

What are regulatory markets and how can they help ensure that AI is safe, fair, and ethical?

By Jovana Jankovic

Misinformation and election interference. Racial bias in the criminal justice system. Self-driving car accidents. Artificial intelligence is changing the world so quickly, it's unclear how regulators can keep up with ongoing developments.

As the COVID-19 pandemic continues to affect our communities and the world, the need for nimble and responsive AI regulation is becoming even more clear. For example, the official coronavirus contact-tracing app in North Dakota was recently found to be sharing people's personal data with third parties like Foursquare and Google—something which is not disclosed in the app's privacy policy.

Enter regulatory markets.

A recent paper from Gillian Hadfield, director of the Schwartz Reisman Institute for Technology and Society, and Jack Clark of AI research lab OpenAI proposes a novel governance and regulatory framework for ensuring that AI can be used safely, ethically, and with the public interest in mind, regardless of industry or purpose.

Hadfield and Clark present the idea of “regulatory markets,” in which market-based incentives like competition, positive industry reputation, and expanded market access can spur innovation, efficiency, and accountability in the ways in which we regulate AI. The model isn't, however, ignorant of the crucial aspects of public safety and government oversight—in fact, it envisions foundational roles for these very important players.

The reality is that public regulation of AI is insufficient—because governments lack the resources and expertise to keep up with lightning-quick developments in AI.

“If we try to regulate 21st century technology with 20th century tools, then we'll get none of the benefits of regulation and all of the downsides,” writes Hadfield in Venture Beat. “This is no small task. If we need to reinvent the rules to keep pace with technological change, where do we start?”

Hadfield's and Clark's proposed three-party regulation model sees governments licensing and overseeing an industry of private regulators who then monitor and keep companies who are using and developing AI in check. Some existing models for hybrid public-private regulation miss the crucial aspect of “regulating the regulators,” in scenarios like, for example, the Boeing 737-MAX crashes of early 2019. But, Hadfield and Clark tout successful regulation examples from which we can learn valuable lessons, such as the private certification providers that oversee medical device quality in a number of countries and the UK's regulation of legal service providers.

The bottom line is that a market-based regulatory framework—with essential government oversight—could make use of the same kind of innovation and technical expertise from which AI itself has grown. The authors even envision a “start-up dynamic” for these kinds of “new entrepreneurial ventures.”

Still, an agile and novel regulatory framework need not be the Wild West. A global competitive market for AI regulators should involve a “formal, publicized procedure for regulating the regulator,” write Hadfield and Clark. “Achieving regulatory competition requires close attention to the design of the market,” and also sees a crucial role for public institutions as overseers.

Overall, regulatory markets would offer “more effective and less burdensome ways to provide a service with hard government oversight to ensure that whatever the regulatory market produces, it satisfies the goals and targets set by democratic governments,” writes Hadfield.

So, how should policymakers and business leaders start thinking about implementing this regulatory framework for AI?

Establishing a government oversight body is a crucial first step. Then, enabling legislation that creates a market for regulatory services, or fostering the development of existing self-regulating frameworks to grow into a regulatory market, follow.

Regulating AI may be uncharted territory, but regulating the safety and security of technology for the public good certainly isn't. In Hadfield and Clark's proposed system, “building safe machine learning systems requires not only technical innovations, it also requires regulatory innovations.”

Want to learn more? [Read “Regulatory Markets for AI Safety,” by Gillian Hadfield and Jack Clark online.](#)

The Hadfield/Clark paper has been covered in [The Observer](#), and Hadfield has written the concept of regulatory markets in [Venture Beat](#).



The [Schwartz Reisman Institute](#) aims to deepen our knowledge of technologies, societies, and what it means to be human by integrating research across traditional boundaries and building human-centred solutions that really make a difference. We want to make sure powerful technologies truly make the world a better place—for everyone. Comprising diverse areas of inquiry, from machine learning, computer engineering, epistemology, systems theory, and ethics to legal design, systems of governance, and human rights, our research agenda and solutions stream cross traditional boundaries and are fundamentally inspired by a commitment to reinventing from the ground up.