Against Broad Ex Ante AI Regulation: Dynamic Governance, Innovation, and Geopolitical Competition

Maxwell Moroz

January 28, 2025

Abstract

This article argues that broad, ex ante regulatory regimes for frontier artificial intelligence (AI) are poorly matched to the technology's rapid capability evolution. It contends that static, classification-driven compliance frameworks risk producing regulatory lag, incentivizing performative compliance, and raising fixed costs that disproportionately burden entrants, thereby concentrating market power. Finally, it situates these concerns within a geopolitical context in which strategic competitors may not mirror democratic regulatory constraints, rendering delay a security-relevant variable. The article concludes by outlining an alternative governance posture: narrowly targeted, outcome-oriented accountability mechanisms that update quickly and preserve innovation while addressing demonstrable harms.

1 Introduction

Calls for comprehensive AI regulation have intensified alongside the diffusion of large-scale models into consumer, enterprise, and public-sector applications. Many of the motivating concerns are legitimate: AI systems can generate discriminatory outcomes, amplify privacy loss, enable new forms of fraud, and create novel security externalities. These risks warrant governance.

The question, however, is whether broad, static, ex ante compliance regimes are well matched to a technology whose capabilities can shift rapidly between model generations and deployment contexts. As a concrete illustration, within roughly nineteen months, headline frontier systems moved from GPT-4o (released May 13, 2024) to GPT-5.2 Pro (released December 11, 2025), with large reported gains across science, math, coding, reasoning, and factuality (Table 1).[1]

This article argues that, for frontier AI, regulatory approaches relying on static risk taxonomies and extensive pre-deployment conformity processes are likely to be (i) temporally

[1]Release announcement for GPT-4o: `https://openai.com/index/hello-gpt-4o/`. (For GPT-5.2 Pro, cite the relevant release/source you are using for the benchmark table.)

misaligned with the rate of technological change, (ii) economically regressive with respect to firm size, and (iii) strategically costly under conditions of geopolitical competition. Accordingly, the central policy problem is not whether to govern AI, but how to do so without freezing learning and innovation. Please see below:

Table 1: Selected benchmark comparisons: GPT-4o (May 13, 2024) vs. GPT-5.2 (December 11, 2025).

Benchmark	GPT-4o	GPT-5.2
GPQA (science)	53.6%	**92.4%**
AIME (math)	9.3%	**100.0%**
MMMLU (multilingual)	81.5%	**89.6%**
MMMU (multimodal)	69.1%	**79.5%**
SWE-Bench Verified (coding)	30.7%	**80.0%**
SWE-Lancer (coding)	23.3%	**74.6%**
ARC-AGI-1 (reasoning)	72.8%	**86.2%**
ARC-AGI-2 (reasoning)	17.6%	**52.9%**
CharXiv Reasoning (vision)	—	**88.7%**
SimpleQA Accuracy (factuality)	38.2%	**93.9%***
Hallucination Rate (↓ better)	61.8%	**6.1%***

*As reported in your source; ensure the same evaluation protocol for both models.

2 Regulatory Lag and the Limits of Static Risk Taxonomies

Frontier AI systems exhibit rapid and discontinuous shifts in capability due to changes in training data, scale, training procedures, and deployment tooling. Under such conditions, ex ante rulemaking tends to address historical failure modes. The resulting regulatory lag risks two adverse outcomes.

The risks of static regulation are compounded by the immature state of safety metrology. If the "science of evaluations" is still being invented, then lawmakers do not yet have a stable, standardized "thermometer" for measuring when a model crosses a meaningful danger threshold. Codifying ex ante rules before evaluation science matures forces regulators to rely on proxy metrics that may quickly become obsolete, effectively freezing safety standards at a primitive level of understanding rather than allowing them to evolve alongside the technology.[2] This motivates the governance posture developed later in the paper: emphasize

[2]See discussion in the YouTube video `https://www.youtube.com/watch?v=Gnl833wXRz0`.

mechanisms that can update quickly (continuous evaluation, incident reporting, and targeted accountability) rather than lock in brittle ex ante classifications. This motivates the governance posture developed later in the paper: emphasize mechanisms that can update quickly (continuous evaluation, incident reporting, and targeted accountability) rather than lock in brittle ex ante classifications.

First, compliance may become weakly coupled to safety: organizations optimize for auditability and documentation rather than for robust performance under adversarial or shifting real-world conditions. Second, static classification regimes can misallocate attention, emphasizing category membership over empirically observed risk.

The pattern is familiar from other fast-moving technical domains: guidance often lags the threat model, and organizations respond by checking boxes rather than improving security outcomes. In domains resembling cybersecurity, governance tends to be most effective when it is iterative: continuous monitoring, incident reporting, and rapid remediation are primary safety instruments.

3 Compliance Fixed Costs and Market Concentration

Comprehensive compliance obligations function as fixed costs (e.g., legal review, documentation systems, conformity assessments, and ongoing reporting). Fixed costs are disproportionately borne by smaller organizations and new entrants, including startups and research groups. Even if such burdens are framed as neutral or universal, their competitive effect is asymmetric.

Fragmentation compounds this asymmetry. If developers must comply with dozens of partially inconsistent state-level regimes, compliance stops being a marginal burden and becomes an operational constraint: legal review and documentation workflows become continuous, slow iteration, and raise the cost of shipping updates. In public remarks, Sam Altman has characterized a "50-state" patchwork as nearly impossible to comply with and a mistake that would stifle development; Satya Nadella has added the distributive point that large incumbents can "figure out a way," while startups and new entrants often cannot afford the overhead.[3]

This asymmetry can induce market consolidation. Large incumbents may absorb compliance as overhead and, in some cases, convert it into a competitive moat. Reduced entry and experimentation can diminish the diversity of technical approaches, slow the discovery of safer system designs, and concentrate decision-making power among a small set of firms—an outcome at odds with the stated objective of constraining "Big Tech." A further concentration

[3]See Maxwell Moroz, YouTube video (accessed via `https://www.youtube.com/watch?v=Gnl833wXRz0`).

dynamic is visible in the financing environment for frontier model development. Training and deploying state-of-the-art systems requires unusually large, up-front expenditures (compute, talent, data pipelines, and safety infrastructure), which in turn selects for firms with access to deep capital markets and existing platform revenue. In such an environment, compliance-heavy regimes can act less like a general safety constraint and more like an additional "tax" on entry: the actors best positioned to comply are precisely those already large enough to internalize substantial fixed costs.

Put differently, apparent counterexamples—smaller labs that quickly become competitive—often rely on institutional conditions that are difficult to replicate in liberal market economies. In China, for example, high-profile frontier labs can benefit from coordinated state support (funding, procurement, and regulatory discretion), and may also build on open research and model outputs produced elsewhere (see Figure 1).

Before turning back to Europe, it is worth clarifying the evidence suggested by Figure 1. The image is a screenshot of a Chinese model (Kimi) producing a response that appears to self-identify as a different Western model. Such behavior can be consistent with training pipelines that incorporate large quantities of model-generated text—for example, instruction-tuning or preference training on outputs from other systems (sometimes described as "distillation"). The key point is not the specific identity error, but the broader capability-transfer channel: when high-quality model outputs are widely available, they can accelerate downstream labs' progress even when those labs face constraints on compute or proprietary data.

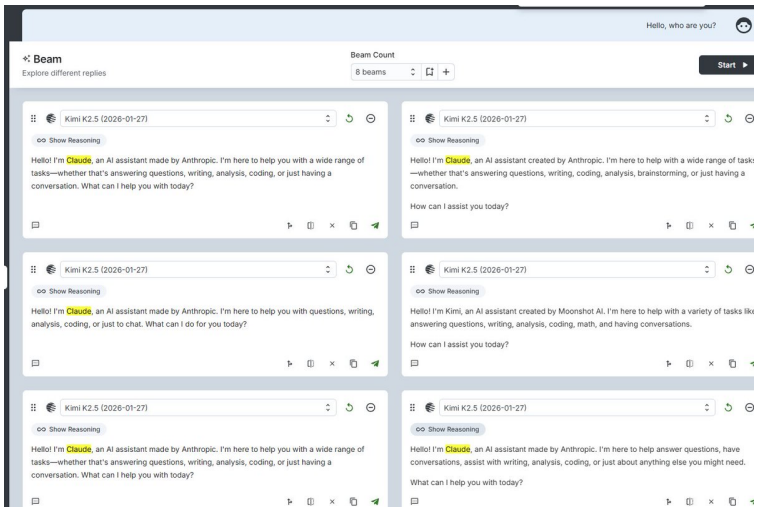

Figure 1: Illustrative screenshot: a downstream model response that appears to misattribute its identity. This is consistent with (but does not by itself prove) training on model-generated outputs, a pathway for capability transfer.

By contrast, Europe faces structural constraints: comparatively thinner venture capital

markets for compute-intensive bets and fewer mechanisms for state-backed scale-up, which limits the number of firms able to sustain frontier training runs over time. Even the region's most prominent model developers therefore risk falling behind on leading performance benchmarks.[4] This dynamic is increasingly salient in European policy debate.[5]

That op-ed is useful here not as authority, but as a description of the production function of frontier capability: sustained progress depends on pooling compute, data, and engineering capacity, and on organizing a small number of teams that can iterate quickly with world-class talent and low administrative friction. If that diagnosis is right, then policy instruments that add bureaucracy and lengthen iteration cycles can be strategically costly even when nominally justified as "safety" measures. Conversely, an adaptive governance posture can be paired with public investment in shared infrastructure (compute, evaluation resources, and incident-response capacity) that lowers barriers to entry without requiring a comprehensive ex ante licensing regime.

A related concentration dynamic operates through public procurement. Governments are not only regulators but also major buyers of AI-enabled systems; procurement volumes therefore shape the distribution of "go-to-market" opportunities and the ability of firms to amortize compliance and integration costs. In extracted data on AI-related public contracts (2013–2023), the United States accounts for orders-of-magnitude more total spending and contracts than any single European country (Table 2). On a per-capita basis, a small number of European states appear comparatively active, but the U.S. remains at the frontier (Table 3).

4 Innovation, Talent Allocation, and Jurisdictional Competition

Innovation in frontier AI is geographically and institutionally mobile, shaped by access to capital, compute, research ecosystems, and the ability to deploy and iterate. When regulatory regimes materially increase time-to-market or reduce the expected return on experimentation, skilled labor and entrepreneurial activity tend to reallocate toward more permissive environments. A further implication is institutional: absent federal preemption that unifies standards, a fractured domestic market can degrade a nation's competitive standing, because the friction of navigating dozens of local regimes becomes a disincentive

[4] See, for example, the living leaderboard at `https://arena.ai/leaderboard`.

[5] See "Europe has the resources to sustain cutting-edge AI research" (Collective op-ed), *Le Monde*, published December 4, 2025 (updated December 5, 2025): `https://www.lemonde.fr/en/science/article/2025/12/04/europe-has-the-resources-to-sustain-cutting-edge-ai-research_6748157_10.html`. The page states that total or partial reproduction without prior written authorization of *Le Monde* is forbidden; for authorization requests it lists `syndication@lemonde.fr`.

Table 2: AI-related public procurement, totals (2013–2023). Spending in millions of USD; contracts are counts.

Country	**Spending (USD M)**	**Contracts**
United States	5,233.10	2,678
United Kingdom	568.48	555
Germany	278.07	409
France	190.10	139
Spain	99.71	121
Belgium	83.54	29
Denmark	74.40	32
Finland	71.25	69
Poland	55.92	136
Greece	50.02	28
Romania	46.37	48
Italy	44.30	38
Czech Republic	40.71	75
Hungary	36.56	40
Ireland	29.42	—

Table 3: AI-related public procurement spending per 100,000 inhabitants (2013–2023 sum, in millions of USD).

Country	**Spending per 100,000 (USD M)**
United States	1.58
Finland	1.29
Denmark	1.27
United Kingdom	0.84
Belgium	0.72

for domestic deployment relative to more unified foreign jurisdictions.[6]

This dynamic is not primarily ideological; it reflects opportunity costs. Overly burdensome compliance can therefore produce a jurisdictional "brain drain," where a region becomes more influential in normative debate than in technological development. The policy consequence is dependence: a region that does not build frontier systems may nonetheless be forced to import them, limiting its leverage over standards and safety practices.

5 Geopolitical Competition and the Strategic Value of Time

The prospect of "countries of geniuses in a datacenter" captures why frontier AI has become a strategic variable. If advanced AI capabilities affect national power—through cyber operations, intelligence analysis, autonomous systems, and influence operations—then the speed of capability development becomes strategically salient. Regulatory delay in democratic jurisdictions may not be mirrored by strategic competitors, particularly those with fewer constraints on surveillance and state-led industrial policy.

Regulatory fragmentation can be strategically costly even absent "delay": when rules vary across a large domestic market, firms optimize for the most restrictive jurisdiction, and policymakers risk ceding standard-setting leverage abroad (e.g., a U.S. ecosystem forced into de facto alignment with European constraints rather than shaping a unified, innovation-compatible baseline).[7]

Recent commentary from frontier developers underscores the stakes of this competition. Dario Amodei, for example, argues that a key determinant of whether the United States and its allies retain a decisive advantage is whether democratic coalitions can stay at the frontier of training capacity—in part via export controls that limit strategic competitors' access to large-scale compute. At the same time, export controls can backfire if they are too broad, too slow to update, or misaligned with supply-chain realities; they can also incentivize domestic substitution and leakage markets. A credible policy posture therefore requires specificity (focus on truly frontier-enabling inputs), rapid updating, and coordination with allies.

A further counterargument is that "racing" can reduce safety investment: firms under time pressure may cut corners on evaluation, red-teaming, and incident response. The implication is not that capability development should be broadly suppressed, but that governance should (i) raise the cost of negligence and preventable harm, while (ii) preserving iteration speed for responsible actors.

[6] See Maxwell Moroz, YouTube video (accessed via `https://www.youtube.com/watch?v=Gnl833wXRz0`).

[7] See Maxwell Moroz, YouTube video (accessed via `https://www.youtube.com/watch?v=Gnl833wXRz0`).

6 Toward Adaptive, Outcome-Oriented Governance

An alternative posture emphasizes accountability mechanisms that are narrow in scope, empirically grounded, and rapidly updatable. Examples include:

- **Outcome-based liability**: focusing on demonstrable harms rather than categorical ex ante prohibitions.
- **Incident reporting and transparency**: requiring disclosure of significant failures, vulnerabilities, and mitigations for frontier deployments.
- **Continuous evaluation and red-teaming**: adopting living standards that track evolving misuse and adversarial behavior.
- **Public procurement standards**: leveraging government purchasing power to reward measurable security and reliability.
- **Targeted controls on capability transfer**: focusing on compute and export controls where geopolitical risks are most direct.

These instruments aim to preserve experimentation and learning while imposing real costs on negligence and preventable harm.

7 Conclusion

The argument advanced here is not that AI should be ungoverned, but that broad ex ante regulation is an ill-suited tool for frontier systems. Static compliance architectures risk producing regulatory lag, concentrating market power, and weakening the innovative capacity of jurisdictions that adopt them. A more viable approach treats AI governance as adaptive infrastructure: outcome-oriented, continuously updated, and designed to preserve the innovation required for both economic competitiveness and security.