Ten Principles for Regulation That Does Not Harm AI Innovation

By Daniel Castro | February 8, 2023

Artificial intelligence (AI) has the potential to create many significant economic and social benefits. However, concerns about the technology have prompted policymakers to propose a variety of laws and regulations to create "responsible AI." Unfortunately, many proposals would likely harm AI innovation because few have considered what "responsible regulation of AI" entails. This report offers ten principles to guide policymakers in crafting and evaluating regulatory proposals for AI that do not harm innovation.

As artificial intelligence (AI) continues to improve, opportunities to use the technology to increase productivity and quality of life will flourish across many sectors of the economy, including health care, education, transportation, and more. In response, policymakers have proposed a variety of regulations to address concerns that this coming wave of AI systems may cause harm. Minimizing potential harm from AI systems is an important goal, but so too is maximizing the potential benefits of AI systems. Implementing many of these proposals, especially in their current form, is likely to have serious consequences because many of AI's potential benefits—including opportunities to use the technology both to save lives and to improve living standards—may be delayed or denied with poorly crafted regulations.

Policymakers want AI systems that do not cause harm, but they have not mastered the art of creating regulations that do not harm AI innovation. If policymakers decide that regulation is necessary, then to avoid slowing AI innovation and adoption, they should follow these 10 principles:

1. Avoid pro-human biases.
2. Regulate performance, not process.

3. Regulate sectors, not technologies.
4. Avoid AI myopia.
5. Define AI precisely.
6. Enforce existing rules.
7. Ensure benefits outweigh costs.
8. Optimize regulations.
9. Treat firms equally.
10. Seek expertise.

PRINCIPLES FOR AI REGULATION

The motivations for AI regulations vary. Some are guided by reasonable safety concerns, such as to ensure autonomous vehicles do not cause undue risk to occupants or pedestrians or that AI-based medical devices work as intended. Some are motivated by fears that government will use AI systems to invasively surveil its citizens. And still others are motivated by more nebulous concerns that AI systems may not be sufficiently "ethical," "trustworthy," or "human centric."[1]

Regulation is a means, not an end. Because regulation can limit innovation, policymakers should always explore nonregulatory options to achieve their goals. For example, industry self-regulation or codes of practice may prove equally or more effective than top-down regulation.[2] One reason these "soft law" approaches are often superior to "hard law" is because the private sector can typically move faster than legislators and regulators to create and implement new rules, as well as update those rules in response to changing conditions. Given the fast rate of change in the technology industry, especially with AI, this flexibility and nimbleness allow industry groups to address emerging threats promptly.

As summarized in table 1, there are ten key principles policymakers should keep in mind when creating regulations for AI so as not to stop or deter innovation. Each of these is explained in more detail below.

Table 1: Principles for regulation that does not harm AI innovation

Principle	Description	Rationale
Avoid pro-human biases.	Allow AI systems to do what is legal for humans (and prohibit what is illegal too).	Holding AI systems to a higher standard than for applies to humans disincentivizes the technology's use.

Principle	Description	Rationale
Regulate performance, not process.	Address concerns about AI safety, efficacy, and bias by regulating outcomes rather than creating specific rules for the technology.	Performance-based regulations allow for flexibility in how to meet objectives and does not impose potentially costly and unnecessary rules on AI systems.
Regulate sectors, not technologies.	Set rules for specific AI applications in particular sectors rather than creating broad rules for AI technologies generally.	Context matters. An AI system to drive a vehicle is different than one to automate stock trades or diagnose illnesses, even if they use similar underlying technologies.
Avoid AI myopia.	Address the whole problem rather than fixate on the portion of a problem involving AI.	Many problems need to be solved regardless of whether they involve AI. Focusing only on the AI-portion of the problem often distracts from resolving the bigger issue.
Define AI precisely.	Define AI clearly to avoid inadvertently including other software and systems within the scope of new regulations.	AI covers a broad range of technology and is integrated into many products. Policymakers should not use broad definitions of AI if they only intend to regulate machine learning or deep learning systems.

Principle	Description	Rationale
Enforce existing rules.	Hold AI accountable for adhering to existing regulations.	Many laws already address common concerns about AI, such as those relating to worker safety, product liability, discrimination, and more.
Ensure benefits outweigh costs.	Consider the full potential costs and benefits of regulations.	Costs, including both direct compliance costs and indirect innovation and competitiveness costs, impact the merits of a regulatory proposal.
Optimize regulations.	Maximize the benefits and minimize the costs of regulations.	Policymakers should find the most efficient way to achieve their regulatory objective.
Treat firms equally.	Apply rules equally to firms regardless of their size or where they are domiciled.	Exempting certain firms from regulations creates an uneven playing field and puts consumers at risk.
Seek expertise.	Augment regulatory expertise with technical and industry expertise.	Technical experts can help regulators understand the impact of regulatory options.

1. Avoid Pro-Human Biases

Policymakers should not discriminate against AI; AI systems should be held to the same standard as humans. Generally, if something is legal for a human to do, it should be legal for an AI system. For example, if it is legal for a security guard to verify the identity of someone entering a building, it should be legal to use an AI system to do the same. Prohibiting, or penalizing, firms for using an AI system to complete a task instead of a human disincentivizes the technology's use. Penalties can take different forms, but include holding AI systems to a higher standard or creating additional obligations firms must satisfy before they can use AI. Conversely,

if something is illegal for a human to do, it should also be illegal for an AI system. For example, if selling a certain piece of art or music infringes on someone else's intellectual property rights, it should not matter whether a human or AI system produced that work.

One example of pro-human bias is exclusive requirements for occupational licensing to provide a service, such as in the health care and legal fields, when comparable AI services can be provided. Occupational licenses are intended to protect consumers by limiting who can practice in the field. However, medical licensing boards or state bars—which consist of professionals in the field who may not want to compete with AI systems—can prevent organizations from offering AI-enabled services even if those services perform as well as, or better than, licensed professionals because the AI system cannot receive a license to practice. For example, most state bars restrict nonlawyers from offering legal services, thereby limiting the development of AI-based legal tools.[3]

2. Regulate Performance, Not Process

To address concerns about safety, efficacy, and bias, regulators should regulate the performance of AI systems rather than create prescriptive rules about specific processes and methods firms must follow. Establishing performance-based metrics for AI systems gives consumers, businesses, and government an opportunity to better compare performance across different systems, as well as set minimum performance requirements. For example, rather than create stringent compliance-based rules, such as requiring lenders to use diverse datasets to train their credit scoring models, regulators should create performance-based rules, such as requiring lenders to validate that their credit scoring models accurately assess risk across all protected classes of individuals. Allowing firms to identify the best way to achieve the desired goal gives them the flexibility necessary to comply most efficiently. Moreover, performance-based regulations can ensure firms meet the desired goals rather than simply check the box on a list of compliance measures.

3. Regulate Sectors, Not Technologies

Since policymakers cannot anticipate all future uses of AI, some have proposed regulating the technology itself, rather than specific uses. But AI is a general-purpose technology with many potential applications. Just as a knife is different in the hands of a chef, a soldier, and a surgeon, so too do the risks and benefits of AI depend on how it is being used.[4] Regulators treat knives differently in different sectors, such as creating unique workplace safety standards for scalpels used in hospitals, knives used for food preparation, and knife blades attached to power tools in industrial applications.[5] Likewise, if there is a need for rules, policymakers should create narrow rules for specific AI applications in particular sectors, such as health care and transportation, rather than for AI itself. An AI system to navigate a vehicle should be treated differently than one to automate stock

trades or diagnose illnesses, even if they use similar underlying technologies. Forcing all sectors to use the same rules for AI will likely impose excessive or duplicative requirements on some while providing insufficient requirements on others. Creating rules for specific AI applications allows regulators with deeper expertise about particular industries to set appropriate rules for AI applications. For example, insurance regulators may already have considered how to address risks from inscrutable credit scoring models, so whether an insurer uses machine learning models is irrelevant.[6]

4. Avoid AI Myopia

Many of the concerns motivating calls to regulate AI are not actually about AI. For example, the White House's "Blueprint for an AI Bill of Rights" lists concerns such as bias in hiring practices and credit scoring, lack of recourse when unfavorable outcomes occur, and insufficient consumer privacy.[7] Yet, none of those issues are unique to AI. Regulators should focus on fixing the broader problem, not just the part of the problem involving AI. The goal for regulators should never be to address, for example, only biased hiring practices involving AI, but rather all biased hiring practices. Moreover, focusing narrowly on AI as a problem ignores opportunities to use AI as part of the solution, such as considering how the use of AI could inject more objectivity into existing hiring practices and decrease human bias.

Unfortunately, many fixate on AI and ignore the bigger picture. This shortsightedness distracts from ongoing efforts to address the whole problem and can sideline those who have been working on the broader issue while focusing limited public and media attention on what is often a relatively small part of the problem. For example, concerns about wrongful arrests from facial recognition technology have dominated many news headlines and policy debates, including congressional hearings, ignoring the broader issue of police reforms that might decrease the number of wrongful arrests in the country.[8] Similarly, initiatives to ban law enforcement agencies from using facial recognition would certainly prevent the technology from being used for any wrongful arrests, but it would do nothing to tackle the larger problem of wrongful arrests. By treating AI as separate from the rest of the problem, it ignores the potential role for AI to be part of the solution.

5. Define AI Precisely

Policymakers should carefully define AI to avoid inadvertently including other software and systems within the scope of new regulations. AI covers a broad range of technology and is integrated into many hardware and software products. Policymakers should not use broad definitions of AI in regulation if they only intend to regulate uninterpretable machine learning or deep learning systems. For example, the European Union's AI Act initially included such a broad definition of AI that basic spreadsheet software such

as Microsoft Excel would likely have fallen within its scope.[9] Using too broad a definition of AI in regulations would impose substantial costs on those developing or deploying products and services that integrate basic analytics or automation. Using a clear and unambiguous definition of AI avoids market uncertainty.

6. Enforce Existing Rules

AI does not exempt organizations from following rules. Many laws and regulations, such as those addressing worker safety, product liability, discrimination, and more, apply regardless of whether AI is involved. In these cases, new regulations for AI are often not needed. For example, firms must adhere to laws prohibiting discrimination in employment decisions regardless of whether they involve a human or a computer in a decision to hire someone. Likewise, lenders must abide by fair lending rules regardless of whether they use an AI system to assess credit risk. To respond to concerns about AI, regulators should explain how they will enforce existing regulations for use of emerging AI products and services, provide guidance to those adopting these tools, and seek public feedback on any areas of potential concern. Indeed, agencies such as the U.S. Equal Employment Opportunity Commission and the Consumer Financial Protection Bureau have announced recent efforts to address concerns about algorithmic fairness using their existing authority and mission.[10]

7. Ensure Benefits Outweigh Costs

To have a net positive impact, the benefits of any regulatory intervention should outweigh the costs. When considering costs, policymakers should not only look at direct compliance costs but also indirect productivity, innovation, and competitiveness costs—such as reduced opportunities to use emerging technologies for social and economic benefit and a decrease in domestic firm investment in bringing these technologies to market. But these indirect costs of technology regulation can be substantial, even amounting to more than the direct compliance costs.[11] Considering these costs might seem obvious, but policymakers often pay little attention to the costs of technology regulations, especially when they believe, often wrongly, that regulation necessarily spurs innovation and thus an analysis of costs is superfluous. For example, a European Commission spokesperson deflected questions about the AI Act's economic costs by arguing with virtually no evidence that the law "will enhance the uptake of AI by increasing users' trust, hence also increasing the demand, and providing legal certainty for AI providers to access bigger markets."[12]

One reason some policymakers ignore the costs of regulations is because they believe costs are irrelevant compared with certain fundamental rights. For example, the EU's impact assessment of the AI Act dismisses concerns that the law might keep certain products out of the market. In considering the impact on a hypothetical AI-based recruitment software, the impact assessment notes that "the respect of the fundamental right in question

(in this case: non-discrimination) prevails over the loss of economic activity."[13] But this shallow assessment ignores an obvious possibility: that the use of AI-enabled recruitment software might reduce existing discrimination in hiring by presenting more objective evaluations of job candidates. When considering all costs, the net impact of the regulation can be negative, especially if poorly designed.

8. Optimize Regulations

Even if a regulation has a net positive impact, policymakers should still work to maximize the benefits and minimize costs. There are often multiple ways to achieve the same outcome, so the goal for policymakers should be to find the most efficient way to achieve their regulatory objective. Unnecessary regulatory costs cause firms to reallocate funding away from other business activities and toward compliance, leaving consumers worse off.[14]

Those championing new laws and regulations are often reluctant to discuss their costs because acknowledging those costs can discourage support for their proposals. However, ignoring the costs or directing hostility toward those raising them creates a policy environment antithetical to constructive dialogue about how to improve policy proposals.[15] Moreover, legislation often goes through multiple rounds of revisions, yet policymakers often do not update their impact assessments.

9. Treat Firms Equally

Policymakers should treat all firms the same to create a level playing field. Policymakers often suggest exempting smaller businesses from their legislative proposals because they recognize that their proposals present high compliance burdens that smaller firms might be unable to bear, and they would have an easier time passing their proposal by only applying their rules to larger firms. But the solution to that problem should be to reduce these burdens across the board, not impose them only on larger firms. Moreover, if certain obligations are necessary to protect consumers from certain products or services, then all firms should adhere to the rules regardless of their size.[16] Likewise, policies should not treat firms differently based on the country where they are domiciled. Again, if the purpose of certain rules is to protect consumers, those rules should be applied equally to all firms. Exempting domestic firms from some rules is usually intended more for protectionism than protecting consumers.

10. Seek Expertise

Relevant technological and industry expertise is necessary to craft effective regulations. Unfortunately, policymakers do not, and cannot, know everything. As the European Commission itself noted in its impact assessment of the AI Act, "[G]iven the complexity and rapid speed of AI development, competent authorities often lack the necessary resources,

expertise and technological tools to effectively supervise risks posed by the use of AI systems to safety and fundamental rights."[17] To address this shortcoming, policymakers should involve AI experts, particularly from sectors they seek to regulate, in the regulatory process.

One way to involve experts is with regulatory sandboxes, which are meant to provide an opportunity for businesses that do not fit into existing regulatory frameworks to bring their products to market. In exchange for temporary exemptions from specific rules, businesses provide detailed information about their operations for regulators, which they and businesses use to both provide effective oversight and protection for consumers and develop permanent new rules to address successful business models.[18] Regulatory sandboxes benefits firms, regulators, and consumers: Firms can bring products and services to market, regulators can receive the information they need to design appropriate rules, and consumers can access innovative businesses.

Regulators should also ensure they have AI and data literacy skills on their teams to understand new technologies. For example, they may seek technical training for their staff or hire a chief technologist to provide technical advice. However, regulators should avoid hiring technologists with preconceived biases about AI or the companies creating the technology.[19] Ideally, regulators should seek technical expertise that is independent and objective. For example, the U.S. National Institute of Standards and Technology (NIST) conducts technical reviews for accuracy of facial recognition systems and other biometrics.[20] These types of independent technical reviews can help regulators craft rules based on trustworthy assessments of AI.

CONCLUSION

AI has the potential to have a substantial positive impact on the economy and society. But policymakers should not take technological progress for granted. Poorly crafted laws and regulations could delay or stall the adoption of technologies that could save lives, increase wages, and improve quality of life. Therefore, policymakers should proceed with caution and be guided by these core principles so that their quest for responsible AI does not result in the creation of irresponsible regulation.

REFERENCES

1. "Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS," April 21, 2021, EUR-Lex, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206.
2. Daniel Castro, "Benefits and Limitations of Industry Self-Regulation for Online Behavioral Advertising," Information Technology and Innovation Foundation. December 2011, https://www2.itif.org/2011-self-regulation-online-behavioral-advertising.pdf.
3. Morgan Stevens and Daniel Castro, "State Policymakers Should Change Their Rules to Support the Development of Tech-Enabled Legal Services," Center for Data Innovation, August 8, 2022, https://datainnovation.org/2022/08/state-policymakers-should-change-their-rules-to-support-the-development-of-tech-enabled-legal-services/.
4. Patrick Grady, "The AI Act Should Be Technology-Neutral," Center for Data Innovation, January 31, 2023, https://datainnovation.org/2023/01/the-ai-act-should-be-technology-neutral/.
5. "Young Worker Safety in Restaurants eTool," Occupational Safety and Health Administration (OSHA), U.S. Department of Labor, n.d., https://www.osha.gov/etools/young-workers-restaurant-safety/food-prep; "Machine Guarding eTool," OSHA, n.d., https://www.osha.gov/etools/machine-guarding/plastics-machinery/thermoforming-machines.
6. Bethan Moorcraft, "The EU's AI Act: Is it unfair to insurers?" Insurance Business America, December 6, 2022, https://www.insurancebusinessmag.com/us/news/columns/the-eus-ai-act-is-it-unfair-to-insurers-429780.aspx.
7. "Blueprint for an AI Bill of Rights," The White House, October 2022, https://www.whitehouse.gov/ostp/ai-bill-of-rights/.
8. Kashmir Hill, "Another Arrest, and Jail Time, Due to a Bad Facial Recognition Match," *New York Times*, December 29, 2020, https://www.nytimes.com/2020/12/29/technology/facial-recognition-misidentify-jail.html.
9. Mikołaj Barczentewicz and Benjamin Mueller, "More than Meets the AI: The Hidden Costs of a European Software Law," Center for Data Innovation, December 1, 2022, https://www2.datainnovation.org/2021-more-than-meets-the-ai.pdf; Insurance Europe, "Response to EC proposal for a Regulation on AI," n.d., https://www.insuranceeurope.eu/publications/2413/response-to-ec-proposal-for-a-regulation-on-ai/.
10. "Artificial Intelligence and Algorithmic Fairness Initiative," U.S. Equal Employment Opportunity Commission, n.d., https://www.eeoc.gov/ai; "Fair Lending Report of the Consumer Financial Protection Bureau," Consumer Financial Protection Bureau, May 2022, https://files.consumerfinance.gov/f/documents/cfpb_2021-fair-lending_report_2022-05.pdf.
11. Ashley Johnson and Daniel Castro, "Maintaining a Light-Touch Approach to Data Protection in the United States," Information Technology and

Innovation Foundation, August 8, 2022, https://itif.org/publications/2022/08/08/maintaining-a-light-touch-approach-to-data-protection-in-the-united-states/.

12. Luca Bertuzzi, "Study warns of compliance costs for regulating Artificial Intelligence," *Euractiv*, July 26, 2021, https://www.euractiv.com/section/digital/news/study-warns-of-compliance-costs-for-regulating-artificial-intelligence/.
13. "Impact Assessment of the Regulation on Artificial intelligence," European Commission, April 21, 2021, https://digital-strategy.ec.europa.eu/en/library/impact-assessment-regulation-artificial-intelligence.
14. Luke Stewart, "The Impact of Regulation on Innovation in the United States: A Cross-Industry Literature Review," Information Technology and Innovation Foundation, June 2010, https://www2.itif.org/2011-impact-regulation-innovation.pdf.
15. Bertuzzi, "Study warns of compliance costs for regulating Artificial Intelligence."
16. Robert D. Atkinson and Michael Lind, *Big Is Beautiful: Debunking the Myth of Small Business* (MIT Press, 2018).
17. "Impact Assessment of the Regulation on Artificial intelligence," European Commission.
18. Patrick Grady, "EU's AI Regulatory Sandboxes Need Fixing," Center for Data Innovation, October 4, 2022, https://datainnovation.org/2022/10/eus-ai-regulatory-sandboxes-need-fixing/.
19. Mitchell Clark, "Googler who helped lead 2018 walkout is joining the FTC," *The Verge*, November 3, 2021, https://www.theverge.com/2021/11/2/22759776/google-meredith-whittaker-ftc-nomination-ai-ethics-regulation.
20. "Face Recognition Vendor Test (FRVT)," National Institute of Standards and Technology (NIST), November 30, 2020, https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt.

ABOUT THE AUTHOR

Daniel Castro is the director of the Center for Data Innovation and vice president of the Information Technology and Innovation Foundation. He has a B.S. in foreign service from Georgetown University and an M.S. in information security technology and management from Carnegie Mellon University.

ABOUT THE CENTER FOR DATA INNOVATION

The Center for Data Innovation studies the intersection of data, technology, and public policy. With staff in Washington, London, and Brussels, the Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as open data, artificial intelligence, and the Internet of Things. The Center is part of the Information Technology and Innovation Foundation (ITIF), a nonprofit, nonpartisan think tank.

Contact: info@datainnovation.org

datainnovation.org