September 2026

Independent Verification:
What Makes a Country Ready to Govern AI?

A roadmap to the five factors that shape where third-party evaluation can take root

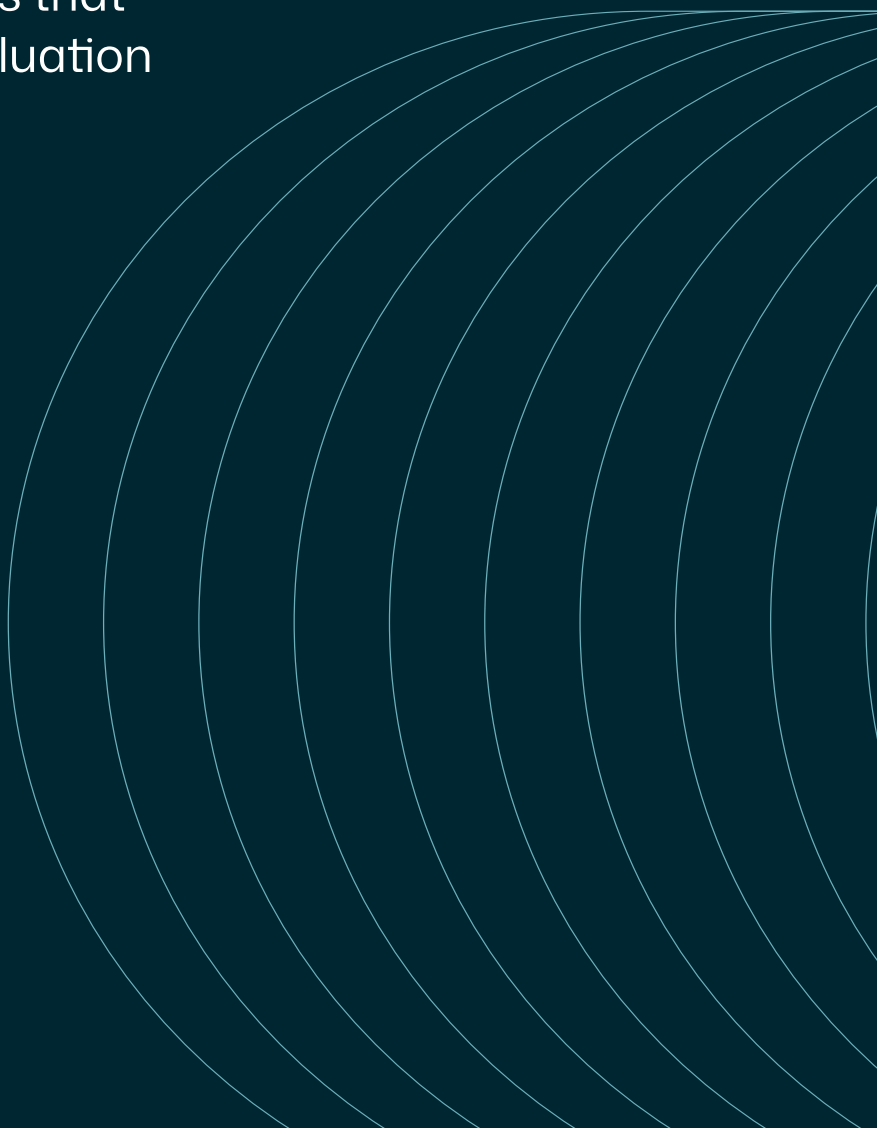

Fathom

Introduction

An Editorial Note from Fathom

Independent evaluation is moving from the margins to the center of the AI governance conversation. Anthropic, OpenAI, and Google DeepMind have each, in their own way, pointed toward third-party evaluation of models. Independent Verification Organizations (IVOs) are explicitly named in the FRONTIER Act, the first federal blueprint for independent AI evaluation, and in enacted state laws, most notably California's SB 813, which Fathom sponsored and which establishes the state's framework for independent verification organizations. Policymakers and industry stakeholders have arrived at the same insight: companies cannot grade their own homework.

At Fathom, we've spent a lot of time making the case for IVOs in the United States. But the idea doesn't stop at the US border, and neither does AI, or the risks and opportunities it brings. Independent evaluation is far more valuable when the systems that provide it are interoperable across countries. The US conversation matters, but it needs to go global.

There is broad consensus that advanced AI needs more oversight than exists today, and jurisdictions are experimenting with different models, mechanisms, and institutions. But what does flexible, interoperable AI governance, including independent evaluation, actually look like, and what determines whether a model can succeed from one country to the next?

At the highest levels of the global AI policy debate, independent research reveals that there are clear opportunities for governance, but such opportunities are being overshadowed by false dichotomies. Too often, potential governance mechanisms have died trying to thread the needle between economic development, innovation, and risk-fit governance. Strong conceptual buy-in, in other words, is not the same as readiness: converting that buy-in into on-the-ground adoption is complicated by a consistent set of political conditions that vary enormously from one capital to the next.

What we can say is that those conditions follow a pattern consistent enough to assess across jurisdictions. This report maps that pattern: the five factors that determine whether a government sees independent evaluation as an opportunity or an unnecessary constraint, what policymakers themselves told us about the constraints they work under, what determines if the framework travels well between countries, and how AI governance leaders can turn it from a diagnosis into a roadmap.

Methodology

In early 2026, Fathom commissioned an independent research study consisting of 50 semi-structured interviews with policy influencers and experts in six geographies: Australia, the EU (Brussels-based influencers and experts), Canada, France, Singapore, and the United Kingdom. The research set out to do the following:

- Investigate AI governance models and conditions outside the US
- Map current policymaker perceptions across a series of tech policy agenda-setting geographies
- Assess feasibility within each for interoperable governance mechanisms like an independent evaluation marketplace

Able to speak candidly and without attribution, policy elites, regulators, civil society leaders, and private sector actors were frank about the challenges across multiple constituencies and capitals.

What emerged from these conversations is that the path forward for innovative forms of AI governance is open, but will require careful coordination to navigate a complex political and geoeconomic environment that can be boiled down as one recurring set of variables.

The five factors that determine AI governance readiness

The first three factors below are determinative: “green lights” in each are paramount, and a blocker in just one is disqualifying in the current geopolitical context.

1. **Economic growth narrative:** how central AI is to a country’s economic story, and whether it is viewed as essential for growth and prosperity.
2. **Geopolitical standing and aspiration:** how a country sees its position in the world, and its willingness to assert independence or differentiate from global powers.
3. **Governmental self-perception:** how a government understands its role, its risk tolerance, and its openness to partnering with independent oversight bodies.

The latter two factors moderate the speed and difficulty of advocacy, but do not serve as immediate disqualifiers.

4. **Political architecture and elections:** the formal structures, veto points, political volatility and likelihood of power shifts, and institutional pathways that shape the speed and feasibility of policy adoption.
5. **Relationship with industry:** how government-industry dynamics influence buy-in, compliance, and the practical ease of implementing AI governance.

Factors

FACTOR 1:

How closely is AI linked to a country's growth agenda?

Our research revealed that most political systems view AI governance models through an economic lens. If AI is understood to be a core element of that country's national development plan, new regulatory frameworks were considered much less viable across every geography. As one UK government official put it, the government "is committed to the AI infrastructure build and committed to having it as a keystone of the economy," but the priority "is just economic growth at all costs."

Nations seeking additional AI-driven growth, even if they did not have a robust domestic industry, want to be seen as conducive to investment and are therefore wary of regulatory frameworks that chill investment potential. The imperative for many, as a Singapore-based head of AI governance noted, is "creating a safe space for people to innovate. [We don't] want to apply brakes to the innovation and stop them from launching."

However, the interplay between regulation and innovation is understood differently across geographies. Binary thinking of "innovation and economic gain" versus "regulation and safety" pervades some countries but not others. Stakeholders in countries with robust regulatory histories believe strongly that regulation is an important element of growth. A director of operations for a French AI company even went as far as to say that "regulation can even encourage innovation and the development of" domestic AI capabilities.

In this context, a sweet spot emerges: absent significant political or economic shifts, countries with stable economic growth, strong regulatory muscle memory, and the lack of strong domestic AI champions are most likely to advance innovative AI governance.

FACTOR 2:

AI governance cannot be advanced in a vacuum

Geopolitical standing, and even more importantly, geopolitical ambition, shapes a country's tolerance for novel governance models. Local governments are trying to strike a delicate balance on the scope, scale, and timeline for any governance efforts. On one hand, governance models that are too heavy-handed risk incurring geopolitical consequences or disinvestment from foreign companies. Another UK government official emphasized that the potential for regulation to be seen as "risking a fight" is a "big barrier to being particularly aggressive or forward-leaning." On the other, governance that is ill-suited to the models or too light risks breaking already-fragile public trust.

Countries aspiring to greater AI sovereignty, or seeking to differentiate themselves from the dominant powers, are most willing to explore different types of oversight. These geographies view governance leadership as a source of strategic leverage. An AI and digital regulation official in the EU noted that "the discourse on sovereignty around AI is very much emotionally [and] politically loaded, it's about [our] identity." In these instances, governance may not be seen as a constraint, but instead as a tool to shape global norms and reinforce a position as rule-setter.

FACTOR 3:

Institutional self-perception and precedent directly influence openness

Hard factors like economics and geopolitics matter most. However, our research showed that softer factors like institutional culture also substantially affect views around feasibility. Some governments and institutions like to act as first movers and pride themselves on experimentation and setting global standards, whereas others are more such as provinces, states, or regions, provide vari. As one Australian official put it, tending to "watch and see what others adopt and then follow," while others are more rigid.

Regulatory precedent matters as well. Does a government historically like to issue rules and standards itself? Or does it have experience delegating some technical validation to independent bodies? For some countries, as an AI and cybersecurity leader in France said, the "administrative and political model is based on compliance and regulation," whereas others are comfortable, as an Australian business leader noted, delegating "testing, reviewing, and analyzing" responsibilities to third parties who "then provide advice back to the government."

Digging further, the current US momentum around IVOs reflects a philosophy of governance that may translate most naturally to governments that are less centralized, in cultures with strong civil societies and a history of public-private partnerships. Prior experience with third-party certification, standards bodies, or hybrid public-private governance significantly affects receptivity. Where these models are familiar and proven, independent evaluation feels like a natural evolution rather than a revolution.

The level of trust in government, industry, and existing oversight mechanisms is also critical, as is how top-of-mind major societal issues are. Where distrust is elevated, or where people are primed by major AI issues, there is increased appetite for near-term action versus a preference for going through more traditional (longer, more iterative) regulatory approaches. A Canadian AI expert further noted that the urgency to regulate would likely increase "with more incidents surfacing that can hardly be ignored."

FACTOR 4:

Political architecture shapes the pathway for adoption, not the possibility

Beyond conceptual fit, formal political architecture determines the practical pathways for adoption. Different systems offer varying degrees of friction, bottlenecks, and possibilities for engagement.

High executive popularity, single-party control, or strong parliamentary dominance can smooth the path, whereas fragmented coalitions, weak mandates, and political volatility, including anticipated upcoming elections and potential changes in power, make action more challenging. In governments that move slowly, imminent elections can mean, as a UK government official noted, the perpetual "chance that the government that comes in is significantly less interventionist."

Subnational governments, such as provinces, states, or regions, provide varied potential entry points that allow for incremental adoption even where national-level momentum is constrained. An Australian government official emphasized that decentralized government systems, such as a subnational jurisdiction, "with a smaller footprint might be willing to experiment because there's less risk."

FACTOR 5:

A government's posture toward industry shapes the practical ease of governance

Practical viability is also shaped in part by a country's relationship with its AI and technology sector. While not as determinative as political, geopolitical, or cultural factors, industry buy-in can influence adoption, and our research found that meaningful regulation, as a UK government official said, would likely "need to have very strong buy-in from the tech industry." Resistance from key players may slow implementation, especially in countries where AI and tech are central to economic growth.

This dynamic is affected by the presence of sovereign AI and whether the government views it as in the country's best interest to truly champion those models, whether frontier AI companies are viewed as partners or adversaries, and the broader ecosystem of government-industry collaboration. Skeptical or adversarial relationships can create hurdles, and there is a fine line to walk between engendering public trust and provoking industry pushback. Governance mechanisms are more likely to be adopted in countries where evaluators are viewed as helpful third parties who can ensure both parties' needs are met.

Next Steps

Governance as a common language

These dynamics are not unique to the six geographies in our research. Rather, they offer a practical framework for assessing readiness across a much broader set of jurisdictions.

That broader applicability may ultimately prove to be the framework's greatest value. As countries increasingly pursue distinct approaches to AI governance, understanding why governments make different choices becomes as important as understanding what choices they make. The five-factor model developed through this research offers policymakers, civil society organizations, and industry leaders a common language for identifying where momentum exists, where barriers remain, and where investment in independent oversight is most likely to succeed.

Readiness, in other words, turns out to hinge less on a country's technological sophistication and more on its political conditions. Economic narrative, geopolitical ambition, institutional culture, political architecture, and relationships with industry, all taken together, determine whether a government perceives independent evaluation as an opportunity or as a constraint.

Turning the framework into a roadmap

Diagnosis is only useful if it changes what comes next. The five factors are not a scorecard for ranking countries against one another; they are a way for AI governance leaders to read the conditions they are actually working with, and to choose the argument that fits.

> **Where AI sits at the center of the growth story,** evaluation can be framed as a way to build the trust that sustains innovation rather than a barrier to it: the mechanism that allows adoption to scale in high-priority sectors like high-risk systems, child safety, and regulated industries, where our research found appetite to be highest.

> **Where a government wants a distinctive role on the world stage,** it becomes a form of strategic leverage: a way to shape global norms and reinforce a position as rule-setter, rather than inherit standards set elsewhere.

> **Where third-party certification and public-private partnership are already familiar,** independent evaluation can be introduced as an extension of assurance functions the country already runs, such as standards bodies, conformity assessment, sectoral auditors, and not as a new layer of oversight requiring new authority.

> **Where national action is slow or contested,** subnational governments can offer lower-risk ground to test what works, as can international coalitions. Since such forums shape the policy environment in which national governments and regulators operate, they are a valuable and light touch tool for demonstrating momentum to reluctant government and industry stakeholders.

> **Where relationships with industry are adversarial,** the first task is positioning rather than persuasion: establishing evaluators as helpful third parties who can ensure both parties needs are met, rather than as an enforcement arm in waiting.

None of these approaches requires a government to have green lights across all five factors at once. The work is figuring out which factor is binding, and using the ones that are available.

Conclusion

AI capabilities are evolving fast, and policymakers, regulators, and deployers know governance needs to catch up. Whereas models are released in weeks, governments, by contrast, legislate at a slower pace. There is a clear call for flexible, adaptive oversight capable of incorporating new technical standards, responding to emerging capabilities, and maintaining public confidence. There is a pathway to flexible and interoperable governance, but only if governments, civil society, deployers, and developers choose to buy in. With AI becoming increasingly dynamic, the institutions that govern it must become equally responsive.

The window for shaping global AI governance remains open, but it will not remain so indefinitely. The choices made over the next several years will help determine whether international oversight develops as an interoperable ecosystem capable of matching the global nature of AI itself. That won't happen everywhere at once, and it doesn't need to. It starts with recognizing which governments are ready to move, and building roadmaps suited to each, so the AI transition ahead is one shaped by more than just a handful of countries.