceptance criteria and retention
Review and trigger	Revision cycle; change of use, customer, process, law or risk
Confidentiality	Public, controlled, trade secret, personal data or authority-only
Status and exception	Compliant / action required / disputed / not applicable with grounds

Hard stop: 'Not applicable' must be justified by a precise scope determination, not by the firm's desire to avoid a cost. A supplier's regulatory opinion is not an official market authorisation.

C. Template 2 - The request-to-evidence concordance matrix

Create a row for every recurring request and a column for the validated evidence assets it can legitimately use. The first four columns are **requester**, **requested fact**, **legal/contractual basis** and **freshness/scientific scope**. The next columns are **existing evidence identifier**, **equivalence confirmed (yes/no/partial)**, **additional information genuinely necessary**, **reviewer**, and **reusable output format**.

This matrix separates three situations. **True duplication:** evidence has identical meaning, boundaries, freshness and authority; reuse is usually technically feasible. **Partial overlap:** one dataset serves multiple requests, but each needs an additional independent field or judgment; modular responses avoid recreating everything. **False equivalence:** products, uses, jurisdictions, methods or test matrices differ as harmonising the files would erase a real distinction and may be unsafe or unlawful.

The objective is not to force every customer to accept one template. It is to allow the organisation to respond correctly from a validated data source, avoiding unnecessary repeated expert interpretation.

D. Template 3 - Evidence passport minimum viable structure

An evidence passport can start as a controlled set of documents and metadata, without buying a new software platform. A minimum data structure consists of:

- **Identity:** evidence ID, legal entity, site, product/grade, document version, owner and reviewer.
- **Claim:** exactly what the document establishes, and what it does not establish.
- **Applicability:** target jurisdictions, legal/contractual basis and any product-status decisions.
- **Science:** method identity, scope, measurement uncertainty if relevant, sampling and limitations.

- **Integrity:** who approved it, signature record, hash/version, source of original results and supersession history.
- **Lifetime:** issue date, review trigger, expiry or review date, change log and archived prior versions.
- **Sharing:** permitted recipients, confidentiality tier, redaction rule and record of external delivery.

An example of an acceptable scoped claim is “*Analytical result for batch X, method Y, matrix Z, sample obtained on date D, with stated measurement uncertainty.*” A prohibited extrapolation would be “*All future batches and all intended uses are safe or lawful*” based on that isolated certificate.

The passport should link evidence requests to source content using stable identifiers. Artificial intelligence can assist with matching, but only a qualified reviewer should authorise scope equivalence for high-consequence documents. Secrets should be disclosed only to the degree necessary and legally required.

E. KPI definitions with denominators

Evidence rework rate: requests returned for missing or inaccurate information ÷ requests completed; report by source and reason. **Evidence reuse share:** accepted responses supplied from a controlled pre-existing asset ÷ all eligible responses; eligibility must exclude materially different cases. **Interpretation lead time:** median days from fully specified question to approved regulatory determination, with a high-percentile measure for difficult cases. **Protected technical capacity:** recorded qualified hours reserved for direct hazard control, corrective action, training and validation ÷ planned qualified hours. **Overdue critical CAPA share:** overdue high/critical corrective actions ÷ all due high/critical corrective actions; never interpret zero overdue as safety in isolation. **Source currency:** evidence objects checked against current authoritative source ÷ objects scheduled for source reconciliation. **Documentary false-positive rate (AI):** inapplicable alerts ÷ all reviewed alerts, supported by an adjudicated sample. **Protection outcome indicators:** process-specific microbiological, chemical, allergen, traceability, labelling or equipment verification endpoints, chosen before intervention.

KPIs must not be optimised in isolation. An increase in evidence reuse that also increases rejection rates is not automatically successful. A reduction in CAPA count that results from reclassification of major findings is not improvement. Use a balanced scorecard with at least one process-reliability measure, one assurance-integrity measure and one material hazard indicator.

F. Four decision gateways for new products and ingredients

Gate 0 - Concept and lawful intended use. Identify the proposed ingredient, processing function, target food categories, intended claims, jurisdictions and customer applications. Determine who is responsible for regulatory classification. Do not start with the assumption that a designation such as 'natural' or 'food grade' guarantees legal suitability.

Gate 1 - Classification and regulatory pathway. Decide whether existing provisions, specific authorisations, novel-food rules, additive requirements, feed rules or national conditions are relevant. Document alternatives and unresolved issues. Where an official view is required, record the authority and submission route; private interpretations do not replace binding decisions.

Gate 2 - Scientific validation. Specify required data on identity, composition, relevant contaminants, stability, safety and demonstrated technological effect. Assess study design, methods, representative matrices and uncertainties. This is not a cost-reduction gate: it prevents poorly targeted experiments and invalid extrapolation.

Gate 3 - Industrialisation and market release. Reconcile validated specifications, change control, batch evidence, labelling, supplier documentation, certificates and customer contractual requirements. Release only after lawful market status and required safety controls have been demonstrated. Use a single approval package to avoid conflicting declarations to different departments.

For a demanding novel-food pathway, consult the EU novel-food legal instruments and the current EFSA scientific and administrative guidance. Even where several trials share a dataset, the evidentiary requirements are not automatically identical. The purpose of gates is to prevent costly late-stage surprises while maintaining the highest relevant protection standards.

G. Worked example - a supplier questionnaire redesign (assumed data)

Assume a supplier receives 42 customer evidence requests per month, 25 request facts demonstrably identical to material already independently verified, nine have partial overlaps and eight require genuinely new scope or review. Under the current workflow each request takes, on average, 70 qualified minutes to complete. A controlled workflow would aim to reduce work on eligible requests while maintaining expert checks for partial and novel requests.

The baseline is $42 \times 70 = 2,940$ qualified minutes/month, or 49 hours/month. If the redesigned process takes 25 minutes on each of the 25 true-duplicate requests, 55 minutes for each of nine partial cases and the unchanged 70 minutes for each of eight distinct cases, total time becomes $25 \times 25 + 9 \times 55 + 8 \times 70 = 1,680$ minutes/month, or 28 hours/month. The assumed difference is 21 qualified hours per month, before system maintenance and quality-assurance time.

A pilot must still verify that the 25 requests truly share scope, that response accuracy does not worsen and that any reduced time is genuinely realised. It must also track corrections, customer acceptance, document version failures and the impact on staff time available for hazard control. The example demonstrates an evaluation method, **not observed efficiency at any named company**.

H. Practical governance: responsibilities and stop criteria

Board/general management: approves scope, resources, acceptable risk within legal limits, and escalation of unresolved critical obligations. **Regulatory affairs/legal:** decides applicability and maintains authoritative legal mappings. **Quality/food safety:** owns hazard controls, protective outcomes and scientific evidence adequacy. **R&D/process engineering:** documents intended function, design changes, validation and practical production constraints. **Procurement/sales:** classifies customer information requests and negotiates sensible formats without misrepresenting legal status. **IT/data security:** owns versioning, access control, backups and reproducible changes. **Independent reviewer:** verifies protection is preserved before simplification goes live.

The pilot stops or rolls back if it causes a missing mandatory record, inaccurate product-status determination, uncontrolled disclosure of intellectual property, degradation of a critical process control, unresolved high-severity data integrity defect or a statistically/technically important rise in accepted evidence errors. A minor speed gain is never a reason to conceal such a result.

I. A 90-day roadmap, with a 12-month validation horizon

Days 1-30: choose scope, capture actual sources, requests and technical workload, map process hazards and existing protective metrics, define baseline and review governance. Deliverables: source register, evidence inventory, no-change protection baseline and pilot protocol.

Days 31-60: classify true and false duplicate requests, build controlled evidence assets, agree release permissions and ownership, configure internal change alerts, test representative customer cases. Deliverables: approved evidence lineage, first reusable data packs, source reconciliation and an intervention risk assessment.

Days 61-90: run bounded parallel or controlled rollout, check rejection rates, qualified minutes and protected controls, close discrepancies, document remaining uncertainties. Deliverables: accepted-evidence dataset, qualified-hours analysis and an independent protective-floor review.

Months 4-12: measure persistence, version drift, high-severity findings, corrective-action timeliness and meaningful protective trends, broaden only after evidence supports doing so. For rare hazards, a 12-month observation period may still be insufficient. The time horizon must be justified by the consequence and incidence of the events being monitored.

J. Guidance for public authorities and standards owners

The same model suggests a regulator-facing agenda: coherent obligation IDs, publication of consolidated current requirements and authoritative interpretations, formal consideration of cumulative-impact pathways, secure once-only submission where legislation allows, measured response times stratified by case complexity, recognised independent assurance where legally relevant, and review of whether additional information requests actually improve decisions. Standards owners and large buyers could publish minimum common information schemas and precise reasons why their requirements depart from equivalent validated evidence.

This does not imply fewer inspections where risk is high. An effective supervisory regime will sometimes **increase** inspection frequency, independent sampling or documentation where failures or fraud make that necessary. Simplification is conditional on equivalent or better assurance, not a presumption of reduced scrutiny.

K. Research transparency and non-claims

The diagnostic framework, evidence-passport architecture, protected-capacity measures and pilot designs are original management proposals derived from the accompanying conceptual analysis. They have not been subjected to controlled industrial validation in this document. No real customer name, identifiable supplier data, proprietary manufacturing procedure or trade-secret formulation is disclosed. The figures are worked hypothetical examples, not reported industrial savings. Independent scientific and regulatory review remains necessary before implementation in critical applications.

Author's concluding operational principle

The most useful question for a technical director is not, ***“How many compliance documents can we eliminate?”*** It is, ***“Which qualified activities directly secure lawful, safe outcomes, what verified evidence demonstrates them and which avoidable transactions are preventing our specialists from doing their most valuable work?”***

That question aligns commercial efficiency with the legal and scientific purpose of regulation.