Appendix for 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. https://doi.org/10.1145/3375627.3375827

A YOUGOV SAMPLING AND WEIGHTS

YouGov interviewed 2387 respondents who were then matched down to a sample of 2000 to produce the final dataset. The respondents were matched to a sampling frame on gender, age, race, and education. The frame was constructed by stratified sampling from the full 2016 American Community Survey (ACS) one-year sample with selection within strata by weighted sampling with replacements (using the person weights on the public use file).

The matched cases were weighted to the sampling frame using propensity scores. The matched cases and the frame were combined and a logistic regression was estimated for inclusion in the frame. The propensity score function included age, gender, race/ethnicity, years of education, and geographic region. The propensity scores were grouped into deciles of the estimated propensity score in the frame and post-stratified according to these deciles.

The weights were then post-stratified on 2016 U.S. presidential vote choice, and a four-way stratification of gender, age (four-categories), race (four-categories), and education (four-categories), to produce the final weight.

B TEXT OF THE QUESTIONS

Below, we present the survey text as shown to respondents. The numerical codings are shown in parentheses following each answer choice.

B.1 Should AI and/or robots should be carefully managed

Please tell me to what extent you agree or disagree with the following statement.
[Respondents were presented with one statement randomly selected from the list below.]

- AI and robots are technologies that require careful management.
- AI is a technology that requires careful management.
- Robots are technologies that require careful management.

ANSWER CHOICES:

- Totally agree (2)
- Tend to agree (1)
- Tend to disagree (-1)
- Totally disagree (-2)
- I don't know

B.2 Support for developing AI

[All respondents were presented with the following prompt.]
Next, we would like to ask you questions about your attitudes toward artificial intelligence.
Artificial Intelligence (AI) refers to computer systems that perform tasks or make decisions that usually require human intelligence. AI can perform these tasks or make these decisions without explicit human instructions. Today, AI has been used in the following applications:
[Respondents were shown five items randomly selected from the list below.]

- Translate over 100 different languages
- Predict one's Google searches
- Identify people from their photos
- Diagnose diseases like skin cancer and common illnesses
- Predict who are at risk of various diseases
- Help run factories and warehouses
- Block spam email
- Play computer games
- Help conduct legal case research
- Categorize photos and videos
- Detect plagiarism in essays
- Spot abusive messages on social media
- Predict what one is likely to buy online
- Predict what movies or TV shows one is likely to watch online

QUESTION: How much do you support or oppose the development of AI?
ANSWER CHOICES:

- Strongly support (2)
- Somewhat support (1)
- Neither support nor oppose (0)
- Somewhat oppose (-1)
- Strongly oppose (-2)
- I don't know

B.3 AI governance challenges

We would like you to consider some potential policy issues related to AI. Please consider the following:
[Respondents were shown five randomly-selected items from the list below, one item at a time. For ease of comprehension, we include the shorten labels used in the figures in bold.]

- **Hiring bias**: Fairness and transparency in AI used in hiring: Increasingly, employers are using AI to make hiring decisions. AI has the potential to make less biased hiring decisions than humans. But algorithms trained on biased data can lead to lead to hiring practices that discriminate against certain groups. Also, AI used in this application may lack transparency, such that human users do not understand what the algorithm is doing, or why it reaches certain decisions in specific cases.
- **Criminal justice bias**: Fairness and transparency in AI used in criminal justice: Increasingly, the criminal justice system is using AI to make sentencing and parole decisions. AI has the potential to make less biased hiring decisions than humans. But algorithms trained on biased data could lead to discrimination against certain groups. Also, AI used in this application may lack transparency such that human users do not understand what the algorithm is doing, or why it reaches certain decisions in specific cases.
- **Disease diagnosis**: Accuracy and transparency in AI used for disease diagnosis: Increasingly, AI software has been

used to diagnose diseases, such as heart disease and cancer. One challenge is to make sure the AI can correctly diagnose those who have the disease and not mistakenly diagnose those who do not have the disease. Another challenge is that AI used in this application may lack transparency such that human users do not understand what the algorithm is doing, or why it reaches certain decisions in specific cases.

- **Data privacy**: Protect data privacy: Algorithms used in AI applications are often trained on vast amounts of personal data, including medical records, social media content, and financial transactions. Some worry that data used to train algorithms are not collected, used, and stored in ways that protect personal privacy.
- **Autonomous vehicles**: Make sure autonomous vehicles are safe: Companies are developing self-driving cars and trucks that require little or no input from humans. Some worry about the safety of autonomous vehicles for those riding in them as well as for other vehicles, cyclists, and pedestrians.
- **Ditigal manipulation**: Prevent AI from being used to spread fake and harmful content online: AI has been used by governments, private groups, and individuals to harm or manipulate internet users. For instance, automated bots have been used to generate and spread false and/or harmful news stories, audios, and videos.
- **Cyber attacks**: Prevent AI cyber attacks against governments, companies, organizations, and individuals: Computer scientists have shown that AI can be used to launch effective cyber attacks. AI could be used to hack into servers to steal sensitive information, shut down critical infrastructures like power grids or hospital networks, or scale up targeted phishing attacks.
- **Surveillance**: Prevent AI-assisted surveillance from violating privacy and civil liberties: AI can be used to process and analyze large amounts of text, photo, audio, and video data from social media, mobile communications, and CCTV cameras. Some worry that governments, companies, and employers could use AI to increase their surveillance capabilities.
- **U.S.-China arms race**: Prevent escalation of a U.S.-China AI arms race: Leading analysts believe that an AI arms race is beginning, in which the U.S. and China are investing billions of dollars to develop powerful AI systems for surveillance, autonomous weapons, cyber operations, propaganda, and command and control systems. Some worry that a U.S.-China arms race could lead to extreme dangers. To stay ahead, the U.S. and China may race to deploy advanced military AI systems that they do not fully understand or can control. We could see catastrophic accidents, such as a rapid, automated escalation involving cyber and nuclear weapons.
- **Value alignment**: Make sure AI systems are safe, trustworthy, and aligned with human values: As AI systems become more advanced, they will increasingly make decisions without human input. One potential fear is that AI systems, while performing jobs they are programmed to do, could unintentionally make decisions that go against the values of its human users, such as physically harming people.
- **Autonomous weapons**: Ban the use of lethal autonomous weapons (LAWs): Lethal autonomous weapons (LAWs) are military robots that can attack targets without control by humans. LAWs could reduce the use of human combatants on the battlefield. But some worry that the adoption of LAWs could lead to mass violence. Because they are cheap and easy to produce in bulk, national militaries, terrorists, and other groups could readily deploy LAWs.
- **Technological unemployment**: Guarantee a good standard of living for those who lose their jobs to automation: Some forecast that AI will increasingly be able to do jobs done by humans today. AI could potentially do the jobs of blue-collar workers, like truckers and factory workers, as well as the jobs of white-collar workers, like financial analysts or lawyers. Some worry that in the future, robots and computers can do most of the jobs that are done by humans today.
- **Critical AI systems failure**: Prevent critical AI systems failures: As AI systems become more advanced, they could be used by the military or in critical infrastructure, like power grids, highways, or hospital networks. Some worry that the failure of AI systems or unintentional accidents in these applications could cause 10 percent or more of all humans to die.

QUESTION: In the next 10 years, how likely do you think it is that this AI governance challenge will impact large numbers of people in the U.S.?
ANSWER CHOICES:

- Very unlikely: less than 5% chance (2.5%)
- Unlikely: 5-20% chance (12.5%)
- Somewhat unlikely: 20-40% chance (30%)
- Equally likely as unlikely: 40-60% chance (50%)
- Somewhat likely: 60-80% chance (70%)
- Likely: 80-95% chance (87.5%)
- Very likely: more than 95% chance (97.5%)
- I don't know

QUESTION: In the next 10 years, how likely do you think it is that this AI governance challenge will impact large numbers of people around the world?
ANSWER CHOICES:

- Very unlikely: less than 5% chance (2.5%)
- Unlikely: 5-20% chance (12.5%)
- Somewhat unlikely: 20-40% chance (30%)
- Equally likely as unlikely: 40-60% chance (50%)
- Somewhat likely: 60-80% chance (70%)
- Likely: 80-95% chance (87.5%)
- Very likely: more than 95% chance (97.5%)
- I don't know

QUESTION: In the next 10 years, how important is it for tech companies and governments to carefully manage the following challenge?
ANSWER CHOICES:

- Very important (3)
- Somewhat important (2)
- Not too important (1)

- Not at all important (0)
- I don't know

B.4 Trust of actors to develop AI

QUESTION: How much confidence, if any, do you have in each of the following to develop AI in the best interests of the public?
[Respondents were shown five items randomly selected from the list below. We included explainer text for actors not well known to the public; respondents could view the explainer text by hovering their mouse over the actor's name. The items and the answer choices were shown in a matrix format.]

- The U.S. military
- The U.S. civilian government
- National Security Agency (NSA)
- Federal Bureau of Investigation (FBI)
- Central Intelligence Agency (CIA)
- North Atlantic Treaty Organization (NATO); explainer text for NATO: NATO is a military alliance that includes 28 countries including most of Europe, as well as the U.S. and Canada.
- An international research organization (e.g., CERN); explainer text for CERN: The European Organization for Nuclear Research, known as CERN, is a European research organization that operates the largest particle physics laboratory in the world.
- Tech companies
- Google
- Facebook
- Apple
- Microsoft
- Amazon
- A non-profit AI research organization (e.g., OpenAI); explainer text for OpenAI: Open AI is an AI non-profit organization with backing from tech investors that seeks to develop safe AI. University researchers

ANSWER CHOICES:

- A great deal of confidence (3)
- A fair amount of confidence (2)
- Not too much confidence (1)
- No confidence (0)
- I don't know

B.5 Trust of actors to manage AI

QUESTION: How much confidence, if any, do you have in each of the following to manage the development and use of AI in the best interests of the public?
[Respondents were shown five items randomly selected from the list below. We included explainer text for actors not well known to the public; respondents could view the explainer text by hovering their mouse over the actor's name. The items and the answer choices were shown in a matrix format.]

- U.S. federal government
- U.S. state governments
- International organizations (e.g., United Nations, European Union)
- The United Nations (UN)
- An intergovernmental research organization (e.g., CERN); explainer text for CERN: The European Organization for Nuclear Research, known as CERN, is a European research organization that operates the largest particle physics laboratory in the world.
- Tech companies
- Google
- Facebook
- Apple
- Microsoft
- Amazon
- Non-government scientific organizations (e.g., AAAI); explainer text for AAAI: Association for the Advancement of Artificial Intelligence (AAAI) is a non-government scientific organization that promotes research in, and responsible use of AI. Partnership on AI, an association of tech companies, academics, and civil society groups

ANSWER CHOICES:

- A great deal of confidence (3)
- A fair amount of confidence (2)
- Not too much confidence (1)
- No confidence (0)
- I don't know

C ADDITIONAL TABLES AND FIGURES

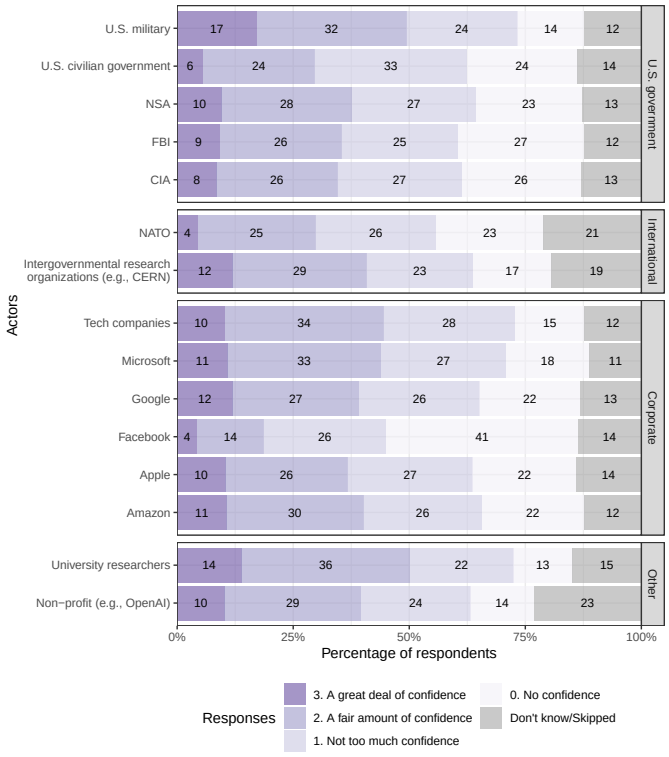

Figure 5: Trust in various actors to develop AI: distribution of responses

Demographic subgroup	Unweighted sample size
Age 18-37	702
Age 38-53	506
Age 54-72	616
Age 73 and older	176
Female	1048
Male	952
White	1289
Non-white	711
HS or less	742
Some college	645
College+	613
Not employed	1036
Employed (full- or part-time)	964
Income less than $30K	531
Income $30-70K	626
Income $70-100K	240
Income more than $100K	300
Prefer not to say income	303
Republican	470
Democrat	699
Independent/Other	831
Christian	1061
No religious affiliation	718
Other religion	221
Not born-again Christian	1443
Born-again Christian	557
No CS or engineering degree	1805
CS or engineering degree	195
No CS or programming experience	1265
CS or programming experience	735

Table 3: Size of demographic subgroups

Variable	Coefficient (SE)
Age 38-53	0.11 (0.07)
Age 54-72	0.35 (0.06)***
Age 73 and older	0.44 (0.07)***
Male	0.02 (0.05)
Non-white	-0.01 (0.05)
Some college	0.03 (0.07)
College+	0.15 (0.07)*
Employed (full- or part-time)	-0.09 (0.06)
Income $30-70K	0.09 (0.08)
Income $70-100K	0.13 (0.10)
Income more than $100K	-0.01 (0.10)
Prefer not to say income	0.04 (0.08)
Democrat	0.13 (0.07)
Independent/Other	0.14 (0.07)
No religious affiliation	-0.04 (0.06)
Other religion	-0.05 (0.08)
Born-again Christian	0.07 (0.07)
CS or engineering degree	-0.35 (0.10)***
CS or programming experience	-0.01 (0.07)
Criminal justice bias	0.05 (0.13)
Disease diagnosis	-0.06 (0.14)
Data privacy	0.16 (0.13)
Autonomous vehicles	-0.07 (0.14)
Digital manipulation	-0.14 (0.15)
Cyber attacks	0.05 (0.14)
Surveillance	<0.01 (0.15)
U.S.-China arms race	0.04 (0.13)
Value alignment	-0.06 (0.13)
Autonomous weapons	0.06 (0.14)
Technological unemployment	-0.12 (0.14)
Critical AI systems failure	-0.27 (0.15)
Intercept	2.25 (0.11)***
N = 10000 observations, 2000 respondents	$F(259,1999) = 3.36$; p-value: <0.001

Table 4: Results from a saturated regression predicting perceived issue importance using demographic variables, AI governance challenge, and interactions between the two types of variables; the coefficients for the interactions variables are not shown due to space constraints

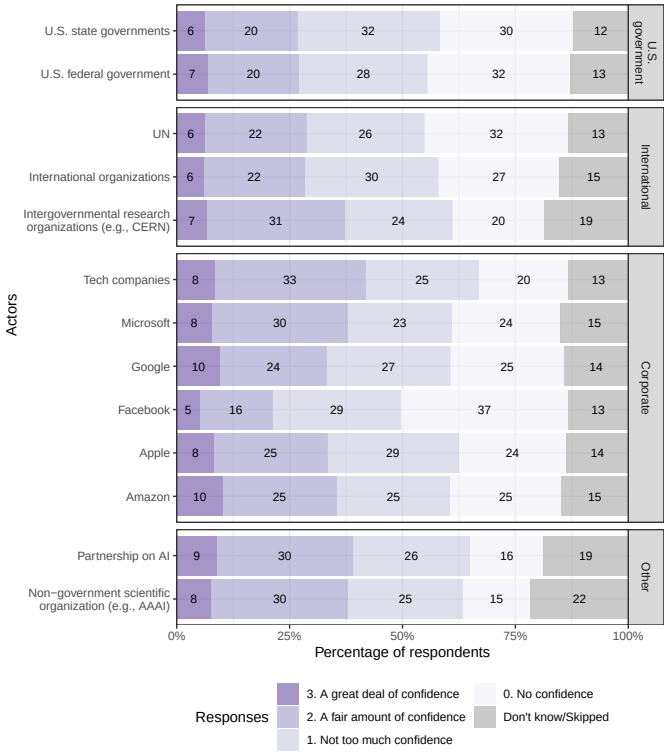

Figure 6: Trust in various actors to manage AI: distribution of responses

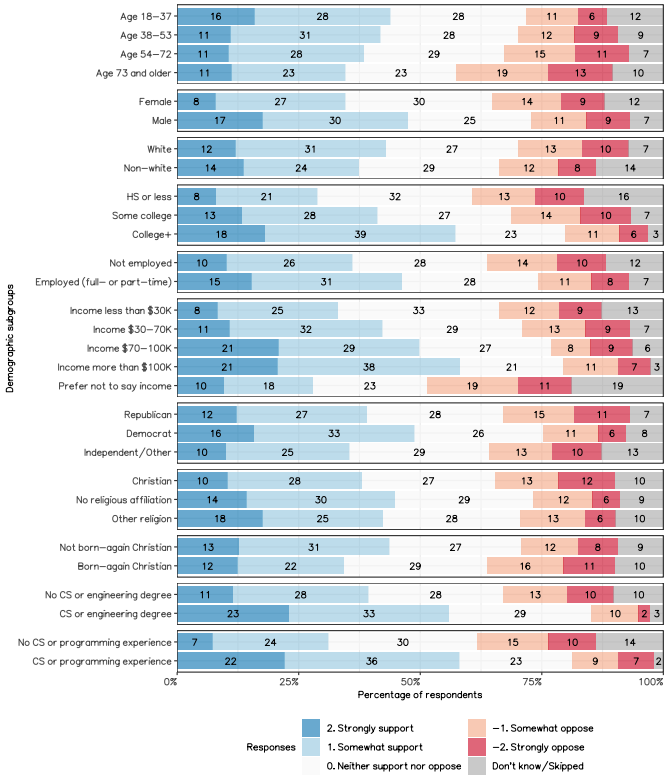

Figure 7: Support for developing AI across demographic characteristics: distribution of responses