NO PREP

Should AI Be Regulated? Arguments on Both Sides

Includes: Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

Should AI Be Regulated? Arguments on Both Sides

The Central Debate

Name: **Period:** **Date:**

Should AI Be Regulated? Arguments on Both Sides

The Central Debate

Few policy debates of the current era are as complex and consequential as the question of how, and whether, to regulate artificial intelligence. The question sits at the intersection of economic policy, national security, civil liberties, consumer protection, and the long-term trajectory of human civilization. Not surprisingly, it generates passionate disagreement from people with fundamentally different values, different assessments of the technology's risks, and different intuitions about what institutions and what actors are best positioned to manage those risks. What makes the AI regulation debate particularly challenging is that neither side is obviously wrong. The case for aggressive regulation reflects genuine, documented harms. The case for restraint reflects genuine risks that overregulation could slow beneficial development, shift power to less safety-conscious actors, and fail to achieve its goals anyway. Understanding both sides rigorously, rather than caricaturing the opposition, is essential for anyone who wants to think clearly about this issue.

The Case for Stronger AI Regulation

Those who favor stronger AI regulation begin with documented harms. Facial recognition AI has led to wrongful arrests of innocent people: a 2024 ACLU analysis documented over two dozen cases in the United States where individuals were wrongfully arrested based on AI facial recognition matches. AI hiring systems have been shown to discriminate against women and people of color when trained on historically biased data. AI-generated disinformation has demonstrably influenced elections in multiple countries. These are not hypothetical future risks; they are present, happening-now realities. Proponents of regulation also argue from the logic of market failure. Left to their own devices, technology companies have strong financial incentives to deploy AI quickly and widely, but relatively weak incentives to

spend heavily on safety, because the costs of unsafe AI are typically externalized onto users and society rather than borne by the companies themselves. This is precisely the condition that justifies regulatory intervention in other industries: pharmaceutical companies are regulated because individual patients cannot conduct drug trials; banks are regulated because depositors cannot independently assess systemic risk. The same logic applies to AI. A third argument concerns power concentration. Powerful AI capabilities are currently concentrated in a very small number of large technology companies. Without regulation, this concentration is likely to deepen as AI improves. Regulation (including interoperability requirements, data access rules, and antitrust enforcement) can prevent a small number of private entities from exercising undue control over technology that affects everyone.

The Case for Regulatory Restraint

Those skeptical of aggressive AI regulation begin with a competence objection: regulators typically lack the technical expertise to write effective rules for rapidly evolving AI systems. Well-intentioned regulations written today may be obsolete by the time they take effect, or may inadvertently prohibit beneficial applications while failing to address the actual harms they were designed to prevent. The history of technology regulation includes numerous examples of rules that protected incumbents from competition without meaningfully improving safety. A second argument concerns the innovation cost. The United States and its democratic allies currently lead the world in AI development. Significant regulatory burdens could shift AI development to China or other countries with fewer safety-conscious cultural norms, less democratic accountability, and less concern for the values embedded in the EU AI Act and similar frameworks. If the choice is between regulated-but-slower democratic AI and unregulated-but-faster authoritarian AI, some argue the global outcome of the former could be worse than the latter. A third argument is that voluntary industry commitments and market competition can achieve safety goals without coercive regulation. Companies that deploy unsafe AI face reputational consequences, legal liability, and loss of user trust. Competitive pressure creates incentives for companies to distinguish themselves on safety. Some researchers argue that 'safety washing' (fake safety commitments) is a real risk, but that it can be addressed through transparency requirements and liability rules rather than prescriptive mandates.

Where the Disagreement Actually Lives

It is worth noting that most serious participants in the AI regulation debate do not actually occupy the extreme positions. Almost nobody argues for zero AI oversight of any kind. Almost nobody argues for banning AI entirely. The real disagreements are about degree, mechanism, and timing: How much regulatory oversight is appropriate for which applications? Should regulation be prescriptive (specific rules) or principles-based (flexible standards)? Should safety testing be mandatory before deployment or voluntary after? Who should be liable when AI causes harm? These are questions that admit of reasonable disagreement. They require weighing real tradeoffs between values that most people care about — innovation, safety, equity, liberty, security. Understanding the strongest versions of different positions, rather than the weakest, is the prerequisite for reaching good policy conclusions. Public opinion on AI regulation has been evolving rapidly as AI becomes more visible in everyday life. Surveys conducted between 2022 and 2025 consistently show strong public support for mandatory AI disclosure in high-stakes decisions like hiring, lending, and law enforcement. Support for regulation is high across political affiliations and demographic groups — the AI regulation debate is notably less partisan than most technology policy issues. This public support creates political conditions for legislation that might otherwise face significant industry opposition. A significant practical constraint on AI regulation is the scarcity of technical expertise in regulatory agencies. Writing effective rules for AI systems requires deep understanding of how those systems work, what failure modes they have, and what technical requirements would meaningfully reduce risk. Most government regulatory agencies have limited numbers of AI-literate staff. This creates both the competence problem highlighted by regulation skeptics and a policy imperative: governments that want to regulate AI effectively must invest heavily in building regulatory capacity. Several countries have recognized this and launched programs to recruit AI experts into government service with competitive compensation. The liability question — who is legally responsible when AI causes harm — is one of the most contested and consequential unresolved issues in AI regulation. Current law is poorly suited to AI: products liability doctrine was developed for physical products, not software systems that learn and change over time. Tort law assumptions about identifying a negligent party are challenged by AI systems where harmful outputs may result from complex interactions between training data, model architecture,

deployment context, and user prompts that no single actor fully controlled. Resolving the liability question requires either adapting existing legal frameworks or developing new ones that account for AI's novel characteristics. Consumer protection regulators have entered the AI debate through a different angle: deceptive practices law. The Federal Trade Commission has already brought actions against companies that made misleading claims about their AI systems — overstating accuracy, fabricating safety testing, or concealing limitations. This approach does not require comprehensive AI legislation; it applies existing consumer protection authority to AI contexts. Several legal scholars have argued this is the most viable near-term regulatory strategy in the U.S. precisely because it does not require Congress to pass new laws.

Key Information: AI Regulation Debate

Detail	Information
Pro-regulation: documented harm 1	Wrongful arrests from facial recognition AI (24+ U.S. cases per 2024 ACLU analysis)
Pro-regulation: market failure argument	Companies externalize safety costs; same logic as pharmaceutical/banking regulation
Pro-regulation: power argument	AI concentrated in few companies; regulation prevents undue private control
Anti-regulation: competence objection	Regulators lack technical expertise; rules may be obsolete or counterproductive
Anti-regulation: innovation cost	Heavy regulation could shift AI development to less safety-conscious countries
Anti-regulation: market alternative	Reputational costs, liability, and competition can drive safety without mandates
Where real debate lives	Degree, mechanism, and timing — not zero vs. total regulation

Detail	Information
Key unresolved question	Who is liable when AI causes harm?

Case Study: Two Students, One Debate

Ms. Torres's class is preparing for an AI policy debate. Aaliyah is arguing for stronger federal AI regulation. She has found data on facial recognition wrongful arrests, evidence of biased hiring AI, and FTC reports on AI-enabled fraud losses. She believes the market has failed and government intervention is necessary. Her classmate Deshawn is arguing for regulatory restraint. He has found evidence that the EU's early privacy regulations (GDPR) created enormous compliance costs for small businesses without proportional safety benefits, and research suggesting that U.S. AI leadership depends partly on its regulatory flexibility. He worries that prescriptive AI rules will be written by regulators who don't understand the technology. Both students believe they are genuinely trying to protect people — they just disagree about which approach does that better. Aaliyah also notes that this debate matters beyond policy class: real hiring algorithms have already harmed real job seekers, and real facial recognition errors have resulted in real wrongful arrests — these are not hypothetical future scenarios. Deshawn concedes the point but argues the question remains whether regulation or innovation would address those harms faster. Both students agree the answer matters enormously for the people affected.

1. What is the strongest argument Aaliyah can make using evidence from the reading, and what is the strongest counter-argument Deshawn could make?

2. Based on both the case study and the reading, where do Aaliyah and Deshawn actually agree, even though they're on opposite sides of the debate?

Multiple Choice

1. What is the 'market failure' argument for AI regulation?

    a) AI regulation failed in the past, creating a market for unregulated AI products used to cause measurable damage and harm

    b) AI companies are currently losing money and need government subsidies to continue developing safe and responsible AI

    c) Companies have incentives to deploy AI quickly but weak incentives to spend on safety because harm costs fall on others

    d) The market for AI products has too many competitors, preventing any one company from investing in safety

2. What is the 'innovation cost' argument against aggressive AI regulation?

    a) The cost of innovation in AI already includes safety — regulation would double-count those costs for companies

    b) AI regulation would prevent companies from innovating new products and developing advanced AI solutions for industry

    c) Innovation costs money that could be better spent on safety research and protective measures for vulnerable populations

    d) Heavy regulation could drive AI development to less safety-conscious countries, potentially worsening global outcomes

3. According to the reading, where does the real disagreement about AI regulation actually live?

a) Between those who want zero regulation and those who want an outright ban on every AI system

b) In questions about degree, mechanism, and timing rather than extreme all-or-nothing positions

c) Between American and European approaches to technology policy and regulatory standards

d) Between technology companies and governments over who profits most from AI and market advantages

Matching

Write the letter of the correct definition on the line next to each term.

Term	Definition
1. ___ Market failure	A) Fake or superficial safety commitments made by companies to satisfy public pressure without genuine safety investment
2. ___ Competence objection	B) Flexible regulatory standards that define goals without mandating specific compliance methods
3. ___ Safety washing	C) Rules that specify exactly what companies must do — specific requirements rather than general standards
4. ___ Prescriptive regulation	D) A condition where market incentives do not produce socially optimal outcomes, typically justifying regulatory intervention
5. ___ Principles-based regulation	E) The argument that regulators lack technical expertise to write effective rules for rapidly evolving AI

Fill-in-the-Blank

1. A 2024 ACLU analysis documented over two _____ cases in the United States where individuals were wrongfully arrested based on AI facial recognition matches.

2. The argument that companies face reputational damage and legal _____ for unsafe AI is used by those who favor market solutions over government regulation.

3. The question of who is _____ when AI causes harm is one of the key unresolved questions in the AI regulation debate.

Short Answer

Answer each question in 2-4 complete sentences.

1. Explain the market failure argument for AI regulation in your own words, using the pharmaceutical or banking industry as an analogy.

_____

_____

_____

2. Why does the reading argue that understanding the strongest versions of opposing arguments is necessary for good policy conclusions?

_____

_____

_____

3. Where do you personally come down on AI regulation, and which argument from the reading most influences your position? Explain.