Fathom on Private AI Governance

According to many of the sharpest minds in artificial intelligence, we are just a few short years away from truly transformative AI applications that will disrupt many aspects of modern life. At Fathom, it's our mission to find, build, and scale the solutions needed to help society transition to a world with AI and navigate this disruption.

As we confront this massive societal transition, the choices we make now about how we develop, deploy, and govern AI will determine whether we harness its extraordinary potential or allow it to unfold haphazardly – or even not at all.

Over the past year, Fathom has convened Americans from all walks of life – technologists, parents, teachers, business leaders, factory workers, artists, academics, national security experts, church parishioners, and civil society advocates – to understand their views on our collective AI future.

Fathom conducted multiple polls, spent hundreds of hours conversing with stakeholders from across society, and hosted The Ashby Workshops, which brought together over 180 AI experts from business, government, academia, and civil society.

Across all our work, we continue to hear one common refrain: **society needs an AI governance model to preserve and deepen our American values, while allowing innovation to flourish.**

AI governance requires creative thinking beyond government systems, which cannot keep pace with the rate of innovation, and Big Tech should not be the only ones who decide how this transformative technology shapes our society.

To that end, we believe we've uncovered a legally, technically, and politically viable governance solution: Private AI Governance.

To be successful, Fathom believes Private AI Governance will require the formation and licensing of private AI standards-setting regulatory bodies. These bodies would be composed of a wide range of stakeholders to ensure their effectiveness, fairness, and balance – for the sake of simplicity, we'll call them Multi-stakeholder Regulatory Organizations or MROs for short.

How Private AI Governance Works

Given today's political environment and the relative legislative gridlock at the federal level, enabling MROs will require state legislatures to act.

The following will enable an effective Private AI Governance structure and the creation and deployment of impactful MROs:

1. First, a state legislature authorizes its Attorney General to license MROs. Licenses will only be granted to organizations that are sufficiently independent from industry and have both technical and legal credibility.
2. Licensed MROs will then govern AI platforms that *voluntarily* seek certification. MROs will identify, mitigate, and monitor risks, collecting the necessary information from AI platforms to do so, and certifying these AI platforms' compliance with best practices for the prevention of personal injury and property damage.
3. Certification will provide AI platforms with protection of a statutory affirmative defense against tort liability for personal injuries and property damage—a significant boon to these platforms. However, if an AI platform acts in ways legally deemed reckless, deceitful, or grossly negligent, the affirmative defense tort liability protection would not apply.
4. MROs can also revoke any AI platform's certification for non-compliance, ending their tort liability protection.
5. Just as MROs can revoke an AI platform's certification, so too can the Attorney General revoke an MRO's license if their behavior is found to be negligent – while the MRO is a private entity, it is still subject to public oversight.

The Benefits of Private AI Governance

By marshaling private-sector expertise and technical innovations, MROs will enable an agile, public-private regulatory model designed to promote innovation, ensure the security and transparency of AI platforms, reduce regulatory uncertainty, and build societal trust.

MROs also provide AI platforms with legal certainty about the requirements they must satisfy to responsibly develop AI technologies and prevent unnecessary harm. Reducing litigation risk encourages investment, fosters innovation, and ensures regulatory adaptability. It also promotes heightened care and adherence to best practices while maintaining accountability and effectively balancing risk mitigation with consumer and public protection.

Therefore, MROs transform regulation from a compliance burden into a competitive advantage by proactively setting clear standards, creating tailored pathways for established companies and startups, and offering concrete legal and economic incentives for compliance.

It's not just about managing risks; it's about accelerating responsible growth and empowering businesses to innovate and thrive confidently in an AI-driven economy.

Importantly, the nimble nature of MROs as private governance organizations—coupled with their *public* oversight—allows AI governance to adapt at the speed of technological relevance while still answering to the public's stated policy goals.

What's Next

Fathom exists to help society navigate the transition to a world with AI—it's already coming. In fact, the transition is taking place around us daily, and we don't have any time to waste. In this political environment, it's clear that federal legislation governing AI adoption is unrealistic. Therefore, it's incumbent on the states to function, as they often have, as the laboratories of democracy.

While hundreds of AI bills have been introduced and passed across the states, none have yet dedicated their focus to private AI governance. It is our sincere hope that one or more states will step up and lead the way towards the creation of MROs as we envision them, capable of providing certainty to businesses, protections for consumers, and a significant benefit for society as a whole.

Andrew Freedman is co-founder and Chief Strategy Officer of Fathom.org, helping find, build, and scale the solutions needed for our transition to a world with AI.