

U.S. Public Opinion on the Governance of Artificial Intelligence

Baobao Zhang*

Allan Dafoe*

baobaozhangresearch@gmail.com

allan.dafoe@politics.ox.ac.uk

Centre for the Governance of AI, Future of Humanity Institute, University of Oxford
Oxford, United Kingdom

ABSTRACT

Artificial intelligence (AI) has widespread societal implications, yet social scientists are only beginning to study public attitudes toward the technology. Existing studies find that the public’s trust in institutions can play a major role in shaping the regulation of emerging technologies. Using a large-scale survey ($N=2000$), we examined Americans’ perceptions of 13 AI governance challenges as well as their trust in governmental, corporate, and multistakeholder institutions to responsibly develop and manage AI. While Americans perceive all of the AI governance issues to be important for tech companies and governments to manage, they have only low to moderate trust in these institutions to manage AI applications.

CCS CONCEPTS

• **Social and professional topics** → **Governmental regulations.**

KEYWORDS

public opinion; AI governance; public trust

ACM Reference Format:

Baobao Zhang and Allan Dafoe. 2020. U.S. Public Opinion on the Governance of Artificial Intelligence. In *Proceedings of the 2020 AAAI/ACM Conference on AI, Ethics, and Society (AIES '20)*, February 7–8, 2020, New York, NY, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3375627.3375827>

1 INTRODUCTION

Advances in artificial intelligence (AI) could impact nearly all aspects of society, including the labor market, transportation, health-care, education, and national security [22]. AI’s effects may be profoundly positive, but the technology entails risks and disruptions that warrant attention. While technologists and policymakers have begun to discuss the societal implications of machine learning and AI more frequently, public opinion has not shaped much of these conversations. Given AI’s broad impact, civil society groups

*B.Z. and A.D. designed the survey. B.Z. analyzed the survey data and wrote the first draft of the article. Both provided critical input to the writing and results and approved the final version of the article.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

AIES '20, February 7–8, 2020, New York, NY, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7110-0/20/02...\$15.00

<https://doi.org/10.1145/3375627.3375827>




Figure 1: Agreement with statement that robots and AI require careful management (EU data from 2017 Special Eurobarometer #460)

argue that the public, particularly those underrepresented in the tech industry, should have a role in shaping the technology [37].

In the U.S., public opinion has shaped policy outcomes [4], including those concerning immigration, free trade, international conflicts, and climate change mitigation. As in these other policy domains, we expect the public to become more influential over time in impacting AI policy. It is thus vital to have a better understanding of how the public thinks about AI and the governance of AI. Such understanding is essential to crafting informed policy and identifying opportunities to educate the public about AI’s character, benefits, and risks.

Using an original, large-scale survey ($N=2000$), we studied how the American public perceives AI governance. The overwhelming majority of Americans (82%) believe that AI and/or robots should

be carefully managed. This statistic is comparable to survey results from respondents residing in European Union countries, as seen in Figure 1. Furthermore, Americans consider all of the 13 AI governance challenges presented in the survey to be important for tech companies and governments to manage carefully.

At the same time, Americans have only low to moderate levels of trust in governmental, corporate, and multistakeholder institutions to develop and manage AI in the public's interest. Public trust varies greatly between organizations. Broadly, the public puts the most trust in university researchers (50% reporting "a fair amount of confidence" or "a great deal of confidence") and the U.S. military (49%); followed by scientific organizations, the Partnership on AI, tech companies (excluding Facebook), and intelligence organizations; followed by U.S. federal or state governments, and the UN; followed by Facebook. Contrary to existing research on attitudes toward other emerging technologies, our study finds that individual-level trust in various actors to responsibly develop and manage AI does not predict one's general support for developing AI.

2 BACKGROUND

2.1 Public Opinion and AI Governance

Studying public opinion allows us to anticipate how electoral politics could shape AI governance as policymakers seek to regulate applications of the technology. Over the past two years, corporations, governments, civil society groups, and multistakeholder organizations have published dozens of high-level AI ethics principles. These documents borrow heavily from core principles in bioethics or medical ethics [10]. But unlike medical ethics, which is guided by the common aim of promoting the health of the patient, AI development does not share one common goal but many goals [19]. Furthermore, many of these AI ethics principles are at tension with each other [38]. For instance, how does one improve the accuracy of algorithmic predictions while ensuring fair and equal treatment of those impacted by the algorithm? As corporations and governments attempt to translate these principles into practice, tensions between principles could lead to political contestation.

In the U.S., voters are divided on how to regulate facial recognition technology and algorithms used to display social media content, with some of these divisions reflecting existing disagreements between partisans [39]. For instance, Democrats, compared with Republicans, are much more opposed to law enforcement using facial recognition technology [34]. While several states and cities that are heavily Democrat have adopted or proposed moratoriums on law enforcement using facial recognition technology, progress to enact federal regulation has been slow. Industry self-regulation is not immune to partisan politics. For example, Google's AI ethics board dissolved after the company's employees and external civil society groups protested the inclusion of Heritage Foundation president Kay Coles James and drone company CEO Dyan Gibbens on the board. Recognizing the public's divergent policy preferences on AI governance issues is necessary to have a productive public policy deliberation.

2.2 How Trust in Institutions Affects Regulation of Emerging Technology

The public's trust in various institutions to develop and manage AI could affect the regulation of the technology. General social trust is negatively correlated with perceived risk from technological hazards [31]. One observational study finds that that general social distrust is positively correlated with support for government regulation, even when the public perceives the government to be corrupt [2]. A follow-up study using more extensive survey data suggests that the public prefers governmental regulation if they trust the government more than they do major companies [24]. This latter finding is supported by how the public has reacted to genetically modified (GM) foods and nanotechnology. These two emerging technologies are similar to AI in that the public has to rely on those with scientific expertise when evaluating potential risks.

Distrust in institutions producing and regulating GM foods is a compelling explanation for the widespread opposition to GM foods in developed countries. Those with a high-level trust of scientists and regulators are more accepting of GM foods; in contrast, distrust of the agricultural/food industry, contrasted with trust in environmental watchdogs, predicts opposition to GM foods [15, 18, 23, 28, 29]. Although scientists are among the most trusted group in the U.S., Americans have cynical views toward scientists when considering GM foods. Only 19% thinks that scientists understand the health consequences of GM foods very well, even though scientists have formed a consensus that GM foods are safe to eat. Furthermore, the American public believes that scientists are more motivated by concerns for their industry than concerns for the public [12].

Nanotechnology, though less salient than GM foods in the media, is the subject of extensive public opinion research. A meta-analysis of 11 surveys conducted in developed countries finds that the public perceives that the use of nanotechnology has greater benefits than risks; nevertheless, a large subset of the public is uncertain of the consequences of nanotechnology [25]. As in perceptions of GM foods, trust in institutions seems to play a significant role in shaping attitudes toward nanotechnology. Americans who have lower confidence in business leaders within the nanotechnology industry also perceive the technology to be riskier [6]. A higher level of distrust in government agencies to protect the public from nanotechnology hazards is associated with a higher perceived risk of the technology [17, 32]. Consumers who are less trustful of the food industry indicate they are more reluctant to buy foods produced or packaged using nanotechnology [30].

2.3 Existing Survey Research on AI Governance and Trust in Tech Companies

Our survey builds on existing public opinion research on attitudes toward AI and trust in tech companies. Past survey research related to AI tends to focus on specific governance challenges, such as lethal autonomous weapons [13], algorithmic fairness [26], or facial recognition technology [1, 34]. A few large-scale surveys have taken a more comprehensive approach by asking about a range of AI governance challenges [5, 8, 35]. We improved upon these surveys by asking respondents to consider a variety of AI governance

challenges using the same question-wording. This consistency in question-wording allowed us to compare respondents' perceptions toward major issues in AI governance.

Another innovation of our research project is that we connect AI governance with trust in tech companies. Trust in tech companies has become a growing topic in survey research. While Americans perceive tech companies to have a more positive impact on society than other institutions, including the news media and the government, their feelings toward tech companies have declined dramatically since 2015 [7]. Nevertheless, there exists heterogeneity in how the public views individual tech companies; one consistent finding in non-academic surveys is that the public strongly distrusts Facebook, particularly in its handling of personal data [3, 14, 20]. In our survey, we examined whether this general distrust of tech companies extend to their management of AI development and applications.

3 SURVEY DATA AND ANALYSIS

We conducted an original online survey ($N=2000$) through YouGov between June 6 and 14, 2018. YouGov drew a random sample from the U.S. adult population (i.e., the target sample) and selected respondents from its online respondent panel that matched the target sample on key demographic variables. The details of YouGov's sample matching methodology can be found in the Appendix (available at <https://doi.org/10.7910/DVN/BREC5M>).

We pre-registered nearly all of the analysis on Open Science Framework (pre-analysis plan URL: <https://osf.io/7gqvm/>). Pre-registration increases research transparency by requiring researchers to specify their analysis before analyzing the data [21]. Doing so prevents researchers from misusing data analysis to come up with statistically significant results when they do not exist, otherwise known as p -hacking. Survey weights provided by YouGov were used in our primary analysis. We followed the Standard Operating Procedures for Don Green's Lab at Columbia University when handling missing data or "don't know" responses [16]. Heteroscedasticity-consistent standard errors were used to generate the margins of error at the 95% confidence level. We report cluster-robust standard errors whenever there is clustering by respondent. In figures, each error bar shows the 95% confidence intervals.

4 PERCEPTIONS OF AI GOVERNANCE CHALLENGES

4.1 Methodology

We sought to understand how Americans prioritize policy issues associated with AI. Respondents were asked to consider 13 AI governance challenges, randomly selected from a set of 13 (see the Appendix for the full text); the order these were given to each respondent was also randomized. After considering each governance challenge, respondents were asked how likely they think the challenge will affect large numbers of people both 1) in the U.S. and 2) around the world within 10 years using a seven-point scale that displayed both numerical likelihoods and qualitative descriptions (e.g., "Very unlikely: less than 5% chance (2.5%)"). Respondents were also asked to evaluate how important it is for tech companies and governments to carefully manage each challenge presented to them using a four-point scale.




Figure 2: Perceptions of AI governance challenges in the U.S. and around the world

4.2 General Results

We use scatterplots to visualize our survey results. In the top graph in Figure 2, the x-axis is the perceived likelihood of the problem happening to large numbers of people in the U.S. In the bottom graph, the x-axis is the perceived likelihood of the problem happening to large numbers of people around the world. The y-axes on both graphs in Figure 2 represent respondents' perceived issue importance, from 0 (not at all important) to 3 (very important). Each dot represents the mean perceived likelihood and issue importance. The correspondent ellipse represents the 95% confidence region of the bivariate means assuming the two variables are distributed multivariate normal.

Americans consider all the AI governance challenges we present to be important: the mean perceived issues importance of each governance challenge is between "somewhat important" (2) and "very important" (3), though there is meaningful and discernible variation across items.

The AI governance challenges Americans think are most likely to impact large numbers of people, and are important for tech companies and governments to tackle, are found in the upper-right quadrant of the two plots. These issues include data privacy as well as AI-enhanced cyber attacks, surveillance, and digital manipulation. We note that the media have widely covered these issues during the time of the survey.

There are a second set of governance challenges that are perceived on average, as about 7% less likely, and marginally less important. These include autonomous vehicles, value alignment, bias in using AI for hiring, the U.S.-China arms race, disease diagnosis, and technological unemployment. Finally, the third set of challenges are perceived on average another 5% less likely, and about equally important, including criminal justice bias and critical AI systems failures.

We also note that Americans predict that all of the governance challenges mentioned in the survey, besides protecting data privacy and ensuring the safety of autonomous vehicles, are more likely to impact people around the world than to affect people in the U.S. While most of these statistically significant differences are substantively small, one difference stands out: Americans think that autonomous weapons are 7.6 percentage points more likely to impact people around the world than Americans (two-sided p -value < 0.001).

We want to reflect on one result. “Value alignment” consists of an abstract description of the alignment problem and a reference to what sounds like individual level harms: “while performing jobs they could unintentionally make decisions that go against the values of its human users, such as physically harming people.” “Critical AI systems failures,” by contrast, references military or critical infrastructure uses, and unintentional accidents leading to “10 percent or more of all humans to die.” The latter was weighted as less important than the former: we interpret this as a probability-weighted assessment of importance, since presumably the latter, were it to happen, is much more important. We thus think the issue importance question should be interpreted in a way that down-weights low probability risks.

4.3 Subgroup Analysis

We performed further analysis by calculating the percentage of respondents in each subgroup who consider each governance challenge to be “very important” for governments and tech companies to manage. In general, differences in responses are more salient across demographic subgroups than across governance challenges. In a linear multiple regression predicting perceived issue importance using demographic subgroups, governance challenges, and the interaction between the two, we find that the stronger predictors are demographic subgroup variables, including age group and having CS or programming experience.

Two highly visible patterns emerge from our data visualization. First, a higher percentage of older respondents, compared with younger respondents, consider nearly all AI governance challenges to be “very important.” In another part of the survey (see Figure 7 in the Appendix), we find that older Americans, compared with younger Americans, are less supportive of developing AI. Our results here might explain this age gap: older Americans see each AI governance challenge as substantially more important than do younger Americans. Whereas 85% of Americans older than 73 consider each of these issues to be very important, only 40% of Americans younger than 38 do.

Second, those with CS or engineering degrees, compared with those who do not, rate all AI governance challenges as less important. This result could explain another finding in our survey




Figure 3: AI governance challenges: issue importance by demographic subgroups

that shows those with CS or engineering degrees tend to express greater support for developing AI. In Table 4 in the Appendix, we report the results of a saturated linear model using demographic variables, governance challenges, and the interaction between these two types of variables to predict perceived issue importance. We find that those who are 54-72 or 73 and older, relative to those who are below 38, view the governance challenges as more important (two-sided p -value < 0.001 for both comparisons). Furthermore, we find that those who have CS or engineering degrees, relative to those who do not, view the governance challenges as less important (two-sided p -value < 0.001).

5 TRUST IN ACTORS TO DEVELOP AND MANAGE AI

5.1 Methodology

Respondents were asked how much confidence they have in various actors to develop AI. They were randomly assigned five actors out of 15 to evaluate. We provided a short description of actors that are not well-known to the public (e.g., NATO, CERN, and OpenAI). Also, respondents were asked how much confidence, if any, they have in various actors to manage the development and use of AI in the best interests of the public. They were randomly assigned five out of 15

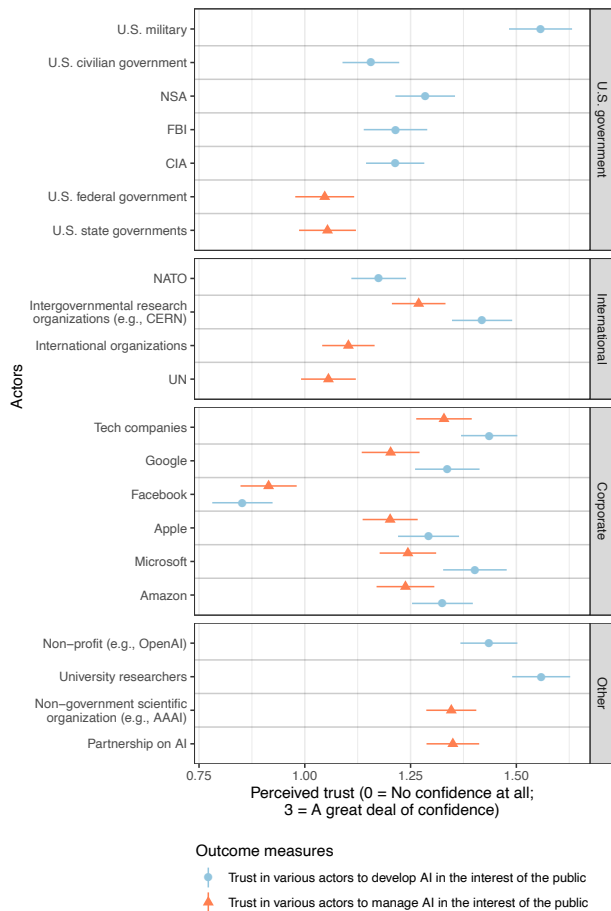


Figure 4: Trust in various actors to develop and manage AI in the interest of the public

actors to evaluate. Again, we provided a short description of actors that are not well-known to the public (e.g., AAAI and Partnership on AI). Confidence was measured using the same four-point scale described above. The two sets of 15 actors differed slightly because, for some actors, it seemed inappropriate to ask one or the other question. See the Appendix for the exact wording of the questions and descriptions of the actors.

5.2 Results

Americans do not express great confidence in most actors to develop or to manage AI, as seen in Figures 5 and 6 in the Appendix. A majority of Americans do not have a “great deal” or even a “fair amount” of confidence in any institution, except university researchers, to develop AI. Furthermore, Americans place greater trust in tech companies and non-governmental organizations (e.g., OpenAI) than in governments to manage the development and use of the technology.

University researchers and the U.S. military are the most trusted groups to develop AI: about half of Americans express a “great deal” or even a “fair amount” of confidence in them. Americans express slightly less confidence in tech companies, non-profit organizations

(e.g., OpenAI)¹, and American intelligence organizations. Nevertheless, opinions toward individual actors within each of these groups vary. For example, while 44% of Americans indicated they feel a “great deal” or even a “fair amount” of confidence in tech companies, they rate Facebook as the least trustworthy of all the actors. More than four in 10 indicate that they have no confidence in the company.

Our survey was conducted between June 6 and 14, 2018, shortly after the Facebook/Cambridge Analytica scandal became highly salient in the media. On April 10-11, 2018, Facebook CEO Mark Zuckerberg testified before the U.S. Congress regarding the Cambridge Analytica data leak. On May 2, 2018, Cambridge Analytica announced its shutdown. Nevertheless, Americans’ distrust of the company existed before the Facebook/Cambridge Analytica scandal. In a pilot survey that we conducted on Mechanical Turk during July 13–14, 2017, respondents indicated a substantially lower level of confidence in Facebook, compared with other actors, to develop and regulate AI.

The results on the public’s trust of various actors to manage the develop and use of AI provided are similar to the results discussed above. Again, a majority of Americans do not have a “great deal” or even a “fair amount” of confidence in any institution to manage AI. In general, the public expresses greater confidence in non-governmental organizations than in governmental ones. Indeed, 41% of Americans express a “great deal” or even a “fair amount” of confidence in “tech companies,” compared with 26% who feel that way about the U.S. federal government. But when presented with individual big tech companies, Americans indicate less trust in each than in the broader category of “tech companies.” Once again, Facebook stands out as an outlier: respondents give it a much lower rating than any other actor. Besides “tech companies,” the public places relatively high trust in intergovernmental research organizations (e.g., CERN)², the Partnership on AI, and non-governmental scientific organizations (e.g., AAAI). Nevertheless, because the public is less familiar with these organizations, about one in five respondents give a “don’t know” response.

Similar to our findings, recent survey research suggests that while Americans feel that AI should be regulated, they are unsure *who* the regulators should be. When asked who “should decide how AI systems are designed and deployed,” half of Americans indicated they do not know or refused to answer [36]. Our survey results also seem to reflect Americans’ general attitudes toward public institutions. According to a 2016 Pew Research Center survey, an overwhelming majority of Americans have “a great deal” or “a fair amount” of confidence in the U.S. military and scientists to act in the best interest of the public. In contrast, public confidence in elected officials is much lower: 73% indicated that they have “not too much confidence” or “no confidence” [11]. Less than one-third of Americans thought that tech companies do what’s right “most of the time” or “just about always”; moreover, more than half indicate that tech companies have too much power and influence in the

¹At the time the survey was conducted (June 2018), OpenAI was a 501(c)(3) nonprofit organization. In March 2019, OpenAI announced it is restructuring into “capped-profit” company to attract investments.

²We asked about trust in intergovernmental research organizations, such as CERN, because some policy researchers have proposed creating a politically neutral AI development hub similar to CERN to avoid risks associated with competition in AI development between rival states [9].

U.S. economy [33]. Nevertheless, Americans' attitudes toward tech companies are not monolithic but varies by company. For instance, our research findings reflect the results from several non-academic surveys that find the public distrusts Facebooks significantly more than other major tech companies [3, 14, 20].

5.3 Predicting Support for Developing AI Using Institutional Trust

Table 1: Regression results: predicting support for developing AI using respondents' trust in different types of actors to develop AI

Variable	Coefficient (SE)	<i>p</i> -value
Corporate actors	-0.02 (0.03)	0.596
International actors	-0.01 (0.03)	0.830
U.S. government actors	-0.01 (0.03)	0.799
Intercept	0.26 (0.03)	<0.001

$N = 10000$ observations from 2000 respondents (standard errors are clustered by respondent); F -statistic: 3.976 on 79 and 1999 DF, p -value: < 0.001. We controlled for the demographic variables found in Table 3.

Table 2: Regression result: predicting support for developing AI using respondents' trust in different types of actors to manage AI

Variable	Coefficient (SE)	<i>p</i> -value
Corporate actors	-0.03 (0.02)	0.271
International actors	0.02 (0.03)	0.517
U.S. government actors	0.02 (0.03)	0.449
Intercept	0.26 (0.03)	<0.001

$N = 10000$ observations from 2000 respondents (standard errors are clustered by respondent); F -statistic: 3.998 on 79 and 1999 DF, p -value: < 0.001. We controlled for the demographic variables found in Table 3.

The American public's support for developing AI is not predicted by their trust in various actors to develop and manage AI. Our null results stand in contrast to studies of GM foods and nanotechnology, where trust in institutions that develop and regulate the technology is associated with greater support for the technology.

For our analysis, we use multiple linear regression to whether individual-level trust in various types of institutions can predict support for developing AI. Note that this analysis is not included in our pre-registration and is therefore exploratory in nature. There are four categories of actors (the same ones shown in Figure 4): U.S. government, international bodies, corporate, and others (e.g., universities and non-profits). In our regressions, the group "others" is the reference group. Support for developing AI is measured using a five-point Likert scale, with -2 meaning "strongly oppose" and 2 meaning "strongly support." Our regressions controlled for all the

demographic variables shown in Table 3 in the Appendix. Trust in none of the actor types predicts support for developing AI, as seen in Tables 1 and 2.

6 CONCLUSION

To understand what the American public thinks about the regulation of the technology, we conducted a large-scale survey. The survey reveals that while Americans consider all AI governance challenges to have high issue importance, they do not necessarily trust the actors who have the power to develop and manage the technology to act in the public's interest. Nevertheless, as our exploratory analysis from the previous section shows, institutional distrust does not necessarily predict opposition to AI development.

One direction for future research is to examine other factors that shape the public's preferences toward AI governance. Technological knowledge, moral and psychological attributes, and perceptions of risks versus benefits are associated with support for GM foods and nanotechnology [25, 27]. Another direction for future research is to improve the measurement of institutional trust. Our existing survey questions focused on the public interest component of trust; future studies could investigate other components of trust, such as competence, transparency, and honesty [15]. Finally, we plan to investigate whether the public's perceived lack of political power makes them more distrustful of institutions to govern AI. In a recent survey of the British public, a large majority felt that they were unable to influence AI development, and many felt tech companies and governments have disproportionate influence [5]. Research on public perceptions of AI is nascent but could have a broad impact as AI governance moves from the realm of abstract principles into the world of mass politics.

ACKNOWLEDGMENTS

The research was funded by the Ethics and Governance of Artificial Intelligence Fund and Good Ventures. For useful feedback we would like to thank: Miles Brundage, Jack Clark, Kanta Dihal, Jeffrey Ding, Carrick Flynn, Ben Garfinkel, Rose Hadshar, Tim Hwang, Katelynn Kyker, Jade Leung, Luke Muehlhauser, Cullen O'Keefe, Michael Page, William Rathje, Carl Shulman, Brian Tse, Remco Zwetsloot, and the YouGov Team (Marissa Shih and Sam Luks). In particular, we are grateful for Markus Anderljung's insightful suggestions and detailed editing. We acknowledge Steven Van Tassel for his copy-editing and Will Marks and Catherine Peng for their research assistance.

REFERENCES

- [1] Ada Lovelace Institute. 2019. *Beyond face value: public attitudes to facial recognition technology*. Technical Report. Ada Lovelace Institute. <https://perma.cc/M4EG-N5MY>
- [2] Philippe Aghion, Yann Algan, Pierre Cahuc, and Andrei Shleifer. 2010. Regulation and distrust. *The Quarterly Journal of Economics* 125, 3 (2010), 1015–1049.
- [3] Nick Statt Casey Newton and Michael Zelenko. 2017. The Verge Tech Survey. *The Verge* (2017). <https://perma.cc/F75M-ZRYA>
- [4] Devin Caughey and Christopher Warshaw. 2018. Policy preferences and policy change: Dynamic responsiveness in the American states, 1936–2014. *American Political Science Review* 112, 2 (2018), 249–266.
- [5] Stephen Cave, Kate Coughlan, and Kanta Dihal. 2019. Scary robots: examining public responses to AI. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 331–337.
- [6] Michael D. Cobb and Jane Macoubrie. 2004. Public perceptions about nanotechnology: Risks, benefits and trust. *Journal of Nanoparticle Research* 6, 4 (2004),

- 395–405.
- [7] Carroll Doherty and Jocelyn Kiley. 2019. *Americans have become much less positive about tech companies' impact on the U.S.* Technical Report. Pew Research Center. <https://perma.cc/PHQ3-LD8D>
- [8] Eurobarometer. 2017. *Special Eurobarometer 460: Attitudes towards the impact of digitisation and automation on daily life.* Technical Report. Eurobarometer. <https://perma.cc/9FRT-ADST>
- [9] Sophie-Charlotte Fischer and Andreas Wenger. 2019. *A Politically Neutral Hub for Basic AI Research.* Technical Report. Center for Security Studies (CSS) at ETH Zurich. <https://perma.cc/JA6M-RDUS>
- [10] Luciano Floridi and Josh Cows. 2019. A Unified Framework of Five Principles for AI in Society. *Harvard Data Science Review* 1, 1 (23 6 2019). <https://doi.org/10.1162/99608f92.8cd550d1> <https://hdsr.mitpress.mit.edu/pub/10jsh9d1>
- [11] Cary Funk. 2017. Real Numbers: Mixed Messages about Public Trust in Science. *Issues in Science and Technology* 34, 1 (2017). <https://perma.cc/UF9P-WSRL>
- [12] Cary Funk and Brian Kennedy. 2016. *The New Food Fights: U.S. Public Divides Over Food Science.* Technical Report. Pew Research Center. <https://perma.cc/TJ2L-K9JU>
- [13] Michael C Horowitz. 2016. Public opinion and the politics of the killer robots debate. *Research & Politics* 3, 1 (2016). <https://doi.org/10.1177/2053168015627183>
- [14] Chris Kahn and David Ingram. 2018. Americans less likely to trust Facebook than rivals on personal data: Reuters/Ipsos poll. *Reuters* (2018). <https://perma.cc/5KSL-2253>
- [15] John T Lang and William K Hallman. 2005. Who does the public trust? The case of genetically modified food in the United States. *Risk Analysis: An International Journal* 25, 5 (2005), 1241–1252.
- [16] Winston Lin and Donald P Green. 2016. Standard operating procedures: A safety net for pre-analysis plans. *PS: Political Science & Politics* 49, 3 (2016), 495–500.
- [17] Jane Macoubrie. 2006. Nanotechnology: public concerns, reasoning and trust in government. *Public Understanding of Science* 15, 2 (2006), 221–241.
- [18] Mathew D Marques, Christine R Critchley, and Jarrod Walshe. 2015. Attitudes to genetically modified food over time: How trust in organizations and the media cycle predict support. *Public Understanding of Science* 24, 5 (2015), 601–618.
- [19] Brent Mittelstadt. 2019. AI Ethics—Too Principled to Fail? (2019). <https://dx.doi.org/10.2139/ssrn.3391293> Forthcoming in *Nature Machine Intelligence*.
- [20] Rani Molla. 2018. Facebook is the least-trusted major tech company. *Recode* (2018). <https://perma.cc/F75M-ZRYA>
- [21] Brian A Nosek, Charles R Ebersole, Alexander C DeHaven, and David T Mellor. 2018. The preregistration revolution. *Proceedings of the National Academy of Sciences* 115, 11 (2018), 2600–2606.
- [22] OECD. 2019. *Artificial Intelligence in Society.* <https://doi.org/https://doi.org/10.1787/eedfee77-en>
- [23] Anna Olofsson, Susanna Öhman, and Saman Rashid. 2006. Attitudes to gene technology: the significance of trust in institutions. *European Societies* 8, 4 (2006), 601–624.
- [24] Hans Pitlik and Ludek Kouba. 2015. Does social distrust always lead to a stronger support for government intervention? *Public Choice* 163, 3–4 (2015), 355–377.
- [25] Terre Satterfield, Milind Kandlikar, Christian EH Beaudrie, Joseph Conti, and Barbara Herr Harthorn. 2009. Anticipating the perceived risk of nanotechnologies. *Nature Nanotechnology* 4, 11 (2009), 752.
- [26] Nripsuta Ani Saxena, Karen Huang, Evan DeFilippis, Goran Radanovic, David C Parkes, and Yang Liu. 2019. How Do Fairness Definitions Fare?: Examining Public Attitudes Towards Algorithmic Definitions of Fairness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 99–106.
- [27] Sydney E Scott, Yoel Inbar, Christopher D Wirz, Dominique Brossard, and Paul Rozin. 2018. An overview of attitudes toward genetically engineered food. *Annual Review of Nutrition* 38 (2018), 459–479.
- [28] Michael Siegrist. 1999. A causal model explaining the perception and acceptance of gene technology 1. *Journal of Applied Social Psychology* 29, 10 (1999), 2093–2106.
- [29] Michael Siegrist. 2000. The influence of trust and perceptions of risks and benefits on the acceptance of gene technology. *Risk Analysis* 20, 2 (2000), 195–204.
- [30] Michael Siegrist, Marie-Eve Cousin, Hans Kastenholtz, and Arnim Wiek. 2007. Public acceptance of nanotechnology foods and food packaging: The influence of affect and trust. *Appetite* 49, 2 (2007), 459–466.
- [31] Michael Siegrist, Heinz Gutscher, and Timothy C Earle. 2005. Perception of risk: the influence of general trust, and general confidence. *Journal of Risk Research* 8, 2 (2005), 145–156.
- [32] Michael Siegrist, Carmen Keller, Hans Kastenholtz, Silvia Frey, and Arnim Wiek. 2007. Laypeople's and experts' perception of nanotechnology hazards. *Risk Analysis: An International Journal* 27, 1 (2007), 59–69.
- [33] Aaron Smith. 2018. *Public Attitudes Toward Technology Companies.* Survey report. Pew Research Center. <https://perma.cc/KSN6-6FRW>
- [34] Aaron Smith. 2019. *More Than Half of U.S. Adults Trust Law Enforcement to Use Facial Recognition Responsibly.* Technical Report. Pew Research Center. <https://perma.cc/FUV7-5BDJ>
- [35] Aaron Smith and Monica Anderson. 2016. *Automation in Everyday Life.* Technical Report. Pew Research Center. <https://perma.cc/WU6B-63PZ>
- [36] Darrell M. West. 2018. *Brookings survey finds divided views on artificial intelligence for warfare, but support rises if adversaries are developing it.* Survey report. Brookings Institution. <https://perma.cc/3NJV-5GV4>
- [37] Sarah Myers West, Meredith Whittaker, and Kate Crawford. 2019. *Discriminating Systems: Gender, Race and Power in AI.* Technical Report. AI Now Institute. <https://perma.cc/85VP-HL8R>
- [38] Jess Whittlestone, Rune Nyrup, Anna Alexandrova, and Stephen Cave. 2019. The role and limits of principles in AI ethics: towards a focus on tensions. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 195–200.
- [39] Baobao Zhang. 2019. *Public Opinion Lessons for AI Regulation.* Technical Report. Brookings Institution. <https://perma.cc/B3FY-E6JH>