
Ensuring the Consent of the Governed: Surveying the U.S. Public's Attitudes Towards AI Governance Regimes

Matthew Sargent
RAND Corporation
1200 South Hayes St
Arlington, VA 22202
msargent@rand.org

K. MacKenzie Scott
RAND Corporation
1776 Main St
Santa Monica, CA 90401
mkscott@rand.org

James Caridi-Doyle
RAND Corporation
1776 Main St
Santa Monica, CA 90401
jcarididoyle@rand.org

Abstract

The rapid advancement of artificial intelligence (AI) capabilities has generated numerous proposals for oversight and governance frameworks to ensure that the technology is developed safely and responsibly. There has been little focus, however, on the question of who the public would trust to provide this oversight. This paper examines the public's perceived need for oversight, as well as the types of organizational structures they would trust to manage increasingly powerful AI. Using a nationally representative survey of 2,000 U.S. adults, we found that the American public very strongly favors monitoring and oversight of the most powerful AI systems. We then examined the types of institutions that Americans would trust most to manage such systems. Respondents were asked to choose from among six different organizational forms commonly represented in proposals for AI governance: the US government, a foreign government, a private company, a public-private partnership, a multinational partnership, or no oversight. The results were striking. Even though the respondents had earlier rated U.S. dominance in AI as very important and had rated the potential of Chinese dominance very unfavorably, the preferred forms of governance were multinational and public-private partnerships. This suggests that even as the American public views AI development partially through the lens of national competition, many Americans also regard multiparty and international bodies as trusted providers of effective AI oversight. These results provide important empirical evidence for debates over the feasibility and perceived legitimacy of international AI governance alternatives.

1 Introduction

The rapid advancement of artificial intelligence (AI) capabilities has led to numerous proposals for oversight and governance frameworks intended to ensure that the technology is developed and deployed safely and responsibly. While there has been significant analysis of the different

governance mechanisms and enforcements strategies that might be involved (Dafoe 2018; Jasanoff 2016; Floridi et al. 2018; Brundage et al. 2020; Jobin, Ienca, and Vayena 2019), there has been comparatively little empirical research examining public attitudes toward AI regulation or their preferences regarding the institutional structures that should be employed. This study addresses this gap by reporting an empirical investigation of public attitudes toward AI governance structures. Using a nationally representative survey of 2,000 U.S. adults conducted through the RAND American Life Panel (Pollard and Baird, 2017), we examine citizen preferences regarding the institutional arrangements that are most trusted to govern advanced AI systems and assess the factors that shape these preferences.

The public's views are significant because they are likely to shape both the political feasibility and perceived legitimacy of different governance proposals, particularly those that envision international AI governance institutions. If the public strongly supports oversight but distrusts the institutions proposed to provide it, governance regimes may face resistance even when they are technically sound. Conversely, if certain institutional forms are perceived as legitimate, policymakers may have greater room to develop robust governance arrangements before advanced AI systems become more deeply embedded in economic, military, and social life.

The results presented below suggest three central findings. First, Americans strongly support monitoring and oversight of the most capable AI systems, with more than three quarters of respondents rating its importance as a six or more on a ten points scale. Second, although respondents expressed strong preference for U.S. leadership in AI and strong concern about potential Chinese leadership, their preferred collaborative governance arrangements such as multinational partnerships and public-private partnerships. These findings indicate that Americans simultaneously view AI as an arena of high-stakes geopolitical competition and regard cooperative, multi-actor institutions as plausible and legitimate forms of oversight.

1.1 Statement of contributions

This paper presents and analyzes previously unreported survey data in Sections 3 and 4. By examining public preferences regarding different AI governance institutions, the study provides important empirical grounding for discussions of governance alternatives. Understanding public opinion in these areas will be essential for policymakers seeking to design AI governance regimes that are perceived as legitimate by the public.

2 The organizational structures of international AI governance proposals

Effective AI governance ultimately depends not only on the methodological robustness and operational feasibility of a proposed model but also on its broad public acceptance and perceived legitimacy. This paper therefore sets aside questions about the specific authorities and regulatory goals of proposed AI governance regimes and instead focuses on whom the U.S. public would trust to provide such regulation. Public opinion surveys can offer early insight into societal values, expectations, and concerns before technologies and governance frameworks become fully entrenched. Moreover, evidence about which forms of governance the public views as legitimate can inform the design of institutions through a more deliberative and inclusive process, rather than relying on the assumption that governments, experts, or firms alone are best positioned to decide.

2.1 The need for public support

If they are to be viewed as legitimate, governance institutions that focus on emerging technologies must enjoy meaningful public support. Without such support, even well-designed oversight mechanisms may face implementation challenges, political resistance, or outright hostility from affected stakeholders. The history of technology regulation demonstrates that public acceptance is often crucial for the long-term effectiveness of regulatory frameworks (Winner, 1980; Jasanoff, 2005). Research has also shown that institutional trust affects public acceptance of new

technologies and compliance with related regulations (Siegrist, 2000; Earle & Cvetkovich, 1995). Because the effectiveness of any AI governance regime will depend partly on public trust in the institutions responsible for oversight, research into public preferences regarding governance alternatives is crucial for designing institutions that can maintain legitimacy over time.

2.2 A taxonomy of the alternative organizational structures for AI governance

To assess U.S. public attitudes toward AI governance, we developed a taxonomy of institutional formats that could plausibly be responsible for AI regulation and oversight across different scenarios. Rather than asking respondents about specific regulatory mechanisms or policy objectives, we focused on the types of organizations that would implement such regimes. This emphasis on organizational form is important because, as our results demonstrate, the public holds strong views about the desirability of different institutions independent of the policies and governance agendas that they might oversee. Although respondents' knowledge of specific governance mechanisms and objectives may still be emerging, they nevertheless have meaningful beliefs about which institutions they consider appropriate and legitimate providers of AI oversight.

Our taxonomy also pays specific attention to *operator control*, that is, determining which individuals, organizations, or government actors would be permitted to run and deploy the most capable AI models under each governance approach. This focus reflects proposals in the AI policy literature that suggest limiting the operation of frontier AI systems to vetted actors with the technical capacity, ethical commitment, and legal accountability to manage high-risk deployments (Brundage et al., 2020; Anderljung & Maas, 2023).

A summary of the different organizational formats that we included in our survey follows.

National Government Oversight A substantial body of proposals advocates for centralized oversight by national governments, particularly in jurisdictions with significant AI development capacity. In the United States, various scholars and policy analysts have proposed federal AI agencies or expanded mandates for existing regulatory bodies such as the Federal Trade Commission or the National Institute of Standards and Technology (Baum, 2020; Dafoe, 2018). Under such regimes, the authority to operate high-capability AI models would potentially be restricted to licensed entities, with licensing granted only to organizations meeting strict safety, security, and transparency requirements. Such entities could include government research laboratories, contracted private-sector firms, and certain academic institutions designated as "authorized operators." Under some proposals, for instance, frontier models could only be trained or deployed at certified national AI centers, with run-time monitoring and reporting obligations akin to those used in nuclear facility oversight (Zhang et al., 2022).

Foreign Government Oversight Other nations have explored similar governance structures, and it is possible that these foreign regulatory institutions could govern the future use of advanced AI in the United States under certain trade, treaty, or market-access arrangements. The European Union's AI Act (European Commission, 2021), for example, establishes risk-based categories for AI systems with mandatory conformity assessments for high-risk categories. In operational terms, approved model operators could include certified European AI developers, regulated cloud service providers hosting AI systems, and, in some cases, multinational corporations with compliance mechanisms that meet EU regulatory standards. If applied within the U.S., foreign government oversight would mean that only organizations registered with, and audited by, the foreign regulatory body could deploy certain AI systems domestically. This could potentially result in U.S. tech firms having to comply with foreign operational licensing regimes, as already occurs in data protection under the EU's GDPR (Greenleaf, 2018). It could also mean that models with close ties to foreign governments, like China's DeepSeek, are providing advanced AI capabilities that are used by US consumers.

Multinational Organizations and International Coordination The global nature of AI development and deployment has led to proposals for coordination among international governments to promote shared governance. Some analysts argue that unilateral national

approaches may be insufficient given the cross-border nature of AI development and deployment (Polyakova & Meserole, 2019). Beyond the global distribution of AI capabilities, a central rationale for multinational governance lies in the incentive structures surrounding AI innovation. Competitive pressures among firms and states create strong incentives to prioritize speed, capability, and market share over caution, increasing the risk of a regulatory “race to the bottom” in which jurisdictions weaken safeguards to attract investment and talent. Under such conditions, unilateral or purely national regulatory approaches may prove ineffective, as developers can shift activity to less restrictive environments. Multilateral coordination can help mitigate these dynamics by harmonizing baseline standards, reducing opportunities for regulatory arbitrage, and aligning incentives toward safer development and deployment practices.

In operational terms, international coordination might involve a multilateral licensing framework, similar to export controls, in which only sanctioned entities meeting collectively defined safety criteria could operate frontier AI models. For example, an international AI oversight body modeled after the International Atomic Energy Agency (IAEA) could maintain a registry of approved operators, conduct on-site inspections of model training facilities, and require real-time technical reporting from AI systems in deployment (Maas, 2019). This approach would intentionally limit operation of the most advanced models to a smaller set of approved governmental, industrial, and non-governmental research institutions. Related proposals include the creation of a “CERN for AI” that would benefit from joint funding from member states and could provide services such as the operation of cutting-edge AI systems (Sargent et. al, 2025).

Public-Private Partnerships Hybrid governance models combining government authority with private sector expertise have gained attention, given the concentration of AI capabilities within a small number of technology companies. These proposals often envision government-led licensing combined with industry-led operational management. In such a system, powerful AI models could be run jointly by trusted operator consortia bringing together technology companies, universities, and public-sector AI laboratories, all operating under shared security protocols. The “GovCloud” approach, where certain high-security computing infrastructure is jointly operated by government and approved private firms, is sometimes cited as a possible blueprint for ensuring both capacity and accountability (Center for Security and Emerging Technology, 2021).

Private Industry Proposals placing primary responsibility for AI governance in industry typically rely on voluntary compliance frameworks (Partnership on AI, 2016; IEEE, 2019). Under these models, powerful AI systems would be operated by firms participating in industry-led governance schemes, such as adherence to shared AI risk-management standards or safety checklists. Existing examples include Alphabet’s internal AI Governance Council and Microsoft’s Responsible AI Standard, under which operation of high-capability models is subject to internal review boards prior to deployment. Under these models, private firms would continue to operate and provide AI services, even as their capabilities become more advanced.

Minimal or No Oversight Finally, some argue that formal oversight may be unnecessary or counterproductive (Thierer, 2020; Wallach, 2015). In such a regime, any technically capable actor could operate frontier AI models without licensing or centralized vetting. Control would be determined entirely by possession of technical resources including computational capacity, model weights, and skilled personnel, meaning that powerful systems could be deployed by startups, individual researchers, and even hobbyists. This approach maximizes innovation potential but raises significant safety concerns, especially as open-source dissemination of cutting-edge models increases the number of possible operators worldwide (Inserra, 2024).

3 Survey results

To better understand public perceptions regarding AI and possible governance alternatives, we fielded survey questions through RAND’s American Life Panel (ALP). The ALP is a highly regarded survey platform that allows researchers of any affiliation to submit questions that are

included in a national survey. The survey was delivered electronically and was primarily completed on mobile devices. Two thousand adults residing in the United States were interviewed in March and April 2025, and the results were provided along with demographic variables and survey weights, calculated by the ALP, to make the survey results nationally representative. These survey weights are generated using a raking algorithm to match the ALP sample to the adult U.S. population as measured by the Census Bureau’s Current Population Survey. The weights match on four bivariate discrete distributions: gender by age, gender by race/ethnicity, gender by highest education, and family income by adult household size. The recruitment and weighting methodologies of the ALP are described in Pollard and Baird (2017).

3.1 Importance of AI monitoring and oversight

As part of the survey, respondents were asked to evaluate the importance of AI governance. When presented with the question, “How important do you think it is that the most capable AI systems—those that rival or exceed human levels of capability—be monitored for safety and overseen by a responsible entity (such as the U.S. government or a non-profit organization)?,” responses were overwhelmingly positive. Using representative weights, 48 percent of respondents rated the importance of monitoring as a 10, and 78 percent rated it a 6 out of 10 or higher. See Figure 1.




Figure 1: Importance that AI be monitored and overseen

This finding suggests that Americans strongly support AI governance in some form. Subsequent questions therefore assessed the types of organizations they would trust to provide such oversight.

3.2 Who should oversee advanced AI?

The survey then asked, “As increasingly capable AI systems are developed... who would you trust most to manage and operate them?” Respondents were instructed to select a single choice from the following list: A foreign government, a multinational partnership, a public-private partnership, no oversight necessary, none of the above, private firms, or the U.S. Government. This framing was chosen because the public is likely more attuned to questions about who should manage and operate advanced AI systems than to more abstract questions about governance architecture. As discussed in the previous section, this focus on operational control as a component of governance aligns with

much of the existing literature on international AI governance. Choices were presented in random order. The results are summarized in Figure 2.




Figure 2: Who should oversee advanced AI?

These findings suggest that despite the strong interest in AI oversight described in Section 3.1, many Americans have not yet formed clear preferences about the kinds of institutions that they would trust to manage such systems. The most popular single answer was “None of the above,” a choice which, in this context, served here as a no-opinion or prefer-not-to-answer response for those who did not wish to choose one of the other options. Among the substantive governance options, however, more than half of respondents selected one of two responses that included multi-member partnerships. This indicates that collaborative governance is regarded as a favorable option for overseeing next generation AI capabilities. Private firms and the U.S. government also received moderate levels of support.

4 Regression analysis

To assess the factors that shaped respondents’ preferences for different types of governance organizations, we estimated a multinomial logit regression in R using the “nnet” package. We examined which predictors were associated with respondents favoring a specific governance structure. The results are presented graphically below to aid intelligibility. The predictors reported below were part of a single regression model and thus serve as controls for one another. The factors included were: reported importance of AI monitoring (from survey response), reported importance of U.S. leadership in AI (from survey response), reported concern about Chinese leadership in AI (from survey response), sex, age (grouped by generation), and household income (grouped into income bands).

Table 1: Predictors of governance organizational preference

Predictor	Governance preference					
	No oversight necessary	Foreign govt	Multi-national partnership	Public-private partnership	Private firms	US Govt
(Intercept)	-2.539*	-1.62	-3.646***	-3.573***	-4.498***	-4.316***
AI should be monitored	-0.314**	-0.369**	0.208***	0.023	-0.097*	0.111*
Importance of US AI leadership	0.267*	0.242*	0.162***	0.250***	0.172***	0.231***
Concern about Chinese dominance	0.000	0.113	-0.046	0.044	0.108*	0.135**
Generation: Millennial	0.899	-1.635*	0.393	-0.235	2.054***	-0.096
Generation: GenX	-0.971	-1.963*	0.540*	-0.021	1.439**	-0.088
Generation: Baby Boomer and Silent	-0.645	-1.247	0.560*	0.360	1.348**	-0.013
Income: \$25,000 to \$49,999	0.262	-0.093	0.312	0.532	0.061	-0.084
Income: \$50,000 to \$99,999	-2.009	-1.459	0.893***	0.936***	0.486	0.039
Income: \$100,000 or more	-0.93	-1.978	0.810***	1.075***	0.153	-0.509
Sex: Male	-1.201	-0.044	0.321*	0.942***	0.947***	0.297

* p < 0.05, ** p<0.01, *** p<0.001

4.1 How did increased interest in monitoring shape choice of governance model?

Our first question was how differences in the perceived importance of monitoring and overseeing advanced AI, summarized in Figure 1, shaped respondents' preferences for different governance alternatives. Treating this measure as a predictor of preferred governance type allowed us to assess how institutional preferences change as the perceived importance of AI monitoring increases. These results are shown in Figure 3.




Figure 3: The impact of the perceived importance of monitoring AI on preference for governance type

As expected, stronger beliefs about the importance of monitoring were associated with a decreased preference for the view that “no oversight is necessary.” Preferences for governance by private firms and foreign governments also declined, likely reflecting perceptions that these alternatives would provide less effective or less legitimate oversight. By contrast, support for multinational partnerships increased, suggesting that this institutional form is viewed as particularly well suited to providing robust governance.

4.2 How did views about AI competition with China affect choice of governance model?

We then considered how attitudes towards geopolitical competition in AI might shape preferences regarding governance.

Our survey included a pair of questions that probed issues of geopolitical competition. The first asked “How important do you think it is that the United States be the world leader in AI development and use?” The question elicited a strong response: 37 percent of respondents rated the importance of American leadership as 10 out of 10, while 72 percent rated it 6 out of 10 or higher. Similarly, when asked “How concerned would you be if China was the world leader in AI development and use?,” 38 percent rated their concern as 10 out of 10, and 72 percent rated it 6 out of 10 or higher. Despite the similarity in these patterns of response, the correlation between the two measures was 0.56, and additional analysis found different factors shaping the two sets of attitudes (Sargent, Scott, and Caridi-Doyle, 2025).

Responses to each of these questions were significant predictors of governance choice. Higher ratings of the importance of U.S. leadership in AI were positively and significantly associated with all substantive governance options, indicating that respondents who valued U.S. leadership highly were more likely to express a preference for a specific governance structure rather than select “none of the above.” Responses to the question assessing concern about Chinese leadership in AI exhibited a more nuanced pattern of associations. These results are shown in Figure 4.




Figure 4: The impact of the reported concern over Chinese AI leadership on preference for governance type

These results show that individuals who are concerned about Chinese leadership in AI were more likely to prefer U.S.-centric governance options, specifically oversight or service provision by the U.S. government or private firms. This is significant because as Sargent, Scott, and Caridi-Doyle (2025) have shown, public concern about potential Chinese leadership in AI are independent from their familiarity with AI. Instead, these anti-Chinese attitudes appear to reflect broader geopolitical competition between the U.S. and China. This suggests that broader competitive tensions may need to be addressed before the U.S. public is willing to accept governance by non-U.S. institutions.

4.3 How did demographic factors shape opinions about governance alternatives?

To further understand the factors that shaped preferences for different types of governance organizations, we included demographic factors for age (generation), income, and sex in the model. In the interest of space, the primary results are summarized in text.

Generation Prior work has documented strong generational differences in attitudes towards AI, including opinions about AI use (Robinson, 2025) and trust in AI (Barna, 2024). Our analysis showed that, relative to Gen Z (born 1995-2012), Millennials (born 1980-1995) were less likely to favor management of AI by foreign governments and more likely to support governance by private firms. Shifting our analysis to progressively older generations of Americans, we found that Gen X (born 1965-1979) had similar attitudes to Millennials, but also supported AI managed by multinationals. Baby boomers and the silent generation (born 1946-1964 and 1925-1945 respectively) were similar to Gen X, favoring multinational and private firms, but with no significant differences in their views about the desirability of AI managed by foreign governments.

Income Our analysis also considered the impact of respondents' household incomes. Income data was provided according to 17 income ranges, which we reduced to four. Our analysis focused on households making less than \$25,000 per year, those making \$25-50,000 per year, those making \$50-100,000, and those making \$100,000+ per year. As a point of reference, the real median household income in the United States in 2023 was \$80,610 (U.S. Census Bureau, 2024). Using the lowest income bracket as the reference group, we found no significant differences in governance preferences in the next highest income bracket. However, both income brackets above \$50,000 showed the same pattern of statistically significant differences in attitudes. Compared with the

lowest income bracket, respondents in both higher-income brackets were more likely to choose public-private partnerships and multinational institutions.

Sex The final demographic variable that we considered was sex. Relative to women, male respondents favored AI oversight by private firms, public-private partnerships, and multinational institutions at significantly higher rates. Our data included only responses indicating male or female sex.

5 Summary of findings

This study provides new empirical evidence about how the U.S. public views the governance of increasingly capable AI systems, with particular attention to whom Americans trust to manage and operate such technologies. Four central findings emerge.

First, there is overwhelming public support for the monitoring and oversight of advanced AI. Nearly four in five respondents rated oversight as important, and almost half assigned it the highest possible importance rating. This suggests that proposals for AI governance rest on a strong foundation of public legitimacy and that policymakers are likely to face little resistance, in principle, to establishing some form of oversight for frontier AI systems. The challenge is instead likely to involve developing consensus over what form the governance should take

Second, while Americans strongly endorse oversight in general, they display more uncertainty about the specific institutional forms that should provide it. Many respondents selected “none of the above” when asked whom they would trust to manage and operate advanced AI, indicating that institutional preferences are less crystallized than attitudes about oversight itself. Among those who did express a substantive preference, however, collaborative models—especially multinational partnerships and public-private partnerships—were the most frequently chosen options. This pattern suggests that cooperative, multi-actor governance arrangements may be perceived as more legitimate than either unilateral national control or purely private governance.

Third, attitudes toward geopolitics and technological competition shape institutional preferences in systematic ways. Concern about potential Chinese leadership in AI increased support for U.S.-centric governance options, particularly oversight by the U.S. government or private firms. Notably, however, even among respondents who view AI through a lens of national competition, higher perceived importance of monitoring was associated with increased support for multinational governance. This combination of findings highlights a tension: Americans simultaneously frame AI as a geopolitical competition and express trust in international or multi-actor institutions to provide effective oversight.

Finally, demographic factors also shape preferences, though less strongly than attitudes about monitoring and geopolitics. Older respondents, higher-income respondents, and men were more likely to favor collaborative or industry-involved governance arrangements. These patterns suggest that socioeconomic position and generational experience modestly conditions views about who should be entrusted with advanced AI.

Taken together, these findings indicate that the U.S. public strongly supports oversight of advanced AI and is broadly open to governance arrangements that extend beyond any single actor or nation. For policymakers, this implies that the central challenge is not whether to pursue AI governance, but how to design institutions that align with public expectations about legitimacy, cooperation, and accountability. Governance proposals that incorporate visible mechanisms for monitoring, include multiple stakeholders, and address national security concerns are likely to enjoy the greatest public support.

References

- Allen, Greg, and Taniel Chan. 2017. *Artificial Intelligence and National Security*. Cambridge, MA: Belfer Center for Science and International Affairs.
- Anderljung, Markus, and Matthijs Maas. 2023. "Operationalizing AI Governance: Licensing Large-Scale General-Purpose AI Models." *AI & Society*. <https://doi.org/10.1007/s00146-023-01623-6>.
- Barna Group. 2024. *Generations and AI*. March 6, 2024. <https://www.barna.com/research/generations-ai/>.
- Baum, Seth D. 2020. "Risk and Regulation of Artificial Intelligence: The Case of Autonomous Weapons." *AI & Society* 35 (3): 343–54. <https://doi.org/10.1007/s00146-019-00874-1>.
- Brookings Institution. 2019. *Governing Artificial Intelligence: Upholding Human Rights & Dignity*. Washington, DC: Brookings Institution.
- Brundage, Miles, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, et al. 2020. "Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims." *arXiv*, February. <https://arxiv.org/abs/2004.07213>.
- Brundage, Miles, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, et al. 2020. "Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims." *arXiv*, February. <https://arxiv.org/abs/2004.07213>.
- Center for Security and Emerging Technology. 2021. *AI and the Future of Disinformation Campaigns*. Washington, DC: Georgetown University.
- Cihon, Peter. 2019. "Standards for Artificial Intelligence Governance: International Standards to Enable Global Coordination." *Journal of AI Policy* 1 (1): 30–49.
- Dafoe, Allan. 2018. "AI Governance: A Research Agenda." Oxford, UK: Centre for the Governance of AI, Future of Humanity Institute.
- Dafoe, Allan. 2018. "AI Governance: A Research Agenda." Oxford, UK: Centre for the Governance of AI, Future of Humanity Institute, University of Oxford.
- Earle, Timothy C., and George T. Cvetkovich. 1995. *Social Trust: Toward a Cosmopolitan Society*. Westport, CT: Praeger.
- European Commission. 2021. *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*. Brussels: European Commission.
- Floridi, Luciano, Josh COWls, Matt Taddeo, and Thomas Morley. 2018. "An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations." *Minds and Machines* 28 (4): 689–707. <https://doi.org/10.1007/s11023-018-9482-5>.
- Future of Humanity Institute. 2019. *The Windfall Clause: Distributing the Benefits of AI for a Common Good*. Oxford, UK: University of Oxford.
- Government of Canada. 2019. *Canada's National Artificial Intelligence Strategy*. Ottawa: Government of Canada.
- Greenleaf, Graham. 2018. "The Influence of the European Data Protection Regulation beyond Europe: Patching the GDPR and Australian Privacy Principles Together." *International Data Privacy Law* 8 (1): 33–45. <https://doi.org/10.1093/idpl/ix027>.

- Horowitz, Michael C., Paul Scharre, and Gregory C. Allen. 2020. "Artificial Intelligence and International Security: The State of Play." Washington, DC: Center for a New American Security.
- IEEE Standards Association. 2019. *Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems*, First Edition. Piscataway, NJ: IEEE.
- Inserra, David. 2024. "Artificial Intelligence and Liberty: A Case against Government Overreach." Washington, DC: Cato Institute.
- Jasanoff, Sheila. 2005. *Designs on Nature: Science and Democracy in Europe and the United States*. Princeton, NJ: Princeton University Press.
- Jasanoff, Sheila. 2016. *The Ethics of Invention: Technology and the Human Future*. New York: W. W. Norton & Company.
- Jobin, Anna, Marcello Ienca, and Effy Vayena. 2019. "The Global Landscape of AI Ethics Guidelines." *Nature Machine Intelligence* 1 (9): 389–99. <https://doi.org/10.1038/s42256-019-0088-2>.
- Maas, Matthijs M. 2019. "Innovating Governance for AI Safety, Security, and Stability." *International Journal of Technology Policy* 5 (2): 221–47.
- Partnership on AI. 2016. *Tenets of the Partnership on AI*. San Francisco, CA: Partnership on AI.
- Pollard, Michael S. and Matthew D. Baird, The RAND American Life Panel: Technical Description. Santa Monica, CA: RAND Corporation, 2017. https://www.rand.org/pubs/research_reports/RR1651.html.
- Polyakova, Alina, and Christopher Meserole. 2019. "Exporting Digital Authoritarianism: The Russian and Chinese Models." Washington, DC: Brookings Institution.
- Robinson, Bryan. 2025. "Gen Z Trust AI Over Humans in Their Careers, New Study Shows." *Forbes*. February 19, 2025. <https://www.forbes.com/sites/bryanrobinson/2025/02/19/gen-z-trust-ai-over-humans-in-their-careers-new-study-shows/>.
- Sargent, Matthew, Mary Lee, Dennis Murphy, and Sabahat Zafar, Assessing the Value of a CERN-Like Multinational Research Organization for AI. Santa Monica, CA: RAND Corporation, 2025. <https://www.rand.org/pubs/perspectives/PEA4024-1.html>.
- Sargent, Matthew, K. MacKenzie Scott, and James Caridi-Doyle, Majority of Americans See U.S. Leadership in AI as Crucial, Survey Finds. Santa Monica, CA: RAND Corporation, 2025. https://www.rand.org/pubs/research_reports/RRA4363-1.html.
- Siegrist, Michael. 2000. "The Influence of Trust and Perceptions of Risks and Benefits on the Acceptance of Gene Technology." *Risk Analysis* 20 (2): 195–203. <https://doi.org/10.1111/0272-4332.202020>.
- Smart Nation Singapore. 2019. *National AI Strategy: Advancing Our Smart Nation Agenda*. Singapore: Smart Nation and Digital Government Office.
- Thierer, Adam. 2020. *Permissionless Innovation: The Continuing Case for Comprehensive Technological Freedom*. Revised Edition. Arlington, VA: Mercatus Center.
- U.S. Census Bureau. 2024. *Income in the United States: 2023 – P60-282*. September 2024. <https://www.census.gov/library/publications/2024/demo/p60-282.html>.
- UK Government. 2021. *National Artificial Intelligence Strategy*. London: Department for Digital, Culture, Media & Sport.

UK Government. 2023. *Policy Paper: Pro-Innovation Regulation of AI*. London: Department for Science, Innovation & Technology.

Wallach, Wendell. 2015. *A Dangerous Master: How to Keep Technology from Slipping beyond Our Control*. New York: Basic Books.

Winfield, Alan F., and Marina Jirotko. 2018. "Ethical Governance of Autonomous Systems and Artificial Intelligence: A Framework." *Philosophical Transactions of the Royal Society A* 376 (2133): 20180080. <https://doi.org/10.1098/rsta.2018.0080>.

Winner, Langdon. 1980. "Do Artifacts Have Politics?" *Daedalus* 109 (1): 121–36.

Zhang, Baobao, Ethan R. Salley, and Matthew J. Burton. 2022. "Who Should Govern AI? Public Opinion on the Role of Government in Regulating Artificial Intelligence." *Journal of AI Policy* 2 (1): 1–20.