When Compliance Becomes Exclusion:
Proportionality, Liability Legitimacy, and the Governance of Third-Party AI Fine-Tuning

leyi zhang
Independent Researcher
smxjudger@proton.me
ORCID 0009-0006-9191-6285
July 2026

Abstract

Third-party fine-tuning occupies an awkward position in contemporary AI governance. The practice is consequential enough to warrant independent legal scrutiny—fine-tuners intervene in a model's behavioral architecture in ways that foundation-model developers neither control nor fully anticipate—yet the compliance obligations now attaching to this activity have been designed with a different kind of actor in mind. The documentation requirements, dynamic auditing obligations, and ethics-review procedures that have become standard features of AI liability frameworks assume actors with stable legal infrastructure, dedicated compliance personnel, and resources sufficient to absorb recurring governance costs. Small and medium-sized enterprises rarely fit this description.

The problem is not that regulation is too demanding in some abstract sense. The problem is that formally identical obligations produce radically unequal practical burdens. When compliance intensity bears no relationship to the scale or probability of harm that a given deployment generates, the governance system has stopped doing what liability systems are supposed to do. It has begun selecting market participants according to institutional endurance rather than risk exposure.

This article argues that proportionality should serve as the normative foundation for calibrating liability in the AI fine-tuning ecosystem—not as a flexible policy preference that regulators may or may not choose to apply, but as a condition of liability legitimacy itself. Drawing on the EU's risk-tiered framework, the fragmented American model, and China's emerging "regulate through support" approach, the article identifies the differentiated obligation structures that already exist within mature governance systems and asks why they have not been more coherently theorized. The article then proposes a liability architecture built around four elements: duty structures calibrated to deployment risk rather than organizational characteristics; regulatory thresholds anchored in application-side influence indicators rather than development-side proxies; graduated transition mechanisms that allow compliance obligations to track institutional growth; and safe harbor protections conditioned on participation in collective governance structures. The last of these is the most important. Without public compliance infrastructure—shared audit pools, standardized templates, centralized ethics review—proportional liability remains a formal commitment without institutional substance.

Keywords: AI fine-tuning liability; capability reconfiguration; proportionality principle; SME compliance burden; safe harbor design; public compliance infrastructure; regulatory legitimacy; AI governance architecture

I. The Problem That Capability Reconfiguration Theory Left Unresolved

Something significant happened to AI governance theory when fine-tuning became a commercially prevalent practice. Until fairly recently, the liability question in AI was primarily a question about the foundation-model developer: who built the system, what capabilities did they choose to embed, and how should responsibility attach when those capabilities cause harm. Fine-tuning complicated this picture in ways that the existing frameworks were not designed to handle.[1]

The concept of "capability reconfiguration liability" represents the most serious attempt to theorize the fine-tuner's distinctive legal position.[2] Under this approach, fine-tuning is understood as a transformative intervention in a model's behavioral architecture. Fine-tuners do not merely use a tool; they reshape what the tool can do, adjusting its outputs, expanding or restricting its operational range, and introducing application scenarios that the original developer neither anticipated nor approved. This intervention, the argument runs, generates independent legal responsibility. Fine-tuners are not downstream users seeking indemnification from upstream developers; they are co-designers of a system whose behavior they have materially altered.

The argument is persuasive, and this article accepts it. But accepting it creates a further problem that the capability reconfiguration framework has not yet confronted directly. Identifying the right subject of liability says nothing about how intensely that liability should apply. It does not tell us whether a three-person startup fine-tuning a medical inquiry model for a regional hospital should face the same compliance architecture as a large enterprise deploying a general-purpose fine-tuned system to millions of users. The capability reconfiguration framework treats fine-tuners as a legally homogeneous category. The commercial reality of fine-tuning ecosystems is that they are anything but.

---

[1]On the tripartite structure of the contemporary AI supply chain, see generally Andres Guadamuz, "The Artificial Intelligence Supply Chain: Legal Perspectives," 12 Journal of Internet Law 1 (2023). The distinction between foundation models, fine-tuners, and deployers has gained increasing analytical prominence in both academic and regulatory discourse; see also EU AI Act, Regulation (EU) 2024/1689, Recitals 97–101 (addressing responsibility allocation across the AI value chain).

[2]The concept of "capability reconfiguration liability" is developed in Wang Qingsong, Risk Allocation and Liability Structures in Large-Model Applications, *Eastern Law*, No. 4 (2025). Wang's framework treats fine-tuning not as ordinary product use but as a secondary act of capability design, analogous in some respects to the modification of industrial equipment by an aftermarket operator.

The practical consequences of this gap deserve attention.[3] Most AI governance frameworks—including the obligations now commonly associated with capability reconfiguration liability—assume that regulated actors possess stable legal infrastructure, dedicated compliance personnel, and institutional resources sufficient to absorb dynamic auditing requirements, documentation obligations, data-governance protocols, ethics-review procedures, and post-deployment monitoring duties. These assumptions are reasonable for large platform firms and major model providers. They bear increasingly little resemblance to the operational conditions of the small and medium-sized enterprises that account for a substantial and growing share of commercial fine-tuning activity.

When obligations designed for one kind of actor are applied without differentiation to another, the result is a form of regulatory asymmetry: the compliance burden is nominally equal but practically unequal, calibrated to institutional scale rather than to the risk that regulation is supposed to govern. The deeper problem is that this asymmetry is not merely inequitable in some distributional sense. It threatens the normative coherence of the liability regime itself. A system that constrains actors primarily because they lack institutional endurance, rather than because they generate disproportionate harm, has lost its justificatory connection to the purposes that give liability law its legitimacy.

The argument of this article proceeds as follows. Part II examines the mechanisms through which uniform compliance obligations generate structural displacement pressure within SME fine-tuning ecosystems. The aim is not simply to document that compliance is expensive for small firms—this is obvious—but to identify the specific ways in which fixed-cost governance structures distort behavior, concentrate innovation within dominant platform ecosystems, and ultimately reproduce forms of structural inequality that bear no relationship to the risk-governance rationale of AI liability. Part III develops the proportionality principle as the appropriate normative foundation for addressing this problem, focusing on the dimension of proportionality stricto sensu, where the analytical work is most consequential. Parts IV and V examine how the EU's risk-tiered framework and China's emerging "regulate through support" model each incorporate differentiated governance structures, and what their respective limitations reveal about the conditions for proportional liability to function in practice. Part VI proposes a differentiated liability architecture. Part VII argues that this architecture cannot succeed without the construction of public compliance infrastructure capable of transforming compliance from an individual burden into a shared institutional service.

II. How Uniform Obligations Generate Structural Exclusion

---

[3]The distributional consequences of formally neutral regulatory obligations have received sustained attention in the broader regulatory literature. See Robert W. Hahn & John A. Hird, The Costs and Benefits of Regulation: Review and Synthesis, 8 Yale Journal on Regulation 233, 251–58 (1991). In the AI-specific context, the asymmetric impact of compliance costs on smaller actors is documented in Roberto Garrone, Proportionate Cybersecurity for Micro-SMEs: A Governance Design Model Under NIS2, 5 Journal of Cybersecurity and Privacy 4 (2025).

A. The Fixed-Cost Problem

The difficulties SME fine-tuners face are sometimes framed as a simple matter of cost: regulation is expensive, and smaller firms have less money. This framing understates the structural character of the problem. The issue is not that compliance costs are high in general; it is that the cost structure of contemporary AI compliance obligations does not scale with organizational size or deployment risk.

Capability reconfiguration liability typically requires fine-tuners to undertake pre-deployment risk assessment, monitor training data quality and provenance, evaluate the effects of parameter changes, maintain documentation records sufficient for regulatory review, implement post-deployment monitoring systems, and develop remediation procedures for identified harms. Each of these obligations requires institutional infrastructure—legal personnel, technical audit capacity, documentation management systems, governance protocols—that must be established at roughly fixed cost regardless of how small or narrowly scoped the fine-tuning activity is. A hospital deploying a lightly fine-tuned administrative assistant model faces structurally similar compliance requirements to a technology company deploying a general-purpose fine-tuned system at consumer scale. The similarity is a product of the obligation structure, not the risk profile.

For large firms, compliance infrastructure can be distributed across existing legal departments, amortized over multiple deployments, and integrated into compliance systems already maintained for other regulatory purposes. For SMEs, each of these elements represents a new fixed cost that must be absorbed within constrained operational budgets. The consequence—documented across analogous regulatory domains from financial services to cybersecurity—is a displacement effect: resources that would otherwise support research and development are progressively diverted toward governance administration.[4] Experimental iteration slows. Tolerance for failure declines. Innovation pathways narrow. In sectors where the ability to iterate quickly matters enormously, this contraction effect can be decisive.

What makes this distortion particularly troubling is that it does not track the risk distribution that AI liability is supposed to address. SMEs operating narrow, application-specific fine-tuning deployments with limited user populations and highly controlled harm scenarios are not, on any defensible risk-assessment methodology, the primary source of systemic risk in AI ecosystems. The compliance burden they bear is disproportionate not because their governance obligations are especially demanding in isolation, but because uniform obligation structures interact asymmetrically with unequal institutional capacities.

---

[4]The distributional consequences of formally neutral regulatory obligations have received sustained attention in the broader regulatory literature. See Robert W. Hahn & John A. Hird, The Costs and Benefits of Regulation: Review and Synthesis, 8 Yale Journal on Regulation 233, 251–58 (1991). In the AI-specific context, the asymmetric impact of compliance costs on smaller actors is documented in Roberto Garrone, Proportionate Cybersecurity for Micro-SMEs: A Governance Design Model Under NIS2, 5 Journal of Cybersecurity and Privacy 4 (2025).

B. Platform Dependency as a Behavioral Response

Structural cost pressure does not simply reduce innovation activity; it reshapes the organizational strategies through which smaller firms seek to participate in AI ecosystems at all.

When independent compliance becomes prohibitively expensive, SMEs often seek what might be called "compliance shelter"[5]—embedding their fine-tuning activities within larger platform ecosystems that absorb governance responsibility in exchange for licensing fees, data-sharing arrangements, or deployment dependencies. The strategy is individually rational. A small firm that routes its fine-tuning activities through a major platform's API may effectively transfer compliance obligations upward, gaining access to governance infrastructure it could not independently afford.

At the ecosystem level, however, this pattern of behavior is troubling. As independent SME fine-tuning becomes more difficult, innovation chains concentrate within a small number of dominant platform infrastructures. Decentralized experimentation—the mechanism through which AI development has historically generated diversity in application design and technological pathways—weakens. The market structure that emerges is not a product of competitive dynamics operating on the merits; it is a product of compliance asymmetry. Actors become dominant not because they produce better systems but because they possess greater institutional capacity to absorb governance obligations.

This inversion represents a significant governance failure. Risk-based liability is designed to allocate responsibility according to the distribution of harm potential. When it functions primarily to select among market participants on the basis of institutional endurance, it has ceased to operate as a risk-governance instrument and begun operating as a mechanism of structural market concentration.

C. The Legitimacy Dimension

The analytical point just made is not merely about efficiency or economic welfare, though the efficiency implications are real. It concerns the normative foundations of liability as a legal institution.[6]

---

[5]The concept of "compliance shelter" describes a behavioral pattern whereby smaller firms reduce independent governance investment by embedding their operations within larger platform ecosystems that assume de facto compliance responsibility. This dynamic has been observed in connection with cloud computing regulation, financial technology licensing, and more recently, AI deployment through major platform APIs. See generally Wei Wen, Compliance Investment and Innovation Resource Allocation 112–24 (Tsinghua University Press 2022).

[6]The argument that liability legitimacy depends upon the distributive structure of governance obligations, and not merely upon their effectiveness in controlling risk, draws on broader debates in regulatory theory. See generally Cass R. Sunstein, After the Rights Revolution: Reconceiving the Regulatory State 47–72 (Harvard University Press 1990). In the technology context, structural exclusion concerns have been raised in connection with platform regulation, telecommunications licensing, and now AI governance.

Liability regimes claim authority over the actors they regulate on the basis that they serve identifiable governance purposes—deterring harm, allocating loss to those who generated it, providing incentives for appropriate precaution. These justifications are intelligible when the intensity of liability tracks the magnitude and probability of the harm that a given actor's conduct creates. They become increasingly difficult to sustain when liability intensity is determined primarily by factors that bear no relationship to harm generation—such as institutional scale, administrative capacity, or the ability to absorb fixed compliance costs.

Where formally equal obligations generate materially unequal burdens, and where the inequality in burden is unrelated to the inequality in risk, the governance structure has a legitimacy problem that cannot be resolved through more efficient administration or clearer guidance. The problem is structural: the obligation architecture is not calibrated to the governance objective it purports to serve.

III. Proportionality as a Condition of Liability Legitimacy

The argument that AI liability should be calibrated to actual deployment risk and realistic institutional capacity is often framed as a matter of policy pragmatism—regulators should be flexible, should account for the practical difficulties faced by smaller actors, should avoid regulations that inadvertently stifle innovation. There is nothing wrong with this framing as far as it goes, but it understates what is at stake. Proportionality in this context is not a policy preference that responsible regulators might choose to incorporate. It is a condition of the normative legitimacy of liability itself.[7]

The claim is straightforwardly derived from the structure of proportionality analysis as it operates in administrative and constitutional law. Proportionality requires that governance measures maintain a rational connection to legitimate objectives (suitability), that they do not impose burdens exceeding what those objectives require (necessity), and—most critically for present purposes—that the weight of the interference with individual interests does not exceed the importance of the objective being pursued (proportionality stricto sensu). This final dimension is where the analytical work matters most, and it is the dimension that uniform compliance structures most systematically fail.

The suitability and necessity requirements, while important, are relatively tractable in the AI fine-tuning context. Whether a given compliance obligation contributes meaningfully to risk reduction is an empirical question, and most of the standard requirements associated with

---

[7]On the concept of proportionality as a principle of legal justification rather than mere policy flexibility, see Robert Alexy, *A Theory of Constitutional Rights* 66–69 (Oxford University Press 2002). Alexy's account of proportionality stricto sensu—which requires balancing the degree of interference against the importance of the objective pursued—provides the theoretical foundation for the calibration argument advanced in this article.

capability reconfiguration liability—pre-deployment assessment, documentation, monitoring—do contribute something, even if their marginal contribution varies significantly across deployment contexts. The necessity question is somewhat harder, and the analysis of public compliance infrastructure in Part VII takes it seriously: if equivalent governance outcomes can be achieved through shared institutional mechanisms at lower individual cost, the necessity standard requires regulators to pursue those mechanisms actively rather than defaulting to individual compliance mandates.

But the weightiest analytical burden falls on proportionality stricto sensu. Here the question is not whether a compliance obligation serves its purpose, but whether the cost of the obligation to a particular class of regulated actors is justified by the importance of the governance objective relative to the harm probability and magnitude that the actor's conduct actually generates. For a narrowly deployed fine-tuned model serving a limited and identifiable user population in a controlled institutional context, the answer to this question is often negative. The compliance burden imposed is simply not commensurate with the harm scenario that the obligation is designed to address.

Proportionality in this sense does not argue for deregulation or for treating smaller actors as presumptively exempt from governance obligations. Where SME fine-tuners engage in genuinely high-risk deployments—systems operating in critical infrastructure, healthcare decision-support, or law enforcement contexts, at scale and with meaningful autonomy—the full weight of applicable governance obligations is warranted and proportionate. The argument is narrower: that liability intensity must track actual deployment risk, harm potential, and realistic institutional capacity, and that a uniform compliance structure which ignores all three fails the requirements of proportionality stricto sensu in a large and important class of cases.

IV. Differentiation in Practice: European and American Approaches

A. The EU AI Act's Risk-Tiered Framework

The EU AI Act is frequently cited as the most comprehensive AI governance instrument currently in force, and it is worth examining carefully for what it gets right—and what it leaves incomplete—with respect to the proportionality problem.[8]

The Act's fundamental organizing principle is risk stratification. Systems classified as high-risk face extensive obligations concerning documentation, data governance, transparency, human

[8]EU AI Act, Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, Arts. 9–17 (high-risk system obligations), Arts. 50–53 (transparency obligations for limited-risk systems). The Act's risk classification framework is elaborated in the Commission's accompanying guidelines; see European Commission, Guidelines on the Classification of AI Systems Under the AI Act (2024).

oversight, and post-market monitoring. Systems in lower-risk categories face substantially lighter requirements. This structure embeds a form of proportionality into the liability architecture, though it does so through categorical risk classification rather than through individualized deployment assessment.

For SME fine-tuners, the most significant practical implication is that many narrow, application-specific fine-tuning activities do not automatically qualify as high-risk under the Act's classification framework. A fine-tuned model deployed for internal document processing, customer service triage, or specialized technical assistance may fall outside the high-risk categories, with the result that compliance intensity is reduced not through any explicit SME accommodation but through the risk-differentiation logic of the classification system itself.

The Act supplements this with several mechanisms that address the institutional capacity problem more directly.[9] Regulatory sandbox provisions allow smaller actors to experiment under supervised conditions with partial relief from formal compliance obligations—a mechanism that acknowledges the particular importance of low-friction entry for SME innovation, even if the practical implementation of sandbox schemes has been uneven across Member States. The Act also permits the contractual reallocation of compliance obligations along the AI value chain,[10] which in principle allows fine-tuners to distribute governance responsibilities with upstream developers and downstream deployers through negotiated arrangements.

The limitation of the EU approach, from the perspective of this article's argument, is that proportionality operates primarily through categorical risk classification rather than through deployment-specific assessment. The classification framework uses a mix of sector-based and use-case-based criteria that map imperfectly onto the actual distribution of harm potential in fine-tuning ecosystems. A fine-tuned model may fall outside the high-risk classification on the basis of its nominal use case while actually operating in ways that generate significant risk; conversely, it may fall within a high-risk category on the basis of superficial formal characteristics while posing only limited harm potential in its actual deployment context. The architecture is better than uniform compliance, but it is not yet what a fully articulated proportionality principle would require.

B. The American Approach: Flexibility and Its Costs

---

[9]EU AI Act, Art. 57 (regulatory sandboxes); see also European Commission, SME Strategy for a Sustainable and Digital Europe, COM(2020) 103 final, at 8–11 (identifying regulatory complexity as a primary structural barrier for smaller enterprises entering digital markets).

[10]EU AI Act, Art. 25 (responsibilities along the AI value chain, including provisions governing substantial modification and the reallocation of provider obligations through contractual arrangement). This provision is significant because it permits fine-tuners to assume provider-equivalent obligations where they make substantial modifications to foundation models, while also allowing upstream developers and downstream deployers to contractually redistribute compliance duties.

The United States has not enacted a comprehensive federal AI statute. Governance instead relies on executive policy frameworks, voluntary guidance documents, sector-specific agency action, and a growing patchwork of state-level legislation.[11] The NIST AI Risk Management Framework provides an influential analytical vocabulary for thinking about AI governance without imposing mandatory compliance obligations.

This arrangement has obvious advantages from the perspective of innovation flexibility. The absence of a comprehensive mandatory compliance regime reduces immediate institutional burdens, and the voluntary character of most federal guidance allows firms to engage with governance frameworks selectively. SMEs are not automatically subject to the kind of systemic compliance architecture that the EU AI Act creates for high-risk systems.

The costs of this approach, however, are increasingly apparent.[12] Regulatory fragmentation generates substantial uncertainty costs that fall asymmetrically on smaller actors. As state-level AI legislation expands—with a growing number of states enacting AI transparency, algorithmic accountability, and automated decision-making requirements—SMEs operating across multiple jurisdictions face an increasingly complex and inconsistent compliance landscape. For large firms with established legal infrastructure, managing jurisdictional variation is a tractable if burdensome problem. For SMEs, the transaction costs of navigating overlapping and inconsistent requirements can be prohibitive.

The American model therefore does not eliminate the structural asymmetry problem; it transforms it. Compliance complexity shifts from fixed regulatory obligation toward strategic uncertainty. The burden is real but less visible, and it accumulates in ways that are difficult to document or challenge through conventional regulatory channels.

Comparing the two approaches suggests a conclusion that neither model makes fully explicit: differentiated governance structures are not a concession to interest-group lobbying or a departure from principled regulation. They are a structural feature of any liability regime that takes seriously the relationship between obligation intensity and deployment risk. Both the EU and the United States, in different ways and to different degrees, have already moved away from uniform compliance. What is missing is a coherent theoretical account of why differentiation is not merely permissible but required.

---

[11]On the NIST AI Risk Management Framework, see NIST, AI Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (January 2023). For the governance implications of the United States' fragmented approach, see the White House, Fact Sheet: National Policy Framework for Artificial Intelligence (2025); Senator Marsha Blackburn, The American AI Act (2025) (proposing a federal risk-classification framework broadly modeled on the EU approach, though with significant differences in scope and enforcement).

[12]On the strategic uncertainty costs generated by regulatory fragmentation, see OECD, OECD Digital Economy Outlook 2023, at 143–51 (OECD Publishing 2023) (analyzing the compliance challenges faced by SMEs operating across multiple jurisdictions with inconsistent AI governance standards).

V. China's "Regulate Through Support" Model and Its Structural Gaps

China's emerging AI governance model cannot be straightforwardly assimilated to either the EU's regulatory ambition or the American preference for governance flexibility. It reflects a different institutional logic—one in which the state reduces practical barriers to compliance through resource provision while maintaining mandatory governance baselines, rather than calibrating the baselines themselves to institutional capacity.

A. Resource Provision as Governance Instrument

The most distinctive feature of the Chinese approach is the deployment of public resource-support mechanisms as instruments of compliance accessibility.[13] Programs providing compute vouchers, data vouchers, and model vouchers reduce the structural disadvantages SMEs face in accessing the foundational technological resources necessary for AI development. Although these mechanisms are typically framed in industrial-policy terms—as instruments for promoting domestic AI competitiveness—they have a governance dimension that deserves more sustained attention.

By lowering the cost of access to computational resources and high-quality datasets, the state indirectly reduces the practical burden associated with meeting governance obligations such as risk assessment, model evaluation, and deployment monitoring. An SME that can afford to run proper evaluation protocols because compute costs have been subsidized is better positioned to satisfy its compliance obligations than one that cannot. In this sense, resource support and governance compliance are not merely parallel policy tracks; they are institutionally linked, with support provision functioning as a form of compliance enablement.

This represents an important institutional insight that the EU and American frameworks have not fully explored. Regulatory enforceability is not solely a function of obligation design; it also depends on whether regulated actors have realistic capacity to comply. Where that capacity is systematically lacking, obligations remain formal commitments rather than behavioral constraints.

B. Application-Oriented Governance Focus

---

[13]MIIT, Special Action Plan for the Digital Empowerment of SMEs (2025–2027) (establishing compute voucher, data voucher, and model voucher programs as instruments for reducing SME access barriers to foundational AI resources). For analysis of these mechanisms in governance terms, see Liu Shangxi, Governance Innovation and Competition Reconstruction in the Digital Economy Era (2025).

A second notable feature of the Chinese approach is its increasingly explicit orientation toward deployment context rather than development characteristics as the primary regulatory trigger.[14] The Interim Measures for the Administration of Generative Artificial Intelligence Services focus on the characteristics of the service provided to users—its public accessibility, its content-generation capabilities, its potential for harm in deployment—rather than on the technical parameters of model development. Registration obligations attach to public-facing generative AI services, not to fine-tuning activity as such.

This application-oriented framework aligns closely with the proportionality argument developed in this article. If the governance justification for imposing compliance obligations on fine-tuners derives from the harm that deployment generates or risks generating, then the appropriate regulatory trigger is deployment context—user scale, accessibility, criticality—rather than development-side proxies such as parameter count or training expenditure. The Chinese framework has moved toward this position, albeit without a fully articulated theoretical justification.

C. What the Chinese Approach Has Not Yet Resolved

Despite its institutional strengths, the current Chinese model exhibits structural limitations that prevent it from fully realizing the governance logic it implies.[15]

The most significant limitation is the absence of explicit linkage between resource support and liability mitigation. SMEs that benefit from public compute or data subsidies do not thereby gain any formal reduction in their compliance obligations or any recognition of their subsidy-enabled compliance capacity in subsequent liability assessment. The two systems—industrial support and governance compliance—operate in parallel rather than in integration. This gap undermines the potential of the "regulate through support" logic as a coherent governance architecture.

A second limitation concerns the underdevelopment of public compliance infrastructure.[16] Resource subsidies improve access to technological capacity, but they do not address the institutional capacity problem directly. An SME with subsidized compute access still lacks the legal infrastructure, audit expertise, and documentation management systems that compliance

---

[14]CAC et al., Interim Measures for the Administration of Generative Artificial Intelligence Services (2023), Art. 7 (training data obligations), Art. 17 (registration requirements for generative AI services with public-facing functionality). The registration obligation constitutes a mandatory governance baseline applicable regardless of enterprise scale.

[15]On the underdevelopment of formal linkages between public resource-support mechanisms and liability mitigation in the Chinese framework, see Liao Lilin, Gaps in AI-Compatible Supplementary Governance Frameworks, 5 Journal of Cybersecurity & Privacy 4 (2025). Liao identifies the absence of explicit "compliance credit" mechanisms as a structural gap that limits the governance effectiveness of existing subsidy programs.

[16]NDRC et al., Measures to Strengthen the Cultivation of Innovative Enterprises in the Digital Economy (2025); MIIT et al., Measures for Artificial Intelligence Science and Technology Ethics Review and Services (Trial) (2026). The latter instrument establishes a tiered ethics-review structure that distinguishes between simplified and full review procedures depending on deployment context and risk category.

obligations require. Addressing this gap requires not merely resource provision but the development of shared compliance services—the subject of Part VII.

Finally, systematic regulatory sandbox mechanisms for fine-tuning activities remain comparatively limited. The sandbox provisions that exist in the Chinese framework are primarily oriented toward fintech and other established regulatory domains; the application of controlled experimentation frameworks to AI fine-tuning is still at an early stage of institutional development.

VI. Elements of a Proportionate Liability Architecture

Drawing on the preceding analysis, this Part proposes four structural elements for a liability architecture that satisfies the proportionality requirements identified in Part III. These elements are not independent; they form an integrated structure in which each component depends on the others for its effectiveness. The most important of them—safe harbor design conditioned on participation in public compliance infrastructure—is addressed last, because its logic only becomes fully clear once the others are in place.

A. Duty Structures Calibrated to Deployment Risk

The starting point for any proportionate liability architecture is that the intensity of governance obligations must vary with deployment risk, not with organizational characteristics. This is a more demanding standard than it might appear. It is not satisfied by categorical risk classification systems that assign compliance obligations on the basis of nominal use cases or sector labels, because such systems reproduce the proportionality problem at a different level of abstraction. What is required is an obligation structure that responds to the actual harm potential of a specific deployment.[17]

In practical terms, this means distinguishing among at least three deployment contexts. High-risk deployments—systems operating in critical infrastructure, healthcare decision-support, law enforcement, or large-scale consumer-facing applications with significant autonomy—should face the full suite of governance obligations associated with capability reconfiguration liability. Medium-risk deployments, involving consequential but bounded applications with meaningful but not critical effects on identifiable user populations, warrant standard governance obligations calibrated to the deployment context. Low-risk deployments, involving narrow application boundaries, limited and identifiable user populations, and controllable harm scenarios, should

[17]The deployment-side threshold framework proposed here bears some resemblance to the EU AI Act's use of deployment context (rather than technical architecture) as the primary determinant of risk classification. However, the EU framework still employs development-side proxies (such as general-purpose model status and FLOP thresholds) for certain categories of systemic-risk classification; see EU AI Act, Arts. 51–55 (GPAI model obligations). The argument here is that application-side indicators should be primary, with development-side metrics serving at most as secondary triggers.

face simplified obligations sufficient to establish accountability without requiring full institutional compliance infrastructure.

Critically, this differentiation should not operate as an organizational size exemption. SMEs deploying high-risk systems must bear the governance obligations appropriate to those systems. The differentiation is between risk tiers, not between large firms and small firms. What makes proportional calibration beneficial to SMEs is not that they are treated as a privileged category but that the SME fine-tuning activity that generates the most structural exclusion pressure—narrow, specialized, low-to-medium risk deployment—is precisely the category that a properly calibrated risk-tier structure would recognize as warranting reduced compliance intensity.

B. Application-Side Thresholds

The second element addresses the basis on which regulatory intensity is determined. Development-side metrics—parameter count, computational expenditure, training data volume—have been widely used as proxies for AI system risk, partly because they are measurable and partly because they correlate, at least at the extremes, with certain categories of harm potential. But they are increasingly unreliable proxies in the fine-tuning context.

Fine-tuned models inherit the architectural scale of their foundation models regardless of the scope or risk profile of the fine-tuning intervention. A narrowly fine-tuned variant of a large foundation model will have a high parameter count that bears no relationship to the risk profile of the specific deployment. Conversely, collaborative fine-tuning arrangements and distributed development environments allow capable systems to be produced without any single actor incurring the development-side indicators that risk-tier frameworks use as triggers. Application-side thresholds—user scale, deployment criticality, revenue dependence, query volume, accessibility to vulnerable populations—provide a more defensible basis for determining compliance intensity because they measure the relevant variable directly: the actual social risk exposure generated by a given deployment.[18]

C. Graduated Transition Mechanisms

A proportionate liability architecture must also address temporal dynamics. Firms grow, and as they grow, their compliance capacity expands and their risk profile changes. A governance structure that imposes abrupt compliance escalation when firms cross developmental thresholds

---

[18]The deployment-side threshold framework proposed here bears some resemblance to the EU AI Act's use of deployment context (rather than technical architecture) as the primary determinant of risk classification. However, the EU framework still employs development-side proxies (such as general-purpose model status and FLOP thresholds) for certain categories of systemic-risk classification; see EU AI Act, Arts. 51–55 (GPAI model obligations). The argument here is that application-side indicators should be primary, with development-side metrics serving at most as secondary triggers.

—in terms of user base, revenue, or deployment scope—creates strong disincentives for organizational growth, particularly for firms operating close to threshold boundaries.[19]

Graduated transition mechanisms address this problem by allowing compliance obligations to increase progressively rather than discretely. Phased documentation requirements, extended reporting grace periods for newly threshold-crossing firms, staged implementation of monitoring obligations, and access to structured compliance support during transition periods would allow the obligation structure to track genuine increases in institutional capacity and risk profile without generating the cliff-edge effects that binary classification systems produce.

D. Safe Harbor Conditioned on Collective Compliance Participation

The fourth element is both the most important and the most distinctive feature of the architecture proposed here. Existing discussions of safe harbor mechanisms for AI fine-tuners typically condition protection on independent compliance investment: the firm has conducted its own risk assessment, maintained its own documentation, established its own monitoring system.[20] This design is defensible in principle but reproduces the fixed-cost problem in a new form. Safe harbor protection conditioned on individual compliance capacity is accessible precisely to the actors who least need it.

The architecture proposed here conditions safe harbor protection not on independent compliance investment but on meaningful participation in public compliance infrastructure. A fine-tuner that uses a lawfully registered foundation model, completes simplified registration or ethics-review procedures through a collective compliance mechanism, engages with shared audit services where available, and implements reasonable remediation measures in response to identified harms should qualify for supplementary rather than primary liability unless its conduct directly and proximately caused the relevant harm. The safe harbor is earned through governance participation, not through independent compliance capacity.

This design has an important systemic implication: it creates affirmative incentives for engagement with collective compliance mechanisms, transforming public compliance infrastructure from an optional resource into a governance instrument with meaningful legal consequences. The value of the safe harbor depends on the quality and accessibility of the collective compliance systems to which it is linked. This dependency, far from being a weakness,

---

[19]The case for graduated compliance transitions draws support from experience with analogous regulatory escalation problems in other sectors. Telecommunications licensing frameworks have historically employed phased obligation structures to accommodate the transition from small-scale experimental operators to larger commercial providers. See generally Doshi-Velez et al., Accountability of AI Under the Law: The Role of Explanation, Berkman Klein Center Working Paper (2017) (noting the importance of calibrating accountability obligations to institutional capacity).

[20]Safe harbor mechanisms in the AI context are discussed in Wang Qingsong, supra note 2, at 78–83. The proposal here extends Wang's analysis by conditioning safe harbor availability on participation in public compliance infrastructure, rather than solely on independent compliance investment. This linkage is designed to create affirmative incentives for SME engagement with collective governance mechanisms.

is what makes the architecture coherent: it ties the legitimacy of liability mitigation to the actual functioning of governance institutions, rather than to the formal performance of compliance procedures.

VII. Public Compliance Infrastructure: The Institutional Precondition

The differentiated liability architecture described in Part VI cannot function without the public compliance infrastructure on which its safe harbor mechanism depends. This is the hardest institutional problem the article raises, and it deserves direct attention.

The structural difficulties confronting SME fine-tuners arise not only from resource scarcity but from systemic deficits in compliance capacity. An SME may have adequate financial resources to fund a fine-tuning project, access to foundation models, and technically competent personnel—and still lack the legal infrastructure, audit expertise, and governance protocols that contemporary AI compliance obligations require. Resource provision, however generous, does not solve this problem. What is needed is a transformation in the institutional character of compliance itself: from an individualized burden that each actor must bear independently to a shared institutional service that the governance system provides.

Shared audit pools represent perhaps the most practically significant element of the infrastructure required.[21] By allowing multiple firms to access collective compliance resources—audit expertise, risk-assessment methodologies, documentation frameworks—shared audit mechanisms reduce fixed governance costs in ways that individual subsidy programs cannot replicate. The cost of audit expertise does not decrease when a small firm receives a compute voucher. It decreases when audit services are institutionally organized at collective scale.

Centralized registration and ethics-review systems serve a related but distinct function. Beyond reducing repetitive administrative burdens, centralized systems improve the informational quality of governance itself. A centralized registry of fine-tuned systems with documented deployment contexts provides regulators with the distributional data necessary to calibrate risk-tier classifications and update obligation structures as deployment patterns evolve. Ethics-review systems organized at industry or sector level allow the substantive expertise embedded in governance procedures to be developed and maintained collectively, rather than requiring each firm to reinvent it independently.

---

[21] On shared audit platforms as a governance instrument, see China Economic Net, Shared Audit Resource Platforms and Industry Community Governance Mechanisms, Audit Observation (July 2025). The article describes pilot programs in Guangdong and Shanghai involving collective AI audit services funded through public-private partnership arrangements. See also WTO Informal Working Group on MSMEs, Key Conclusions of the Working Group Meetings 2024–2025 (identifying shared compliance infrastructure as a priority mechanism for reducing SME trade costs in digital services).

Standardized compliance templates and guidance systems address a third form of compliance burden that receives less attention than cost barriers: interpretive uncertainty. For SMEs without dedicated legal infrastructure, the challenge of compliance is often not primarily financial. It is the difficulty of understanding what specific obligations require in the context of a particular deployment, and of documenting compliance in ways that will satisfy regulatory review. Standardized templates that translate generic compliance requirements into actionable, deployment-specific guidance can dramatically reduce this burden without reducing the substantive standard.

The broader significance of public compliance infrastructure extends beyond the benefits it provides to individual firms. When compliance becomes publicly accessible, the nature of governance changes. Regulation no longer operates solely through the imposition of individual obligations; it also functions through enabling actors to participate in governance at manageable institutional cost. This shift matters for AI governance specifically because the structural concentration dynamics described in Part II are not merely distributional problems. They represent a genuine threat to the diversity of technological development that makes AI ecosystems epistemically and socially robust.

There is also a cross-border dimension that warrants emphasis.[22] Where compliance obligations are significantly asymmetric across jurisdictions, mobile fine-tuning activities will tend to migrate toward lower-intensity regulatory environments—not eliminating the associated risks but displacing them into less accountable governance contexts. The resulting regulatory arbitrage weakens the long-term sustainability of AI governance architectures globally. Public compliance infrastructure that makes high-standard governance accessible at manageable cost reduces the incentive for compliance migration, strengthening rather than undermining the effectiveness of demanding governance frameworks.

VIII. Conclusion

The argument of this article began with a gap in capability reconfiguration theory. Identifying fine-tuners as independent subjects of AI liability is a genuine conceptual advance, but it resolves only half the institutional problem. The other half concerns how intensely that liability should apply—a question that the capability reconfiguration framework, focused on the correct allocation of responsibility among supply chain actors, has not yet addressed directly.

---

[22]The cross-border risk displacement argument draws on a broader literature analyzing regulatory arbitrage in technology governance. Where compliance obligations are asymmetric across jurisdictions, mobile firms may relocate operational activities to lower-intensity regulatory environments—not necessarily reducing aggregate risk, but redistributing it across governance boundaries. This dynamic has been documented in connection with data-processing activities and cryptocurrency exchange operations. For the AI context, see OECD, OECD Digital Economy Outlook 2023, supra note 10, at 189–97.

The answer this article defends is that liability intensity must be calibrated to actual deployment risk through a proportionality analysis that takes institutional capacity seriously as a governance constraint, not merely as a practical inconvenience. The failure to make this calibration is not a technical deficiency in governance design. It is a normative failure: a liability regime that constrains actors primarily on the basis of institutional endurance rather than harm potential has ceased to function as a risk-governance instrument and begun functioning as a mechanism of structural market concentration.

The comparative analysis suggests that the trajectory of mature governance systems is already toward differentiation—that both the EU's risk-tiered framework and China's "regulate through support" model reflect, in different ways, an institutional recognition that uniform compliance structures are neither effective nor legitimate in the context of highly heterogeneous technological ecosystems. What these systems lack is a coherent theoretical account of why differentiation is normatively required. That account is what the proportionality argument provides.

The differentiated liability architecture proposed in Part VI—deployment-risk-calibrated duty structures, application-side regulatory thresholds, graduated transition mechanisms, and safe harbor protections conditioned on collective compliance participation—is designed to give institutional form to this theoretical commitment. None of its elements is novel in isolation; each has analogues in existing governance frameworks. The contribution is in their integration, and particularly in the linkage between safe harbor availability and participation in public compliance infrastructure, which transforms the relationship between individual liability mitigation and collective governance functioning.

The article does not claim to have resolved all the institutional design questions that this framework raises. How risk tiers should be defined, how application-side thresholds should be set and updated, how shared audit mechanisms should be funded and governed, and how safe harbor protections should interact with existing tort and regulatory liability rules across different jurisdictions—these are questions that require careful empirical investigation and ongoing regulatory elaboration. What the article claims is something more modest but perhaps more fundamental: that these questions cannot be answered well without first establishing that proportionality is not a preference but a condition, and that the legitimacy of AI liability regimes depends on whether they meet it.

References

Robert Alexy, A Theory of Constitutional Rights (Oxford University Press 2002).

CAC et al., Interim Measures for the Administration of Generative Artificial Intelligence Services (2023).

China Economic Net, Shared Audit Resource Platforms and Industry Community Governance Mechanisms, Audit Observation (July 2025).

Doshi-Velez et al., Accountability of AI Under the Law: The Role of Explanation, Berkman Klein Center Working Paper (2017).

European Commission, Guidelines on the Classification of AI Systems Under the AI Act (2024).

European Commission, SME Strategy for a Sustainable and Digital Europe, COM(2020) 103 final.

Roberto Garrone, Proportionate Cybersecurity for Micro-SMEs: A Governance Design Model Under NIS2, 5 Journal of Cybersecurity and Privacy 4 (2025).

Andres Guadamuz, "The Artificial Intelligence Supply Chain: Legal Perspectives," 12 Journal of Internet Law 1 (2023).

Robert W. Hahn & John A. Hird, The Costs and Benefits of Regulation: Review and Synthesis, 8 Yale Journal on Regulation 233 (1991).

Liao Lilin, Gaps in AI-Compatible Supplementary Governance Frameworks, 5 Journal of Cybersecurity & Privacy 4 (2025).

Liu Shangxi, Governance Innovation and Competition Reconstruction in the Digital Economy Era (2025).

MIIT, Special Action Plan for the Digital Empowerment of SMEs (2025–2027).

MIIT et al., Measures for Artificial Intelligence Science and Technology Ethics Review and Services (Trial) (2026).

NDRC et al., Measures to Strengthen the Cultivation of Innovative Enterprises in the Digital Economy (2025).

NIST, AI Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (2023).

OECD, OECD Digital Economy Outlook 2023 (OECD Publishing 2023).

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (AI Act).

Senator Marsha Blackburn, The American AI Act (2025).

Cass R. Sunstein, After the Rights Revolution: Reconceiving the Regulatory State (Harvard University Press 1990).

Wang Qingsong, Risk Allocation and Liability Structures in Large-Model Applications, Eastern Law, No. 4 (2025).

Wei Wen, Compliance Investment and Innovation Resource Allocation (Tsinghua University Press 2022).

White House, Fact Sheet: National Policy Framework for Artificial Intelligence (2025).

WTO Informal Working Group on MSMEs, Key Conclusions of the Working Group Meetings 2024–2025.