



AI RegRisk Think Tank

From Dynamic Architecture to Durable Stewardship: Building on Harvard's Governance Crossroads

A Response to "Governance at a Crossroads" by the AI RegRisk Think Tank

The dialogue on artificial intelligence governance is crowded with false choices. The recent paper from [Harvard Kennedy School, "Governance at a Crossroads,"](#) provides a pivotal contribution by rightly reframing the debate. It moves us beyond the sterile argument of whether to regulate and toward the far more critical question of how to govern a technology that evolves at an exponential pace.

The paper's proposed three-part mechanism—adaptable standard-setting, independent third-party audits, and clear liability frameworks—is a necessary architectural blueprint. But a blueprint is not a foundation. At the AI RegRisk Think Tank, we contend this architecture must be built upon a robust, and often-overlooked, prerequisite: **a defined, measurable, and operational governance program that is woven into the fabric of business.**

The dialogue around AI governance has focused too much on the existence of committees and not enough on what they are actually governing. The Harvard model provides the what. Our work focuses on the how. True, dynamic governance can only be realized when its external architecture is supported by an internal, programmatic commitment to stewardship. The ultimate goal is not that every institution makes perfect decisions, but that its program consistently puts it in a position to make **informed, defensible decisions** that find a responsible balance between the competing interests of shareholders, customers, employees, and society at large.



The Essential Blueprint: Affirming the Need for a Dynamic Framework

We wholeheartedly endorse the core premise of the Harvard paper. Static, top-down regulation is doomed to fail. The proposed dynamic model is the correct path forward for several key reasons:

It Embraces Evolution

By focusing on adaptable standards, the model acknowledges that today's best practices will be obsolete tomorrow.

It Leverages the Ecosystem

The reliance on third-party audits creates a market for trust and verification.

It Anchors Accountability

A clear liability framework creates predictable consequences, which in turn incentivizes proactive risk management.

This blueprint is the right one. However, an architect's plans are only as good as the builder's ability to execute them. The success of this entire model is contingent on the maturity of the very thing it seeks to govern: an institution's AI governance program. Without it, there is nothing of substance to audit and no fair basis for assigning liability.



The Missing Foundation: The Primacy of the Governance Program

The Harvard framework, if implemented without this programmatic prerequisite, risks inadvertently creating the same problems that have plagued cybersecurity assurance for years. The "cyber trap" saw an excessive focus on auditing technical controls and best practices, leading to a culture of compliance theater. It checked for good hygiene but often failed to address the core issue: was the organization's executive leadership at the helm, making informed, risk-based business decisions? There is no zero risk environment.

AI governance cannot afford to fall into the same trap. An over-emphasis on tactical standards for the technology itself will miss the forest for the trees. This is where governance must be understood as an internal, operational discipline first and a subject of external checks second.

□ We define this as **programmatic governance**—an integrated, auditable system that translates high-level principles like "Ethical AI" into practice.

This programmatic foundation is the missing layer that makes the Harvard architecture viable. It is built upon three interconnected imperatives:





Strategic Oversight & Accountability

1 The Leadership Imperative

This is the leadership imperative. It is the board and C-suite formally defining the organization's risk appetite for AI within a defined program, aligning it with core purpose, and establishing clear lines of authority for the program's execution. This ensures accountability begins at the top, rooted in the duty to oversee the program itself.



Integrated Operational Governance

2 The Organizational Imperative

This is the organizational imperative. It is the connective tissue of the program, weaving AI risk considerations into the fabric of the enterprise—from strategic planning and vendor management to the day-to-day processes of model development. This is what creates an auditable program, providing the substance for the "third-party audit" pillar to function effectively.

NOTE: it is the overall program that is owned by the board and executives, not just the operationalized governance, but the oversight and accountability is all about the boards and executives establishing being accountable for establishing and owning the governance program.





Tactical Execution for Unique AI Challenges

3

The Frontline Imperative

This is the frontline imperative. It comprises the specific tools and controls needed to manage novel AI risks. Critically, these activities are not the program itself; they are the outputs of a well-governed program that directs their use in line with the organization's strategic risk appetite.

These imperatives are not a checklist; they are an integrated system. They transform the abstract goals of "Ethical AI" into a single, answerable question:

Is the program fit for purpose?



Auditing the Program, Not Just the Outputs

With this programmatic foundation in place, the role of third-party audits is elevated and clarified. They are no longer a perfunctory check on technical controls but a meaningful validation of a living, breathing governance system.

Without a Program to Audit

The audit degrades into governance theater, verifying a veneer of compliance while missing deep, systemic risks. It incentivizes passing a technical test over building a culture of responsibility.

With a Program to Audit

The audit becomes a powerful capstone exercise. It provides independent assurance that the institution's governance program is not just well-designed but is operating as intended to put the organization in a position to make informed decisions.

This reframes the auditor's primary job: from searching for flaws in an algorithm to verifying the integrity and fitness-for-purpose of the governance program itself. This is the only way to avoid the cybersecurity trap and achieve a more scalable and meaningful form of oversight.



A Communal Contract: Fusing Architecture and Execution

The path forward requires a new, communal contract where all stakeholders recognize their role in building and relying on this programmatic foundation.



For Academia and Standard-Setting Bodies

Lead the sociotechnical research that informs adaptable standards focused on maturing governance through a programmatic approach, not just through technical specifications.



For Industry

Build, own, and execute a defined governance program as a strategic enabler, not a compliance cost. The ultimate competitive advantage will belong to those who can innovate safely at speed because their program provides the necessary guardrails and confidence.



For the Public Sector

Evolve your supervisory role to focus on the integrity of an institution's governance program. Set the expectation that every organization must be able to demonstrate its capacity for repeatable, informed, and defensible decision-making. No new rules need to be established, just a more mature interpretation and expectation.



For the Emerging Audit Ecosystem

Develop the expertise to audit the program as a discrete, executive-owned function. Prepare to validate the systems that produce decisions, not just the technical outputs.

Conclusion: From a Crossroads to a Shared Path

The Harvard paper masterfully leads us to the crossroads. The AI RegRisk Think Tank's mission is to pave that road with a durable foundation. Dynamic architecture and programmatic execution are two halves of the same whole. By fusing Harvard's external framework with an internal, communal commitment to programmatic readiness, we can build a future where governance is not a brake on progress, but the very trust infrastructure that enables it.

