



Pulse

US attitudes and awareness regarding effective giving and philanthropic cause areas

Surveys and Data Analysis
Department

November 2024 • Public Polling Report | USA

AUTHORS

Jamie Elsey — Senior Behavioral Scientist II, Rethink Priorities

MANAGER

David Moss — Principal Research Director, Rethink Priorities

Pulse: US attitudes and awareness regarding effective giving and philanthropic cause areas

Jamie Elsey — Senior Behavioral Scientist II, Rethink Priorities

David Moss — Principal Research Director, Rethink Priorities

Suggested citation: Elsey, J.W.B., & Moss, D. 2024. Pulse: US attitudes and awareness regarding effective giving and philanthropic cause areas. *Rethink Priorities*.

Surveys and Data Analysis at Rethink Priorities

The Surveys and Data Analysis Department at Rethink Priorities conducts rigorous and impact-focused behavioral science research, and provides analytic consultation to impact-driven organizations.

<https://rethinkpriorities.org/surveys>

For all queries, please contact surveys@rethinkpriorities.org

Funding disclosure: We thank Open Philanthropy for funding this research report (Open Philanthropy Project LLC). The views expressed herein are not necessarily endorsed by Open Philanthropy.

© 2024 Rethink Priorities

RETHINK PRIORITIES

530 Divisadero St. PMB #796
San Francisco, California 94117
United States of America



About Rethink Priorities

Rethink Priorities is a research and implementation group that identifies pressing opportunities to make the world better. We act upon these opportunities by developing and implementing strategies, projects, and solutions to key issues. We do this work in close partnership with foundations and impact-focused nonprofits.

www.rethinkpriorities.org

Contents

Editorial note	4
Executive summary	5
Introduction	7
Attitudes towards major cause areas	8
Cause area importance	8
Support for donations to cause areas	10
Awareness of Effective Altruism	14
Population level estimation	14
Demographic subgroup estimates	15
Attitudes towards Effective Altruism	18
Attitudes among those who have already heard of EA	18
Attitudes among those who have not heard of EA	19
Awareness of EA-related terms, organizations, and people	22
Terms	22
Organizations	23
People	24
Charitable giving behavior and attitudes	26
Donation behavior	26
Attitudes towards a ‘10% pledge’	28
Beliefs about own income relative to the rest of the world	30
Attitudes towards Artificial Intelligence	32
AI worry	32
Stances on costs and benefits of AI	33
General discussion	37
Takeaways for impact-focused organizations and philanthropists	37
Takeaways regarding AI risk	38
Concluding remarks	38
Methodological transparency	39



Editorial note

Rethink Priorities is an independent, non-partisan, non-profit 501(c)3 think tank that does polling and policy analysis. Rethink Priorities is not funded by any candidate or political party committee and does not poll on behalf of any political candidate or party.

Executive summary

The scale of charitable giving in the United States is massive, with **half a trillion dollars donated in 2023**. Donors hope that their donations will be used effectively to produce positive results. However, **little is known about public awareness of, and attitudes towards, charitable giving opportunities that focus explicitly on impact and cost-effectiveness**.

Rethink Priorities' *Pulse* project seeks to address this knowledge gap through large scale public polling of US adults to assess awareness and perceptions of effective giving opportunities, impactful cause areas, and the ideas and communities associated with such philanthropic efforts, including 'Effective Altruism' (EA).

This report presents results from the first wave of *Pulse*, fielded between July and September 2024, and is based on the responses of 4,890 US adults, with estimates representative of the US adult population.¹ The core areas covered, and key insights gained, are:

- **Attitudes towards ten major philanthropic cause areas**
 - Relatively traditional causes such as **Cancer Research and Mental Health** ranked **top** in terms of importance and support for donations, and were **rated more highly than several causes often highlighted by advocates of effective giving and EA (Animal Welfare, Pandemic Preparedness, AI Risk, and Existential Risk)**
 - **Global Health and Development** was rated more highly than most other core EA cause areas, but we estimate around 69% of US adults would be more supportive of donating to cancer research than to improve global health/poverty
 - Perceptions of cause areas varied substantially according to political party identification: **Democrats rated most causes more highly than Republicans**, with **high polarization in ratings of Climate Change**, **moderate polarization for causes such as Farmed Animal Welfare and Global Health and Development**, and **low polarization for AI Risk**
- **Awareness of Effective Altruism and a range of related public figures and organizations**
 - **Awareness of EA in the general public was low**, at around 1%
 - EA awareness was **highest among the more highly educated** and those **aged between 25-44**, which have been target groups of EA outreach
 - **Female respondents displayed less awareness than males**, and **Black/African American respondents showed lower awareness than White/Caucasian or Asian/Asian American respondents**
 - **Most people who were aware of EA held positive views of it**, though some associate it negatively with **scandals such as FTX**
 - Perceptions of the core ideas of EA among those who were *not* aware of it were almost all positive or neutral
 - Reported awareness of 'effective giving'-focused charitable organizations and major figures of the EA movement was also low, at less than 10%
- **Donation-related behavior and use of charity evaluators**
 - **54% of US adults reported having given to charity in the past year**
 - **Only a small minority of donors (2%) reported using a recognized charity evaluator**

¹ Estimates account for Age, Sex, Household income, Educational attainment, Racial identity, Political party identification, US Census region, and State.

- **Views of a charitable ‘giving pledge’ and one’s own income relative to the rest of the world**
 - The majority of US adults disagreed with the perspective that, owing to the relative income of people in the US relative to the rest of the world, that it was both possible and morally good for most people in the US to give 10% of their incomes to improve the lives of those in low income countries
 - US adults, across income levels, likely substantially **underestimate how their own income stacks up relative to the households across the rest of the world**
- **Concerns about artificial intelligence**
 - **Worry about the negative effects AI might have on their lives and society more broadly has increased among US adults since 2023, though 60% of people still reported quite low daily worry**
 - Amongst three perspectives on AI development, **US adults most agreed with the view that AI could pose a significant danger to humanity**, requiring research, oversight, regulation, or even a pause to prevent negative outcomes
 - However, **the view that ethical concerns (such as algorithmic biases and privacy harms) should be more of a focus than speculative harms also found support**
 - **US adults tended to disagree that we should be accelerating AI development as much as possible in order to reap its benefits, rather than worrying about risks**

Key Takeaways on Charitable Giving and Effective Altruism (EA):

- **Low Awareness but High Receptivity:** Awareness of EA and charity evaluators among the US public was low, but **many supported core ideas associated with effective giving**, such as evaluating charities for impact and promoting cost-effective charities and neglected cause areas
- **Opportunities for Growth:** While receptive individuals may not become deeply involved in ‘the EA movement’, the findings show the **potential for broader awareness, understanding of, and engagement with effective giving** (e.g., the use of charity evaluators)
- **Demographic Patterns and Tailored Outreach:**
 - Awareness of EA was lower among females and Black/African American respondents, but these demographic subgroups appeared no less receptive to several core EA-related ideas—we can **expand the coalition of people engaged with effective giving** by finding ways to better reach these individuals
 - Republicans showed lower awareness and positivity towards EA and associated causes and ideas compared to Democrats. More work needs to be done to understand why this is and whether it can be improved

Key Takeaways on AI Risk:

- **Increasing Concern and a Focus on Risks:** US public concern about AI risks has grown since 2023, with fewer individuals reporting no concern and more worrying “a fair amount” or “a lot”. The public endorsed perspectives on AI development that emphasized risks (whether existential or ethical), and disagreed with calls to accelerate AI
- **Opportunity for Bipartisan Support:** AI risk as a cause area was viewed as similarly important by Republicans and Democrats, suggesting that **safety-focused regulatory proposals could gain traction across party lines**

Introduction

The scale of charitable giving in the US is massive. In [2023](#), over half a trillion dollars was given to US charities, with \$374.4 billion of this (67% of the total) coming from individuals.² People give to charity for a multitude of reasons: from a feeling of compassion, a sense of moral duty, to generate positive feelings, to show solidarity, or to influence and improve a particular issue. Regardless of motivations, most donors presumably hope that their money will be used wisely and have a positive impact.

Over the past two decades, this focus on impact has risen to prominence among a collection of philanthropists, philosophers, research organizations, and charities, in a movement known as [Effective Altruism](#) (EA)³. The core tenets of EA are to use evidence-based methods and reasoning to find the most cost-effective ways of doing good. In practice, members of the Effective Altruist community often advocate for donating to cost-effective charities working in relatively neglected areas such as global poverty or farmed animal welfare, or for people to take up careers that tackle pressing societal issues or potential threats to the long term future, such as pandemics and advanced artificial intelligence (AI).

Because effective charities may be as much as [100x more impactful](#) than ineffective ones⁴, and given the extent of US charitable giving, a tremendous positive impact could be achieved if people in the US understood and engaged with a more impact-oriented approach to their charitable behavior. Yet, relatively little is known in terms of the US public's awareness of effective giving opportunities, their perception of various potentially impactful cause areas, or their attitudes towards ideas associated with such philanthropic efforts.

Rethink Priorities' *Pulse* project seeks to address this gap through highly-powered polling of the US public. This report presents results from the first wave of *Pulse*, fielded between July and September 2024, and is based on the responses of 4,890 US adults. The report provides insights into the US public's:

- **attitudes towards ten major philanthropic cause areas**, including global health and development, farmed animal welfare, and existential risk
- **awareness of Effective Altruism** and a range of related public figures and organizations
- **donation-related behavior** and use of charity evaluators
- **views of a charitable 'giving pledge'**
- **estimates of their own income relative to that of the wider world**
- **concerns about artificial intelligence**

We used quota-based sampling and multilevel regression and poststratification to generate estimates for the US public, accounting for Age, Sex, Household income, Educational attainment, Racial identity, Political party identification, US Census region, and State. Detailed methodology behind these results can be found in the [Methodological transparency](#) section.

Rethink Priorities will also publish more in-depth investigations into these areas in future reports. Readers can subscribe to our [newsletter](#)⁵ to be kept informed of such releases, as well as analyses of subsequent waves of *Pulse*.

² Lilly School of Philanthropy (2024). *Giving USA 2024*. <https://perma.cc/5K46-LUCZ>

³ Center for Effective Altruism. What is effective altruism? <https://perma.cc/Z6VH-QT8M>

⁴ Caviola, L., Schubert, S., Teperman, E., Moss, D., Greenberg, S., & Faber, N. S. (2020). Donors vastly underestimate differences in charities' effectiveness. *Judgment and Decision Making*, 15(4), 509-516.

⁵ <https://rethinkpriorities.org/newsletter/>

Attitudes towards major cause areas

Opportunities for effective giving are available across many cause areas. It is crucial to understand the extent to which the general public engages with various cause areas that are promoted by impact-focused organizations. To do this, we split respondents such that 60% were asked the extent to which they supported or opposed people giving charitable donations to a range of cause areas, and 40% were asked how important they thought each cause area was. This split was done because the perceived importance of an issue does not necessarily correspond exactly to the extent to which people support it as a charitable cause area.

We provided respondents with several relatively more traditional cause areas (Cancer Research, Civil Rights, and Climate Change), as well as a range of cause areas that have seen a greater relative focus among impact-oriented researchers and philanthropists, particularly in the EA community.⁶

Each cause came with a brief description or example.⁷ While this explanatory information aimed to be simple, informative, and uncontroversial, it should be noted that different framings of cause areas could lead to substantially different prioritization ratings.

Cause area importance

Respondents who were randomized to see the importance items were told:

“Governments, charities, philanthropists, companies, and individuals all make choices about the most important issues of our time.

We would like to know what you think are the most important issues to be tackled in the world.

On the following pages, we will show you several different issues or 'cause areas'. Please rate how important you think each one is.”

Average ratings are shown in Figure 1. The highest two rated causes in terms of importance were both individual ailments: cancer and mental health issues. Readers may be surprised to see Climate Change so far down the list, but this owes to political polarization surrounding the issue: respondents who were Democrats gave Climate Change an average rating of 8.7, whereas Republicans gave it an average of 5.8. In contrast, Republicans, Democrats, and Independents rated cancer research 8.9-9.0.


Global Health and Development—a cause area that has received substantial focus from advocates of effective giving—ranked among the middle of the cause areas. Several other cause areas that have become a focus of organizations related to effective giving and the EA community—namely Pandemic preparedness, Existential Risk, AI risk, and Farmed Animal Welfare—fell in the bottom half of importance ratings. Nevertheless, essentially all causes were rated much closer to being ‘extremely important’ than ‘not at all important’.

⁶ More traditional causes such as cancer research are not argued to be unimportant by effective giving advocates or the EA community. Rather, cases such as Global Health and Animal Welfare are seen to be relatively more neglected, or to have specific opportunities for highly cost-effective donations.

⁷ Exact question wording can be found at: <https://tinyurl.com/rp-pulse-2024-survey>

Figure 1: Most important issues to be tackled in the world, according to US adults

Most important issues to be tackled in the world, according to US adults



Estimates based upon 1925 US adults

Besides looking at means and ordinal rankings of causes, it is possible to compare ratings across cause areas more directly, using an estimate of within subjects *Probability of Superiority* (Figure 2).⁸ These values estimate the approximate likelihood that a US adult would rate one cause as higher than another, with values of .5 indicating no difference (i.e., 50% chance cause A is greater, 50% chance cause B is greater). Values below .5 indicate cause A being rated lower than cause B, and values over .5 indicate cause A being rated higher than cause B. The possible range of values is from 0 to 1 (0 = no one is expected to rate cause A > B; 1 = everyone is expected to rate cause A > B). Based on these estimates, we would expect around 75% of US adults to rate Cancer Research and Mental Health as more important than Farmed Animal Welfare.

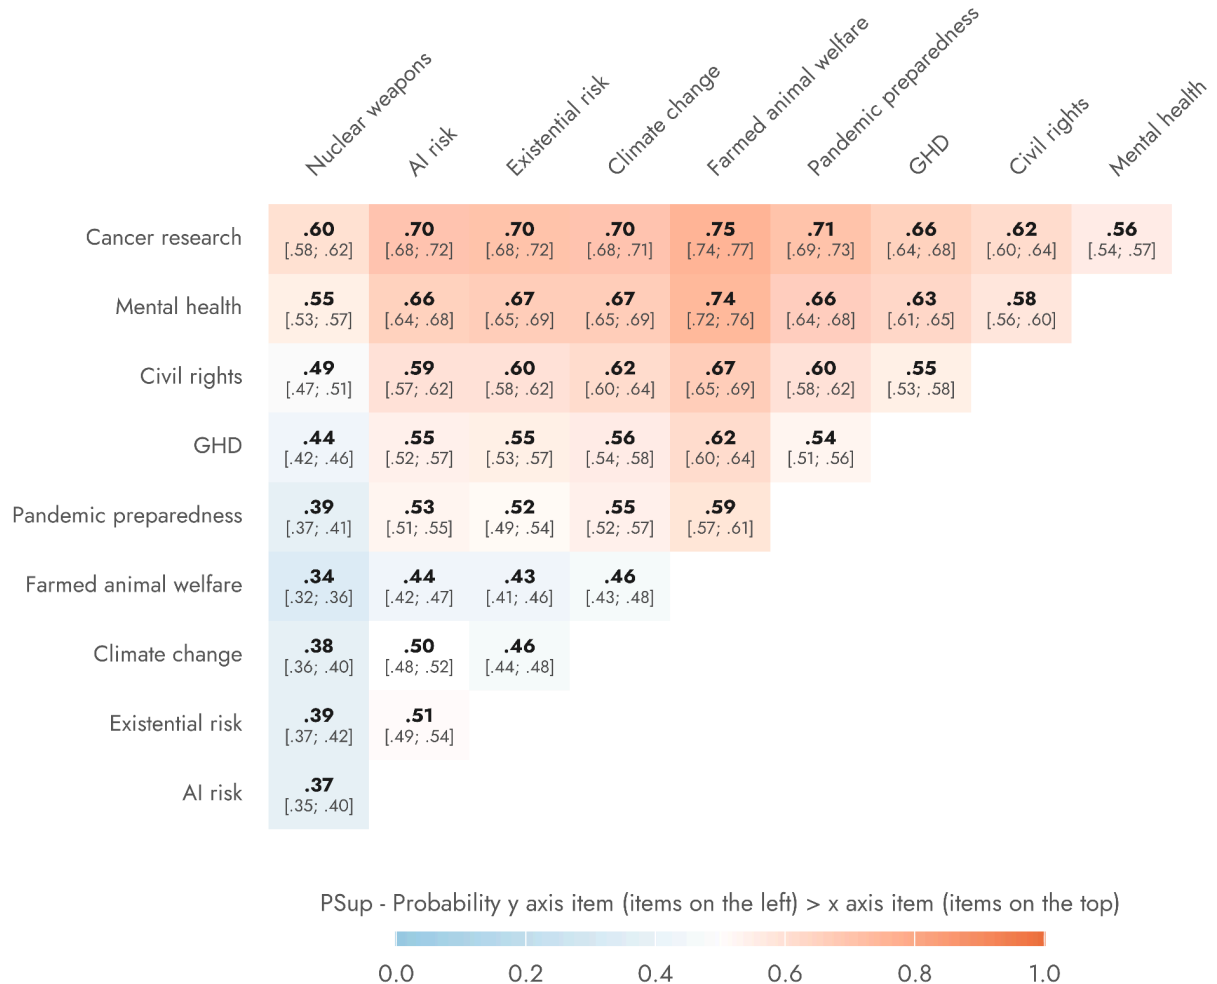
⁸ *Probability of Superiority* estimates for Importance and Donation support items in this section use weighting as opposed to multilevel regression and poststratification. See the *Methodological Transparency* section for details.

Figure 2: Comparisons between importance ratings for different cause areas

Comparisons between importance ratings for different cause areas

Using within subjects *Probability of Superiority (PSup)*

Example interpretation: *PSup* in top-left cell is .60 for *Cancer research* > *Nuclear weapons*.
We expect ~60% of people to rate *Cancer research* more important than tackling *Nuclear weapons*.



Estimates based upon 1941 US adults

Support for donations to cause areas

Respondents who were randomized to see the donation support/opposition items were told:


“On the following pages you will see several ‘cause areas’ or issues.

We would like to know to what extent you support or oppose charitable donors giving money to each of them.”

Average ratings are shown in Figure 3. Similarly to the rankings based on importance, Cancer Research and Mental Health came out on top. However, Farmed Animal Welfare fared relatively better in terms of its placement among the ratings, and Nuclear Weapons fared much worse.

Figure 3: Support for donating to different causes among US adults


Support for donating to different causes among US adults



Estimates based upon 2927 US adults

Figure 4: Comparing importance and support ratings

Comparing importance and support ratings



When comparing ‘Importance’ and ‘Support’ ratings (Figure 4), Nuclear Weapons was something of an outlier, receiving relatively low donation support in comparison to the perceived importance of dealing with this issue. People may, understandably, feel that there is little opportunity or likelihood of success in donations geared towards preventing nuclear

proliferation or war, even while acknowledging the importance of the issue.⁹ For other items, there was a strong coupling between importance and support, in terms of average population ratings.

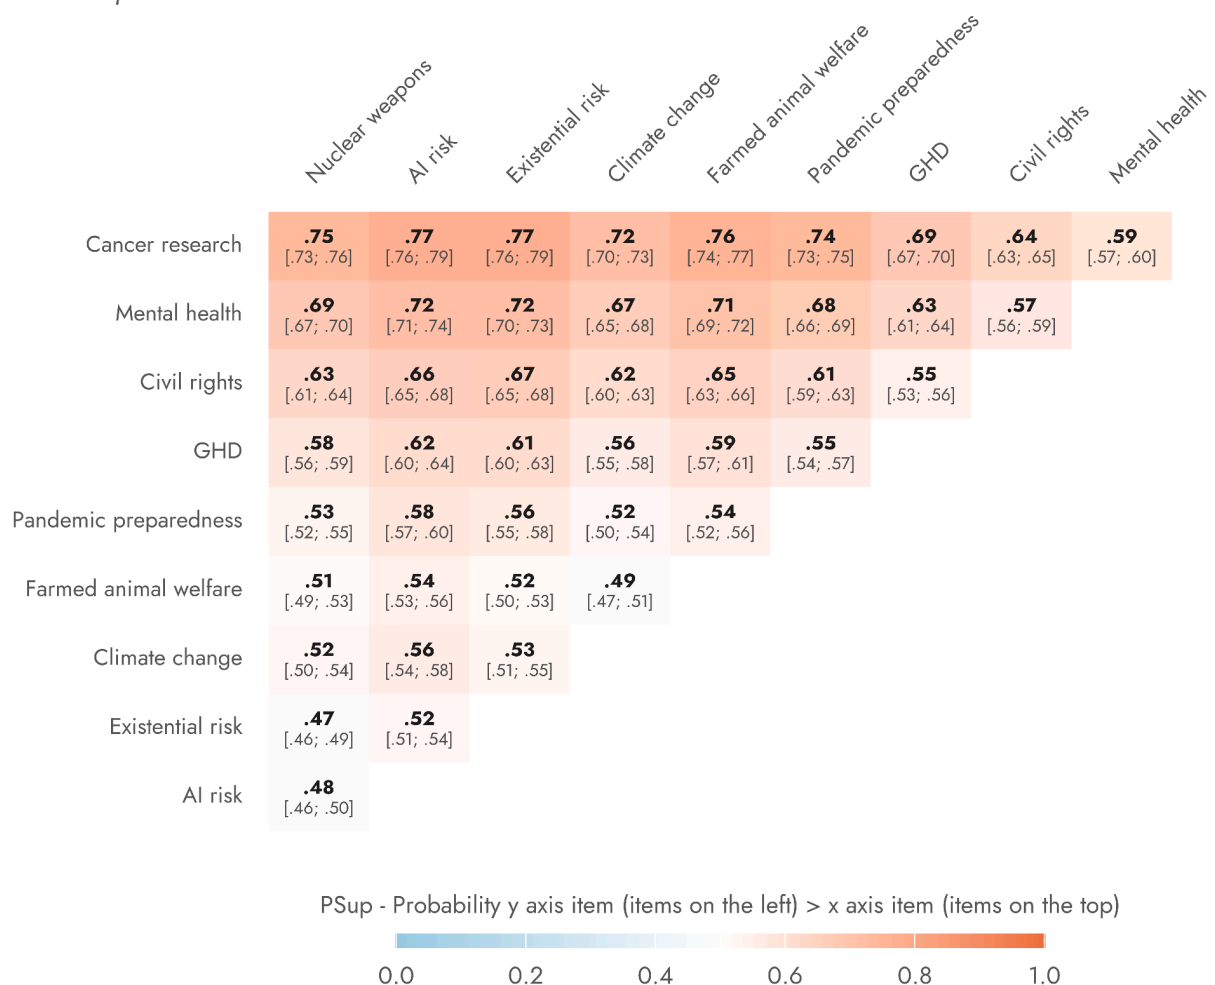
As with importance ratings, it is possible to look at *Probability of Superiority* among the support ratings (Figure 5). Whereas Farmed Animal Welfare was rated lower than every other item in terms of importance, in terms of support for donations, it was rated slightly higher than AI risk, and similarly to Climate Change or Nuclear Weapons.

Figure 5: Comparisons between support/opposition ratings for different cause area donations

Comparisons between support/opposition ratings for different cause area donations

Using within subjects *Probability of Superiority* (PSup)

Example interpretation: PSup in top-left cell is .75 for *Cancer research* > *Nuclear weapons*. We expect ~75% of people to be more supportive of donations to *Cancer research* than to tackling *Nuclear weapons*.



Estimates based upon 2949 US adults


For most causes, with both Importance and Support ratings, respondents who identified as Democrats tended to provide higher ratings than those who identified as Republicans, with Independents intermediate between them (Figure 6). The most sizable difference was for Climate Change, but other causes also showed quite substantial political variation, such as Pandemic Preparedness, Civil Rights, Farmed Animal Welfare, and Global Health and

⁹ Philanthropic organizations are in fact investigating ways to most effectively mitigate nuclear risks, e.g., *Founder's Pledge* (<https://perma.cc/EE64-53GD>) and *Longview Philanthropy* (<https://perma.cc/67D8-DHC5>)


Development. AI Risk appeared to be more uniformly rated than many causes, looking more similar to Cancer Research than to more polarized issues such as Climate Change. Republicans have not been found to donate less or be less supportive of charity than Democrats, but do tend to favor donations to religious causes and especially local religious congregations.¹⁰ Advocates of effective giving may benefit from considering how promotion of high impact causes can better align with values across the political spectrum.

Figure 6: Ratings of cause area importance and support for donations across political party affiliations

Ratings of cause area importance and support for donations across political party affiliations



Support for donations



¹⁰ Margolis, M. F., & Sances, M. W. (2017). Partisan differences in nonpartisan activity: The case of charitable giving. *Political Behavior*, 39, 839-864.

Awareness of Effective Altruism

Population level estimation


Our estimates for the number of people who have ‘heard of EA’ have multiple thresholds. The ‘Claim’ level is simply the percentage of people who would be expected to pick the term as something they have heard of from the list of terms. However, we know this to be a substantial overestimate, likely due to people wrongly assuming they have heard of the term, or claiming they have on the basis of knowing the constituent words (see below, for example, where 8% of people claim to have heard of the invented term ‘Globally Neutral Advocacy’).

Upon selecting Effective Altruism from the list of terms, people were asked to indicate whether they thought they had at least a basic understanding of what it means. If so, they were required to provide a brief explanation of their understanding of the term. Permissive and Stringent levels of awareness are based on these explanations. The Stringent level was reached when the respondent unambiguously referred to one or more of the core ideas of EA, such as maximizing the amount of good done, and doing so through charity or careers, or referring to figures associated with the EA movement (including reference to scandals or controversial figures). These responses were also assessed for copy-pasted/paraphrased definitions, which were not counted. The Permissive level was reached if the respondent was alluding to some core ideas, but with lower specificity than the stringent level. Note that the estimate for the percentage of people reaching the Permissive level includes both those who reached the Permissive and who reached the Stringent levels, whereas, Stringent—being a higher level of attainment—only includes those who reached the Stringent level.

Although we estimate that 12% of US adults might claim to be aware of Effective Altruism, our permissive estimate for awareness of Effective Altruism is 1-2% of US adults, and less than 1% for stringent assessment (Figure 7).

Figure 7: Estimated percentage of US adults who have heard of EA

Estimated percentage of US adults who have heard of EA




Estimates based on 4852 US adults

Demographic subgroup estimates

Estimates of awareness in demographic subgroups provide additional information about how the concept of Effective Altruism permeates across different social and demographic groups. Outreach in the Effective Altruism community has typically been targeted towards younger and more highly educated individuals. This is borne out in patterns of awareness across subgroups, with awareness of Effective Altruism being highest among those who have completed the highest levels of education (Figure 8), as well as peaking in the 25-34 or 35-44 year old age groups (Figure 9).

Figure 8: EA awareness depending on education


EA awareness depending on education



Estimates based on 4852 US adults

Figure 9: EA awareness depending on age

EA awareness depending on age




Estimates based on 4852 US adults

Several other demographic trends are apparent. Awareness of Effective Altruism was higher among respondents with higher household incomes (Figure 10), as well as among male respondents relative to female respondents (Figure 11).

Figure 10: EA awareness depending on household income


EA awareness depending on household income



Estimates based on 4852 US adults

Figure 11: EA awareness depending on sex

EA awareness depending on sex




Estimates based on 4852 US adults

Awareness of Effective Altruism also tended to be lower among Black or African American respondents relative to White or Caucasian and Asian or Asian American respondents, although there was considerable uncertainty around these estimates (Figure 12).

Figure 12: EA awareness depending on racial identity


EA awareness depending on racial identity



Regarding Political Party affiliation, we observed a slight tendency for Republican respondents to be less likely to have heard of Effective Altruism than Democrat and Independent respondents (Figure 13).

Figure 13: EA awareness depending on political party identification

EA awareness depending on political party identification



While these demographic breakdowns convey some sense of the variation across different social and demographic groups, it should be highlighted that awareness among very niche populations may be substantially higher than the estimates provided here. For example, we

might expect young people working in Silicon Valley to be substantially more likely to have heard of Effective Altruism (due to geographic and intellectual overlap with major Effective Altruism hubs and ideas, or via awareness of Sam Bankman-Fried). Likewise, people studying at elite educational institutions that have their own EA groups would be expected to have higher awareness than the generally highly educated.

Attitudes towards Effective Altruism

Attitudes among those who have already heard of EA

In addition to awareness, we assessed attitudes towards Effective Altruism. Among those aware at the Permissive and Stringent levels (Figures 14 and 15, respectively), there was a tendency towards positive and very positive attitudes, with both these levels seeing over 50% of the respondents at least slightly positive towards Effective Altruism. However, a non-negligible minority of respondents feel negative towards Effective Altruism. Part of this may result from the association of Effective Altruism with recent scandals, such as the collapse of the FTX cryptocurrency exchange, which may be the only lens through which some people have encountered Effective Altruism. Nine respondents were flagged as alluding to FTX in their explanation of Effective Altruism, and all of these reported a negative sentiment. These respondents made up over half of the relatively small absolute number of people who expressed negative sentiment towards Effective Altruism, suggesting that FTX could be an important factor contributing to negative perceptions of the movement. Negative perceptions resulting from FTX may be especially important in understanding the relatively greater proportion of negative respondents among the stringently aware, as people who referenced FTX in the explanation of Effective Altruism were coded as meeting the stringent criteria.

Figure 14: Sentiment towards EA among those with a 'permissive' level of awareness

Sentiment towards EA

EA awareness: Permissive

Numeric estimates: 5.18 [4.75; 5.60], SD = 1.82, N = 73





Figure 15: Sentiment towards EA among those with a 'stringent' level of awareness

Sentiment towards EA

EA awareness: Stringent

Numeric estimates: 4.51 [3.84; 5.19], SD = 2.09, N = 39



Attitudes among those who have not heard of EA

At the end of the survey, those respondents who had not heard of Effective Altruism were presented with a definition of the term that focused on its core ideas, and asked to indicate how positively or negatively they felt about it:

“Effective altruism is a philosophy and a community. Effective altruists use reasoning and evidence to find the most effective ways of doing good, and then act upon what they find to try to improve the world. For example, assessing how much impact different charities have for every dollar donated, and then promoting the most cost-effective charities.

It mostly advocates for supporting charities working on global poverty, factory farming, or risks to the long term future, such as artificial intelligence.”


This allows for a population-level assessment of potential attitudes towards Effective Altruism, factoring out those who have already heard of the term. Again, attitudes were broadly positive (Figure 16). Relative to those who have encountered Effective Altruism, there was a greater preponderance of Neutral responses, and fewer Negative responses.

Figure 16: Attitudes towards the idea of Effective Altruism

Attitudes towards the idea of Effective Altruism

Ratings are estimated based upon a model including 4852 US adults, but later processing steps extract estimates specifically for US adults who were not aware of what EA is, and who read a description of the core ideas.

Numeric estimate: 5.0 [4.99; 5.08]




It is of particular interest to understand, among those who do not know of EA, whether any particular demographic subgroups may be more or less receptive to the general principles. This can help advocates understand whether there is untapped potential among particular types of people. Figure 17 displays how different demographic subgroups break out in terms of Positive, Neutral, and Negative attitudes towards the formulation of EA presented above.

The results suggest a general tendency for younger and more highly educated US adults to have perceived the ideas of EA more positively, which aligns with seemingly successful outreach to these target subpopulations and their high relative awareness. There was also a difference in sentiment in terms of political views, with Democrats being more likely to provide a positive rating (73%) than either Independents (61%) or Republicans (59%).

However, compared with the lower awareness among Female (vs. Male) and Black or African American (vs. White or Caucasian and Asian or Asian American) respondents, there was remarkably little variation in terms of how positively the core concepts of EA were viewed.

Figure 17: Sentiments towards the ideas of Effective Altruism among demographic subgroups

Sentiment towards the ideas of Effective Altruism among demographic subgroups



Awareness of EA-related terms, organizations, and people


Terms

Besides Effective Altruism, we asked respondents about several other terms (Figure 18). Longtermism and Effective giving were chosen as being relevant to Effective Altruism. Other items function as reference points (e.g., relatively science-related or intellectual terms such as Evidence-based medicine) or as lures with which to gauge the likelihood of true awareness (e.g., Globally neutral advocacy is a made up term that, similarly to Effective Altruism, people might be tempted to claim they know owing to recognizing the constituent words).

We can see that estimated awareness (or rather, likelihood of claiming awareness) of Longtermism is actually lower than for a lure (Globally Neutral Advocacy). We expect that, like with Effective Altruism, true awareness would be much lower than the implied 5%. Effective giving reaches claimed awareness of 21%. Again, similarly to Effective Altruism, we expect that many respondents would suggest a definition based upon ‘giving that is effective’, with a much smaller minority being aware of the core concepts of formally evaluating giving opportunities and prioritizing amongst them to maximize impact.

Figure 18: Estimated awareness of different terms among US adults

Estimated awareness of different terms among US adults



Estimates based on 4852 US adults


Organizations

Estimated awareness of different organizations is presented in Figure 19. As with awareness of EA-related terms, it is important to stress that awareness of organizations—especially niche ones—is likely lower than the percentages presented below. For example, even though we estimate only around 1% of US adults would claim awareness of 80,000 Hours, fewer still passed a multiple choice question to select the option that most accurately characterized the organization.

For organizations and public figures, we’ve categorized some as EA-related. This does not imply any official affiliation between them and the Effective Altruism community, nor endorsement from organizations/figures of Effective Altruism, and vice versa. The term EA-related is used purely descriptively to indicate a general association (e.g., GiveWell is counted as EA-related because it uses reason and evidence to find and promote the most cost-effective charities, whereas OpenAI—along with current or former members of its executive suite—is counted as EA-related due to Effective Altruism’s links with AI safety, and current or former board members reportedly motivated by EA principles).

Figure 19: Estimated awareness of different organizations among US adults

Estimated awareness of different organizations among US adults



Estimates based on 4852 US adults

Awareness of OpenAI was the highest of organizations associated with Effective Altruism. More people claimed awareness of FTX (the defunct cryptocurrency exchange whose CEO, Sam Bankman-Fried, publicly promoted Effective Altruism) than claimed awareness of Effective Altruism itself, or of any definitely EA-aligned organizations such as GiveWell or Open Philanthropy.

Relative to major philanthropic organizations or charities such as the Gates Foundation or Doctors Without Borders, EA-aligned philanthropic and charitable organizations appeared to be little recognized.


People

Estimated awareness of different organizations is presented in Figure 19. Among figures with high involvement in the Effective Altruism community, only Peter Singer reached awareness levels even close to more well known public intellectuals such as Sam Harris or Jordan Peterson (though Sam Harris notably has promoted Effective Altruism, and encouraged reason-driven charitable donations). In contrast, approximately 1 in 5 US adults were estimated to claim awareness of Sam Bankman-Fried. Figures such as Will MacAskill and Nick Bostrom, though widely recognized within effective altruist spheres, had claimed recognition at levels similar to a fake 'lure' (Jim Spacer) of only around 2%.

When considering the exact percentages, it should again be borne in mind that they represent *claimed* awareness, which could be cases of mistaken recognition. Despite this caveat, these awareness estimates suggest that Effective Altruism and its related organizations and figures remain quite niche in the general public consciousness of the US.

Figure 20: Estimated awareness of different people among US adults

Estimated awareness of different people among US adults



Estimates based on 4852 US adults


Charitable giving behavior and attitudes

Donation behavior

Respondents were asked to indicate whether they had donated to any charities in the past 12 months (Figure 21). If yes, they were asked whether they had used a charity evaluator to aid with their decision. Based on such self-report, we estimate approximately 54% of US adults have donated to charity over the past 12 months (this closely matches estimates from [YouGov](#), which estimated US giving at 54% in 2022¹¹). However, few use a charity evaluator.

Figure 21: Donation behavior among US adults

Over half of US adults report giving to charity in the last year, but a small minority report using a charity evaluator




Estimates based on 4852 US adults

Among those who reported using a charity evaluator, we also looked at the name of the charity evaluator they reported to determine whether it was in fact a formal charity evaluator (i.e., an organization that assesses various charities to determine their likely impact, trustworthiness, or cost-effectiveness; see Figure 22). We found that even fewer people provided the name of a confirmed charity evaluator (in this case, the response was overwhelmingly Charity Navigator). Hence, only a small proportion of people who donate appear to do so using any kind of formal evaluation. An even smaller minority of these—only two respondents—referred to a charity evaluator that focuses primarily on effective giving (in these cases, GiveWell).

Figure 22: Use of a recognized charity evaluator among US adults

Over half of US adults report giving to charity in the last year, but a very small minority indicated using a recognized charity evaluator



Estimates based on 4852 US adults


¹¹ Dumitru, O. (2023). Half of Americans say they have donated money to charity in the past year. *YouGov*. <https://perma.cc/S37D-8EJM>

With increasing educational attainment, the proportion of people who reported donating increased, as did the proportion of people who donated using a confirmed evaluator (Figure 23). Importantly, this trend of more highly educated respondents being more likely to use a confirmed evaluator was not solely due to them being more likely in general to make donations. Among donors, more highly educated respondents were proportionally more likely to report having used a charity evaluator.

Figure 23: Charitable giving behavior depending on education

Charitable giving behavior depending on education

The likelihood of being a charitable donor, and of being a donor who used a recognized charity evaluator, was higher among those with higher educational attainment. Among donors, the likelihood of using a recognized charity evaluator also reliably increased with increasing education.



Estimates based on 4852 US adults

Attitudes towards a ‘10% pledge’

Some organizations in the effective giving space—such as [Giving What We Can](#) (GWWC)—encourage wealthy people (relative to the world population) to pledge 10% of their incomes to effective causes.¹² The giving pledge of GWWC highlights that people earning \$47,000 or more per year after tax are in the top 2% of incomes globally, with \$62,000 being in the world top 1%. This value is slightly below the median US household income after tax of \$64,240 as of 2022 (according to the [US Census Bureau](#)).¹³

We were requested to poll the US public to assess the extent to which US adults would agree with the view that it would be both possible and good for most people in the US to take such a pledge, given their relative wealth. We stress that the statement assessed and reasoning provided differs in both tone and exact content to pledges such as the GWWC pledge. We presented this question to half of the respondents who took part in the survey (with the other half answering the relative income question in the following section). The wording of the perspective was more forcibly phrased than the rationales that accompany real giving pledges:

“It has been noted that, by global standards, even average incomes in the United States reflect a level of wealth and standard of living that is higher than most other people in the world.”

As a result, some people and organizations argue that it is both possible and morally good for most people in the United States to donate 10% of their incomes to charities aimed at improving the lives of people in low income countries.

To what extent do you agree or disagree with this perspective?”


We estimate that more US adults would tend to disagree than to agree with such a view (Figure 24), and in particular estimate more than three times as many US adults would *Strongly disagree* than *Strongly agree*. Moreover, this tendency towards disagreement was evident across income levels (Figure 25), and in every specific subgroup except Democrats, Black or African American respondents, and 18-24 year olds, who showed approximately even agreement and disagreement.

Figure 24: Attitudes towards a ‘10%’ pledge

Attitudes towards a ‘10% pledge’

It has been noted that, by global standards, even average incomes in the United States reflect a level of wealth and standard of living that is higher than most other people in the world.

As a result, some people and organizations argue that it is both possible and morally good for most people in the United States to donate 10% of their incomes to charities aimed at improving the lives of people in low income countries. To what extent do you agree or disagree with this perspective?

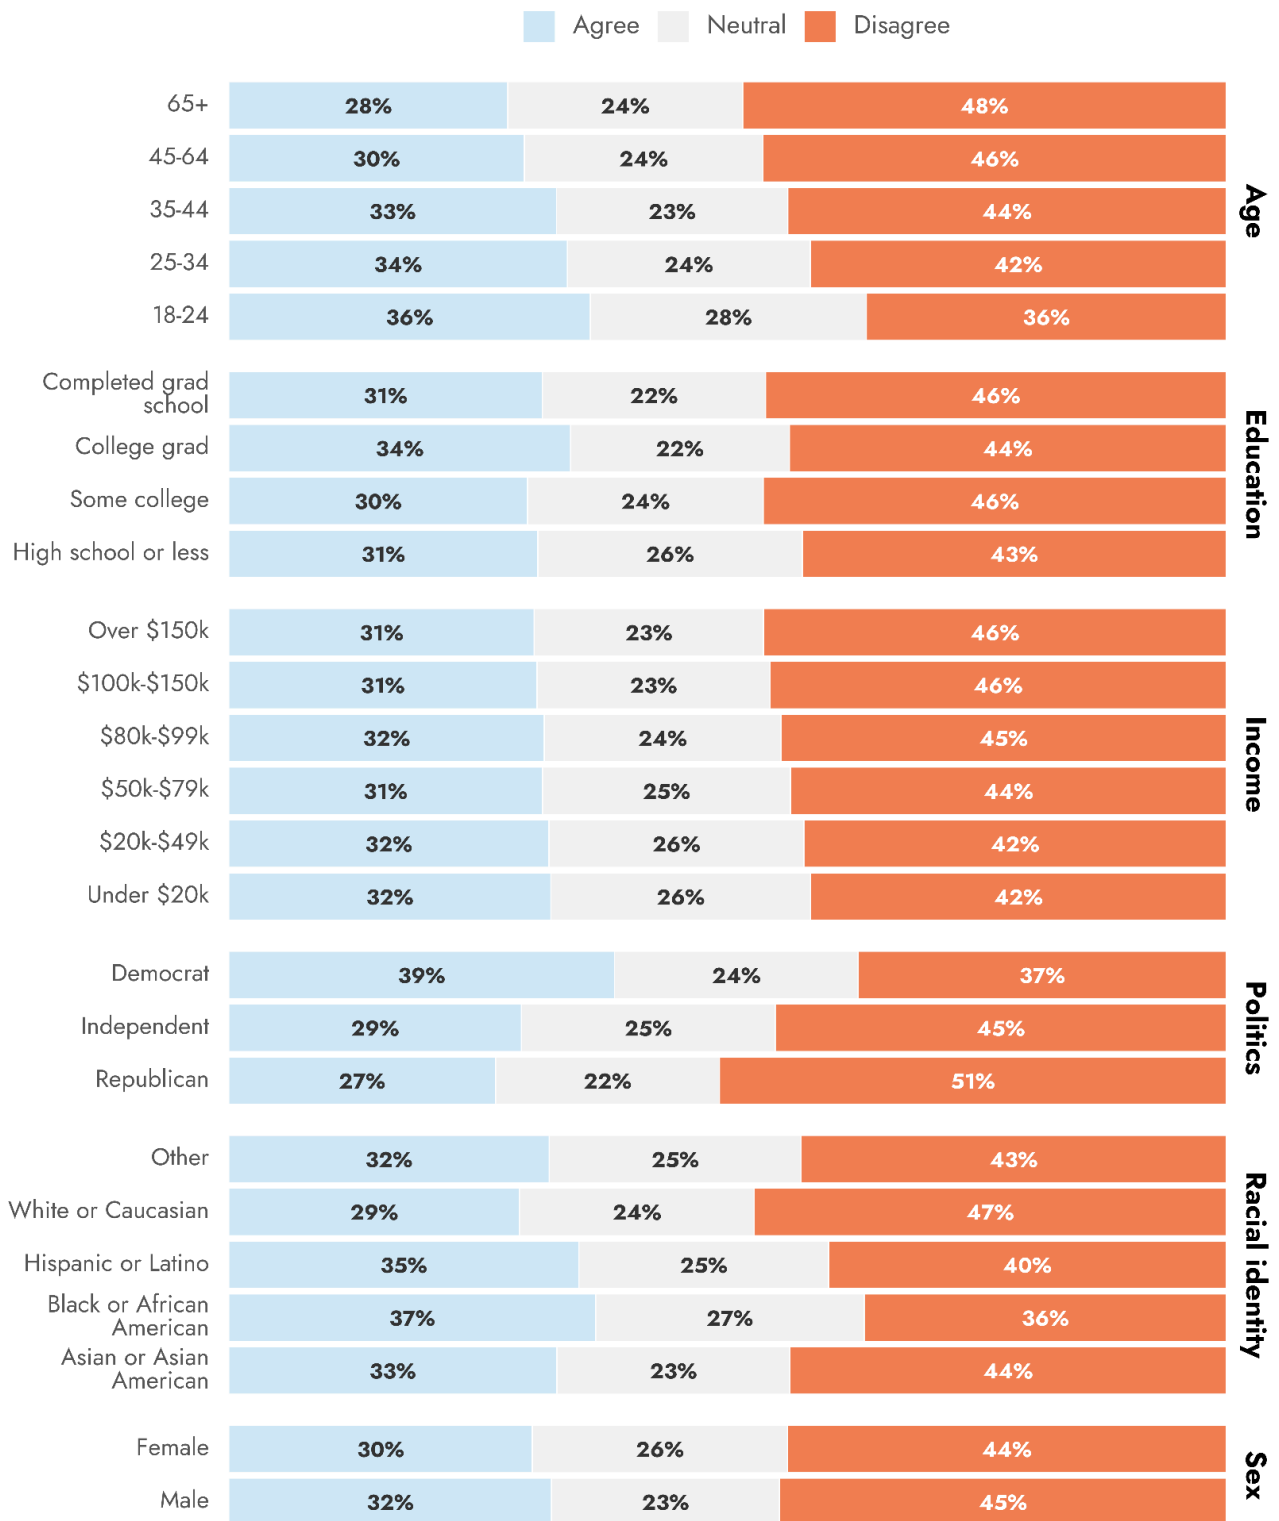


¹² Giving What We Can. *The 10% pledge*. <https://perma.cc/JK62-TALA>

¹³ Creamer, J. & Unrath, M. (2023). *End of Pandemic-Era Expanded Federal Tax Programs Results in Lower Income, Higher Poverty*. <https://perma.cc/3MTN-Z7UA>

Figure 25: Attitudes towards a '10% pledge' across subgroups

Attitudes towards a '10% pledge' across subgroups



Beliefs about own income relative to the rest of the world

As noted in discussion of the 10% pledge in the previous section, part of the reasoning behind people taking such a pledge is the relative wealth of individuals in the US in comparison to lower income countries. However, we expect that US adults might underestimate their relative wealth. We therefore asked half of the survey respondents to provide an estimate of what percentage of world households they thought their household income was greater than (the other half answered the 10% pledge question above).

Although we cannot precisely line up respondents' incomes with other global households (as we do not know their exact income, their household/family structure, nor crucially their *post-tax* incomes), we can get a sense of how people believe their income stacks up relative to the rest of the world—and accordingly, possible misconceptions—by breaking responses down by household income (pre-tax).


US adults likely underestimate their relative wealth. We can see this both when using the estimated means of each income subgroup or their estimated modes (Figures 26 and 27, respectively). The mean reflects the 'average' estimate from each income level, which may be affected by people who provided very high or low numbers relative to most people in that income bracket. The mode reflects the single most common expected response someone might give from that income bracket.

As an example, we estimate that those with a pre-tax household income of \$50,000-\$79,999 on average believe that this is greater than 43% of world households, with the most common response being 30%. In contrast, a relative income calculator from [GWWC](#)¹⁴ suggests that a post-tax income of \$50,000 for a household of 2 adults and 2 children would be among the 83% richest households in terms of purchasing power, or 79% for \$40,000.

Figure 26: Percentage of world households that US adults believe their household income is greater than, broken down by income - using the estimated population mean

Percentage of world households that US adults believe their household income is greater than, broken down by income

Using the estimated mean of the distribution




Estimates based on 2422 US adults

¹⁴ Giving What We Can. *How rich am I?* <https://perma.cc/2TBF-PCML>

Figure 27: Percentage of world households that US adults believe their household income is greater than, broken down by income - using the estimated population mode

Percentage of world households that US adults believe their household income is greater than, broken down by income

Using the estimated mode of the distribution



Estimates based on 2422 US adults

Attitudes towards Artificial Intelligence

In addition to awareness of and attitudes towards effective philanthropic cause areas and altruistic behavior, we asked additional questions related to a recently ascendant cause area: AI Risk.


AI worry

Respondents were asked to indicate the extent to which they worried about the impact AI might have on their lives and society more broadly (Figure 28). As we have asked this same question previously—in [April](#)¹⁵ and in [June](#)¹⁶ of 2023—we were able to assess the extent to which concerns about AI have evolved over time. Relative to 2024, it is estimated that fewer members of the US public do not worry at all, and increasing percentages are expected to worry a lot or nearly all the time.

Figure 28: Worry about AI among US adults is increasing over time

Worry about AI among US adults has been increasing over time

In your daily life, how much do you worry about the negative effects AI could have on your life or on society more broadly?



Estimates based on 4852 US adults in 2024, 2407 in June 2023, and 2444 in April 2023

¹⁵ Elsey, J.W.B., & Moss, D. (2023). US public opinion of AI policy and risk. *Rethink Priorities*. <https://rethinkpriorities.org/research-area/us-public-opinion-of-ai-policy-and-risk/>

¹⁶ Elsey, J.W.B., & Moss, D. (2023). US public perception of CAIS statement and the risk of extinction. *Rethink Priorities*. <https://rethinkpriorities.org/research-area/us-public-perception-of-cais-statement-and-the-risk-of-extinction/>


This impression from the estimated percentages of people picking each option is supported by the probability of superiority over time, comparing our current (July-September, 2024) responses with those of April and June 2023 (Figure 29). Estimates of 56% and 58% relative to June and April respectively suggest a modest but reliable increase in reported worry over time. The increasing awareness and use of AI tools such as ChatGPT, as well as media coverage regarding possible downsides and risks, may have made concerns over the negative effects of AI more salient.

Figure 29: Estimated difference in worry about AI from 2023 to 2024

Estimated difference in worry about AI from 2023 to 2024

Probability of Superiority provides an estimate of the chance that, if we drew someone from the population from July-September of 2024, their reported worry about AI would be higher than someone drawn from the population in April or June of 2023. A value of 50 would indicate no difference between the timepoints.

The estimates of 56% and 58% suggest a modest but reliable increase in reported worry in July-September 2024 relative to June and April of 2023.



Stances on costs and benefits of AI

Besides worry about AI, respondents were asked the extent to which they agreed or disagreed with several AI-related stances (Figure 30). These stances reflect three dominant perspectives regarding possible costs and benefits of AI: the AI risk perspective (which focuses on far-reaching catastrophic or civilization-level risks such as an AI takeover), AI ethics (which focuses more on possible harms from factors such as biases in algorithms), and AI acceleration (the idea that we should be ramping up AI capabilities as fast as possible to derive their benefits). The exact wording for each perspective is presented with the figure below.

We estimate that the US public would agree—and among the stances, agree most strongly—with the idea that advanced AI could come to pose a significant danger to humanity, necessitating research, regulation, or pauses in AI development (i.e., the AI risk stance). However, the public also agreed that ethical risks, such as biased algorithms and privacy harms, should be a focus above more speculative harms (i.e., the AI ethics stance). In contrast, we estimate that the US public tends to disagree with accelerating the development of AI as

quickly as possible to secure its benefits, rather than worrying over harms (i.e., the AI acceleration stance).

Figure 30: US adults' agreement with AI-related attitudes

US adults' agreement with AI-related attitudes

AI poses significant danger:


Advanced artificial intelligence could come to pose a significant danger to humanity. We need to act now through measures like careful research, oversight, regulation and perhaps pauses in AI development in order to prevent this.

Ethics over speculation:

Artificial intelligence poses ethical problems, such as the risk of algorithms being biased against different people, and harms to privacy. We should focus on these kinds of harms, rather than more speculative future harms.

AI acceleration:

Artificial intelligence offers many potential benefits, such as new technologies and boosting the economy. We should promote the development of more advanced AI as quickly as possible, rather than worrying about possible harms.




Estimates based upon 4852 US adults

Although the content of each of these statements might suggest that all three are in tension with one another, we observed a significant positive correlation between AI risk and AI ethics (Spearman's rank = .49), and significant negative correlations between each of these items and AI acceleration (-.18 with AI ethics, and -.34 with AI risk). Hence, although the AI ethics statement emphasized the importance of focusing on ethical concerns *over* speculative harms (which might include the kinds of threats referred to in the AI risk stance), agreement with AI ethics did not preclude and was even positively associated with endorsement of AI risk. Agreement with AI risk and AI ethics might therefore reflect a general appreciation of potential harms from AI development.

Based on within subjects *Probability of Superiority* comparisons, we would expect approximately 73% of the US public to express greater agreement with AI risk and AI ethics than with AI acceleration (Figure 31).

Figure 31: Comparisons between AI attitudes

Comparisons between AI attitudes




Estimates based on 4852 US adults

It might be suggested that communities plausibly at greater risk of facing such bias (for example, racial or ethnic minorities) would show greater endorsement of AI ethics. To the contrary, ratings of the importance of AI ethics were quite constant across different racial identities (Figure 32). However, the mean ratings and estimates of *Probability of Superiority* in each racial subgroup suggest that only White or Caucasian adults reliably rated AI risk more highly than AI ethics (Figure 33). Respondents of all racial identities were on average much more in agreement with each of these risk-oriented statements than with AI acceleration, although White or Caucasian respondents were more negative about AI acceleration than were Asian or Asian American, Black or African American, and Hispanic or Latino adults.

Figure 32: US adults' agreement with AI-related attitudes - broken down by racial identification


US adults' agreement with AI-related attitudes - broken down by racial identification



Estimates based upon 4852 US adults

Figure 33: Comparisons between AI attitudes - broken down by racial identity

Comparisons between AI attitudes - broken down by racial identity




Within subjects probability of superiority

Estimates based on 4852 US adults

Female adults had more negative attitudes towards AI acceleration, and more positive attitudes towards each of the risk-oriented statements, than did males (Figure 34).

Figure 34: US adults' agreement with AI-related attitudes - broken down by sex

US adults' agreement with AI-related attitudes - broken down by sex



Estimated mean agreement [95% HDI]

Estimates based upon 4852 US adults

General discussion

Takeaways for impact-focused organizations and philanthropists

Given the scale of charitable giving in the US, a substantial positive impact could be achieved if US donors understood and engaged with a more impact-oriented approach to their charitable behavior.

At present, awareness of EA—a movement focused on finding and promoting the most effective ways of doing good—among the general public is low. In addition, very few donors use charity evaluators when making donation decisions. Yet, we estimate that the general public is positive towards several of the core ideas of EA/effective giving, such as evaluating the impact of different charities and promoting the most cost-effective ones. While we don't think these findings mean that there are many millions of Americans primed to become highly engaged contributors to the 'EA movement', we do think they suggest there is substantial room to expand the collection of people who consider and act upon EA ideas and effective giving more broadly. Examples might include more carefully considering their charitable behavior, or in thinking about how their career choices might be more or less impactful.

Awareness of EA and use of charity evaluators was highest amongst the highly educated. This may reflect the receptivity of such individuals to the reasoning behind EA and effective giving, but also that advocacy around these ideas has been geared towards the highly educated. Awareness of EA was relatively low amongst females (vs. males) and Black/African American respondents (relative to White/Caucasian or Asian/Asian American respondents), but there was little variation in receptivity to the core ideas of EA across sexes or racial identities. To expand the coalition of people engaged with effective giving, it may be beneficial to consider how to better reach these demographic subgroups.

Advocates of effective giving may also benefit from assessing whether different arguments in favor of particular donations and cause areas might be more or less appealing to different people. US adults broadly disagreed with one version of a perspective that is often provided in favor of donating to lower income countries, namely that most people in the US are relatively wealthy and capable of making quite sizable donations. We estimate that US adults probably underestimate their wealth relative to the rest of the world. Yet, even if people recognize that their incomes place them amongst the globally wealthy, we expect many would be unmoved by this reasoning: people are still greatly influenced by their perceived *local* relative wealth and the real struggles of their own economic situation, even if this situation is good by global standards. However, there are many different reasons that might be given in favor of charitable activities and particular cause areas. Donors need not agree with all aspects of EA, or with the types of arguments that many in the EA community favor, in order to shift their donations in more impactful directions.

One final consideration with respect to demographic subgroups and outreach was that, for every indicator related to awareness of and positive attitudes towards effective or impact-oriented philanthropy—from awareness of EA, to ratings of specific cause areas, to attitudes towards ideas and principles related to effective giving—Republicans showed lower ratings than Democrats. More work needs to be done to determine the reasons for this and whether they can be surmounted. Republicans have not been found to be less charitable, though they may on average give more to religious causes, and especially to local congregations. It is unclear whether the kinds of attitudes or traits that might explain why a

person leans Republican are fundamentally less compatible with the kinds of causes advocated for by the effective giving community, or whether this represents more of an issue with framing and surface-level perceptions.

Several causes deemed to have high potential for impact by the effective giving community were rated by the US public as important and deserving of donor support. However, they tended to rank low relative to more traditional causes such as Cancer Research or Civil Rights. It may prove useful in future research to understand the profiles and motivations of people who support these different causes, perhaps enabling more efficient advocacy to people who would be expected to be receptive to hearing about more neglected but impactful cause areas, or to more effectively communicate about these causes to those who may currently be less positive about them.

With respect to advocacy and promotion, our findings also highlight reputational risks as a consequence of negative publicity. Although most people who had heard of EA were positive about it, amongst those who were negative, a sizable proportion mainly associated EA with the fraudulent activities of Sam Bankman-Fried. Close associations between particular individuals and a 'brand' such as EA can have downsides, given that the brand is not in control of the individual's behavior and may be missing important information. However, it is also possible that the respondents we observed referring to this particular scandal already had negative views about EA, and appreciate having a concrete negative example to point to.

Takeaways regarding AI risk

The extent to which worry about the future of AI plays upon the minds of the US public on a daily basis has increased since 2023. Compared with April 2023, in July-September 2024, 19% (vs. 27%) do not worry at all, 41% (vs. 44%) worry only a little bit, 27% (vs. 21%) worry a fair amount, 11% (vs. 7%) worry a lot, and 3% (vs. <1%) worry nearly all the time. Moreover, among stances on the development of AI, US adults most agreed with a statement emphasizing serious risks from AI and the need for actions such as regulations. Another perspective, which focused on 'ethical' risks such as algorithmic biases and privacy over more 'speculative' harms, was also widely supported. In contrast, the view that development should be accelerated as much as possible to reap the benefits of AI, rather than focusing on risks, was mostly disagreed with.

Taken together, these findings suggest that the US public may be more focused on the risks than the benefits of AI, and are receptive to possible regulatory approaches. We also found that views of AI risk as a cause area were relatively politically neutral (there was little difference between Democrats, Republicans, and Independents in their ratings). This may suggest an opportunity for safety-focused policy proposals, which could receive bipartisan support. Developing bipartisan communication strategies to promote AI safety regulations could be useful.

Concluding remarks

This first wave of *Pulse* provides estimates derived from a large sample regarding philanthropic cause area preferences, awareness of EA and EA-related figures and organizations, donation-related behavior, and perceptions of AI risk. We hope this work has provided useful insights for charitable and philanthropic organizations and the broader effective giving community. Alongside future work with new waves of *Pulse*, we will be releasing additional reports delving into subsections of this report in more detail. Readers can subscribe to our [newsletter](#) to be kept informed of such releases.

Methodological transparency

This document contains information regarding details of the survey design, implementation, and analysis, as proposed by the American Association of Public Opinion Research (AAPOR).

1. **Data Collection Strategy.** This research was based on answers to a cross-sectional online survey.
2. **Who Sponsored the Research and Who Conducted It.** This research was conducted by *Rethink Priorities* (RP), a registered 501(c)3 non-profit organization and think tank, registered in California, US. RP financial disclosures can be found at: <https://rethinkpriorities.org/transparency>. Funding support for this research was provided by Open Philanthropy (Open Philanthropy Project LLC).
3. **Measurement Tools/Instruments.** All questions for the survey are provided in a separate survey appendix, available at <https://tinyurl.com/rp-pulse-2024-survey>.
4. **Population Under Study.** This research aimed to assess US Adult public opinion.
5. **Method Used to Generate and Recruit the Sample.** This research used a non-probability sampling approach, involving respondents opting-in to take part in the study as part of *CloudResearch* online panels. Respondents were not informed about the specific content of the survey until entering the survey, thereby reducing bias from respondents who enter being specifically interested in charity/philanthropy. Respondents were required to be adults (aged 18 and above) and resident in the US. The *CloudResearch* panel administrator followed a guiding quota provided by *Rethink Priorities* to ensure that sample demographics closely matched the *US Census Bureau's American Community Survey* proportions of demographic subgroups, with balancing on Age, Sex, Race/Ethnicity, Household Income, Region, and Education, with an additional quota for political party identification based on *Gallup* polling (quotas were used to ensure sampling did not result in a heavily unrepresentative sample, but results still utilize either multilevel regression and poststratification, or weighting). Compensation for respondents was agreed upon between the panel provider and respondent prior to entry into the survey and could vary depending on the panel sourced by *CloudResearch*.
6. **Method(s) and Mode(s) of Data Collection.** Respondents were recruited via the *CloudResearch* online panel system, and completed the survey using the *Qualtrics* survey software platform. The survey was offered in English.
7. **Dates of Data Collection.** Pilot fielding of the survey began on July 18th, 2024. The main survey was fielded from July 25th, 2024 and September 16th, 2024.
8. **Sample Sizes and Precision of the Results.** 8034 potential respondents began the survey, with a final sample of 4890 following a multistage quality assurance procedure. *Stage 1:* 8034 survey entries. *Stage 2:* 5627 after early termination due to failure to pass initial basic attention checks or self-termination. *Stage 3:* 5207 after screening on *Qualtrics* bot detection, duplicate detection, and fraud detection scoring. *Stage 4:* 4890 after quality assurance screening for repeated failure to pick very high incidence items including a late attention check item and/or repeated picking of fake items. The exact

sample size for each question may vary due to missing data or branching of the sample to see alternative questions.

Margins of error in primary analyses represent the means and 95% highest density intervals (HDIs) of posterior distributions derived from Bayesian Multilevel Regression and Poststratification (MRP). Please see the plots and main text of the report for the uncertainty associated with each result, as there is no single +/- margin of error that applies to all estimates. Margins of error are presented with more decimal places than the central point estimate to ensure that the width of the margin of error is not understated (rounding the upper and lower bounds of the error margin could artificially reduce its apparent width).

MRP is a technique that can be used to estimate outcomes in a specific target population based upon a potentially unrepresentative sample population. The technique involves generating estimates of how a range of features (e.g., education, income, age) are associated with the outcome of interest from the sampled population, using multilevel regression. Based on the known distribution of combinations of these features in the target population, the poststratification step then involves making predictions from the multilevel regression model for the target population. This approach is widely used to make accurate predictions of population level opinion and voting based upon unrepresentative samples, and also allows inferences to be made about specific subgroups within the population of interest¹⁷. Model specification for multilevel regression used respondent State/District, Region, Age bracket, Completed education, Household income bracket, Sex, Racial identification, Sex * Racial identification, Completed education * Age bracket, Political party identification, and the State/District's Republican vote share for the 2020 Presidential election.

Some analyses utilize weighting as opposed to MRP, again with the mean and 95% HDI of the estimated posterior distribution used as summary statistics. The weighting accounts for Region, Age bracket, Completed education, Household income bracket, Sex, and Racial identification, and political party affiliation. Margins of error for weighted analyses incorporate a penalization for the design effect induced by weighting. The following section describes the poststratification and weighting details.

9. **How the Data Were Weighted.** Regression model predictions were poststratified according to the cross-tabulated proportions of the US population with the respective demographic features outlined in the previous section based upon public release of the *Census Bureau's 2022 5 year American Community Survey* for US adults. This poststratification table was extended to include a posterior distribution of expected political party identification based upon multilevel regression linking respondent demographics and political party identification. Weighting procedures similarly used proportions from the *Census Bureau's 2022 5 year American Community Survey* as well as a 3-month moving average for political party identification derived from publicly released data from *Gallup*. Weights were generated using the *anesrake* R package implementation of raking used by the *American National Election Studies* group¹⁸. Raw sample and target percentages across demographics are available at <https://tinyurl.com/rp-pulse-2024-demos>.

¹⁷ Wang, W., Rothschild, D., Goel, S., & Gelman, A. (2015). Forecasting elections with non-representative polls. *International Journal of Forecasting*, 31(3), 980-991; Park, D. K., Gelman, A., & Bafumi, J. (2004). Bayesian multilevel estimation with poststratification: State-level estimates from national polls. *Political Analysis*, 12(4), 375-385.

¹⁸ Pasek, J., & Pasek, M. J. (2018). Package 'anesrake'. *The comprehensive R archive network*.

10. **How the Data Were Processed and Procedures to Ensure Data Quality.** As detailed above, data underwent a multistage quality assurance procedure: *Stage 1:* Survey entry. *Stage 2:* Exclusions after early termination due to failure to pass initial basic attention checks or self-termination. *Stage 3:* Screening on *Qualtrics* bot detection, duplicate detection, and fraud detection scoring. *Stage 4:* Quality assurance screening for repeated failure to pick very high incidence items including a late attention check item and/or repeated picking of fake items.

Respondents who selected Effective Altruism as a term they had heard of underwent assessment by a trained human coder to determine the level and validity of claimed awareness. Such respondents were asked to indicate whether they thought they had at least a basic understanding of what Effective Altruism means and, if yes, to provide a short explanation of the term. Responses were coded as aware at a Stringent level of assessment when the respondent unambiguously referred to one or more of the core ideas of EA, such as maximizing the amount of good done, and doing so through charity or careers, or referring to figures associated with the EA movement (including reference to scandals or controversial figures). Responses were coded as aware at the Permissive level if the respondent was alluding to some core ideas, but with lower specificity than the stringent level. Responses were also assessed for copy-pasted/paraphrased definitions, which were not counted towards either Stringent or Permissive awareness. The assessment of the human coder was supplemented by assessment with the Claude Sonnet LLM using the same criteria. Discrepancies between the human and LLM-assisted code were then re-assessed to determine the final rating.

11. **Limitations of the Design and Data Collection.** Survey data and its analysis and interpretation can be prone to numerous issues. One particular concern for non-probability samples (i.e., opt-in, online surveys) is the potential for biases in recruitment that are not or cannot be counteracted by weighting or poststratification approaches—for example, although we aim to correct for certain demographics, it remains possible that people who come to the online survey panel are unrepresentative of the general population in other ways that are not captured in our weighting or poststratification procedures. Hence, the full uncertainty of estimates is greater than the statistically-derived margins of error. In some cases, important estimates (such as the proportion of people who have heard of Effective Altruism) are based on quite low base rates of awareness. These estimates may be less stable or accurate than estimates for more well known items, because more substantial shifts in the estimate could be produced if sampling for some reason reached a particular pocket of the population with high awareness.

