Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 3

ISSN: 1540-3270

Regulating Innovation: Legal Challenges and Governance Frameworks for AI-Driven Startups

[1]Ms. Alisha Kumar, [2]Ms Vaishali Jain, [3]Dr. Alaknanda Rajawat, [4]Dr. Arti Sharma

[1]Assistant Professor, Jagan Institute of Management studies

[2]Assistant Professor, Jagan Institute of Management Studies

[3]Assistant Professor, Jagannath University, Jaipur

[4]Assistant Professor, Jagannath University, Bahadurgarh

Abstract:

New models of artificial intelligence distribute responsibility on the provider of an AI system, differentiate the responsibilities based on the risk of the intended use, and impose the responsibilities through documentary conformity requirements that are evaluated prior to deployment. Each of these three design decisions plays a game-changing role and creates its own distinct set of issues with the structural nature of AI startups, and the poor outcomes are not the result of transitional friction, but rather foreseeable consequences of the way the regulated object has been designed. We make three assertions. The first is that there are obligations (that are part of the system provider) in the middle of the algorithmic value chain, where the firm does not have access to the information held by the customer and the control over the behaviour of the model is not what the model provider has, since the customer is the one who has the information. Second, "whatever a statute allocates a contract reallocates," by which I mean that the upstream provider, by the terms of service, and the downstream enterprise customer, by the terms of procurement, disclaim, thereby reallocating the incidence down through bargaining power—rather than the intent of the legislature—because startups are squeezed between. Third, the risk-tier classification assumes a fixed intended use, but early-stage companies change their use, which means a firm's risk-tier is not stable in the population that might not be able to handle re-classification. We also differentiate between compliance cost, which can be paid, and compliance capability (which requires organisational structures that are not available in young firms), and propose that this differentiation is why funding-based support instruments are not as effective. Six governance instruments are evaluated for these problems, and seven propositions with an analytical protocol are provided. No artifical statistics are stated.

Keywords: artificial intelligence regulation; startups; regulatory incidence; algorithmic value chain; risk-based regulation; regulatory sandboxes; compliance capability; innovation policy

1. Introduction

The regulation of AI has advanced quicker than ever to become a legally binding document. In a couple of years, voluntary ethics declarations were replaced by risk-management frameworks, management-system standards, and in the European Union, a comprehensive regulation that requires risk-based obligations (European Union, 2024; ISO/IEC, 2023; NIST, 2023). This logic is sound and on the abstract level good; obligations must be commensurate to the risk presented by a system; obligations must be imposed on the party with the greatest control over the risk; compliance must be easily shown beforehand, not in court after the harm.

This paper explores the effect of that logic on the regulated party, in the case of a startup. It's not a rhetorical question. Most of the applied AI development activity is found in companies with less than 50 employees, often on models that aren't their own, in customers' environments they never see, and at a point in time when the intended use that dictates the regulatory classification has yet to be finalised. As has been long recognised in the literature, the timing of action when regulating emerging technology is a trade-off between the capacity to steer a technology and the knowledge of how to steer it (Collingridge, 1980) and the structuring of legal-ethical action is a lag effect that tends to follow technical development (Bennett Moses, 2007; Marchant et al., 2011). The one that is not so obvious is the distributional issue in a regulated population, who is being weighed within the framework?

We contend that three design characteristics are responsible for the answer, and that all three work against small businesses. The first is related to the object that is being regulated. Any framework must have obligations attached to it – and the choices of whom to attach them to (the model, the system, the intended use, the deployer, the outcome) have different implications for who takes what burden. New frameworks have been proposed to attach obligations to a point in the value chain at which information and control are both low, which has been studied in the accountability literature, but not yet explored in terms of its implications for firm size (Cobbe et al., 2023).

The second is related to contract. Statutory allocation doesn't constitute final allocation. The terms of service of upstream model providers exclude liability for the use of their models downstream, and contractual warranties, agreements of indemnity and audit rights by downstream enterprise and public-sector customers reallocate those responsibilities upstream. The two movements merge in the same party. The remaining incidence of an AI obligation therefore depends on the relative bargaining strength and the firm that has the lower bargaining power takes over what the other firms give up. The mechanism is not tied to the statute's substance and, we propose, the most significant factor in determining the location of the burden of AI compliance.

The third one is about the stability of the classification. Risk-tiered frameworks classify by intended purpose and it assumes that intended purpose is known and consistent. It's often neither for an early stage company. A change from a minimal obligation classification to a high-risk classification can occur in a firm without any modification to their technology, just by acquiring a customer in a different sector – and the firm may only find out about the classification change through a procurement process.

The paper will be structured as follows. Next, in Section 2, the candidate regulatory objects and their incidence properties are presented. In Section 3 the value-chain analysis and the thin middle problem are developed. This is followed by a discussion of contractual reallocation in Section 4. Section 5 deals with the issue of classification instability and the distinction between the costs of compliance and the capacity to comply. Section 6 evaluates six governance instruments to these problems. Propositions and an analytical protocol are presented in Section 7 and implications and limitations are discussed in Section 8 and 9. The argument is intentionally structural, does not rely on the specific content of any existing instrument nor is it meant to be applicable to frameworks to be written.

2. What Is Being Regulated?

Every regulatory framework must select an object to which obligations attach, and the selection is consequential in ways that debates about stringency obscure. Table 1 sets out the principal candidates with their incidence properties.

Table 1. Candidate regulatory objects and their incidence

Object	What it captures	What it misses	Who bears the obligation	Position of a startup	Characteristic failure
The model itself	Capabilities and training data of the underlying system	Everything about how it is used	Whoever trained it	Usually outside scope; a beneficiary	Regulates a general capability with unknown application
The system as placed on the market	A bounded product with a declared purpose	Reconfiguration and unforeseen use after supply	The party that supplies it under its own name	Directly in scope, frequently as the named provider	Places obligations where use context is not observable

Object	What it captures	What it misses	Who bears the obligation	Position of a startup	Characteristic failure
The intended use or application	Risk arising from a specific deployment context	General-purpose systems with no single purpose	Whoever declares the purpose	In scope, but the declaration is unstable	Classification can be changed by a sales decision
The deployer	The organisation that puts a system into operation	Design defects fixed upstream of deployment	The customer organisation	Out of scope, but pressured by contract	Deployer may lack technical capacity to assess
The data	Processing of personal or sensitive information	Harms not mediated by personal data	Controller and processor	In scope where personal data is involved	Under-inclusive for many AI harms
The outcome	Realised harm, addressed after the fact	Nothing ex ante; deterrence only	Whoever is found causally responsible	Exposed, and least able to insure	Ineffective where harm is diffuse or unattributable

Two observations follow. The first is that no single object is apparent as correct, and every available framework has picked otherwise, resulting in a certain cross-framework irregularity for firms that fall under several. The second, and more important in the present argument, is that objects are not only different with respect to coverage, but also with respect to incidence. A small number of very large firms are governed by an object defined by the model. An object determined by the result determines who is liable, thus favouring insurers or those who can withstand. The largest number of firms is regulated by an object defined by the system as supplied and they regulate it at the point in the chain where firms are least. This is not a criticism of the choice, because there are real benefits for the bounding of the regulated artefact, but the consequence of distribution that has not generally been recognised as part of the framework design.

A related issue of definition is that of generality. A system trained for no specific purpose cannot be classified by intended purpose; the frameworks have considered this by totally separating general purpose systems, so that it is unclear where the general turns to the specific in a chain. For an applied startup, the answer is yes—they make it specific in their hands, which is why they are named the provider of the system, albeit with a small portion of the capability.

2.1 The prior problem of definition

Delimiting is a prerequisite for regulation and the definitional perimeter of an AI system has been a particularly tough problem to solve. Definitions expressed as a set of system(s) that deduce from inputs how to produce outputs describe much of the common software used in everyday commerce, such as commonly used statistical models, which have been around for decades, and simple re-expression of these narrows the definition so much as to exclude systems of concern. Indeed, this scope question has repeatedly been raised as unresolved but not settled through the drafting process (Veale & Zuiderveen Borgesius, 2021), and the field of techlaw has seen multiple arguments that the scope of technology is often problematic and that it is often better to regulate harms and relationships than regulate artefacts (Calo, 2017; Crootof & Ard, 2021).

The Incidence consequence is specific. A wide definition encompasses companies that are not called AI companies, but have not necessarily done anything to get ready for the legal requirements, and it does so unevenly:

a large firm will spot the overlap in the legal function, while a small firm might not become aware of the overlap until a customer's diligence questionnaire asks. A narrow definition, however, opens the door to reformulation by firms sophisticated enough to describe their systems positively, once again in favour of the savvy firms. Both errors are non-neutral with respect to firm size and the selection between them is consequently a distributional choice made as a technical choice.

One related challenge is the assessment tools that the framework uses. Conformity and impact assessments assume an institutional setting with a person of suitable expertise and independence carrying out the impact assessment and a person who follows up on the results, and algorithmic impact assessment analyses have pointed out that the design of the institutions around the instrument has the potential to make or break the process (Selbst, 2021). An impact assessment is not an assessment if it is carried out by the person who constructed the system, as is typically the case in a firm that does not have a compliance function. This was a broad comment on compliance, which occurred much earlier in time than the advent of AI frameworks, about the need to translate compliance goals into internal organisational practice, where compliance regimes thrive or falter (Bamberger, 2010).

3. Incidence in the Algorithmic Value Chain

3.1 The thin middle

Figure 1 sets out a simplified four-layer chain and compares three attributes across it: information about the deployment context, control over model behaviour, and capacity to bear fixed compliance cost. The pattern is the paper's central structural claim.

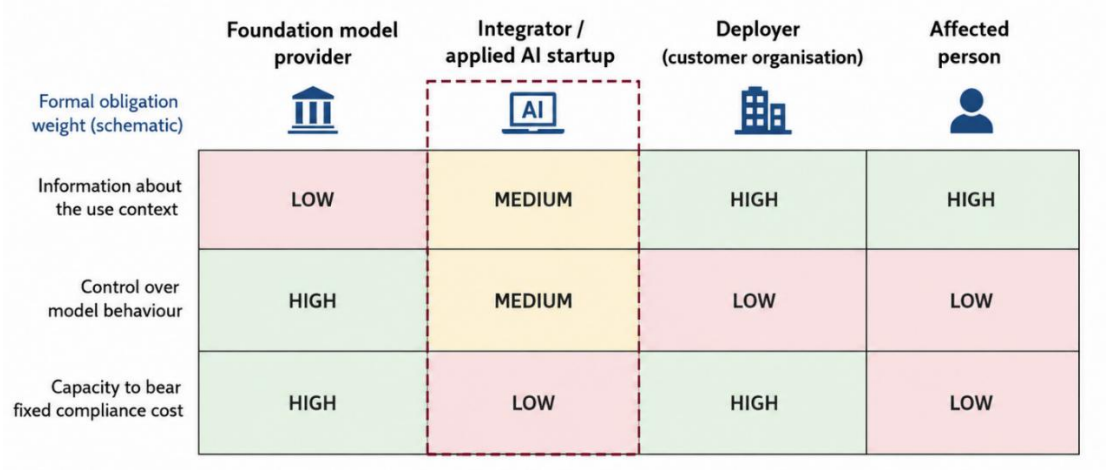

Figure 1. Incidence in the algorithmic value chain. The applied AI startup carries substantial formal obligation while holding neither the deployment information available to the customer nor the behavioural control available to the upstream provider, and it is exposed to contractual reallocation from both directions. Attribute levels are analytic judgements offered for discussion, not measurements.

The deployer is the one who knows the information about the use context, what decisions the system will inform, who will be affected by the system, and what the consequence of the error is. The upstream provider has control over the behaviour of the model, including what training data is used, capabilities, and what guardrails are present that can only be modified by the downstream integrator, which cannot inspect. Either end can have the ability to pay fixed compliance cost, but not the middle. So the applied startup will be at the intersection of maximum obligations and minimum dischargeability in terms of the framework.

It is not an exemption that this type of businesses should be granted. They have the responsibility to make decisions with real consequences, pick their application context, and often are the only ones who know the technical system and domain. It is an argument that duties articulated as if imposed by the provider will be in a substantial number of instances obligations that the designated party cannot meet in full, and that a scheme which

would lead to widespread structural non-compliance of good faith providers has a design problem, not an enforcement problem.

3.2 Three specific discharge failures

Three duties are examples of the challenge to be found in the concrete. Data governance requirements assume that there is a way to describe training data. An integrator can't, because the composition of the training corpus that is the third-party's model is not known and often it isn't known. Technical documentation/explainability requirements assume the provider can explain how the outputs are generated; the integrator can explain how they generate their own outputs but not the one that generates most. Post market monitoring assumes visibility of the performance in deployment; if the customer deploys the system in their own environment, without sharing the outcome data, the provider's monitoring responsibility is based on a cooperation that the contract may not obtain. The answer is right but isn't possible in practice, as the integrator won't be able to get these things for free, by contract, from the vendor it's upstream from since the vendor doesn't spend enough to justify custom contracts. The asymmetry is structural and not a lack of diligence on the part of the parties and is a governance implication of the analysis that obligations that can only be met by the upstream party should be placed on that party or be accompanied by an obligation for the party to disclose to the downstream party.

3.3 The open-weight complication

The chain in Figure 1 has an identifiable upstream provider which has a commercial relationship with the integrator. This is not the case for an increasing number of applied development projects. If the parameters of an open-weight model are downloaded, the parameters are tuned with data from the venture and the model is self-hosted, the upstream contract is missing and there's no provider to disclaim and often no entity within a regulator's reach that could provide the information that the integrator's obligations assume. The startup turns the provider of the system, and takes on the burden of obligations, the chain analysis above reveals it is least capable of fulfilling. This creates a push-pull effect which is not necessarily what the framework designers had in mind. Building on a commercial model via an interface ensures that obligations are formally passed to the integrator but actually shared, as the upstream provider has documentation, publishes evaluations and has an interest in its customers' compliance. When the open weights are built, the entire documentary and evidentiary burden shifts to the smallest party to the arrangement. The regulation, that applies to any named provider, thus punishes the more clear and, to a greater extent, independently verifiable route and encourages smaller providers to rely on a few commercial providers, a policy issue in itself.

The point generalises. If the regulated object can be named by the supplier of a system by that name, the incidence of the framework focuses on corporate structure, not on causality, and corporate structures can be selected. The argument that the design of a system itself represents a form of regulation, preceding and often more controlling than law has been used in the past to analyse the allocation of regulatory effect (Reidenberg, 1998), and this is the reverse, the allocation of the law reshapes the technological architecture, and does so in a way that concentrates.

4. Contractual Reallocation

Statutory allocation will decide who is responsible for a regulator. It doesn't imply who is responsible for the economic cost, as those who will, move by agreement. There are two concurrent and parallel movements in play here – with AI startups.

Upstream, warranties are typically excluded from model access, the responsibility for compliance with applicable law is shifted to the customer and the applicable terms are often such that the customer will be liable in the event of a claim by the provider for use of the model. By contract, the integrator is then responsible for the properties of a component which they cannot inspect. The opposite is true downstream: enterprise and public-sector customers require warranties of regulatory compliance, indemnities, audit and inspection rights, service credits based on the accuracy or availability of the service, and more recently a requirement to demonstrate conformity to a named standard prior to contract award. The customer is likely to be bigger, better informed and will be buying under a template that will not change for a small business.

This results in the reallocation of a unit obligation as shown in Figure 3, (b). The position is statutory, distributed along the chain while under the contractual position, it is concentrated in the middle. The main mechanisms are described in Table 2.

Table 2. Mechanisms of contractual reallocation

Mechanism	Direction	Instrument	Effect on the startup	Feasibility of resistance
Warranty disclaimer	Upstream to middle	Model provider terms of service	Assumes risk for component properties it cannot verify	None at typical spend levels
Use-based indemnity	Upstream to middle	Acceptable use policy with indemnity	Liability for downstream misuse it does not observe	None; terms are standard-form
Compliance warranty	Downstream to middle	Procurement contract	Warrants conformity across regimes it may not have assessed	Low; refusal costs the contract
Audit and inspection rights	Downstream to middle	Procurement contract	Fixed cost of audit readiness regardless of firm size	Low; increasingly standard in public procurement
Standard certification requirement	Downstream to middle	Tender qualification criteria	Certification cost becomes a condition of market access	None; it is a threshold requirement
Unlimited or uncapped liability	Downstream to middle	Procurement contract	Liability exceeding enterprise value	Moderate; sometimes negotiable at the margin

The last row is very important, and creates a result which no regulator would have desired. A liability cap negotiated on a level that is suitable to a large supplier may be more than the startup's total enterprise value which means that the startup is not, in fact, insuring against liability, but rather, it is providing a warranty that it may not be able to pay. The customer receives no economic value, the regulator's interest is not served and the startup has an unquantifiable contingent exposure. It is a simple wrong of contract between parties of very unequal size and is the plainest example of the use of a rule of proportionality in the procurement process rather than altering the substantive rules.

The overall message is methodological. Evaluations of the effects of AI regulation that look at the statutory assignment of duties will systematically underestimate the incidence of regulation, since the actual assignment is contractual and follows the bargaining power. The following is testable, and is given as P2 below.

5. Classification and Capability

5.1 The instability of intended purpose

Risk-tiered frameworks are those that categorize systems based on the risk associated with the system's purpose. It's appropriate for a manufacturer that places a stable product into the market and is not a good fit for manufacturers where the purpose of the product has not been well established. A business can develop a

document-analysis system that has no specific field in mind, sell first to marketing persons, then to a customer in recruitment or lending customers and then the same system comes under a substantially heavier tier. No technical changes have been made. The firm has crossed a regulatory divide by a commercial decision.

The calculations of such under declared pivot rates are presented in Panel (a) of Figure 3. The point is that, regardless of the specific probabilities (which are merely illustrative), the chances that the applicable classification for an early-stage venture has changed at least once get significant within a few years, and are greatest for companies that are actively seeking product-market fit.

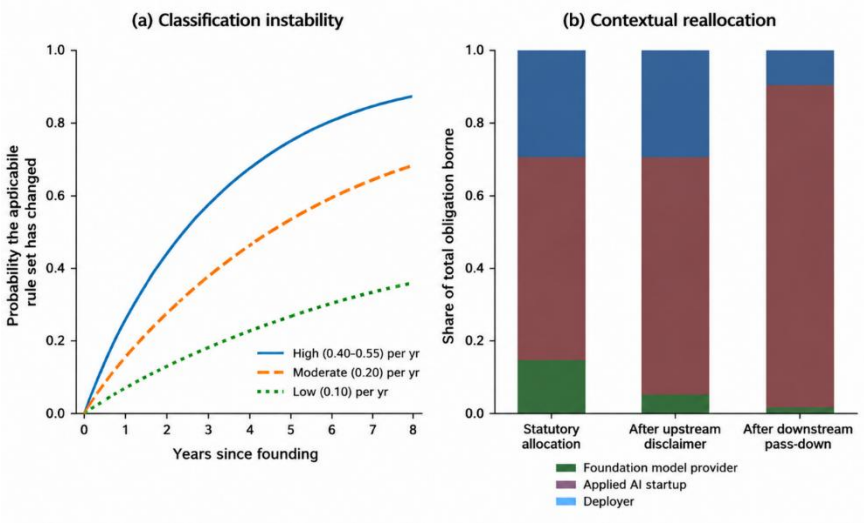

Figure 2. Panel (a) probability that a venture's applicable risk classification has changed, by pivot rate. Panel (b) reallocation of a unit obligation across the value chain from statutory to post-contractual allocation. All values are generated from declared parameters and illustrate model behaviour only; they are not empirical estimates.

Table 3. Sources of classification instability

Source	Mechanism	Typical trigger	Warning available to the firm	Mitigating design
Customer-sector migration	Same system enters a regulated domain	A single sales decision	Often none until procurement diligence	Classification tied to deployment, with deployer duties
Feature extension	Added capability crosses a definitional boundary	Product roadmap	Requires legal review of each release	Guidance with worked examples at the boundary
Model substitution	Upstream component replaced or upgraded	Provider deprecation	Little; deprecation timing is not controlled	Provider notice obligations
Scale threshold crossing	Obligations attach above a user or turnover level	Growth	Predictable if thresholds are transparent	Graduated rather than cliff-edge thresholds

Source	Mechanism	Typical trigger	Warning available to the firm	Mitigating design
Cross-framework divergence	Same system classified differently across jurisdictions	International sale	Requires multi-jurisdiction assessment	Mutual recognition of conformity evidence

5.2 Compliance cost is not compliance capability

The support instruments usually used for small businesses are financial instruments and on the implicit theory that the constraint is cost. When it comes to AI obligations this is only half true, and the half that it doesn't get is the reason for the poor performance of such tools. There are many duties that are not buying, but are organisational states. Organisational structures, defined roles and routine practices are assumed when implementing a risk management system throughout a product lifecycle, a quality management system, technical documentation that is kept up to date, logging with defined retention, post market monitoring and incident reporting. Having a document made by a consultant is not, by itself, an operating quality management system—nor will a document provide evidence of an operating system—so an auditor who reviews evidence of operation will recognize the difference.

Table 4 separates these. This difference is important as it implies which instruments of support are effective. Cost is addressed with grants and vouchers. Capability is addressed through shared compliance infrastructure, template management systems, sector level assurance bodies and staged obligations that are included as they become ready. The literature on internal auditing has consistently highlighted the need for a substrate that extends beyond a mere documentation to an actual organizational process, and that is the case with meaningful algorithmic accountability (Raji et al., 2020), while analyses of conformity assessment for new AI frameworks have also pinpointed practical constraints, such as the availability of auditing capacity (Mökander et al., 2022).

Table 4. Compliance cost versus compliance capability

Obligation type	Example	Nature of the requirement	What money alone buys	What is additionally required	Effective support instrument
Documentary	Technical documentation of the system	Artefact	Substantially the whole obligation	Accuracy and currency	Templates; grants
Procedural	Risk management across the lifecycle	Operating process	A written procedure	Roles, cadence, evidence of operation	Shared infrastructure; staged obligations
Structural	Quality management system	Organisational state	A manual	Sustained practice over time	Sectoral shared services; extended timelines

Obligation type	Example	Nature of the requirement	What money alone buys	What is additionally required	Effective support instrument
Evidentiary	Post-market monitoring and logging	Data infrastructure plus contractual access	Storage	Customer cooperation; instrumentation	Mandated data-sharing terms
Assurance	Third-party conformity assessment	External capacity	A fee, if capacity exists	An available and accredited assessor	Public investment in assurance capacity
Continuous	Incident reporting and corrective action	Standing readiness	A policy	Detection capability; escalation path	Sector-level reporting utilities

5.3 Uncertainty as a cost in its own right

The costs mentioned to date are compliance costs. Legal uncertainty comes with its own price, and this price is borne more heavily by early-stage firms, and works through three channels which are not often included in impact assessments.

Capital is the first one. Regulatory classification is also becoming commonplace in investment diligence and a product for which a venture cannot confidently classify can carry an element of risk that is hard to factor in the valuation process. The rational investor's answer is not "no" but "discount," "delay," or the setting of "conditions"– all costs incurred beforehand which never become obligations. The other is insurance. AI-related exposure is underdeveloped, exclusions are wide and the "no limits" contractual liabilities outlined in Section 4 often remain uninsurable with the residual risk falling on the balance sheet of the party with the fewest means to bear it. The third is roadmap distortion, when a venture does not choose to adopt a feature because it is not sure if it falls inside or outside a definitional boundary, the forgone product cannot be detected in an evaluation of the effects of the framework.

Such costs shouldn't be interpreted as proof of the need for less stringent frameworks. They are all evidence of the fact that clarity is in itself a valuable policy tool, that worked examples at classification boundaries are policy tools, not administrative courtesies, and that classification opinions, when requested, are substantive policy tools. Interpretive elaboration, not the text of the obligation, has been at the centre of the debates on algorithmic accountability, and the literature's caution against a wide gap between the nominal right and the content in its practice also holds true for a nominal safe harbour (Wachter et al. 2017; Kaminski 2019, 2023). An unresolved interpretive question is essentially a prohibition for a firm making a decision on building.

6. Governance Instruments Assessed

6.1 Timing and the instrument set

Figure 2 places the timing problem that must be solved by any instrument. When intervention is possible early in the diffusion of a technology it is tractable, but the consequences are not known; later, the consequences are known, but the technology is embedded and costs of change are high (Collingridge, 1980). The diagram starts with startups on the left and statutory frameworks on the right, so that the firms most impacted by a statutory framework are those that started under a different one or none.

Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 3

ISSN: 1540-3270

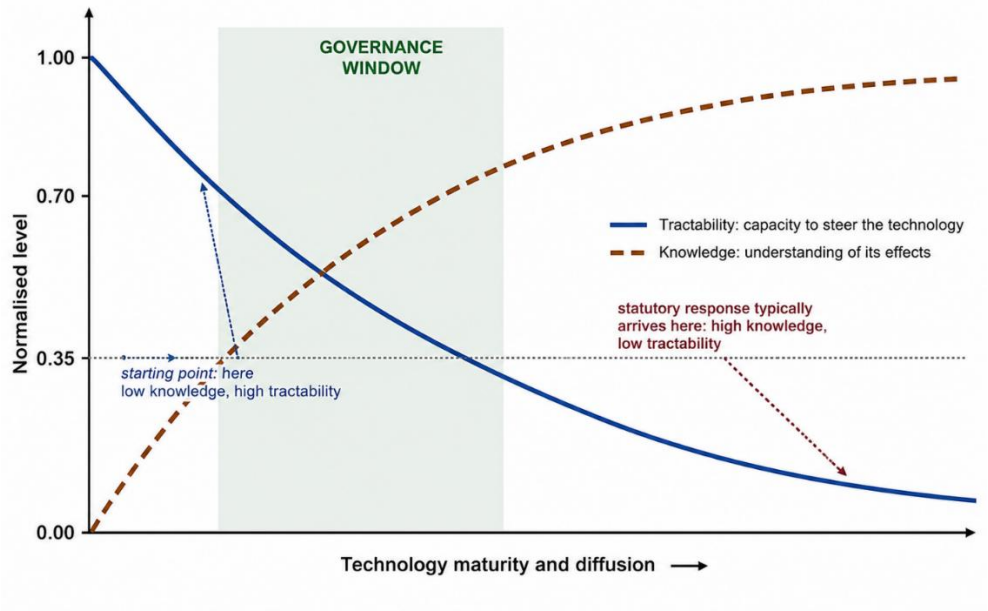

Figure 3. The governance timing problem. Tractability declines and knowledge of effects rises with technological maturity, leaving a limited window in which both are adequate. Startups enter early; statutory responses typically arrive late.

Table 5 assesses six instruments against the problems identified above. Three observations are worth drawing out.

Table 5. Governance instruments assessed against the identified problems

Instrument	Mechanism	Problem addressed	Evidence and experience	Principal limitation	Suitability for startups
Regulatory sandbox	Supervised operation with relaxed requirements and regulator dialogue	Timing; classification uncertainty	Established in financial regulation and extended to AI frameworks; some evidence of improved funding outcomes for participants	Capacity-limited; selection favours firms already able to apply	High for those admitted; irrelevant to the majority
Management-system standards	Certifiable organisational framework	Compliance capability	Widely adopted model in other domains	Certification cost and audit capacity; risks becoming a procurement gate	Moderate; valuable if subsidised, exclusionary if not
Assurance and audit ecosystem	Independent third-party verification	Evidentiary and structural obligations	Nascent; methods and accreditation still developing	Assessor capacity is the binding constraint	High if capacity exists;

Instrument	Mechanism	Problem addressed	Evidence and experience	Principal limitation	Suitability for startups
					currently thin
Safe harbours	Defined conduct confers protection from liability	Uncertainty; disproportionate exposure	Familiar from intermediary liability regimes	Requires specification precise enough to rely on	High; among the most valuable for small firms
Staged and graduated obligations	Duties attach with scale or maturity	Capability; fixed-cost regressivity	Used in financial and data regimes via thresholds	Cliff-edge effects; incentives to remain below thresholds	High if graduated rather than stepped
Procurement conditions	Public buyers impose requirements as purchase terms	Effective enforcement without statute	Already the operative constraint for many suppliers	Unaccountable, unappealable, and not proportionality-tested	Low; a principal source of the reallocation in Section 4

The first being sandboxes, the most common instrument advocated and the least discussed of all. Financial regulation has been evaluated as a capital access improved for participants, a result of the work functioning as a part certification signal (Cornelli et al., 2024) and the mechanism has been analysed as a wider transformation to experimental, adaptive regulation (Zetzsche et al., 2017). The critical literature has, however, offered significant challenges: sandboxes are capacity constrained; admission to the sandbox is discretionary; and there is a lack of procedural constraints normally attached to allocative decisions when a supervisor gives individualized relief (Allen, 2019). The most important condition in the current argument is the arithmetic condition. A sandbox that allows a handful of companies to participate is not the cure for a population of thousands of companies who make mistakes in start-up regulation.

The second relates to procurement – it is the most powerful tool in the table, and least analyzed. The operative regulator for many of the AI startups is the procurement function of a big customer rather than a supervisory authority: the requirements are imposed at an earlier stage, enforced more effectively by the contract award, and do not have to comply with the same proportionality, consultation and appeal requirements as public regulators. If the analysis in section 4 is correct, for a significant proportion of the market, procurement policy is a form of AI regulation, and the proportionality requirements, as well as liability caps on suppliers below specific limits, should be incorporated into procurement policy.

The third is the concept of safe harbours that get less attention and are significantly more scalable than sandboxes. Safe harbour does not depend on admission, on the capacity of the firm to supervise, and on an ongoing relationship: If a firm satisfies certain conditions, it has a specific protection. The design challenge is to be able to describe the conditions precisely enough to be able to rely on the condition without nullifying the underlying obligation, yet not so demanding that only large firms are able to meet it. This is not easy at all, but it is not an instrument capacity problem since it is a drafting problem.This is a non-easy one, but not an instrument capacity problem because it is a drafting problem - make sense of this?

6.2 Standards as private regulation

An instrument in Table 5 requires special treatment due to the systematic understatement of its governance implications. Horizontal AI frameworks tend to release their technical information in the form of standards created by private standards setting organisations and the acceptance of a conforming harmonised standard provides a presumption of compliance. The result is that the operative rules, that is, the detailed test procedures, formats for the documentation, and performance criteria a firm in fact has to satisfy, are drafted by technical committees, not by legislatures.

You have to pay to be on those committees, and they are costly. It takes years of continuous attendance, time from technical staff, and in many cases a paid membership in the body. While these are normal expenses for a large company, they are not feasible for a start-up and are likely to have predictable consequences: people writing these rules will not be among people who will be most affected by the fixed-cost incidence of these rules. This isn't a charge of bad faith, but the committee members generally work towards constructive standards. It is a structural phenomenon, that is, whose working assumptions become the default, and it is one that is identified in the literature of capture, although it is not the one that involves influence over legislators (Carpenter & Moss, 2014; Dal Bó, 2006; Stigler, 1971).

This trend towards a decentred and privately drafted form of rulemaking has been discussed in great detail in the regulatory literature, both for its efficiency benefits and its accountability drawbacks (Black, 2001; Brownsword & Goodwin, 2012). The implication in this argument is also practical: subsidising or facilitating small firms' involvement in standards is not a tangential industrial-policy move, but is a direct factor in the distribution of the burden of the framework. It costs a lot less than alternatives too, as it works at the drafting phase, not afterwards when remediation is needed.

7. Propositions and Analytical Protocol

Table 6 states seven propositions with the strategy that would test each. Several are testable from contract documents, tender records, and registry data that already exist.

Table 6. Propositions and testing strategy

Proposition	Statement	Testing strategy	Parameter of interest
P1	Formal obligation weight is highest at the middle of the algorithmic value chain, where information and control are both lowest.	Code obligations by chain position across frameworks	Position profile [w1]
P2	Post-contractual incidence differs systematically from statutory allocation, in the direction of the party with least bargaining power.	Content analysis of upstream terms and downstream procurement contracts	Reallocation share [b2]
P3	Liability caps in procurement contracts are uncorrelated with supplier enterprise value, producing uninsurable exposure for small suppliers.	Sample of contracts matched to supplier financials	Correlation [r3], expected near zero

Proposition	Statement	Testing strategy	Parameter of interest
P4	Applicable risk classification changes at least once for a substantial share of ventures within their first five years, driven by customer-sector migration rather than technical change.	Longitudinal tracking of ventures and their customer sectors	Hazard [h4]
P5	Financial support instruments improve documentary compliance but not structural compliance, because the binding constraint is organisational rather than monetary.	Comparison of compliance outcomes by support type	Contrast [b5d], [b5s]
P6	Sandbox admission is predicted by prior regulatory capacity, so the instrument reaches firms least in need of it.	Admission records regressed on pre-application firm characteristics	Coefficient [b6]
P7	Procurement requirements precede statutory obligations in binding on suppliers, so effective regulatory onset is earlier than statutory commencement.	Tender documents compared against statutory timelines	Lead time [t7]

Four measurement cautions apply. Contract terms are confidential in the private sector, so P2 and P3 are testable primarily on public-sector procurement, which is systematically different and may understate the severity of private-sector terms. Chain position is not recorded in any register and must be coded from product documentation, which introduces classification error that should be assessed by double coding. Survivorship is severe, because firms that failed to meet procurement conditions do not appear in supplier records, which biases any sample of contracted suppliers toward those that could comply. And regulatory classification is often not determined until a dispute or a procurement process forces the question, so observed classifications are endogenous to precisely the events P4 is meant to predict.

8. Discussion

8.1 Implications for framework design

There are four design implications to be drawn. The first is one of discharging obligation allocation, as opposed to the position in the supply relationship. An obligation that can only be fulfilled by the information of the person up stream should be attached to the person up stream or be accompanied by a disclosure obligation down stream; otherwise the obligation has created a duty that its addressee can not fulfil. The second is that statutory allocation should be accompanied by restrictions on contractual reallocation, as a no-regrets approach to contractual reallocation will undermine the distributional decisions made through statutory allocation, if one party to a contract has less bargaining power. The absence of these kinds of proportionate liability caps or prohibitions on flowing down obligations on which the recipient is unable to satisfy is an omission, not a considered choice – these are conventional tools known from consumer and commercial law.

The third is that classification should be graduated and not stepped where possible and reclassification should be accompanied by a transition period. Where a firm crosses a boundary as a result of a sales decision, then they are

not non-compliant on the day of the sale. The fourth is that support should match the type of the obligation, as set out in Table 4: money for documentary obligations, shared infrastructure and time for structural obligations, and public investment in the capacity to provide assurance for obligations that require it, for otherwise it is a prohibition.

The fifth implication is not a legal one, but rather an enforcement one. The idea of responsive regulation is that enforcement should be based on persuasion, and only escalate when the cooperation does not occur, and that the credibility of the escalation is more important than its routine use (Ayres & Braithwaite, 1992). Subsequent developments highlight that responsiveness should not just be limited to how the regulated firm behaves, but rather to its capacity and institutional structure (Baldwin & Black, 2008). This means that a non-compliant small business should be remediated on first offence, with a specific timeframe, and only be penalized on persistence or recklessness of non-compliance; and that supervisors should be provided with resources in order to provide advice to the small business, not just to sanction. This is not leniency, this is recognition that a company with twelve employees who fail a documentary requirement and a company with twelve thousand employees who fail a documentary requirement are two different things and require two different responses. This is akin to the argument by commentators on AI regulation for an adaptive and experimentalist approach to supervision, which is also adduced on the basis of technology's dynamism and the need to keep learning and adapting (Black & Murray, 2019; Guihot et al., 2017; Veale et al., 2023).

8.2 The innovation trade-off, stated honestly

It would be nice to say these changes are free, but these changes do not come free. All proportionality mechanisms establish a dividing line where some companies will operate towards, some staged obligations will give leeway for protection to those to whom it does not apply, and a safe harbour protects some behaviour that would otherwise be actionable. The theory of regulation offers plenty of justification for skepticism about arguments for lighter regulations because they are made with most vigor by those with least to lose from them, and the capture literature abounds with examples of the flexibility with which regulatory design adapts to the interests of incumbents (Dal Bó, 2006; Stigler, 1971).

It's not about regulating AI less, here. It's that its incidence ought to be purposeful. Whether or not it was deliberate, a framework that sets the same structural requirements for a 12-person firm and a 12,000-person firm has made a distributional choice and it benefits incumbents. If there really is a risk that only organisations with a much bigger capacity can deal with, then that is a defensible outcome, and in some cases with high stakes it is quite clearly true. It can't be defended as an untouched consequence of drafting for the paradigm large manufacturer case. The difference between these two is the practical content of proportionality which is formulated as "incidence is not to be assumed", but measured.

8.3 Limitations

There are three restrictions which must be explicitly stated. These are a set of conceptual proposals: the numbers are hypothetical and come from the parameters declared, not from experimentation; the numbers in Figure 1 are not empirical, but rather the structure of an argument; the levels of the attributes in Figure 1 are analytic judgements, not evidence of any kind. Second, this is an area of rapid change, and reliance on characterisations of frameworks, standards and their implementation dates or timings should be checked against current instruments to ensure they are accurate; guidance and delegated acts are also being issued and have a material impact on the points above. Third, the analysis is based largely on the European and Anglo-American regulatory environment, and on software-focused applied AI companies; its transferability to sectoral supervisors' attention-based regimes and to hardware-intensive or safety-critical industries that have well-established conformity regimes is not taken for granted.

9. Conclusion

Discussions on AI regulation have focused on the degree of regulation: whether the requirements are too strict or too loose, and whether innovation will be hindered. What has been forgotten in that framing is the previous and more tractable question: on whom do obligations fall, given a set of obligations? The question of what it is has

been answered in this paper, based on three design parameters. Coupling responsibilities to the provider of a system puts them in the middle of the algorithmic value chain – the place where information and control are minimal. Contract then transfers whatever the statute provided in the direction of the bargaining power and not the direction of the legislature. And classification by intended purpose is unstable in a population, which is unstable with regard to intended purpose.

All of these are arguments for regulating AI, not for exempting it from regulation. Each is addressed by design: the allocation of duties to the party who can discharge them, limit the contractual reallocation of duties that cannot be discharged, grade classification and consequences, and matching of support instruments to whether an obligation is a buy or an organisational state. Most important of all, and least talked about, is that for a large section of this market the regulating entity is not a regulator but a procurement function that is not subject to proportionality rules that have been developed by public regulators over a century. The seven propositions given here and the analytical protocol that they provide are meant to be testable, rather than just asserted, and hence refutable.

References

[1] Allen, H. J. (2019). Regulatory sandboxes. *George Washington Law Review, 87*(3), 579-645.

[2] Ayres, I., & Braithwaite, J. (1992). *Responsive regulation: Transcending the deregulation debate*. Oxford University Press.

[3] Baldwin, R., & Black, J. (2008). Really responsive regulation. *Modern Law Review, 71*(1), 59-94.

[4] Bamberger, K. A. (2010). Technologies of compliance: Risk and regulation in a digital age. *Texas Law Review, 88*(4), 669-739.

[5] Bennett Moses, L. (2007). Recurring dilemmas: The law's race to keep up with technological change. *University of Illinois Journal of Law, Technology & Policy, 2007*(2), 239-285.

[6] Black, J. (2001). Decentring regulation: Understanding the role of regulation and self-regulation in a "post-regulatory" world. *Current Legal Problems, 54*(1), 103-146.

[7] Black, J., & Murray, A. D. (2019). Regulating AI and machine learning: Setting the regulatory agenda. *European Journal of Law and Technology, 10*(3).

[8] Brownsword, R., & Goodwin, M. (2012). *Law and the technologies of the twenty-first century*. Cambridge University Press.

[9] Calo, R. (2017). Artificial intelligence policy: A primer and roadmap. *UC Davis Law Review, 51*(2), 399-435.

[10] Cobbe, J., Veale, M., & Singh, J. (2023). Understanding accountability in algorithmic supply chains. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1186-1197).

[11] Carpenter, D., & Moss, D. A. (Eds.). (2014). *Preventing regulatory capture: Special interest influence and how to limit it*. Cambridge University Press.

[12] Collingridge, D. (1980). *The social control of technology*. St. Martin's Press.

[13] Cornelli, G., Doerr, S., Gambacorta, L., & Merrouche, O. (2024). Regulatory sandboxes and fintech funding: Evidence from the UK. *Review of Finance, 28*(1), 203-233.

[14] Crootof, R., & Ard, B. J. (2021). Structuring techlaw. *Harvard Journal of Law & Technology, 34*(2), 347-417.

[15] Dal Bó, E. (2006). Regulatory capture: A review. *Oxford Review of Economic Policy, 22*(2), 203-225.

[16] European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.

[17] Guihot, M., Matthew, A. F., & Suzor, N. P. (2017). Nudging robots: Innovative solutions to regulate artificial intelligence. *Vanderbilt Journal of Entertainment & Technology Law, 20*(2), 385-456.

[18] ISO/IEC. (2023). *ISO/IEC 42001:2023 Information technology - Artificial intelligence - Management system*. International Organization for Standardization.

[19] Kaminski, M. E. (2019). Binary governance: Lessons from the GDPR's approach to algorithmic accountability. *Southern California Law Review, 92*(6), 1529-1616.

[20] Kaminski, M. E. (2023). Regulating the risks of AI. *Boston University Law Review, 103*(5), 1347-1411.

[21] Marchant, G. E., Allenby, B. R., & Herkert, J. R. (Eds.). (2011). *The growing gap between emerging technologies and legal-ethical oversight*. Springer.

[22] Mökander, J., Axente, M., Casolari, F., & Floridi, L. (2022). Conformity assessments and post-market monitoring: A guide to the role of auditing in the proposed European AI regulation. *Minds and Machines, 32*(2), 241-268.

[23] NIST. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. National Institute of Standards and Technology.

[24] Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33-44).

[25] Reidenberg, J. R. (1998). Lex informatica: The formulation of information policy rules through technology. *Texas Law Review, 76*(3), 553-593.

[26] Selbst, A. D. (2021). An institutional view of algorithmic impact assessments. *Harvard Journal of Law & Technology, 35*(1), 117-191.

[27] Stigler, G. J. (1971). The theory of economic regulation. *Bell Journal of Economics and Management Science, 2*(1), 3-21.

[28] Veale, M., Matus, K., & Gorwa, R. (2023). AI and global governance: Modalities, rationales, tensions. *Annual Review of Law and Social Science, 19*, 255-275.

[29] Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the Draft EU Artificial Intelligence Act. *Computer Law Review International, 22*(4), 97-112.

[30] Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law, 7*(2), 76-99.

[31] Zetzsche, D. A., Buckley, R. P., Barberis, J. N., & Arner, D. W. (2017). Regulating a revolution: From regulatory sandboxes to smart regulation. *Fordham Journal of Corporate & Financial Law, 23*(1), 31-103.