



PROGRAM FOR PUBLIC CONSULTATION
SCHOOL OF PUBLIC POLICY, UNIVERSITY OF MARYLAND



ARTIFICIAL INTELLIGENCE

A National Survey of Registered Voters

Steven Kull
PRIMARY INVESTIGATOR

Evan Scot Fehsenfeld
Evan Charles Lewitus
RESEARCH ANALYSIS

March 2024

ARTIFICIAL INTELLIGENCE

A National Survey of Registered Voters

OVERVIEW

Over the last few years, Artificial Intelligence (AI) technology has advanced rapidly, and its usage throughout society has greatly increased. Developments in AI have brought numerous benefits, as well as many concerns about the effects of its current usage, and future potential for harm.

What role, if any, the Federal government should have in regulating the development and deployment of AI programs in the US, and internationally, has been the subject of much debate.

One set of proposed federal regulations target AI programs which make decisions in areas that can have significant impacts on people's lives, including in healthcare, banking, housing, education, employment, legal services, and utilities. Calls to regulate these AI programs have been in response to evidence that some have violated regulations or best practices, made biased decisions (e.g. by race, gender or age), and been hacked or misused.

Proposals have been put forward, in the Algorithmic Accountability Act by Sen. Wyden and Rep. Clark, as well as the European Union's Artificial Intelligence Act, which would regulate these decision-making AI programs, by:

- Requiring they pass a government-run test before they can be deployed, to ensure they will not violate regulations and best practices, make biased decisions, or have security vulnerabilities.
- Allowing the government to audit those currently in use, and if any of the above problems are found, require they be fixed.
- Requiring AI companies to disclose information to the Federal government on how these AI programs were trained.

Another set of proposed regulations target "deepfakes", which are fake images, video or audio generated using AI that are intended to look and sound like real people or events:

- Prohibit political campaign advertisements from using deepfakes. (Federal Elections Commission)
- Make it illegal to create and publicly distribute any deepfake which depicts a real person engaged in a pornographic act, without the consent of the person being portrayed. (Preventing Deepfakes of Intimate Images Act by Rep. Morelle; and DEFIANCE Act of 2023 by Sen. Durbin)
- Require all deepfakes be clearly labeled as such. (AI Labeling Act by Sen. Schatz and Rep. Kean; AI Disclosure Act by Rep. Torres; DEEP FAKES Accountability Act by Rep. Torres)

A broader proposal for government oversight and regulation of AI that has been put forward, is to create a new Federal agency devoted just to issues involving AI, rather than having such responsibilities dispersed throughout the various Federal departments. (Digital Platforms Commission Act by Sen. Bennett)



As AI technologies have spread globally, and become more able to cause harm across borders, the calls for international treaties have grown.

One international treaty being considered would prohibit the development and use of weapons that use AI to fire autonomously on targets, and create an international organization to oversee and enforce that prohibition. This proposal has been discussed by members of the UN, and has been advocated for by the Campaign to Stop Killer Robots.

Another international proposal that has been put forward relates to the potential development of powerful, large-scale and highly autonomous AI programs in the future, that could cause significant damage internationally – intentionally or unintentionally. These concerns have spurred calls, by UN Secretary-General António Guterres and OpenAI CEO Sam Altman, for the creation of an international organization to oversee and regulate all large-scale AI projects, similar to how the International Atomic Energy Agency regulates nuclear energy and weapons.

Survey Design

Unlike standard polls that rely on respondents' existing impressions and information, PPC took respondents through a process called a 'policymaking simulation' that seeks to put respondents in the shoes of a policymaker. Respondents were:

- given a briefing on policy options under consideration;
- evaluated strongly stated arguments both for and against each option; and only then
- made their final recommendation.

Respondents also evaluated how acceptable each proposal is to them using a 0-10 scale, with 0 as not all acceptable, 10 as very acceptable, and 5 as just tolerable.

The entire text of the survey was reviewed by experts on each side of the debate, to ensure that the briefings were accurate and balanced, and that the arguments presented were the strongest ones being made. Changes were made in response to their feedback.

The Flesch-Kincaid readability test was used to verify that the survey material would be comprehensible to a person with a high school education.

Briefing: Artificial Intelligence Developments, Benefits and Concerns

Before evaluating any proposals, respondents were presented with an overview of AI. First, they were told generally what AI is, what the new developments have been, and the benefits they have brought:

In brief, AI programs have the ability to complete tasks, make recommendations or even make decisions in a way that would have otherwise required human intelligence.

So what is new and different? The first computer programs were simply a set of instructions that a computer followed in an automatic and rigid manner.

Later, with developments in the field of Artificial Intelligence, computer programs were gradually taught to learn from examples, and even be somewhat autonomous. As a result, AI programs can increasingly come up with their own solutions to complex problems, and also communicate in ways that are increasingly human-like. You may have heard about or experienced the new ChatGPT program.

AI can also do things like create images and videos of people or events that appear very real even though they are not.



The use of AI programs has increased greatly over the last several years, by companies, governments, and individuals. They have been used in a wide range of areas including:

- *healthcare,*
- *engineering,*
- *transportation,*
- *consumer services,*
- *government services,*
- *banking,*
- *education,*
- *criminal justice*

and more. They have increased efficiency in many industries and improved scientific research.

It is widely believed they also have the potential to do such things as create new life-saving drugs, detect tumors earlier than doctors can currently, and reduce traffic injuries and deaths. Experts estimate that the use of AI will grow the global economy by several trillion dollars.

Respondents were then informed of the concerns with AI programs:

While, as discussed above, there are many positive results associated with AI, there are also widespread concerns about negative effects. Those with such concerns include people who are directly involved in developing AI, as well as people in government, people who represent workers, and so on.

Some of these concerns are long-term or hypothetical. There are concerns that as AI programs become more powerful, they could get into the hands of bad actors who could use AI in their efforts to do things like taking down energy grids or taking control of military weapons systems.

Among some AI experts there are also concerns that large-scale AI programs could be created that are highly intelligent with advanced capabilities, and, perhaps most significantly, have a high level of autonomy. According to these experts, these AI programs could become uncontrollable by humans and engage in dangerous behavior that causes massive harm.

Other experts think these concerns are exaggerated and an overreaction.

Later in the survey, they were presented briefings on what the immediate concerns with AI are, as well as what the potential long-term concerns are with increasingly powerful AI programs.

Respondents were also briefed on the preventative approach to regulation, in which the government (or an international body such as the United Nations) takes actions to help ensure no harm is done in the future; and evaluated arguments for and against that approach, including how it can hamper innovation and development.

Methodology

Fielding

The survey was fielded Feb 16-23, 2024 online with a national sample of 3,610 registered voters provided by Precision Sample from its larger sample, using a non-probability-based sample. The sample was provided by Precision Sample from its partner network of proprietary online, mobile, and social panels exceeding 9 million in the United States via an opt-in methodology. Panelists were randomly invited to participate in this study and balanced to the necessary demographics through quota management. Every panelist is validated by Precision Sample's proprietary 20-step validation process during the panel recruitment, registration and panel management phases. The sample has a confidence interval that varies from +/- 1.3 to 1.8%.



Weighting

Responses were weighted by age, gender, education, income, race and geographic region. Benchmarks for weights were obtained from the US Census' American Community Survey and the Current Populations Survey of Registered Voters. The sample was also weighted by partisan affiliation. The sample has a median weight of 1.5 and a mean weight of 2.0.

Analysis of Congressional Districts

A further analysis was conducted by dividing the sample six ways, depending on the PVI Cook rating of the respondent's Congressional district. This enabled comparison of respondents who live in very red, somewhat red, leaning red, leaning blue, somewhat blue, and very blue districts. Congressional districts were determined for 72% of respondents, using their self-reported 5-digit zip code, which was cross-checked against their reported state of residence.

SUMMARY OF KEY FINDINGS

Regulating Decision-Making AI Programs

Mandatory Government Pre-Tests

A very large bipartisan majority of nearly eight-in-ten favor requiring that any new AI program which will be used to make decisions that can significantly impact people's lives pass a government-run test before it can be deployed, to help ensure it will not violate regulations or best practices, make biased decisions, or have security vulnerabilities.

Government Audits

A very large bipartisan majority of three-quarters favor the government having the authority to audit AI programs that make decisions which significantly impact people's lives, and require any problems that are found to be fixed (e.g. regulatory violations, biases, security vulnerabilities).

Disclosure of Training Data

A very large bipartisan majority of seven-in-ten favor requiring that AI companies provide the government with information about how the AI was trained.

Regulating Deepfakes

Require Labels

An overwhelming bipartisan majority of over eight-in-ten favor requiring that all deepfakes which are publicly distributed be clearly labeled as such, excluding those used in entertainment to impersonate a real person with their consent.

Prohibit Deepfakes in Political Campaign Ads

An overwhelming bipartisan majority of over eight-in-ten favor making it illegal for political campaigns to use deepfakes in their campaign advertisements.

Prohibit Sharing of Pornographic Deepfakes Without Consent

An overwhelming bipartisan majority of over eight-in-ten favor making it illegal to publicly distribute a pornographic deepfake without first getting the consent of the individuals being depicted.



Federal Agency for AI

A very large bipartisan majority of over seven-in-ten favor the creation of a new federal agency to enforce regulations on AI, provide guidance on policy, and oversee developments in AI.

International Ban on Lethal Autonomous Weapons

An overwhelming bipartisan majority of eight-in-ten favor the US actively working to create an international treaty to ban the development and use of weapons that can fire on targets autonomously, with a UN agency in charge of monitoring and enforcement.

International Agency to Monitor and Regulate AI

A very large bipartisan majority of three-quarters favor the US actively working to create an international organization with intrusive powers to monitor and regulate large-scale AI projects that have the potential to cause international harm – intentionally or unintentionally.



BRIEFINGS

New Development in AI

An explanation of what Artificial Intelligence (AI) is, and what the recent developments have been, was presented as follows:

In brief, AI programs have the ability to complete tasks, make recommendations or even make decisions in a way that would have otherwise required human intelligence.

So, what is new and different? The first computer programs were simply a set of instructions that a computer followed in an automatic and rigid manner.

Later, with developments in the field of Artificial Intelligence, computer programs were gradually taught to learn from examples, and even be somewhat autonomous. As a result, AI programs can increasingly come up with their own solutions to complex problems, and also communicate in ways that are increasingly human-like. You may have heard about or experienced the new ChatGPT program.

AI can also do things like create images and videos of people or events that appear very real even though they are not.

The use of AI programs has increased greatly over the last several years, by companies, governments, and individuals. They have been used in a wide range of areas including:

- *healthcare*
- *engineering*
- *transportation*
- *consumer services*
- *government services*
- *banking*
- *education*
- *criminal justice*

and more. They have increased efficiency in many industries and improved scientific research.

It is widely believed they also have the potential to do such things as create new life-saving drugs, detect tumors earlier than doctors can currently, and reduce traffic injuries and deaths. Experts estimate that the use of AI will grow the global economy by several trillion dollars.

Immediate Concerns and Federal Proposals

The immediate concerns that AI programs present, and which can be regulated domestically by the government, were presented to respondents as follows:

First, we will address immediate concerns about AI programs that are already being used.

For example, some AI programs have:

- *violated regulations, though they were not instructed to do so*
- *provided incorrect information*
- *made flawed recommendations or decisions*
- *unintentionally treated some groups in a biased way (e.g. by race or gender)*



AI programs have also been purposely used to:

- create misinformation very quickly and on a large scale,
- create fake videos of people or events that appear very real which have misled people or damaged reputations
- steal private data

AI programs have also been hacked and used for harmful purposes. Some of these concerns can be addressed at the national level, by the federal government. We will explore proposals for what the government might do.

Preventative Approach to Regulation

Several of the proposals for regulating AI are based on the government taking a preventative approach, and so respondents were introduced to that idea:

As mentioned, there is debate about what role the government should play in regulating AI companies and AI programs.

There are two general approaches that the government can take:



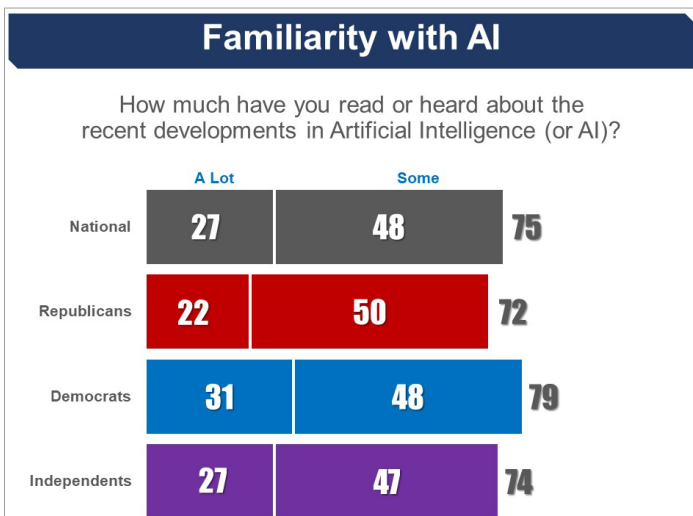
One approach is for the government to take action only after a company has sold a product or service, something has gone wrong, and the product has harmed consumers in some way.

Another approach is for the government to more actively intervene in advance to try to prevent harm from happening. This is called a preventative approach. This approach is used by the government in some areas, such as in healthcare, whereby the government requires new drugs to pass a series of tests before they can be put on the market.

Familiarity with AI

After being presented a briefing on the recent developments in AI, respondents were asked how much they “have read or heard about those developments. Three-quarters said a lot (27%) or some (48%), 21% said a little, and just 4% said nothing at all.

Respondents were also asked how much they had read or heard about concerns with AI, after being presented a description of the immediate and long-term concerns. Over seven-in-ten said a lot (27%) or some (46%), 22% said a little, and just 6% said nothing at all.



FINDINGS

REGULATING DECISION-MAKING AI PROGRAMS

Mandatory Government Pre-Tests

A very large bipartisan majority of nearly eight-in-ten favor requiring that any new AI program which will be used to make decisions that can significantly impact people's lives pass a government-run test before it can be deployed, to help ensure it will not violate regulations or best practices, make biased decisions, or have security vulnerabilities.

Respondents were first introduced to the idea of pre-testing new AI programs as follows:

One way that the government can take a preventative approach with AI is to require that new AI programs pass a series of tests before they can be put into general use. This is called "pre-testing". This would be similar to how the government requires testing new drugs.

The proposal was then introduced, including a description of the types of AI programs that the proposal is targeted at: AI programs that are used to make decisions that can have significant impacts on people's lives:

There is now a proposal to require pre-testing of new AI programs that are going to be used to make decisions that can have significant impacts on people, including in healthcare, banking, housing, education, employment, legal services, and utilities like electricity.

For example, this would include AI programs used:

- *by banks to determine who gets accepted for a loan,*
- *by government agencies to determine whether a person is eligible for government benefits, such as food stamps*
- *by health insurance companies to determine whether a person's medical treatment is covered*
- *by companies to determine whether a person should be hired*
- *by utility companies to determine how to allocate resources, like electricity when there is a shortage*

The tests would try to ensure that the AI program:

- *follows regulations to reduce the chances that it will break the law*
- *follows best practices established by professionals, to reduce the chances it will cause harm*
- *has security protections for data privacy and against hacking*
- *does not have unintended biases that result in it treating some groups worse than others, based on their race, gender, religion, age, sexual orientation, or nationality*

These tests would be run by the government, or by an independent third-party verified by the government. If the AI program does not pass the tests, it will not be approved for general use.


The arguments in favor were found convincing by very large and bipartisan majorities. Four of the five arguments against were also found convincing by majorities, but much smaller majorities than the pro arguments and less bipartisan. More Republicans than Democrats found the con arguments convincing.



Require Pre-Testing of AI Programs

ARGUMENT IN FAVOR


AI programs have the potential to cause harm to millions of people, even in ways we can't yet anticipate. Because of the widespread use and power of AI, the government has a duty to regulate it. The government shouldn't just react after the harm has been caused. The government already takes a preventative approach with lots of products that can cause mass harm, such as new drugs and chemicals. AI should be treated the same.



Require Pre-Testing of AI Programs

ARGUMENT AGAINST


The government should only intervene when there is clear evidence that a problem has occurred. These AI programs are nothing like new drugs; in most cases they are just doing tasks that humans used to do, in many cases substantially better. If the government were to try to anticipate all possible harms of new technologies, it would slow down or bias the direction of development. It would cost money for AI businesses that would then pass the costs onto consumers. Being so cautious will hurt innovation and we could lose out on many possible benefits of AI.



Require Pre-Testing of AI Programs

ARGUMENT IN FAVOR


These technologies are advancing quickly, and the corporations creating them are often reckless in their pursuit of profit. To get to market ahead of their competitors, they cut corners on safety testing. It's better to be cautious with this new technology, even if it means slowing down some innovation, than to find ourselves cleaning up a huge mess later. It would rightly strengthen confidence in US-made AI products. If we had taken a more cautionary approach with new technologies, like the internet and social media, we might not have some of the problems we have now.



Require Pre-Testing of AI Programs

ARGUMENT AGAINST


The private sector can move faster than the government to address risks from rapidly advancing technologies. The market is developing certifications and standards to test for safety and reliability without government intervention. The AI industry has already voluntarily committed to testing their products. The government getting involved would just slow down this whole process, and there's no guarantee they would do a better job at pre-testing than the industry itself.



Require Pre-Testing of AI Programs

ARGUMENT IN FAVOR


AI programs have frequently made errors that have caused widespread and irreversible harm. An AI program that was developed for state governments to detect fraud in unemployment insurance ended up wrongfully accusing thousands of people of fraud. Those people had their wages taken to repay benefits they had received. Some ended up being evicted from their homes. These AI programs are being used all over the country. They must be tested before they are put into use.



Require Pre-Testing of AI Programs

ARGUMENT AGAINST


With any new technology there will be some hiccups, but that doesn't mean we should overreact. We already have laws to deal with possible problems. If, in the unusual case, an AI program unintentionally violates some regulations or causes harm, then the company that made it, or the organization that used it, will be held legally liable.



Require Pre-Testing of AI Programs

ARGUMENT IN FAVOR


Companies have been using AI programs to increase profits by purposely harming their customers. For example, AI programs used by health insurance companies were designed to reject as many claims as possible, rather than make the most accurate judgment. Hundreds of thousands of people were wrongfully denied coverage for needed medical treatment. And when people pointed this out, the company just blamed the AI and denied responsibility. We need to make sure that companies can't exploit AI for their own benefit.



Require Pre-Testing of AI Programs

ARGUMENT AGAINST


It is not in the interest of the company to harm their customers, who will then take their business elsewhere. It is clearly in the company's interest to anticipate any problem in advance and to fix them as quickly as possible. This is a problem that the market will solve. We don't need the government inserting itself into the market with a whole pre-testing bureaucracy.



Require Pre-Testing of AI Programs

ARGUMENT IN FAVOR


A serious flaw that many AI programs have is bias. Research has shown that many AI programs treat some groups of people worse than others—especially racial minorities, older people and women. For example, an AI program used by banks to decide who gets accepted for a loan found that it rejected these groups more than human loan officers would. Thus, AI programs can make inequalities worse. AI has the potential to make life better for everyone, no matter who they are, but that will only happen if we actively steer it in the right direction.



Require Pre-Testing of AI Programs

ARGUMENT AGAINST

This proposal will give Federal bureaucrats excessive power in shaping AI programs. It gives them the ability to insert their own bias into these pre-tests and could require that AI programs give priority to certain values – liberal or conservative – over being good at the actual task it's supposed to accomplish. This will result in AI programs that are worse at their jobs, which will harm everyone.



Asked for their final recommendation, 81% were in favor, including 76% of Republicans and 88% of Democrats


Majorities in all types of Congressional districts were in favor, from very red (79%) to very blue (84%).



Require Pre-Testing of AI Programs

FINAL RECOMMENDATION

Do you favor or oppose the proposal to require pre-testing of AI programs?



CONGRESSIONAL DISTRICTS



Government Audits

A very large bipartisan majority of three-quarters favor the government having the authority to audit AI programs that make decisions which significantly impact people’s lives, and require any problems that are found to be fixed (e.g. regulatory violations, biases, security vulnerabilities).

Respondents were told that, “some AI programs are already in use and have not been pre-tested. AI programs can also change over time as they learn more or are updated by the company.

They were presented a proposal to preventatively regulate decision-making AI programs that are already in use:

Give the government the authority to audit AI programs, or to contract independent third-parties to audit them, that are already in use and that make decisions which have significant impacts on people’s lives.

The audits would include tests on whether the program follows regulations and best practices, has data privacy and security protections, and does not have unintended biases. If the audit finds that the AI program has problems in any of those areas, then the company who owns the AI program would have to fix them and redistribute the corrected version.

Asked for their final recommendation, 77% were in favor, including 74% of Republicans and 82% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (74%) to very blue (78%).



Disclosure of Training Data

A very large bipartisan majority of seven-in-ten favor requiring that AI companies provide the government with information about how the AI was trained.

Respondents were told that, “if the government does pre-test or audit AI programs, another question is how much access the government would have, to see how the AI companies develop their programs.”

They were then presented the following proposal:

Require that AI companies provide the government with information about how the AI was trained, when the government requests it. This would include a summary of the data used to train the AI, and a description of how the data was obtained. This would not include any sensitive information about individuals, such as medical or financial records.


All of the arguments in favor and against were found convincing by bipartisan majorities, but the pro arguments did better overall, and among Democrats and Republicans.



Require Disclosure of Training

ARGUMENT IN FAVOR


Without access to training data, the government could never know whether an AI company has been illegally collecting data unless a whistleblower inside the company decides to come forward. These data could include private health or financial information obtained without consent, or even images or videos of minors. One AI company recently settled a lawsuit for collecting face scans of people without their consent to create an AI-powered surveillance program.



Require Disclosure of Training

ARGUMENT AGAINST


The government should not have the power to force a private company to hand over information about how its AI program was developed. Any time a company has to disclose its data it increases the risk that it is leaked, and the government is known to have many data breaches. This could compromise intellectual property. If the AI program is working fine, then it shouldn't matter what is happening under the hood; and if there is a problem, then the AI company can look into the training data itself.



Require Disclosure of Training

ARGUMENT IN FAVOR


Pre-tests and audits will not be able to see all the problems with an AI program because they are only looking at the outputs, and they won't be able to look at every possible output. Being able to look inside the AI program – the data that it was trained on and how it was programmed – will give the government an important tool for catching problems before they can cause harm.



Require Disclosure of Training

ARGUMENT AGAINST

It is not consistent with the principles of the American constitution for the government to be able to effectively search your property (including your AI programs), whenever it wants to, looking for something it believes might be illegal. If a government agency has reason to believe there is something wrong, it can go to a judge, make its case, and the judge can decide whether to require the company to disclose the training data.




Asked for their final recommendation, 72% were in favor, including 67% of Republicans and 81% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (69%) to very blue (72%).

Require Disclosure of Training

FINAL RECOMMENDATION

Do you favor or oppose the proposal to require that AI companies provide the government with information about how the AI was trained, when the government requests it? This would include a summary of the data used to train the AI, and a description of how the data was obtained.



CONGRESSIONAL DISTRICTS



REGULATING DEEPFAKES

Before evaluating any proposals to regulate deepfakes, a description of what deepfakes are was first presented, as follows:

As you may know, some AI programs can create fake images, audio, or videos, of real people or events, that look completely real. These images or videos are known as “deepfakes”.

Programs like Photoshop have already made it possible to make fake images, but with AI programs the deepfakes are more realistic, harder to detect, and can be readily applied to video as well as fixed images.

Require Labels

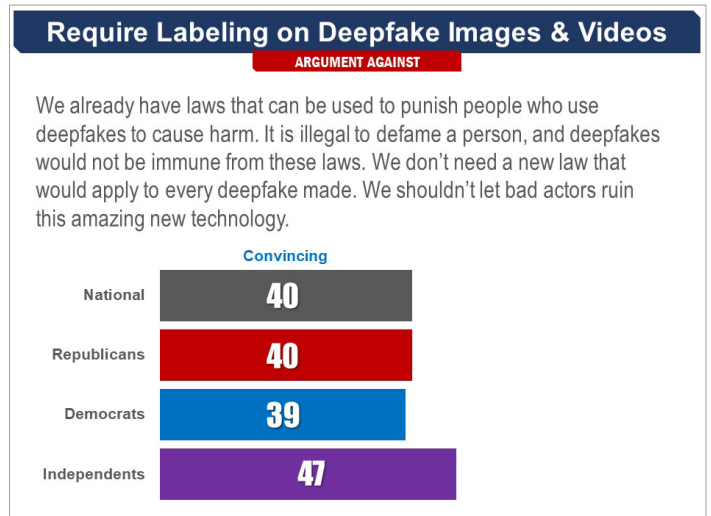
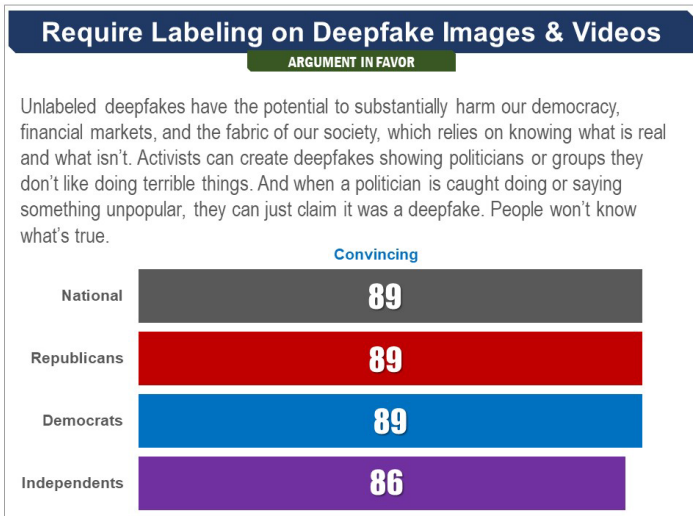
An overwhelming bipartisan majority of over eight-in-ten favor requiring that all deepfakes which are publicly distributed be clearly labeled as such, excluding those used in entertainment to impersonate a real person with their consent.

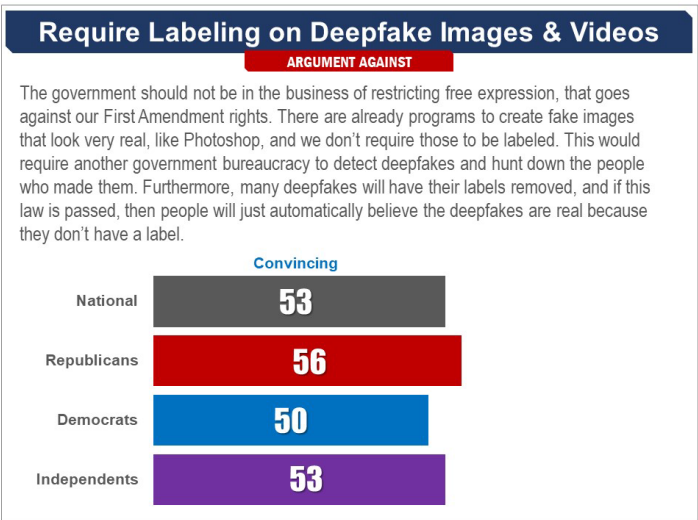
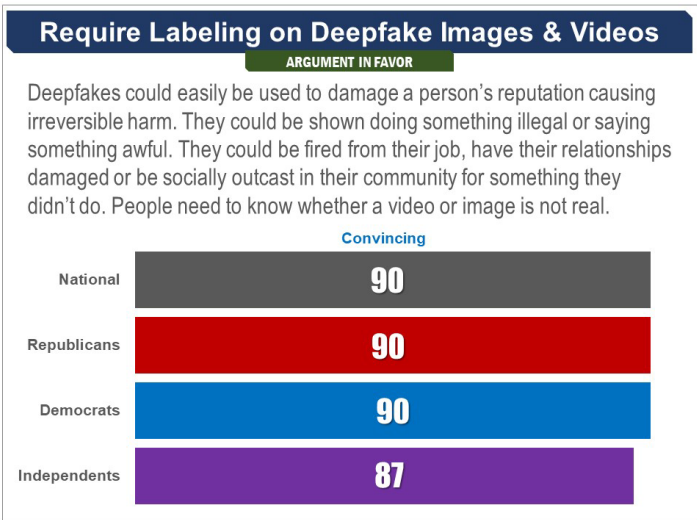
Respondents were presented the following proposal to regulate deepfakes, by requiring they be labeled:

Require that any deepfake image or video distributed publicly – e.g. posted online or shown on TV – must have a label that states that it is not real and was generated by AI. For videos, this label would need to be present the entire time the deepfake is on the screen. For audio deepfakes, they would be required to have a verbal statement at the beginning.

Deepfakes that are used for entertainment purposes to impersonate a real person (such as portraying a movie actor as younger), would not be required to have a label, as long as the person being portrayed has given their consent.


The arguments in favor were found convincing by very large bipartisan majorities. The arguments against did not do as well, with one found convincing by less than half overall, including less than half of Republicans and Democrats; and the other by just over half.





Asked for their final recommendation, 83% were in favor, including 83% of Republicans and 85% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (81%) to very blue (85%).



Prohibit Deepfakes in Political Campaign Ads

An overwhelming bipartisan majority of over eight-in-ten favor making it illegal for political campaigns to use deepfakes in their campaign advertisements.

Respondents were told that, “there have already been campaign advertisements that have used deepfakes depicting politicians doing or saying things they have not, and events that have not happened.”

They were then introduced to the following proposal: “Make it illegal for political campaigns, including PACs, to use deepfakes in their campaign advertisements.”


The arguments in favor were found convincing by overwhelming bipartisan majorities, while the arguments against did substantially worse, with each found convincing by less than half, including less than half of Republicans and Democrats.



Prohibiting Deepfakes on Campaign Ads

ARGUMENT IN FAVOR


Defamation laws will not solve this problem. It is nearly impossible to convict someone of defaming a public figure like a politician and lawsuits take years. Also, defamation laws do not apply to many political deepfakes such as fake videos of terrible things that the ad falsely claims happened when their opponent was in office. We need a new law that explicitly prohibits any use of deepfakes in political ads.



Prohibiting Deepfakes on Campaign Ads

ARGUMENT AGAINST


There are already laws in place that can apply to politicians creating fake videos with the purpose of harming their political opponent's reputation. That's called defamation and it's illegal. We do not need to make a whole new law.



Prohibiting Deepfakes on Campaign Ads

ARGUMENT IN FAVOR


We have already seen the damage that comes from politicians lying and spreading disinformation. Deepfakes can make this so much worse. It can further divide the country: People who oppose a politician will tend to believe the fake videos that make them look bad, and their supporters won't. Even if the news reports that a video is fake, it is often too late because the video has been viewed by millions.



Prohibiting Deepfakes on Campaign Ads

ARGUMENT AGAINST

There are many good uses for this technology in campaigns. They can be used to show people what a politician's policies could achieve, for example by creating videos of new hospitals being built in a rural town; or to show people what risks the politician is worried about, such as a local bridge collapsing because of their opponent's refusal to invest in infrastructure. We should not simply ban this form of expression.




Asked for their final recommendation, 84% were in favor, including 83% of Republicans and 86% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (78%) to very blue (78%).

Prohibit Deepfakes in Campaign Ads

FINAL RECOMMENDATION

Do you favor or oppose the proposal to make it illegal for political campaigns, including PACs, to use deepfakes in their campaign advertisements?



CONGRESSIONAL DISTRICTS



Prohibit Sharing of Pornographic Deepfakes Without Consent

An overwhelming bipartisan majority of over eight-in-ten favor making it illegal to publicly distribute a pornographic deepfake without first getting the consent of the individuals being depicted.

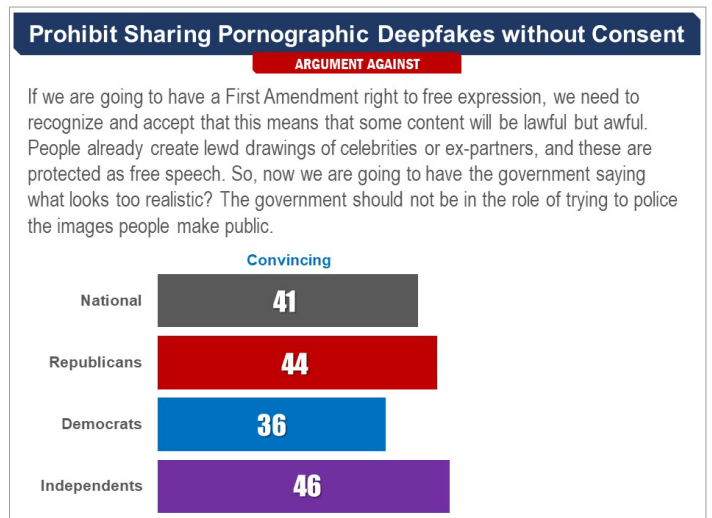
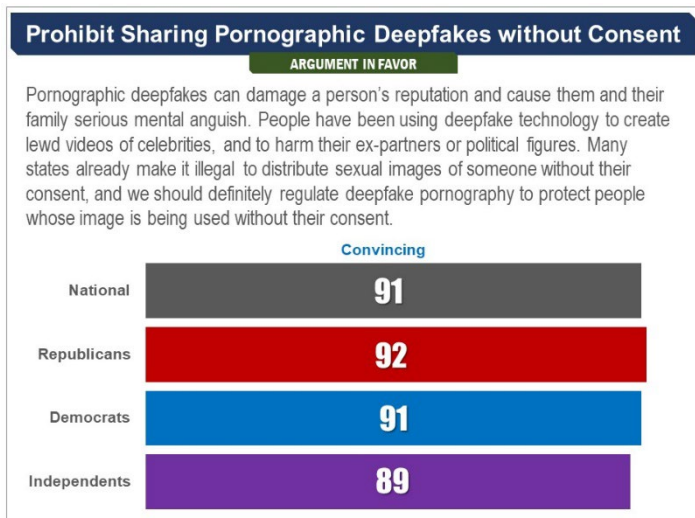
Respondents were introduced to the topic of pornographic deepfakes as follows:

As you may know, people have created deepfake images and videos of individuals engaging in sexual activities without that person's consent. For example, people's faces have been put on images and videos of other people engaging in sexual acts. These deepfakes have then been posted publicly online.

The proposal was then presented:

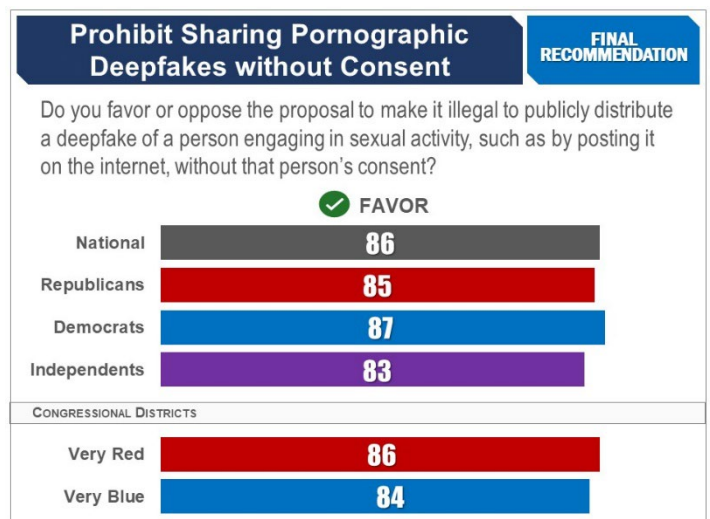
Make it illegal to publicly distribute a deepfake of a person engaging in sexual activity, such as by posting it on the internet, without that person's consent. It would not apply to people who only make such deepfakes for their personal use and do not make them public.

The argument in favor was found convincing by an overwhelming bipartisan majority, while the argument against did quite poorly, with less than half finding it convincing, including less than half of Republicans and Democrats.



Asked for their final recommendation, 86% were in favor, including 85% of Republicans and 87% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (86%) to very blue (84%).



Federal Agency for AI

A very large bipartisan majority of over seven-in-ten favor the creation of a new federal agency to enforce regulations on AI, provide guidance on policy, and oversee developments in AI.

Respondents were introduced to the idea of a new Federal agency for AI as follows:



So far, we have been talking about some specific problems with AI. We are now going to explore a more general proposal for having a Federal agency for AI. Currently, a variety of federal agencies are responding to specific concerns with AI programs that are arising in their own area of expertise.

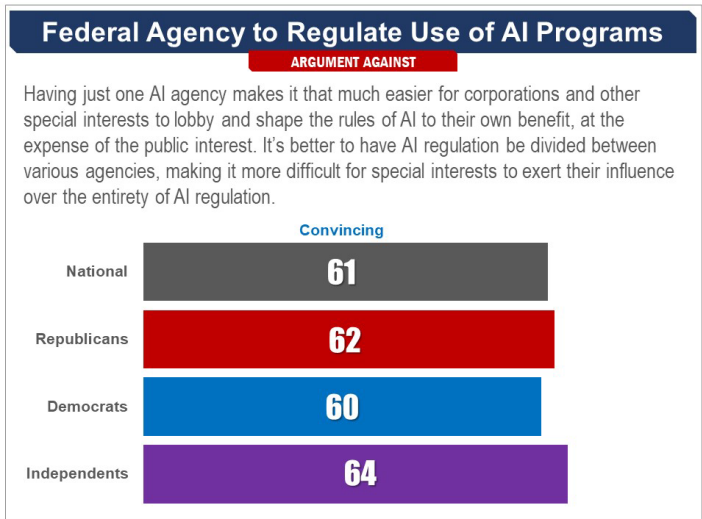

The proposal was then presented:

This proposed agency for AI would take a preventative and comprehensive approach to overseeing and regulating the development and use of AI programs. The agency would:

- *closely monitor the state of AI programs and their uses, and try to anticipate potential problems*
- *define best practices for developing and using AI programs, based on input from AI experts, industry leaders, and other professionals*
- *make recommendations for AI regulations to Congress and the Executive Branch*
- *enforce AI regulations that have been adopted*


The arguments in favor were found convincing by very large bipartisan majorities. The arguments against did not do as well, but were still found convincing by majorities overall.





Asked for their final recommendation, 74% were in favor, including 68% of Republicans and 81% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (71%) to very blue (75%).



International Regulation

Before evaluating specific proposals for regulating AI on the international level, respondents were told:

As you may know, AI programs are being developed in many different countries, and are sold internationally. Currently, there are no global treaties or agreements regulating the development and uses of AI programs.

International Ban on Lethal Autonomous Weapons

An overwhelming bipartisan majority of eight-in-ten favor the US actively working to create an international treaty to ban the development and use of weapons that can fire on targets autonomously, with a UN agency in charge of monitoring and enforcement.

A description of lethal autonomous weapons and concerns about their future development were presented as follows:

As you may know, AI programs have been put into weapons to assist with finding and locking onto targets. There is a concern that the weapon will not only be programmed to find a certain type of target (enemy combatant or military site), but also to make the decision whether to fire on a target, independent of any human choice at the time. These types of weapons are known as lethal autonomous weapons.



The reason that militaries would build lethal autonomous weapons is that they can be more efficient and effective than weapons which require some human control: Thousands of them can be deployed at the same time without the need for an equivalent amount of humans controlling them or making the final decision to attack targets.

There is a concern that these weapons may not always accurately distinguish the target, and may end up firing on civilians or non-military sites.

They were then presented the following proposal:

A proposal has been put forward for an international treaty that would prohibit lethal autonomous weapons. Weapons could use AI to find and lock onto a target, but a human would have to decide whether it fires on that target.

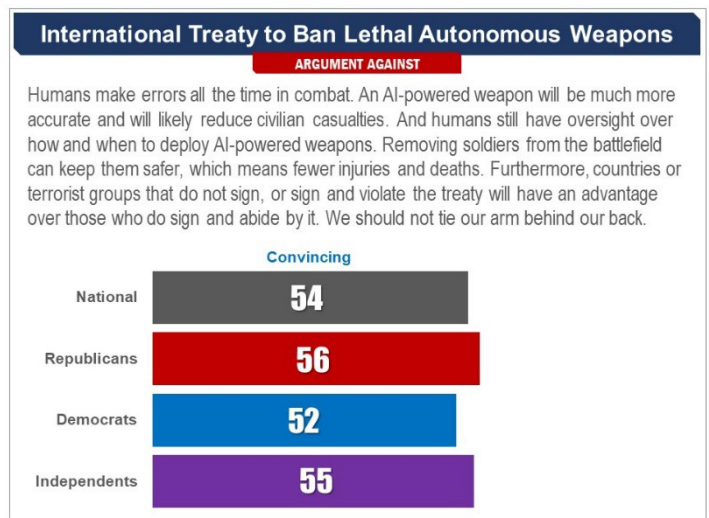
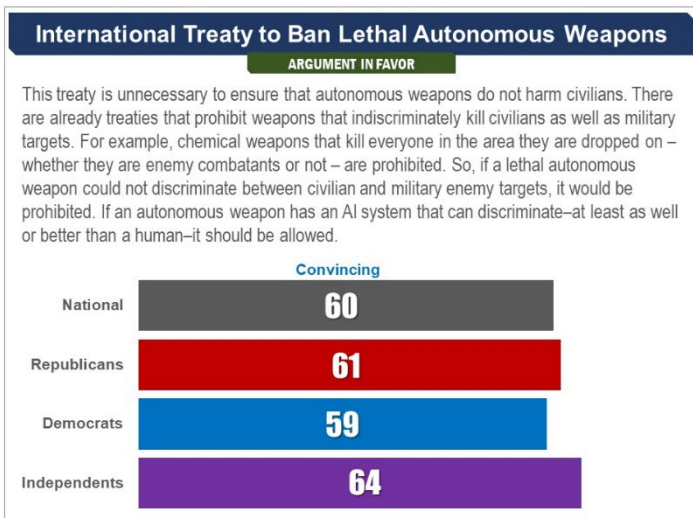
The treaty would also have a UN agency enforce this requirement. Member nations would have to disclose information about the use of AI in their weapons systems and allow the UN agency to inspect their weapons systems.

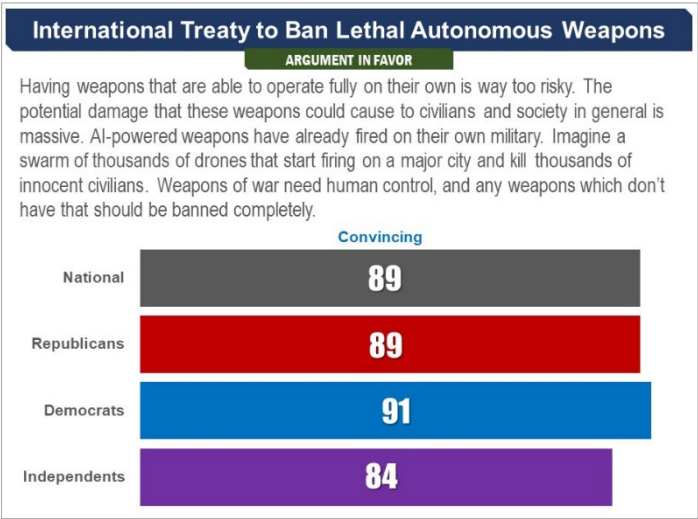
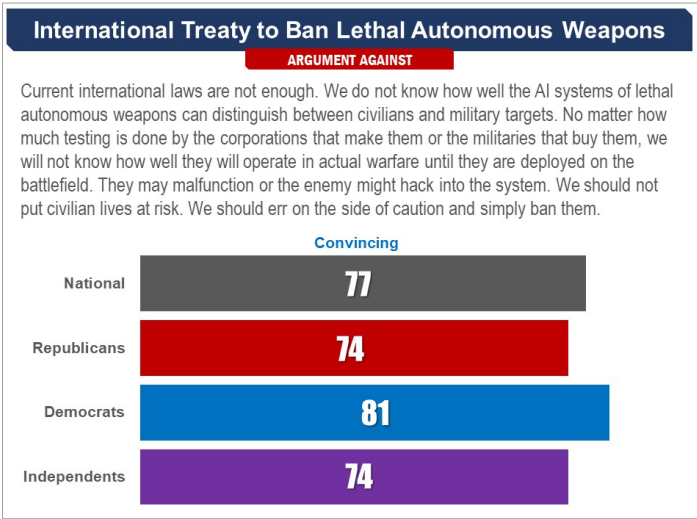
Non-member nations would be pressured to ban lethal autonomous weapons as well.

This proposal is modeled after other international treaties for monitoring and regulating potentially dangerous technologies, such as nuclear and biochemical weapons.

So, the question is whether the US should actively work with other nations to create an international treaty to ban lethal autonomous weapons.


All of the arguments were found convincing by bipartisan majorities, but the arguments in favor did substantially better, overall and among Republicans and Democrats.





Asked for their final recommendation, 81% were in favor, including 77% of Republicans and 85% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (77%) to very blue (83%).



International Agency to Monitor and Regulate Large-Scale AI Projects

A very large bipartisan majority of three-quarters favor the US actively working to create an international organization with intrusive powers to monitor and regulate large-scale AI projects that have the potential to cause international harm – intentionally or unintentionally.

Respondents were briefed on the concerns about future developments of large-scale AI programs that have the potential to cause international damage:

Now let's turn to a proposal for dealing with large-scale AI programs.

Among some AI experts, there is a concern that large-scale AI programs could be created that are highly intelligent, have advanced capabilities, and, perhaps most significantly, have a high level of autonomy. According to these experts, these AI programs could become uncontrollable by humans and engage in dangerous behavior that causes massive harm.

On the other hand, some AI experts have said that these fears of an AI program becoming so powerful and destructive independent of human control are neither realistic nor based on any evidence.



A recent survey of AI experts found that more than half believe there is at least a five percent chance that AI could be developed to the point that it could cause extremely bad outcomes, even possibly human extinction.

In addition to concerns about AI acting autonomously, there are also broad concerns that highly powerful AI programs could be hacked or misused to cause massive harm.

The proposal for regulating large-scale AI programs was then presented:

A proposal has been put forward for an international treaty for regulating large-scale AI programs. This treaty would have two parts:

1. Member nations (those that signed the treaty) would establish a set of regulations for the development and use of large-scale AI programs, with the goal of ensuring that they:
 - can always be shut down by human operators in case they get out of control
 - have robust security measures to protect them from being hacked or misused
 - do not cause major unintended and problematic consequences

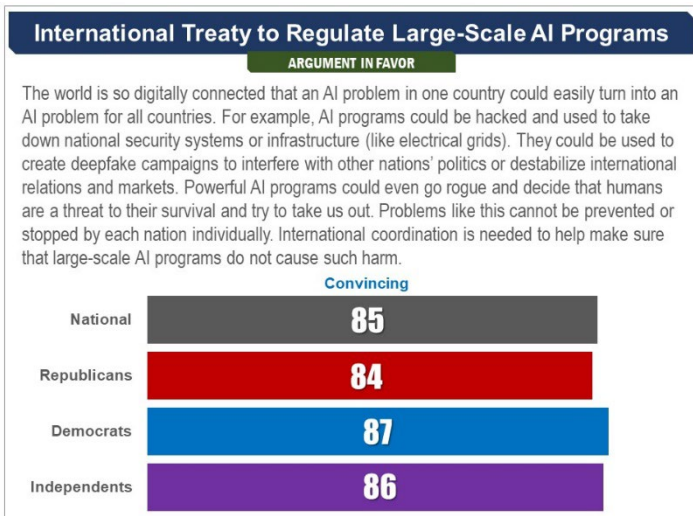
As AI technology advances and changes, member nations could establish new regulations.

2. An international agency would be created to monitor and inspect whether nations' large-scale AI projects are following the agreed-on regulations, and help fix any problems that arise. Member nations would be required to disclose information about their large-scale AI programs and agree to inspections and non-member nations would be pressured to do so as well.

This proposal is modeled after previous international treaties for monitoring and regulating potentially dangerous technologies, such as nuclear and biochemical weapons.

So, the question is whether the US should actively work with other nations to create such an international treaty to establish an agency to regulate large-scale AI programs.

The arguments in favor and against were both found convincing by bipartisan majorities, but the argument in favor did much better, overall and among Republicans and Democrats.



Asked for their final recommendation, 77% were in favor, including 71% of Republicans and 84% of Democrats.

Majorities in all types of Congressional districts were in favor, from very red (70%) to very blue (79%).



ARTIFICIAL INTELLIGENCE

A National Survey of Registered Voters

ACKNOWLEDGEMENTS

STEVEN KULL

Primary Investigator

EVAN SCOT FEHSENFELD | EVAN CHARLES LEWITUS

Research Analysis

JP THOMAS | KATRINA FULLER

Communications & Media Outreach

ALLISON STETTLER

Design & Production



PROGRAM FOR PUBLIC CONSULTATION

SCHOOL OF PUBLIC POLICY, UNIVERSITY OF MARYLAND

The Program for Public Consultation seeks to improve democratic governance by consulting the citizenry on key public policy issues governments face. It has developed innovative survey methods that simulate the process that policymakers go through — getting a briefing, hearing arguments, dealing with tradeoffs — before coming to their conclusion. It also uses surveys to help find common ground between conflicting parties. The Program for Public Consultation is part of the School of Public Policy at the University of Maryland.

PROGRAM FOR PUBLIC CONSULTATION

School of Public Policy, University of Maryland

www.publicconsultation.org

© PROGRAM FOR PUBLIC CONSULTATION 2024