

How Fathom’s Research Led Us to Propose Multistakeholder Regulatory Organizations

By Andrew Freedman, Chief Strategy Officer, Fathom

Fathom’s research into AI governance began with recognizing that while AI presents a formidable governance challenge, our difficulties are not entirely novel. All emerging technologies represent and demand the exploration and navigation of a new frontier – an endeavor characterized by radical uncertainty, risk, and the desire to promote innovation.

We’ve been here before – the United States has faced many a frontier – and our track record of success would suggest that somewhere along the way we discovered critical insights for the governance of emerging technologies. Reviewing the history of formerly frontier technologies including the internet, financial sector, and railroads, we rediscovered one key to their respective success: private governance.

Private Governance Through the Ages

For millennia prior to the early 1880s, time was measured by the placement of the sun, with midday determined by when the sun was highest in the sky over a particular village or town. This resulted in at least 144 different time zones in North America,¹ which by the mid-1800s was creating major headaches for the railroad industry. Passengers often missed their trains owing to “miscommunications” between stations, and train collisions became increasingly frequent. Inspired by the successful adoption of “Railway Time” (now known as “Greenwich Mean Time”) by all British railway companies in 1880, in 1883, all U.S. railroad companies came together to create time zones across America and Canada. These four time zones remain relatively unchanged to this day, and yet the federal government only codified them 35 years after the fact, in 1918.

More recently, the core architecture of the internet – from the assignment of internet domain names and IP addresses to the standardization of protocols and traffic routing – emerged from rules created by public-private, multistakeholder organizations. Established in 1986, the Internet Engineering Task Force (IETF) has coordinated the operation, management, and evolution of the Internet ever since. IETF comprises a large international community of

¹ United States Department of Transportation (January 2023), History of Time Zones and Daylight Saving Time (DST). Available at <https://www.bts.gov/explore-topics-and-geography/geography/geospatial-portal/history-time-zones-and-daylight-saving>

network designers, operators, vendors, and researchers, who together make decisions on the basis of “rough consensus and running code” across more than 100 working groups. Its successes include the development and management of the Internet Protocol Suite, which includes the Transmission Control Protocol (TCP) and Internet Protocol (IP) that facilitate communication between devices on the Internet.

Time and time again, we’ve seen private governance in emerging industries result in the best decisions about best practice, driving an ecosystem of smarter choices.

The Opportunity for Private AI Governance

Driven by the belief that effective governance requires broad multistakeholder input, we began to circulate and solicit feedback on a private governance approach to AI. We conducted multiple polls,² which found that 65% of voters would trust a public-private coalition made up of AI companies, scholars, and policy experts to develop proper guardrails for AI, compared to only 50% who would trust the companies themselves and 45% who would trust the federal government. We then took the idea to hundreds of stakeholders over several months, who echoed support for private governance. Different people had different reasons, of course, but many felt that a private AI governance model could bring the right people to the table in a way that public governance currently cannot. Encouraged by the sense that we were onto something, we decided to push onward.

Just as Fathom’s theory of change begins with listening, so too do its solutions. Drawing heavily on the expert input we’d received, we began designing a private governance solution to create national standards for AI development.

Stakeholders at [The Ashby Workshops](#), which brought together over 180 leaders from business, government, academia, and non-profits, expressed the concern that conversations about governance remain tied to current-day model capabilities, with insufficient consideration for how these capabilities may evolve over time.³ The risk and analogous concern is that any near-term regulation might overindex on today’s capabilities, risks, and opportunities, quickly becoming outdated as the frontier moves on. This became a key concern for us as we were developing our model, and the extent of our collective uncertainty came into focus: we don’t even know what we don’t know yet.

² Fathom.org, Fathom’s Inaugural Report (July 2024). Fathom.org, AI at the Crossroads: Public Sentiment and Policy Solutions (September 2024). Both available at <https://fathom.org/resources>.

³ Fathom.org, The Ashby Workshops 2025: Report Highlights (February 2025). Available at <https://fathom.org/resources>.

No one knows for certain how frontier systems will evolve, all the ways they'll be deployed, and what their impacts will be. Even more challenging, AI will rework the conditions in which markets and governance play out. For instance, many systems depend upon the assumption that there is substantial friction associated with procuring cognitive labor: it is expensive and time-consuming to hire advanced lawyers, civil engineers, biologists, and other specialists. But AI may fundamentally change this reality, and with that change, the assumption of friction will diminish or vanish. AI will also rework *how* markets and governance play out. AI is *itself* a governance instrument—a tool that can be used to accomplish the ends of policymakers. Whether and how, the government uses AI will determine not just the capabilities of government but the nature of policy itself.

The enormity of these challenges reinforces the need for flexible and adaptive governance; not only to allow it to evolve in tandem with the technology it seeks to shepherd, but to allow us the freedom to make mistakes. The need for patience and a willingness to get things wrong came through as a key insight from The Ashby Workshops, but such a willingness requires a governance model that allows us to fix our mistakes.⁴

These reflections sent us in the direction of public-private regulatory markets. The regulatory markets approach to governance would have governments require that industry developers obtain regulatory services from a private regulator, which would itself be overseen by the government, preserving democratic oversight and accountability. This allows the government to set a direction for AI governance, and the market to pioneer the regulatory methods that best achieve that vision. Crucially, the private regulator can experiment and iterate on its methods, allowing it to rectify its mistakes and respond dynamically to an emerging and evolving industry.

The extent of our uncertainty also led us to decide against proposing a new, tailored liability regime for AI until we know what we're dealing with, and to work within the confines of existing law until then. Several of our advisors recommended we start with tort law for two reasons. Tort exists in every state as part of the common law. This means that defining the standard of care owed to consumers by industry in one state helps to shape the standard nationally, driving national standards in the process. Tort law also maps to some of our most significant near-term concerns regarding AI development: physical harm and property damage. This is a rough mapping – physical harm and property damage don't cover all known and conceivable AI harms – but it's a strong starting point.

⁴ Fathom.org, The Ashby Workshops 2025: Report Highlights (February 2025). Available at <https://fathom.org/resources>.

The Fathom Solution: Multistakeholder Regulatory Organizations

Bringing together the ideas of private governance, regulatory markets, and the existing tort regime, we arrived at the concept of a multistakeholder regulatory organization (MRO). MROs would be licensed by a state to audit and certify AI companies. MROs would then identify, develop, and iterate best practices for AI development and deployment. Developers could opt into a certification process and, if successful, earn legal clarity in cases concerning liability for negligence resulting in personal injury and property damage. In doing so, MROs provide developers with legal clarity, empowering them to continue to push at the frontier, while providing consumers with better, safer AI products and services.

We aren't alone in viewing MROs as the solution to the AI governance problem. In California, Senator Jerry McNerney recently introduced a bill to authorize MROs, which Fathom supports. This legislation – [SB 813](#) – has also won the backing of a wide swath of scholars, researchers, and thought leaders, who signed an [open letter to California legislators](#) urging support for the bill.

The signers, who have dedicated their careers to the research of artificial intelligence, technology policy, and governance, represent different perspectives and have historically held varying views on AI. Yet they all "agree that SB 813 stands out as the most responsive, well-designed model yet, able to adapt and evolve over time with the underlying technology."

Taking and tailoring lessons from the past, we feel strongly that MROs are the solution to the formidable governance challenge that is AI – providing a conservative yet innovative approach to governance that ensures our technologies preserve U.S. competitiveness and shape the world for the better.