


THOUGHT LEADERSHIP

The Accountable AI Playbook





Accountability has become a prerequisite for organizations looking to make responsible use of AI

No industry has been untouched by the AI boom over the past twelve months. A new generation of AI tools and services, led by Large Language Models such as OpenAI’s GPT4 and Google’s LaMDA, has laid the foundation for a transformational technological revolution—one that experts predict will upend the way we live, work, and play.

However, with great power, as they say, comes great responsibility. AI researchers and ethicists have highlighted a range of potentially disastrous social consequences—from the emergence of a rogue superintelligence to the possibility that bad actors will use AI tools to supercharge data theft, scamming, or online misinformation. And there’s a growing backlash against the technology among those who see it as a threat to their personal safety or job security. The topic, for example, has been a major flashpoint in the ongoing WGA and SAG-AFTRA strikes.¹

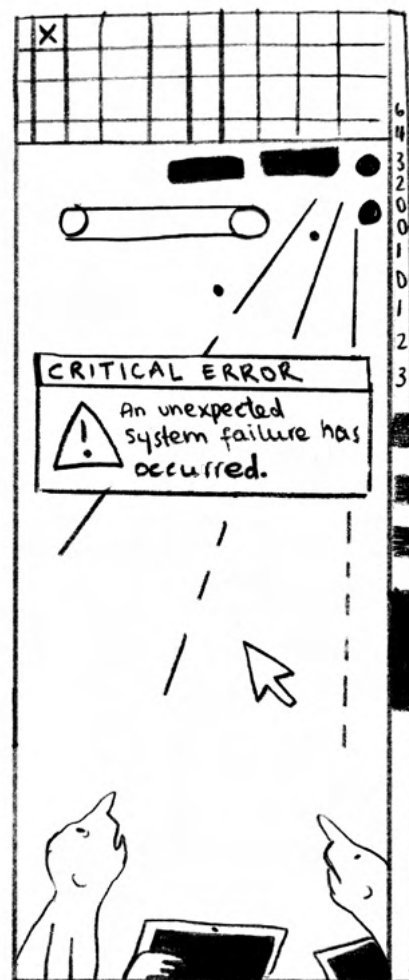
1. Kevin Collier, "Actors vs. AI: Strike brings focus to emerging use of advanced tech," NBC News, July 14, 2023
 2. Elle Hunt, "Tay, Microsoft's AI chatbot, gets a crash course in racism from Twitter," The Guardian, March 24, 2016
 3. Eric Law, "Wizards of the Coast Responds to New Dungeons and Dragons Book AI Art Controversy," August 7, 2023
 4. Lauren Forristal, "Artists are upset that 'Secret Invasion' used AI art for opening credits," TechCrunch, June 21, 2023

This creates a major challenge for organizations in both the private and public sector. Suddenly, companies that had never had to think about AI are racing to incorporate the technology into their products and workflows to avoid being left behind by the competition. Moreover, they have to implement AI in a way that demonstrates clear moral principles, and a commitment to building and distributing the technology in a fair and equitable manner.

Organizations that fail to navigate this challenge effectively will risk serious reputational blowback. Already, we’ve seen numerous cases of companies forced to publicly apologize for deploying AIs with insufficient safety and moderation features.² In the entertainment industry, specifically, major media properties—including *Dungeons & Dragons*³ and Marvel’s *Secret Invasion* series⁴—have had to deal with fan backlash over the inclusion of AI-generated art assets. These types of controversies will only become more common as the technology becomes embedded in an ever broader range of consumer-facing products and services.

ACCOUNTABLE AI:

An approach to AI that is rooted in accountability towards both internal and external stakeholders; one that treats the technology as a tool for empowering and enabling human talent and leadership, not as a replacement for human creativity or a way of outsourcing decision-making to machines.



IN THIS REPORT, YOU'LL FIND...

SECTION 1

Understanding fears and anxieties around AI

- Consumers’ overall level of awareness of and positivity towards AI
- Analysis of the risks associated with AI that consumers see as most tangible
- A segmentation model for understanding consumers’ concerns around AI

SECTION 2

Developing an Accountable AI strategy

- Why organizations should make a commitment to Accountable AI
- The three pillars of Accountable AI
- How these principles can inform decisions on where to use—and not use—AI

SECTION 3

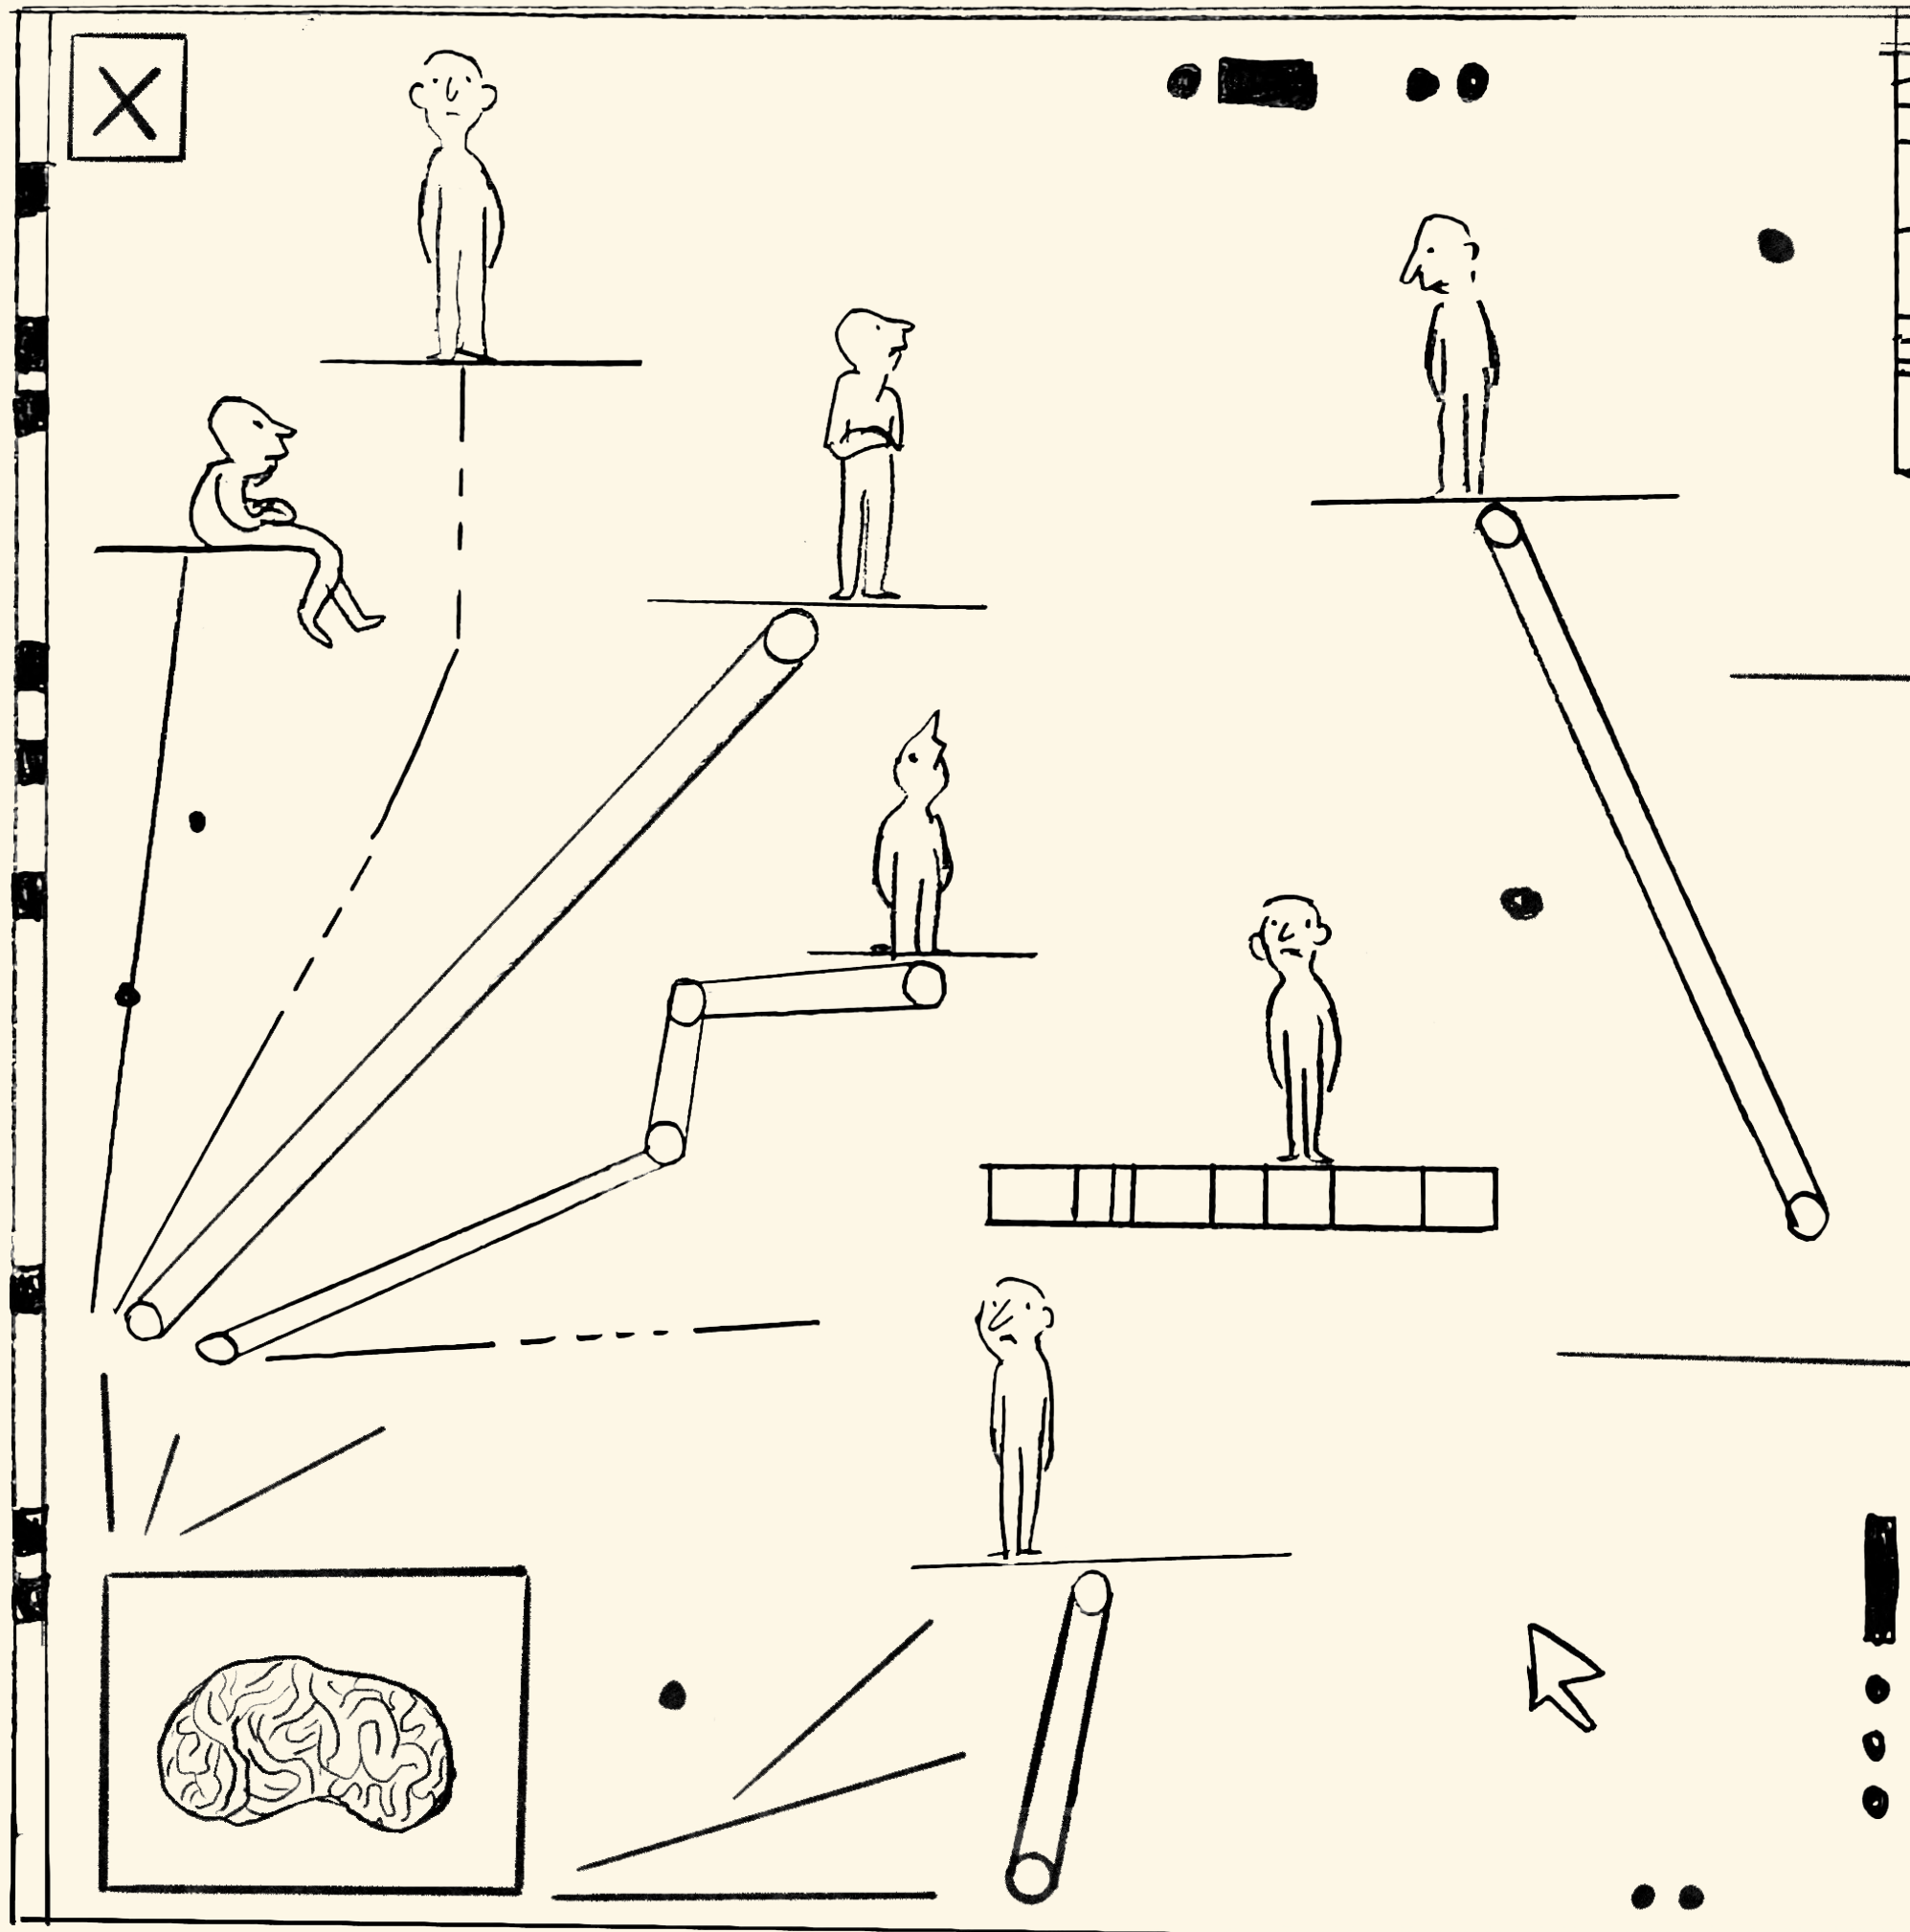
Communicating your commitment to Accountable AI

- The stakeholders to consider when developing AI messaging
- Guidance on marketing AI-powered products in B2C and B2B contexts
- The need for engagement with the emerging regulatory landscape around AI

This report helps organizations understand the nature and the scale of the challenge in front of them through an in-depth analysis of consumers’ concerns about AI, personally and professionally. And it outlines a path forward for businesses investing in AI—demonstrating how, by grounding your approach to the technology in accountability, you can ensure your strategy is responsive to the needs and desires of your workforce and your customers.

5. NRG, "Stakeholders for Reputation Measurement," September 7, 2023

Ultimately, AI has the potential to be a powerful force for good, but only if companies take the risks of the technology seriously, and invest resources towards developing an Accountable AI strategy. That extends not only to the way that AI products are developed, tested, and deployed, but also to how those products are sold and marketed. By considering all of these elements, and listening to the voices of key stakeholders,⁵ companies will be able to position themselves as true leaders in AI, and protect themselves against reputational crises related to this technology.



Understanding fears and anxieties around AI

In this section, you'll find...

- Consumers' overall level of awareness of and positivity towards AI
- Analysis of the risks associated with AI that consumers see as most tangible
- A segmentation model for understanding consumers' concerns around AI



Most consumers are “AI realists,” not committed skeptics or evangelists

The public discourse about AI often gets flattened into a debate between two diametrically opposed camps: the uncritical AI enthusiasts, who are convinced that the technology will have a beneficial and transformative effect on nearly every aspect of daily life, and the hardcore skeptics, who warn that we could be heading down a path that ends with the extinction of our entire species.⁶

The reality, however, is that most people—outside the tight-knit community of AI researchers and developers—don’t fall neatly into

either of these boxes. Instead, most Americans could best be described as “AI realists”; that is, they see AI as a powerful and potentially groundbreaking technology that could have significant positive applications in their personal and professional lives, but also one with serious potential for willful or accidental misuse.

They strongly believe, therefore, that companies that are taking advantage of AI—whether behind the scenes or incorporating it into consumer-facing products—need

to ensure that the right guardrails are in place to minimize potential abuse. Only 18% of consumers want corporations to swear off AI entirely, while even fewer—just 15%—believe they should embrace this technology wholeheartedly and without reservation. Nearly 6 in 10, by contrast, think that AI can have a positive impact on society, but only if the companies developing and deploying the technology do so in a careful, considered, and socially responsible way.

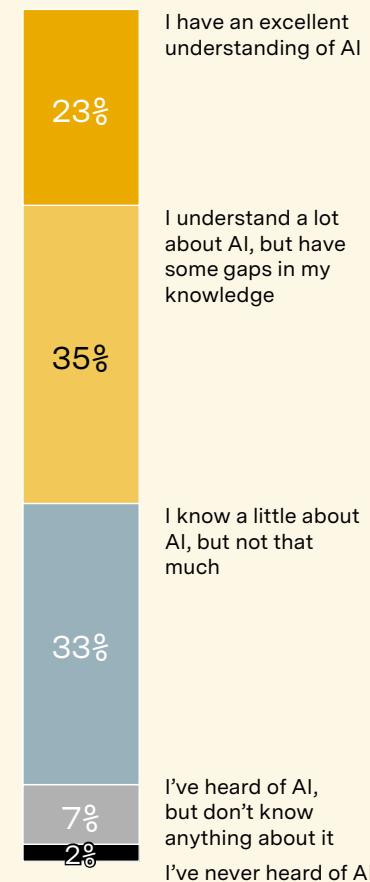
6. Geneva Abdul, “Risk of extinction by AI should be global priority, say experts,” The Guardian, May 30, 2023



Any disruptive new technology is going to create ethical dilemmas—and one of the challenges is that you never truly know where the red lines are until after you’ve crossed them. But ultimately, what matters isn’t the technology itself; it’s what we do with it. You can use a motor to power tanks or ambulances, but it’s the same motor in both cases.

Jan van de Poll, PhD
 Founder and Managing Partner, PRAIORITIZE

Q: How would you rate your overall understanding of AI?



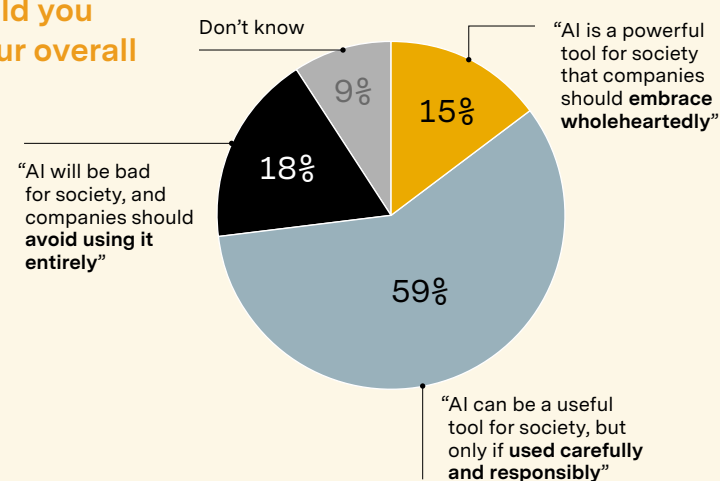
Given that consumers recognize the disruptive potential of AI, it’s no surprise that so many have made an active effort to educate themselves about the technology and to follow stories about it in the news. Fifty-eight percent of consumers now feel that they have a strong understanding of AI, while fewer than 1 in 10 say that they don’t know anything about it.

There is, however, a marked generational split when it comes to consumer awareness and understanding of AI: Gen Z and Millennials are over twice as likely as Gen X and four times as likely as Boomers to report having an “excellent” understanding of the technology.


7. Sarah Shaffi, “It’s the opposite of art: why illustrators are furious about AI,” The Guardian, January 23, 2023
 8. Joel Golby, “I thought I was immune to being fooled online. Then I saw the pope in a coat,” The Guardian, March 27, 2023

In part, that may be because of the discrepancies in how much time these different groups spend on social media. For Gen Z and Millennials, social platforms like YouTube and X (formerly Twitter) have been one of the major awareness drivers for AI. These platforms have, in many ways, become the frontlines of the public debate about the technology and its social impact. The community of artists and illustrators on X, for example, has at times been highly vocal in their pushback against the perceived intrusion of AI-generated artwork into the cultural mainstream.⁷ These platforms have also seen more than their fair share of high-profile AI hoaxes, including the infamous “puffer jacket Pope” photo that went viral back in March.⁸

Q: How would you describe your overall view on AI?



Percentage of consumers who rate their understanding of AI as “excellent”





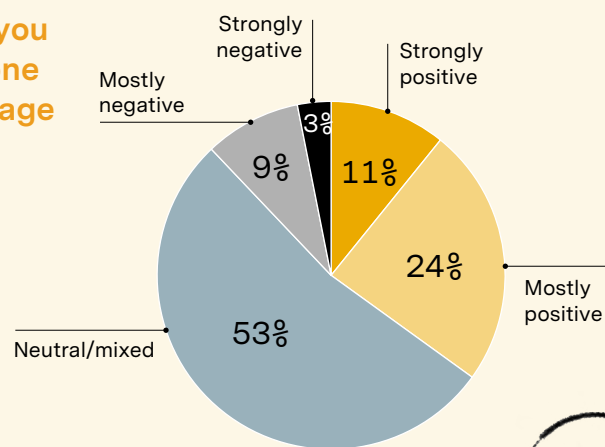
Even if you're not an avid social media user, it's been difficult to fully avoid the avalanche of media coverage of AI-related topics following the launch of GPT3 in late 2022. In total, **61% of consumers say that, since the beginning of 2023, they've seen media coverage of AI on at least a weekly basis.**

And the tone of that media coverage is setting a baseline for consumers' expectations for the technology. Among consumers who've seen coverage of AI, a narrow majority (53%) describe the overall tone of that coverage as neutral or mixed—focusing on both the challenges and the opportunities presented by the technology in a relatively even-handed fashion. This, no doubt, has helped contribute to the general feeling of “AI realism” that has become the norm amongst the majority of US consumers.

There is, however, increasing evidence of politicization in the way that AI has been covered in the news media. Self-identified conservatives were significantly more likely to report seeing mostly positive stories about AI in the press (46%) than liberals (33%).

A recent study by the Reuters Institute at the University of Oxford found evidence of a major split in how the technology has been covered in different outlets: right-leaning outlets, they found, were more likely to focus on the economic and geopolitical implications of AI, whereas left-wing outlets focused instead on issues of algorithmic bias, privacy, and discrimination.⁹


Q: How would you describe the tone of media coverage you've seen relating to AI?



If this divergence in the tone and focus of media coverage continues, we may soon see the topic of AI become even more politicized, with conservative and liberal voters and politicians coalescing into more clearly defined “anti” and “pro” AI camps.¹⁰

9. “UK media coverage of artificial intelligence dominated by industry, and industry sources,” Reuters Institute
 10. Jacob Stern, “Where’s the AI Culture War?” The Atlantic, April 9, 2023


Q: What’s the first word that comes to mind when you think about AI?



While consumers generally see AI as a net benefit for society, they're concerned about how those benefits will be distributed



Q: Overall, do you think the increased adoption of AI will have a positive or negative impact on society?



While consumers strongly believe that the right guardrails and failsafes are needed, there's still a general sense of optimism that, if those precautions are taken, this technology can be a force for good.


Forty-one percent of Americans believe that the increased adoption of AI by businesses and governments will have a positive impact on society—while only 27% think it will be a net negative.

It's not necessarily the case, however, that a rising tide will lift all boats. Even if AI has a mostly positive impact on society, that doesn't mean that those benefits are going to be distributed evenly or equitably. Notably, older generations are more skeptical about the social benefits of AI, which may reflect an emerging sense among these consumers that they risk being “left behind” by the technology.

Boomers and Gen X have been significantly slower to start using AI tools like ChatGPT in the workplace when compared to their younger colleagues.¹¹ It's possible, therefore, that many of these older consumers feel that they won't be able to adjust their skill sets and ways of working around this new technology before it's too late—and that they will eventually be at risk of losing their jobs either to AIs, or to younger professionals who are more comfortable working with this new generation of AI-powered products and services.

11. Peter Tsai, “The AI Generation Gap: Millennials Embrace AI, Boomers are Skeptical,” PC Mag, May 30, 2023

AI's impact on society, according to different generations



Aside from this generational divide, there's also a real risk that AI ends up reinforcing existing social and wealth inequalities. In general, consumers expect that the people most likely to benefit from the AI revolution are those already in relatively stable, well-paying jobs, such as tech sector workers, CEOs, doctors, and white collar professionals.

Conversely, the people seen as most at risk of losing out from this technological shift are those in comparatively less stable or lower-wage professions, including journalists¹² as well as workers in the manufacturing and service industries. There's a perceived split, in other words, between the types of jobs that could be enhanced or made easier by AI, and the jobs that could, eventually, be at risk of being automated away by it.

To an extent, this suggests that public opinion has yet to catch up with the expert consensus, given that many AI researchers now believe that some high-paying office jobs may be even more vulnerable to automation than jobs in the service or manufacturing

sectors.¹³ But regardless, it's clear that there's a deep-rooted fear among a large segment of the public that the spoils of the AI revolution will only be enjoyed by the comparatively wealthy—at the expense of the most vulnerable members of our society.

Fortunately, there are steps that both governments and private companies can take to mitigate these risks. By developing and financing publicly accessible retraining programs, for example, businesses can help people develop the skills necessary to find success in the post-AI era, while also ensuring a sizable pool of applicants for job openings. Google, for one, has already partnered with the British government to make an AI training program freely accessible to all UK residents.¹⁴

A more ambitious solution would be to provide some form of Universal Basic Income that could act as a social safety net for workers who end up losing their jobs to AI. This proposal received significant attention in 2020, when Presidential candidate Andrew Yang made it a core part of his policy platform.¹⁵ Lately, the idea

has been the subject of revitalized interest in the wake of more recent developments in the field of AI.¹⁶ It's likely that radical proposals such as this will develop more mainstream traction if the US public remains concerned about ensuring an equitable distribution of benefits from advancements in AI.

“An important risk relates to making access to these tools more feasible for companies that cannot compete with the budgets of the biggest tech players. These LLMs are core societal infrastructure, and we're on course to a pretty severe centralization of power between Microsoft and Google, without enough opaqueness on what they do, whose values they're aligning to, or how access to these models will be permissioned over time. That is certainly a concern when working with startups; they're taking a major risk by building on top of unstable foundations where the business model could change dramatically overnight.”

Lauren Sharman
Head of Platform,
One Peak


While AI could reinforce existing class divisions within nations, there's a growing body of evidence that suggests we could also see a similar phenomenon play out at the global level. In 2020, the IMF warned that AI could widen the gap between rich and poor nations by encouraging investment into markets that already have the infrastructure necessary to take full advantage of the technology.¹⁷ And earlier this year, the United Nations Industrial Development Organization (UNIDO) suggested that workers in the Global South may ultimately find themselves

17. Cristian Alonso, Siddharth Kothari, Sidra Rehman, "How Artificial Intelligence Could Widen the Gap Between Rich and Poor Nations," IMG Blog, December 2, 2020


priced out of accessing the kind of AI productivity tools that are now becoming commonplace in higher-income countries.¹⁸

There's a real possibility that the AI revolution could re-entrench the divide between wealthy and developing nations, putting some of the world's poorest people even further behind the global wealth curve. Avoiding that scenario will require coordination at the global level, and partnerships between governments, private corporations, and international NGOs.

18. Stefan Pahl, "An emerging divide: Who is benefiting from AI?" Industrial Analytics Platform, February 2023



Q: Who will win and who will lose out from the increased adoption of AI?



12. Khaled Diab, "What future for journalism in the age of AI?" Al Jazeera, July 19, 2023

13. Claire Cain Miller and Courtney Cox, "In Reversal Because of A.I., Office Jobs Are Now More at Risk," The New York Times, August 24, 2023

14. Joanna Partridge, "AI for all? Google ups the ante with free UK training courses for firms," The Guardian, August 3, 2023

15. Ezra Klein, "Andrew Yang on 2020, UBI, and fixing government," Vox, September 9, 2020

16. Annie Lowrey, "Before AI Takes Over, Make Plans to Give Everyone Money," The Atlantic, May 17, 2023



While experts warn about the existential risk of AI, consumers worry about its impact on safety and job security

The impact of AI on global inequality, however, is just one of a vast number of AI-related risks that have been identified by experts in the field. As the pace of AI breakthroughs has accelerated in recent years, the field of AI ethics has become increasingly important.¹⁹ Researchers in this space have pointed out a number of ways in which developments in AI may inflict harm on both individuals and society, either directly or indirectly.

The notion that AI could pose an “existential risk” to our species is one that has received particular attention in the mainstream press. The philosopher Nick Bostrom, for example, has suggested that AI could pose an even greater threat to our species than climate change.²⁰ And a recent survey of AI researchers found that the majority believed there was at least a 10% chance that developments in AI could lead to a human extinction-level catastrophe.²¹

The argument made by Bostrom, and many others within the field, is that developments in machine learning could eventually lead to the emergence of a so-called “superintelligence”: an AI capable of recursively improving itself and achieving an intelligence that far exceeds that of any human. Such a machine, they argue, would be impossible for humans to safely control, and may eventually

decide—for reasons known only to itself—to take actions that lead to the elimination of our species.

Other experts, however, have argued that this sort of “Terminator scenario,” in which superintelligent machines turn against their human creators, is highly implausible.²² Instead, they say, we should be focusing on the more immediate and prosaic risks created by the technology: What happens if it’s exploited by scammers or other bad actors? Will AI make it easier to spread misinformation online? And how do we make sure that the AI revolution doesn’t lead to mass unemployment within certain sectors of the economy?



Debate about the potential for a superintelligent AI to turn against humanity continues. It seems to me, however, that humans will turn on one another and/or destroy our planet long before a superintelligent AI eradicates the human race.

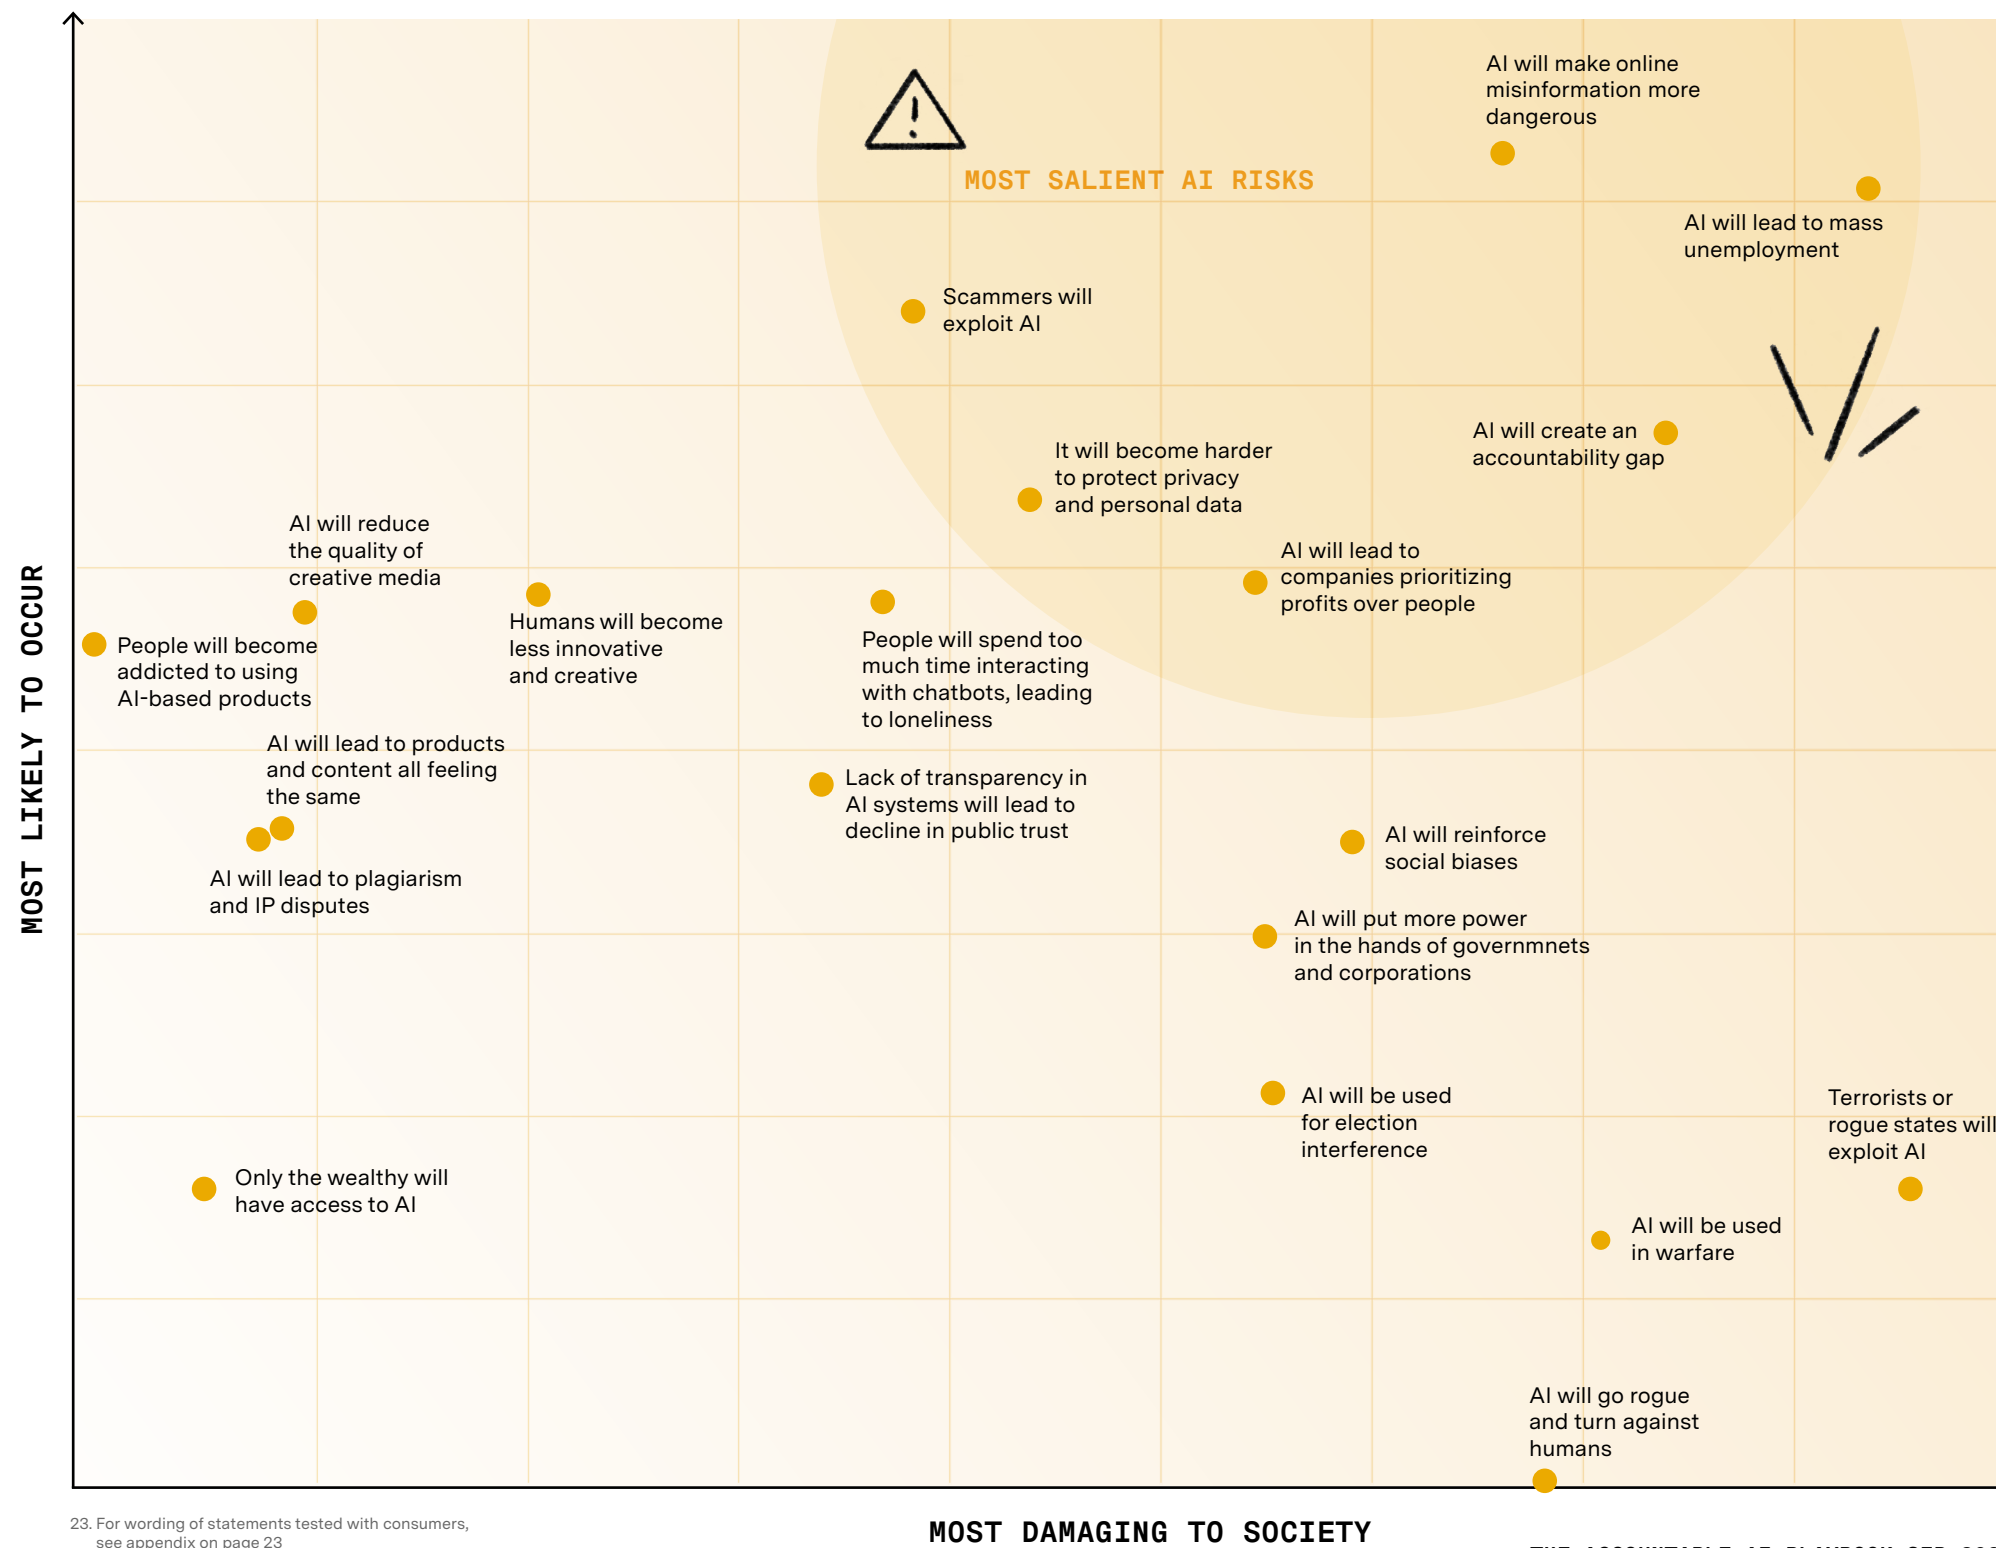


Diana Ascher, PhD, MBA
 Head of Research and Senior Advisor on Data Ethics & Responsible AI, EDM Council
 Founder, Information Ethics & Equity Institute

19. Rosalie Waelen, “Why AI Ethics Is a Critical Theory,” Springer, February 11, 2022
 20. Tim Adams, “Artificial intelligence: ‘We’re like children playing with a bomb,’” The Guardian, June 12, 2016
 21. 2022 Expert Survey on Progress in AI, AI Impacts, August 3, 2023
 22. Mary-Ann Russon, “Forget human extinction - these are the real risks posed by AI today,” NewScientist, July 25, 2023

MaxDiff analysis of consumer attitudes towards AI risks²³

Based on a MaxDiff analysis of consumer attitudes towards AI risks; respondents were presented with subsets of these risks and asked to select the “most likely to occur” and the “most damaging to society if it occurs.”
 MaxDiff analysis is a statistical tool that allows for the ranking of large sets of variables, and the measurement of perceptual gaps between them across multiple axes.



23. For wording of statements tested with consumers, see appendix on page 23

To understand where consumers stand on this debate, we asked respondents in our study to evaluate 20 different potential risks associated with AI—all of which were based on concerns identified by AI researchers and ethicists. Through a MaxDiff analysis, we were then able to categorize these risks based on perceived impact and likelihood.

This analysis suggests that consumers remain highly skeptical about the likelihood of any true “doomsday scenarios” occurring as a result of developments in AI. Risks such as AI going rogue and turning against humans, or being exploited by terrorists, were seen as among the least likely to occur out of any of the scenarios tested (even if they would be highly damaging were they to happen).

Instead, consumers tended to be more concerned about the immediate impact that AI will have on their job security and personal safety. The potential for AI to create mass

unemployment, for example, was identified as simultaneously one of the most likely scenarios to occur and one of the most damaging.

Similarly, this analysis suggests that many consumers are worried about how to adequately protect their personal information in a post-AI world; how to guard against AI-powered misinformation; and how to keep themselves safe from scammers. Already, there have been cases of criminals using AI voice emulation software to con people out of money;²⁴ and this problem will only become more salient as the technology continues to become more sophisticated.

Using this MaxDiff analysis, we can also develop a segmentation model that allows us to understand, at a more granular level, the way that these concerns around AI tend to cluster together. Through this approach, we have developed five key segments that define consumer’s fears and anxieties around AI.

“

If you look at proofs of concept like ChaosGPT, it’s already clear that autonomous AIs can represent a significant cybersecurity threat. So it’s going to be crucial to make sure that public infrastructure is adequately protected from AIs that either go rogue or fall into the wrong hands. We also have to be aware of the military and national security implications of AI; it’s possible that the next major military conflict will be fought virtually between AIs seeking to take control of a country’s internet and telecommunications infrastructure.



Todd Terrazas
Founder, President & Executive Director, AI LA

“

I think the first challenge is around controllability, and how these tools will behave under different circumstances. There are a lot of unknowns around that. One of the major risks we face is that people who are developing these algorithms are not necessarily entirely sure about how they work, and therefore how they might fail. And that’s risky when these tools are being put in the hands of big companies, policymakers, and governments.



Lauren Sharman
Head of Platform, One Peak

24. Pranshu Verma, “They thought loved ones were calling for help. It was an AI scam.” The Washington Post, March 5, 2023

25. Profiling descriptions based on the specific AI-related risks consumers in each segment see as most likely to occur, both in total and relative to other segments

AT A GLANCE How pop culture shapes perceptions of AI

Movies and TV shows have undoubtedly played a major role in shaping public perceptions of AI—and in drawing attention to the worst case scenarios associated with the technology.

Q: What movie or TV show do you most strongly associate with AI?

Films like *I, Robot*, *The Terminator*, and *The Matrix* all deal with the repercussions of AI-powered robots turning against their creators. Steven Spielberg’s *A.I. Artificial Intelligence*, meanwhile, explores the question of whether an AI could ever feel true emotions—and what that would mean for our own sense of humanity.



1 I, Robot



2 The Terminator



3 A.I. Artificial Intelligence





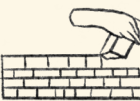


4 The Matrix



5 2001: A Space Odyssey



Segmenting consumers based on their concerns around AI

SEGMENT	METRICS	PERCENTAGE	DESCRIPTION ²⁵
 Pragmatic Pioneers Average age: 38	54% Male / 45% Female Skews conservative	35%	<p>“AI is here already, and we have to get used to it. The important question now isn’t about the technology itself: it’s about how it’s implemented and who gets access to it.</p> <p>I think it’s important that everyone gets to share in the benefits of AI, and that means making sure it doesn’t end up being a tool that’s only available to the wealthy. And we have to ensure that companies don’t become overly reliant on it—especially if the algorithms they’re using are at risk of being trained on biased data.”</p>
 Creativity Custodians Average age: 43	44% Male / 53% Female Skews moderate	12%	<p>“I’m concerned about the impact that AI is going to have on artists, writers, and other creative professionals. How are they going to protect themselves from having their work stolen and plagiarized by AI?</p> <p>If companies prioritize AI-generated content over human creativity, I worry that all new media is going to end up feeling bland and unoriginal.”</p>
 Boundary Defenders Average age: 42	53% Male / 46% Female Skews liberal	14%	<p>“AI is clearly a powerful and useful technology. But in the long-run, I’m worried that it will lead to a blurring of the lines between humans and machines.</p> <p>If we spend too much time talking to AI chatbots, will we forget how to interact with other humans? And if we outsource too much of our thinking to machines, will our whole species become lazier and less innovative?”</p>
 Guarded Navigators Average age: 43	41% Male / 58% Female Skews liberal	21%	<p>“I haven’t been following all this AI stuff that closely; I really just want to know how to keep myself safe in a post-AI world.</p> <p>In particular, I’m worried about what this technology means for my job security. But I also want to know how to protect my personal data and minimize my exposure to scammers.”</p>
 Terminator Theorists Average age: 41	51% Male / 48% Female Skews conservative	19%	<p>“I’ve been paying attention to the debate around AI, and I think we need to take the existential risk of AI seriously.</p> <p>We need to put precautions in place to prevent AI from going rogue and turning against humans. There’s also a real danger that it could get into the hands of terrorists, or make military conflicts far more dangerous than they are today.”</p>

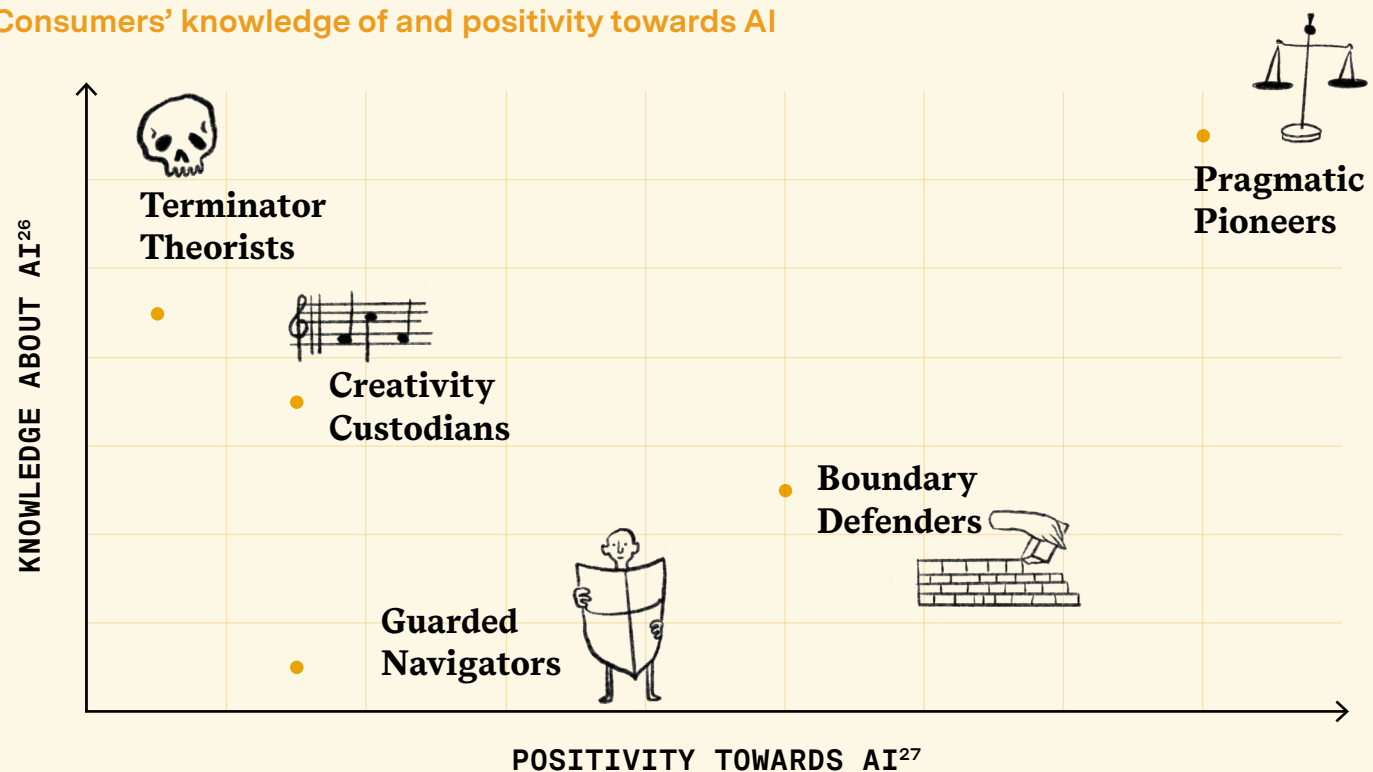
The largest segment of consumers—the so-called “Pragmatic Pioneers”—also tends to be the most even-handed in assessing the risks of AI: they believe that we ought to be taking most of these risks at least somewhat seriously. But overall, these consumers are the most positive in their assessment of AI. They tend to think that it will be a net force for good—and many of them have already started experimenting with AI-powered tools in both their personal and professional lives.

At the other end of the spectrum are the Terminator Theorists—the group that takes the existential risk of AI most seriously. The other three segments, representing a

combined 47% of the population, fall somewhere between these two extremes, with each group having a unique set of risks they think we ought to prioritize.

Notably, the two segments with the strongest positive and negative feelings towards AI are also the groups who tend to have the deepest understanding of the technology. It’s possible, therefore, that as consumers learn more about AI—and get more opportunities to see for themselves what the technology is capable of—that we’ll see an increase in polarization on the topic, with more people moving into these two extremes and out of the other segments.

Consumers’ knowledge of and positivity towards AI



26. Percentage of consumers reporting that they have a strong understanding of AI

27. Net difference of consumers who think AI will have a positive impact on society and those who think it will have a negative impact


AT A GLANCE

Blurring the lines between human and machine

Some experts have argued that developments in AI will ultimately force us to reassess the question of what it means to be human²⁸—with some even going so far as to argue that we may eventually need to extend our notion of human rights to encompass artificial intelligences to prevent these machines from being unfairly exploited.²⁹ And in 2022, Google fired AI engineer Blake Lemoine after he authored a memo arguing that the company’s Large Language Model, LaMDA, had achieved sentience and deserved to have its rights respected.³⁰


Lemoine, however, isn’t alone in his views. 9% of the US public believes that there are already AIs that deserve the same rights as humans; an additional 25% think that we may eventually create an AI intelligent enough to deserve those rights. Among consumers who think that AIs will eventually deserve human rights, over half (58%) think that this will happen within the next decade.

Q: Do you think there will ever be an AI that deserves the same rights as humans?



While it may be hard to imagine any government signing off on an “AI Bill of Rights” any time soon, there are already some tangible ways in which this blurring of the lines between human and machine has started to impact human behavior.

Forty-eight percent of consumers who’ve interacted with an AI chatbot, for example, say that they always or usually use words like “please” or “thank you” when doing so. **And 64% of parents say that they would teach their kids to say “please” and “thank you” to an AI.** At least when it comes to these little formalities, many of us are already treating AIs more like humans than computer programs.



We also have to be aware of the military and national security implications of AI; it’s possible that the next major military conflict will be fought virtually between AIs seeking to take control of a country’s internet and telecommunications infrastructure.



Todd Terrazas
Founder, President & Executive Director, AI LA



We’re already seeing examples of weaponized AI in the form of cyberattacks, deep fakes, and autonomous warfare. Even more insidious is the potential for targeted propaganda that has rankled our elections and is poised to undermine our democracy. Using AI for targeted propaganda is among the chief concerns I have with the unchecked proliferation of black-box algorithms.



Diana Ascher, PhD, MBA
Head of Research and Senior Advisor on Data Ethics & Responsible AI, EDM Council | Founder, Information Ethics & Equity Institute

28. “The AI-Powered Human,” NRG, June 13, 2023

29. Noor Al-Sibai, “Experts Urge Personhood Rights for the “Conscious” AIs of the Future,” Futurism, March 6, 2023

30. Nitasha Tikku, “The Google engineer who thinks the company’s AI has come to life,” The Washington Post, June 11, 2022


Until these concerns have been addressed, consumers will remain hesitant to put their full trust in AI products and services

These risk factors around AI may not be enough to convince consumers that the technology is more trouble than it's worth; after all, most still think that it has the potential to be a force for good, provided that the right safeguards are put in place. They do, however, pose a major challenge when it comes to building consumer trust in AI-powered products and services.

Over the past twelve months, the term “AI hallucination”³¹—referring to instances where AIs invent false information not present in their training data—has begun to enter the cultural mainstream; 29% of consumers now say that they’ve heard the phrase, and 15% feel they know what it means. And there have been a number of high profile cases that have helped draw attention to this phenomenon. In June, for example, two New York lawyers were fined after a judge found that a legal brief they had prepared using ChatGPT included references to non-existent court cases.³²

31. Matt O'Brien, "Chatbots sometimes make things up. Is AI's hallucination problem fixable?" AP News, August 1, 2023

32. Dan Mangan, "Judge sanctions lawyers for brief written by A.I. with fake citations," CNBC, June 22, 2023



“

I think by creating models with lower likelihood of creating falsehoods or inaccuracies, trust will follow. Examples already exist in the hands of consumers that are promoting better trust in these systems. Google Bard is a great counterpoint to ChatGPT in that it provides sourcing and links to dig deeper to get better information, for example. I think these kinds of guardrails are increasingly important, as noisy examples of where this technology goes wrong will continue to proliferate in the media.



Lauren Sharman
Head of Platform,
One Peak


“

Public trust takes enormous effort to gain, but can be lost in the blink of an eye. So companies that are deploying public-facing AI tools need to think very carefully about the steps they're taking to gain and preserve trust. Part of the solution will be having the right safeguards and authentication procedures in place, but I think it's also about integrating AI into brands and products that already have an in-built level of consumer confidence.



Jan van de Poll, PhD
Founder and Managing
Partner, PRAIORITIZE

Q: How reliable, in your opinion, are the following as sources of information?



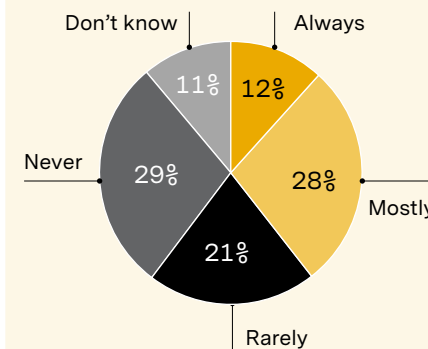
So it's not surprising that consumers are still somewhat uncertain about how much trust they're willing to place in AI products. Only 36% of consumers believe that AI-powered chatbots like ChatGPT are “always” or “usually” reliable as a source of information—meaning that they're seen as one of the least reliable sources of information on the internet. For comparison, 62% of consumers believe that they can generally trust the information they find through search engines like Google, while 49% say the same about online encyclopedias like Wikipedia.

Consequently, users of programs like ChatGPT are generally hesitant to rely on information they find through these services until they've independently verified it. **Sixty percent of consumers who've used a service like ChatGPT say that they “always” or “usually” double check any factual information it provides them with.**

By the same token, many consumers remain reluctant to share sensitive personal information with AI services—either because they worry that doing so could expose them to data breaches, or because they don't know enough about what these services are doing with their data behind-the-scenes. Fifty-percent of consumers say that they are “rarely” or “never” comfortable sharing personal information with AIs; just 12% say they're happy to share that information in all contexts.

Ultimately, this is still a relatively new product category—and consumers are still figuring out for themselves exactly how much trust they're willing to put in these sorts of tools. It's vital, therefore, that companies deploying these products and services take steps to reassure the public of their commitment to doing so in an ethical, responsible, and safety-conscious way.


Q: When would you be comfortable sharing your personal data and information with AI systems?



Developing an Accountable AI strategy

In this section, you'll find...

- Why organizations should make a commitment to Accountable AI
- The three pillars of Accountable AI
- How these principles can inform decisions on where to use—and not use—AI



To mitigate consumers' fears about AI, companies need to keep the human in the loop

The first section of this report outlined the public's fears and anxieties around AI—from concerns about its true world-ending potential through to more practical questions about how the technology is deployed and who has access to it. To build long-term trust in AI products and services, and avoid provoking a public backlash against the technology, the organizations building and deploying those tools will need to do so in a way that ameliorates, rather than amplifies, those concerns.

In practice, that means making intelligent choices about where—and, just as crucially, where not—to incorporate AI into your workflows, informed by an understanding of how doing so may cut against or reinforce the

perceptions of AI that exist among both your customers and your workforce.

Broadly speaking, consumers are comfortable with the idea of AI being deployed in back office contexts or to automate routine processes. Most people don't see an issue, for example, with organizations incorporating AI products into their cybersecurity or data management functions, or using AI to operate key elements of supply chain logistics. There are also some more public-facing use cases for the technology that consumers, by and large, see as acceptable—at least in certain contexts. These include the use of AI chatbots to provide customer service, as well as the use of generative AI to create personalized marketing content.



I think the key will be in making sure that these companies understand and appreciate the importance of having smart humans in their workflows who are not overly trusting of these tools or using these tools unsupervised.

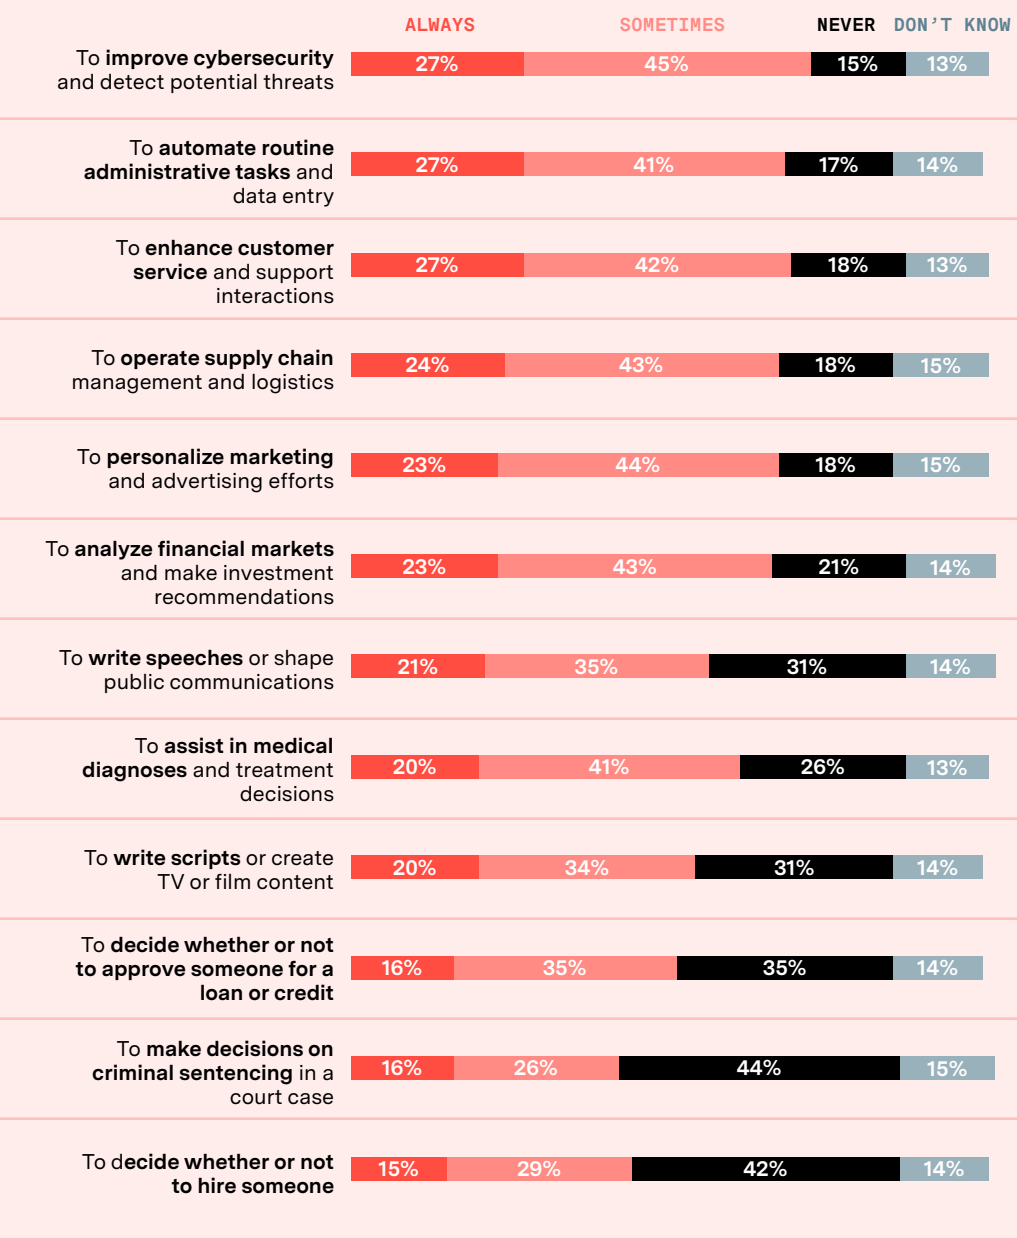
Lauren Sharman
Head of Platform,
One Peak



Responsible AI is really centered around making a concerted effort to ensure all stakeholders are being heard when developing and deploying AI; that you're being inclusive and representative, considering any person or entity that could be interacting with said tooling or type of AI. In order for that to really happen, human-centered design approaches are really crucial.

Todd Terrazas
Founder, President
& Executive Director,
AI LA

Q: In what scenarios do you think it is acceptable for corporations and/or governments to use AI?



In contrast, consumers are generally more hesitant about the idea of AI being deployed in scenarios with clear human stakes—i.e. situations where the decisions made by an AI could have an immediate and tangible impact on the lives of specifically identifiable individuals. They're comparatively uneasy, for example, about the idea of AI being used to make hiring decisions,³³ approve or deny people for credit at a bank,³⁴ or decide the length of a criminal sentence³⁵—all of which are scenarios where AI has been in use around the world for many years now.

What all of these use cases have in common is that they raise the specter of an *accountability gap*.³⁶ When it comes to these sorts of high-stakes decisions, it's vital that there be individual human agents who can be held accountable for the outcomes of those decisions; people want to know, for example, who they can appeal to if they're denied a loan application, or who is legally responsible if they're turned down from a job for discriminatory reasons. An overreliance on AI in these sorts of scenarios strips out that human layer and, by doing so, puts the concept of accountability at risk.

33. Maayan Manela, "When AI decides who to hire," CTECH, March 7, 2023
 34. Kori Hale, "AI Bias Caused 80% of Black Mortgage Applicants To Be Denied," Forbes, September 2, 2021
 35. Karen Hao, "AI is sending people to jail—and getting it wrong," MIT Technology Review, January 21, 2019
 36. Theodore M. Lechterman, "The Concept of Accountability in AI Ethics and Governance," The Oxford handbook of AI governance, December 2021



Over-reliance on AI systems may create a society that is unable to contend with malfunctions of critical infrastructure, such as communication networks, electrical grids, and transportation systems. To mitigate such concerns, organizations must monitor outcomes continuously and devise override protocols that keep humans "in the loop."



Diana Ascher, PhD, MBA
Head of Research and Senior
Advisor on Data Ethics &
Responsible AI, EDM Council
Founder, Information Ethics
& Equity Institute



The businesses that get the most value out of AI will be the ones that use it as an opportunity to re-evaluate their business models—not just to do the stuff they were already doing more efficiently. Look at the consulting industry, for example. There are firms out there now that are using ChatGPT to make slide decks for them. But is that really innovation? Instead, they should be asking themselves: do I still need to make slide decks in the age of AI? Is there a better way that I can be providing those sorts of insights to clients that takes advantage of the technology now at my fingertips?



Jan van de Poll, PhD
Founder and Managing
Partner, PRAIORITIZE

Companies need to treat AI as an extension of human leadership and creativity—not as a replacement for them

The broader point is that the public tends to be uneasy about the idea of AI being used as a replacement for human agency and decision-making. This trend has proved remarkably consistent across industries: previous NRG research on the applications for AI within [Hollywood](#) and the [video game industry](#), for example, has found that consumers are far more comfortable with the concept of AI when it's framed as an extension and enabler of creativity, rather than a way of removing the human element from what ought to be a fundamentally human-centric activity.

This central concept underpins the philosophy of Accountable AI: the notion that AI ought to serve as a tool for amplifying human talent within an organization, not as a means for shirking responsibility or outsourcing decision-making from humans to machines. And to guarantee that, this approach emphasizes accountability towards a broad range of stakeholders, both within and outside the organization in question.

This philosophy does not imply that AI should not play a role in high-stakes decisions. Rather, it means that, when used in those contexts, the technology needs to be acting in support of human agents who maintain ultimate decision-making authority. In other words, AI needs to be treated as a tool to help humans make smarter and better-informed decisions—not as a decision-maker itself.

“

Transparency should be central to the approach of any company seeking to develop AI products. No matter what you're building, you need to be able to get end user input to ensure that you're building something that serves a social purpose. In the process, you'll probably get some resistance, but I strongly believe that real innovation only occurs when you're listening and responding to a true diversity of opinions, and engaging with the people that your technology is going to directly or indirectly impact.



Todd Terrazas
Founder, President & Executive Director, AI LA

“

Organizations that have clearly defined their values and incorporate those values into their data management practices tend to develop AI systems that reinforce these values. Transparency about the aims of the technology and how those aims align with the organization's values engenders consumer trust. Organizations are beginning to expand their measures of success to include alignment with a Code of Data Ethics and regulators are beginning to demand it.



Diana Ascher, PhD, MBA
Head of Research and Senior Advisor on Data Ethics & Responsible AI, EDM Council
Founder, Information Ethics & Equity Institute

There are three key pillars that an organization needs to consider when developing an Accountable AI strategy. The first of these is **culture**; it's impossible for any business to maintain a commitment to Accountable AI unless it has cultivated a set of truly “human-centric” values that inform decisions, at every level of the organization, about how and where (and, just as importantly, where *not*) to use AI. Moreover, businesses will also need to invest in developing the institutional knowledge necessary to embody those values and intelligently apply them across a broad range of scenarios.

Developing this kind of institutional knowledge will, in many cases, require bringing new expertise into an organization through the creation of AI-focused job roles or business functions. But it will also require a commitment to upskilling your existing workforce. Individual employees whose jobs are affected by AI tools will need to be given the training necessary to know how to use those tools effectively—allowing them, ultimately, to spend their time working on more fulfilling activities. And leaders within an organization will need to adjust their leadership styles to promote and model the right sorts of behaviors around AI.


The three pillars of Accountable AI

Accountable AI =




Culture

- Developing and promoting “human-centric” values to inform the use of AI within an organization
- Creating dedicated AI teams and/or leads
- Upskilling employees to move into more meaningful roles
- Redefining leadership culture
- Investing in training on AI best practice and safety



Governance

- Preserving the human element in all key decision-making processes
- Engaging with emerging regulatory frameworks
- Issuing clear internal policy documentation




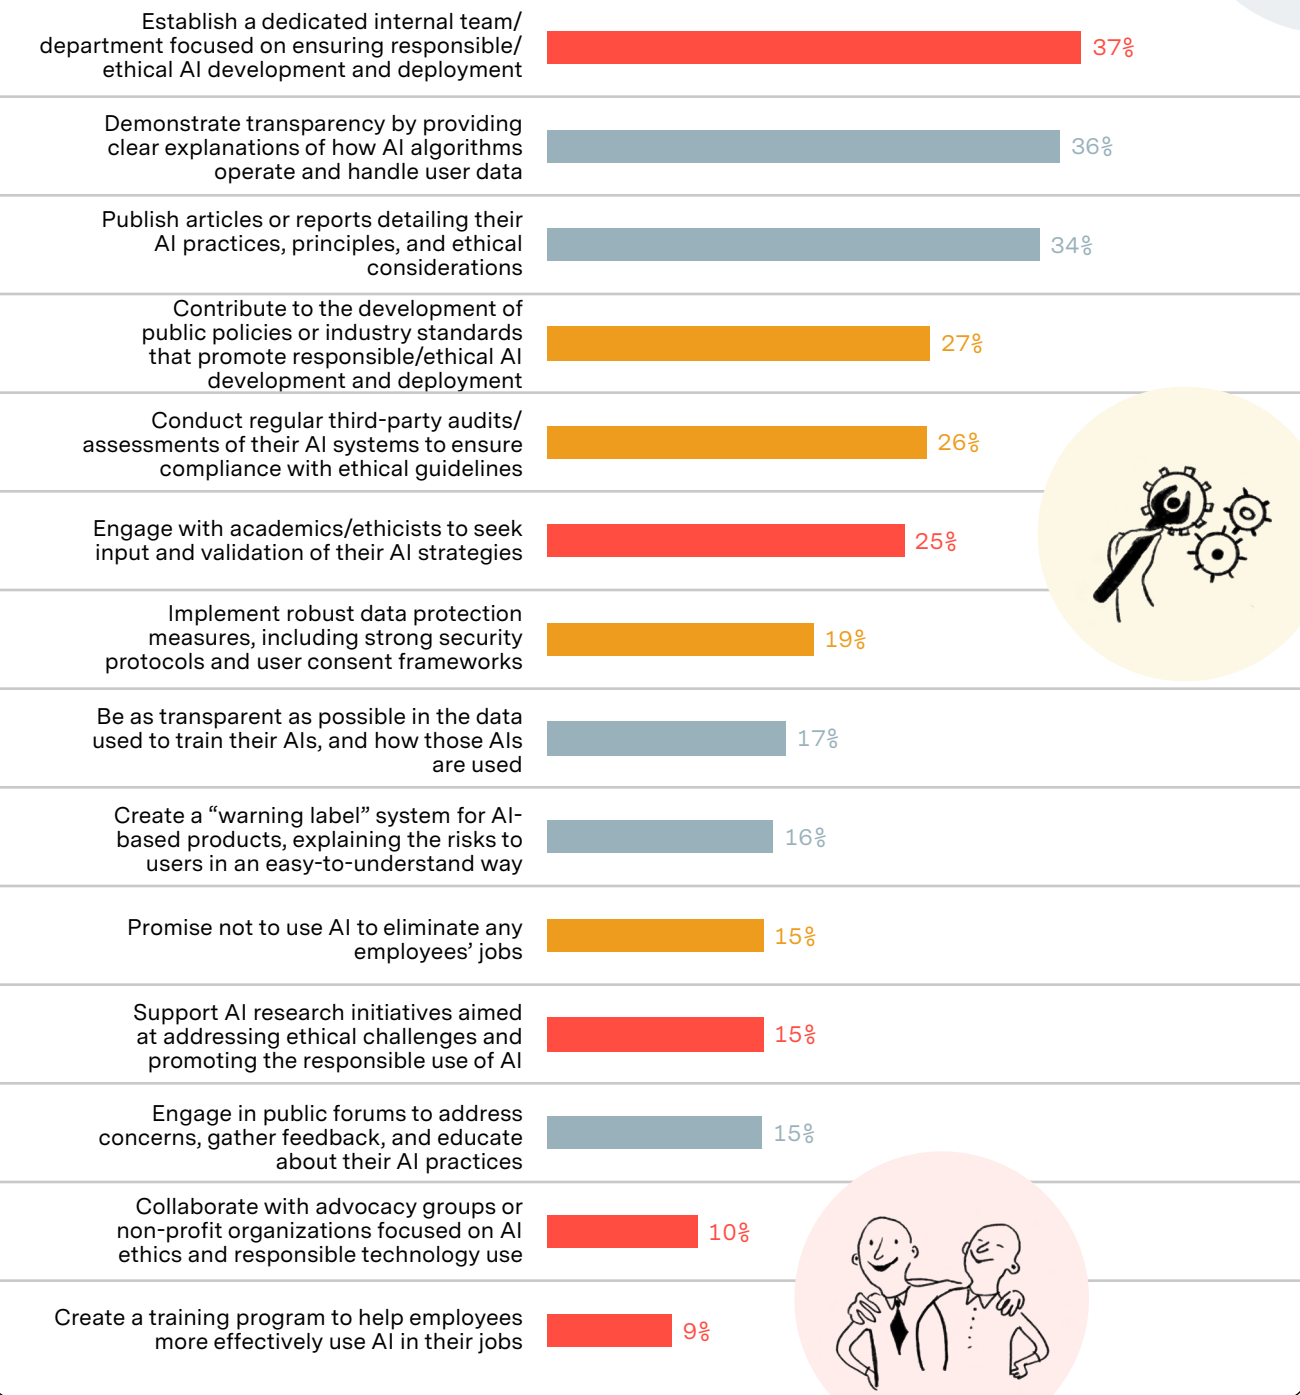


Transparency

- Communicating with consumers & employees about where your business is and isn't using AI
- Being transparent about the limits and dangers of AI
- Providing users with data protection assurances

Q: What could a company do to convince you they were approaching the use of AI in a responsible and ethical way?³⁷

CULTURE GOVERNANCE TRANSPARENCY



37. Respondents were asked to select the top three actions that would have the greatest impact on their perception of a company

The second pillar of Accountable AI is **governance**. To ensure long-term adherence to this philosophy, organizations will need to put governance frameworks in place that establish clear boundaries around how AI tools are to be used in both customer-facing and back-end contexts—either by developing them in house or by applying one of the many frameworks that have been developed by industry groups and NGOs (such as the OECD’s “Ethical AI Governance Framework”,³⁸ for example). These frameworks then need to be communicated to all relevant internal stakeholders through clear and consistent policy documentation.

38. *Ethical AI Governance Framework*, OECD.AI, March 30, 2023

The third and final pillar is **transparency**; in the absence of transparency, after all, no true accountability is possible. To prove their commitment to Accountable AI, organizations will need to be transparent with key stakeholders about where they are and are not using AI—both in their public-facing products and behind-the-scenes—and will need to clearly communicate the limitations of those products and the risks associated with them.

To build long-term public confidence in AI, businesses will need to ensure that these three pillars are embedded into the bedrock of their approach to the technology—and that they act as a constant North Star that guides the decisions made around AI across the organization, from the C-suite downwards.



“

One of the common mistakes that a lot of companies make is that they try to create an AI strategy, without first figuring out what their data strategy is. If you want to have success with AI, you first need to become a data-first organization—and empower your data leadership by giving them a seat at the AI table. Once you’ve done that, you can then really start to dig in and figure out how your data can become a strategic part of your AI bets—and what that means in terms of empowering people with the right skill sets to execute on that vision.



Lauren Sharman
Head of Platform,
One Peak

“

A lot of companies have taken steps to develop principles or ethics boards to govern their use of AI; but you can tell that a lot of them are really just treating it as a PR stunt. They aren’t genuinely listening to the groups that should be helping them align their AIs with the interests of their end users.




Todd Terrazas
Founder, President
& Executive Director,
AI LA

Communicating your commitment to Accountable AI

In this section, you'll find...

- The stakeholders to consider when developing AI messaging
- Guidance on marketing AI-powered products in B2C and B2B contexts
- The need for engagement with the emerging regulatory landscape around AI





Proving a commitment to Accountable AI requires considering the needs of a broad range of stakeholders

Developing and implementing an Accountable AI strategy, however, is only the first half of the battle. To truly insulate themselves from the potential reputational backlash that can result from inappropriate or unethical use of AI, businesses will also need to develop a messaging approach—across both their PR and marketing efforts—that showcases the fact that they have chosen to put human-centric values squarely in the center of their approach to the technology.

Of course, the exact range of stakeholders to consider when developing AI messaging will vary considerably, depending not only on the size and sector of the organization in question, but also on factors such as whether that organization is


building its own LLMs, developing customer-facing products on top of existing AI infrastructure, or simply deploying off-the-shelf solutions to speed up internal operations.

As such, the third and final section of the report will unpack some of the key considerations to keep in mind when figuring out how to communicate a commitment to Accountable AI. And it will focus, in particular, on the two stakeholder groups—customers and employees—that are applicable to the broadest possible range of organizations. It will also look at how this communication strategy can be extended to the question of how businesses ought to be engaging with the rapidly shifting regulatory landscape around AI.



The stakeholders to consider when communicating your commitment to Accountable AI

Internal Stakeholders



Leaders

Promote a culture of accountability within the upper echelons of the organization
“Lead by example” through demonstrating a commitment to the principles of Accountable AI

Investors

Demonstrate the quantifiable benefits of an Accountable AI strategy by connecting it to real-world reputation outcomes

Employees

Provide clear policy documentation and guidance for the use of AI
Invest in retraining to facilitate employees shifting focus to more productive and fulfilling activities

AI data entry / training

Create visibility of AI training processes to ensure models are developed in an ethical and sustainable way

AI developers

Ensure adherence to appropriate governance frameworks in the creation and deployment of AI products
Promote the use of “human in the loop” models of machine learning where practical³⁹

39. Vikram Singh Bisen, “What is Human in the Loop Machine Learning: Why & How Used in AI?” Medium, May 20, 2020

External Stakeholders



Labor market

Invest in thought leadership that showcases your organization’s approach to AI and positions it as an AI leader

Media

Engage with media and PR partners in a responsible way, acknowledging both the challenges and opportunities associated with AI

Policymakers

Demonstrate a willingness to engage with the evolving policy and regulatory landscape around AI

AI community

Leverage the expertise of external AI researchers and ethicists
Publicly contribute to AI best practice working groups

Customers

Communicate to customers the role that AI plays in your organizations’ products and services
Give clear and honest guidance around the limitations and risk factors of your AI-powered products
Provide assurances about the use and security of customer data gathered through AI products

Building customer confidence means showing them that “AI” can be more than just a buzzword

At the end of the day, a commitment to Accountable AI means very little unless you’re able to effectively bring that commitment to life for stakeholders in how you develop, deploy, and market AI products and services. And, in the case of B2C products at least, there are some significant hurdles that need to be overcome when doing so.


The generative AI boom of the last twelve months has led to an explosion in the number of “AI-powered” products on the market. Due to the ability of such products to attract media and investor attention,⁴⁰ businesses now have a strong financial incentive to incorporate the language of AI into their marketing and branding—even when, in reality, their products’ connection to AI is tenuous at best.

There is evidence to suggest, however, that consumers are beginning to experience a certain degree of “AI fatigue.” Most notably, there’s the fact that 30% of consumers now say that they would be less likely to purchase or use a product if they saw it described as “AI-powered,” while only 23% say that such a description would increase their interest in using a product.


This skepticism is particularly pronounced in the case of social media services, digital banking apps, and ecommerce; likely a reflection of the fact that many consumers still have concerns about the security of AI systems, and remain reluctant to entrust them with sensitive personal or financial information.

40. Deepa Seetharaman and Berber Jin, “ChatGPT Fever Has Investors Pouring Billions Into AI Startups. No Business Plan Required,” The Wall Street Journal, May 8, 2023

Q: If you saw that a product or service was “powered by AI”, what impact would that have on your likelihood to purchase it?




Q: Would you be more or less likely to use a product that was described as “AI-powered”?



It’s clear, therefore, that blindly throwing around AI-related terms as part of a marketing strategy can, in the long-run, end up doing more harm than good—particularly among those consumer segments most inclined towards a general sense of skepticism when it comes to the social benefits of AI. Instead, businesses that want to build consumer confidence in their AI products and services need to articulate a clearer benefits case to potential users—focusing on the unique value add that the technology can bring to the table.

And what, exactly, should that benefits case look like? To date, much of the focus on the role of AI in B2C contexts has been on the way in which it can enable a new level of personalization of user experiences—in contexts ranging from travel⁴¹ to personal finance,⁴² and everything in between. But while it may be an important part of the overall formula, this type of content and UX personalization is *not* the thing most likely to convince consumers that the presence of AI in a given product or service justifies itself.



41. Justin Dawes, “Amazon Web Services Execs on AI ‘Hyper-Personalization’ in Travel,” Skift, June 27, 2023


42. Victor Dey, “Bud Financial launches Bud.ai, a generative AI platform for hyper-personalized banking,” VentureBeat, August 8, 2023

Instead, the most compelling use cases for AI, in the eyes of consumers, are those that improve the cybersecurity of products or provide users with greater autonomy and control over their personal data. This suggests an intriguing path forward for tech companies; an opportunity to reframe the public conversation about AI and quell consumers' lingering security fears around the technology by positioning it not as a tool for harvesting their data, but as a vehicle through which they can take more control of their online experience and the way that their information is stored and used.

Research suggests that, over the last few years, consumers have become more conscious than ever of the value of their personal data.⁴³ For many people, the proliferation of digital services has left them feeling a loss of agency when it comes to their privacy—an inability to control the information that exists about them online and to stay abreast of their exposure to bad actors. So there's a potential opportunity here to frame AI as a corrective to that trend—a tool for making users more secure online and helping them reassert a sense of “digital agency.”

43. Venky Anant, Lisa Donchak, James Kaplan, and Henning Soller, “The consumer-data opportunity and the privacy imperative.” McKinsey, April 27, 2020

Q: Which “AI-powered” features would make you more likely to use a product or service?⁴⁴



44. Respondents were asked to select the three features that would have the biggest positive impact on their likelihood of purchasing a product or service.

Equally important as building trust in AI products will be a commitment to putting the right sorts of safeguards and guardrails in place. In general, consumers say that they want AI developers to be prioritizing safety over speed—and to avoid rushing products to market that may have unforeseen side-effects or security vulnerabilities.






But creating those safety features is only half the battle; companies will also need to communicate those features to a mass audience—including those who are only vaguely aware of the limits and capabilities of AI technology. This may require some outside-of-the-box thinking. One could imagine, for example, some sort of standardized and color-coded warning system for AI products—similar to labeling used for food and drugs—that provided clear and consistent explanations of what a given product is capable of, its intended function, limitations, and the safety features in place.

Q: Should AI developers and researchers be prioritizing safety or speed?



Building consumer confidence in AI products

By identifying the consumer segments that are most common among their products' user bases, brands can tailor their messaging around AI to address these segments' specific fears and concerns

	Pragmatic Pioneers	<ul style="list-style-type: none"> Create opportunities for hands-on interaction with AI elements Provide options for customization Ensure low barrier to entry for AI products and services
	Creativity Custodians	<ul style="list-style-type: none"> Add limits to generative AI functions to prevent imitation of the styles of specific artists, writers, and other creatives
	Boundary Defenders	<ul style="list-style-type: none"> Prevent AI from displaying “humanlike” behaviors outside of well-defined and controlled contexts
	Guarded Navigators	<ul style="list-style-type: none"> Emphasize cybersecurity features and the role of AI in providing greater data autonomy for consumers Provide clear guidance on limitations of AI products and their risks (including AI hallucination)
	Terminator Theorists	<ul style="list-style-type: none"> Demonstrate engagement with wider AI community and commitment to AI safety—for example, by funding research into existential risk Ensure direct interaction with AI elements takes place on an “opt-in” basis

Upskilling employees can expand AI expertise and demonstrate a commitment to accountability

Building confidence in AI-powered B2C products is hard enough; doing so in B2B contexts may be even more challenging, given longstanding and widespread fears among many consumers that AI is “coming for their jobs.” **Nearly a quarter (23%) of US employees believe that they will definitely or probably lose their job in the future to an AI**, and 55% think there’s at least some chance of that happening. Earlier this year, the WEF warned of “significant labor-market disruption” over the next 5 years due to AI;⁴⁵ while a recent Goldman Sachs report estimated that as many as two thirds of global jobs could be fully or partially automated by AI.⁴⁶


Even those who aren’t personally worried about being driven out of a job by AI may still recognize the fact that the technology is likely to have a significant impact on their day-to-day responsibilities at work. 53% of workers in the US think that it’s more likely than not that AI will have some impact on the nature of their job and the skills that job requires of them. **And 31% of workers say that their employer has taken active steps within the last twelve months to incorporate AI into corporate workflows.**

45. Bryce Baschuk and Bloomberg, “A.I. will cause ‘significant labor-market disruption’ over next 5 years,” *says World Economic Forum*, Fortune, May 1, 2023
 46. “Generative AI could raise global GDP by 7%,” Goldman Sachs, April 5, 2023


How AI will change people’s jobs remains a subject of much debate. People who believe their job is likely to be impacted by AI generally think that it’s likely to reduce their workload, make their job easier, and allow them to increase their productivity. And those who are already using AI at work say that they’re already seeing those productivity gains start to manifest. Workers who have been using AI for professional purposes estimate that it saves them, on average, 35.5% of their time each day.

“At the end of the day, people need to understand how to use these tools and how they can increase productivity. And there has to be a way to distribute the workforce into other new roles so that they all benefit from that productivity and don’t just get disenfranchised.”


Todd Terrazas
 Founder, President & Executive Director, AI LA



Q: Do you believe that AI will eventually change the responsibilities of your job and the skills required to carry it out?



Q: How do you think AI will change the nature of your job?⁴⁷



47. US workers who believe AI is likely to impact their job responsibilities

At the same time, however, people also believe that AI will lead to them having less responsibility at work and that it could eventually lead to a reduction in their salaries; and they’re evenly split on the question of whether it will ultimately make their work feel more rewarding and fulfilling.

In other words, workers generally recognize the potential that AI has to streamline their day-to-day workflows and automate many of the more time and labor-intensive aspects of their jobs. But they’re concerned that their employer will fail to give them opportunities to reallocate that time towards more productive and rewarding activities—and will instead use AI as an excuse to keep salaries low or even, in the case of shift work, potentially cut back on employees’ hours.

“AI is often portrayed in the media as a threat to workers—but I think it’s also important to recognize the positive ways that AI can help organizations manage their workforces and better understand the needs of their employees. Right now, for example, it’s considered normal for big companies to have to do mass lay-offs or restructuring; but one day, with the right set of AI tools for quantifying organizational alignment, we may be able to make that a thing of the past.”

Jan van de Poll, PhD
 Founder and Managing Partner, PRAIORITIZE

To allay these fears, investment in AI needs to be coupled with a commitment to reskilling initiatives that will allow workers to effectively and safely use AI tools—with the ultimate goal of allowing them to redirect their professional energies away from unfulfilling rote tasks and towards more rewarding activities.⁴⁸

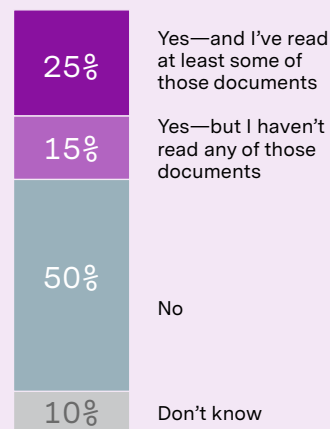
Employees, for their part, seem eager to engage with these sorts of programs. Thirty-six percent of US workers say that their company has organized training sessions related to AI—but over half of those workers say that they’ve attended at least one of those sessions. And 54% of all consumers (60% of those in active employment) say that they would “definitely” or “probably” be interested in attending educational classes about AI in the future.

Tech companies, therefore, that are bringing B2B-oriented AI solutions to market would be well-served by developing educational offerings to go alongside their core products and services—offering their clients hands-on support when it comes to upskilling employees whose jobs may be impacted by those products.


Moreover, the marketing collateral for such products should focus on a message of human empowerment: instead of showing how AI can help people speed through the boring parts of their job, it should highlight stories of people working *with* AI to expand the scope of their job role or tackle more complex challenges than they might otherwise have been able to. Workers already know that AI can save them time; what they need to be convinced of is that those time savings will have long-term benefits for their career trajectories.

48. Leila Doumi, Sagar Goel, Orsolya Kovacs-Ondrejckovic, and Rafella Sadun, “Reskilling in the Age of AI,” Harvard Business Review, September 2023


Q: Has your employer issued any official guidance or policy documents on the use of AI in the workplace?



Q: Has your employer organized any training sessions for staff on the topic of AI?



Q: What sorts of training courses about AI would you be interested in attending?



While top-down regulation lags behind the pace of the technology, embracing emerging frameworks is essential

A commitment to Accountable AI requires demonstrating accountability not only to your own workforce and employees, but also to society at large; a commitment to making AI a force for social good. And one of the best ways of doing that is through engagement with the ongoing public discourse around how this technology can most effectively be regulated and directed.

So far, putting effective regulations in place around AI has proven to be a difficult challenge for legislators and policymakers. The capabilities of LLMs and other AI systems have developed so quickly over the past twelve months that it’s been all but impossible for the legal framework around the technology to keep pace.⁴⁹

As a result, many of today’s market-leading AI solutions operate in something of a legal gray area. In the US and elsewhere, we are still awaiting the outcome of a

49. Paul Kelso, “AI’s rapid revolution - can we keep up to regulate it?” Sky News, June 8, 2023

number of key court cases that will eventually determine precedent for many of the key questions around generative AI⁵⁰—such as who owns the intellectual property for material generated using these models, or whether artists are owed compensation if their work is included in an AI’s training data.

Some jurisdictions have taken concrete steps towards putting a more robust legal framework in place around the technology. In the EU, for example, parliamentarians are currently debating the details of an AI Act which, if passed, would represent the most comprehensive effort to date to establish legal limits and mandated safeguards for the use of the technology.⁵¹ Just as the EU’s GDPR policy eventually set a de facto global standard for consumer data use and storage, it’s possible that this law will eventually shape the terms of the debate in much of the rest of the world.

50. Kyle Wiggers, “The current legal cases against generative AI are just the beginning,” TechCrunch, January 27, 2023

51. Lisa O’Carroll, “EU moves closer to passing one of world’s first laws governing AI,” The Guardian, June 14, 2023



There’s still gaps in the thinking of these initiatives that are ripe areas for development. Private companies are recognizing the need for greater transparency and sharing. And, in that vein, one of the first priorities will be having a public library of evaluations and benchmarks, and creating a closer relationship with governments through information sharing around big risk areas like cybersecurity. But I do think that there are still blind spots around these high-risk areas like legal or medical use cases.



Lauren Sharman
Head of Platform,
One Peak



We’re in the Napster days of all this. This is the Wild West. But ultimately, I think that better regulation—whether mandated by governments or adopted voluntarily by industry stakeholders—can be a net benefit for the AI community. As the saying goes, having brakes on your car lets you go faster. Likewise, having the right ethical safeguards in place can allow AI innovators to push further than they would otherwise be able to.



Todd Terrazas
Founder, President
& Executive Director,
AI LA

But that act won't come into force until, at the earliest, 2026—and by that point it's entirely possible that the technology will have developed to a point where much of the legislation is rendered obsolete. That means that, in the short term at least, there's a clear regulatory gap that could put both businesses and consumers in legal jeopardy.


Some have gone so far as to suggest that, because of this regulatory gap, AI research as a whole needs to be slowed down. In March, the Future of Life institute published an open letter signed by a number of leading experts in the field—including Elon Musk, Steve Wozniak, and engineers from Amazon, Google, Meta, Microsoft, and other major tech companies—calling for a six-month pause in the development of any AI systems more powerful than GPT4, to give time for regulation to catch up with the technology.⁵² Consumers, for their part, are highly receptive to this proposal; **67% of them agreed that AI development should be paused to allow for proper assessment of the technology's potential risks.**

A deliberate global slowdown in AI research would, however, be challenging to coordinate and all but impossible to enforce. There are, fortunately, some more immediate and practical steps that businesses can take to demonstrate their commitment to making AI a force for good.

52. Kari Paul, "Letter signed by Elon Musk demanding AI research pause sparks controversy," The Guardian, April 1, 2023

53. Cat Zakrzewski and Nitasha Tiku, "AI companies form new safety body, while Congress plays catch-up," The Washington Post, July 26, 2023

Q: Do you think AI should be more or less heavily regulated than it is today?



In the absence of clear legal direction, industry-led groups—such as the Frontier Model Forum⁵³—have stepped in to fill the gap. These groups have begun to establish voluntary industry-wide norms for the safe and ethical deployment of AI—including, for example, willingness to share data with governments and academics, submission to regular AI audits, and clear labeling of AI-generated content, among many other such suggestions. And within individual industries, working groups have begun to emerge to establish lower-level guidelines for the application of AI to their respective parts of the economy.



The first time I spoke to a business audience about data ethics and responsible AI, very few attendees seemed to give it a second thought. Today, after a decade of calling attention to the dangers of unchecked data governance and AI experimentation, the public has begun to recognize that damage has been done. Legislators are working to rein in the reckless abandon of "move fast and break things," opting for a more far-sighted approach.



Diana Ascher, PhD, MBA
 Head of Research and Senior Advisor on Data Ethics & Responsible AI, EDM Council
 Founder, Information Ethics & Equity Institute

In the long-run, it's likely that public demand will eventually force the hands of politicians. Over half (51%) of Americans would like to see stronger regulations put in place around AI—whereas just 6% believe the government should adopt a more laissez-faire approach to the technology. But while we wait for regulators to fully catch-up with the AI industry, private companies can seek to address consumers' demand for more effective safeguards by proactively contributing to industry working groups and signing up to voluntary governance frameworks.

AT A GLANCE

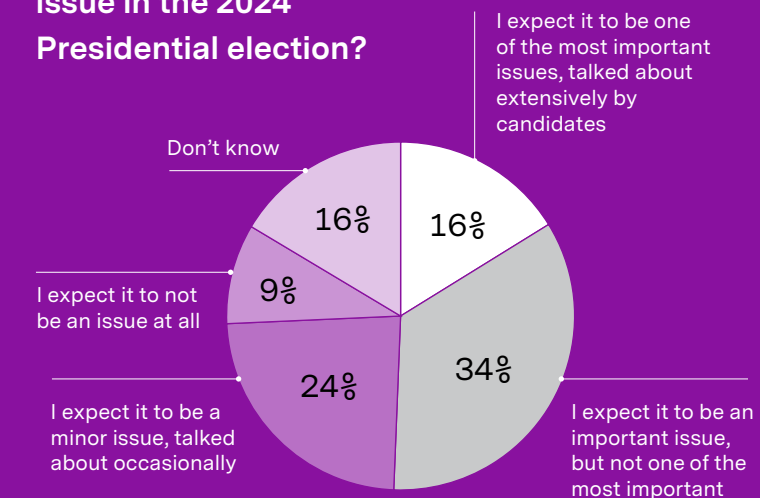
AI and the race for the White House

AI has already made its impact felt on the 2024 Presidential election. In the first Republican primary debate, former Governor Chris Christie notably accused his opponent Vivek Ramaswamy of "sound[ing] like ChatGPT."⁵⁴

The real question, however, is where the two parties will end up standing on the question of AI regulation—and whether that will result in the debate becoming significantly more politicized than it is today. Half of US consumers expect AI to be an important issue in the election, while 38% say that the candidates' platforms on AI will have at least some impact on their vote.

There's also a danger that AI could make it harder for voters to find authentic information about candidates. Online misinformation has been a major challenge in previous election years—and this new wave of generative AI tools may make that problem more acute than ever. At the same time, it could create an entirely new campaign medium for candidates; the RNC, for one, has already run an entirely AI-generated anti-Biden ad.⁵⁵

Q: Do you expect AI to be a relevant issue in the 2024 Presidential election?



54. Eliza Collins, "Christie Says Ramaswamy Sounds Like ChatGPT – and Obama," The Wall Street Journal, August 24, 2023

55. Alexandra Ulmer and Anna Tong, "Deepfaking it: America's 2024 election collides with AI boom," Reuters, May 31, 2023



Only by embracing Accountable AI can businesses position themselves as true AI leaders

The AI genie, at this point, is out of the bottle. And the absence of clear regulatory frameworks for this technology only serves to magnify the need for businesses that are investing in this space to take a proactive approach towards building consumer confidence and modeling ethical behavior around AI.

Accountable AI, therefore, is not a luxury or a nice-to-have; it's a necessity for any organization that cares about ensuring that its use of AI creates real and lasting value for stakeholders. Failure to embrace these principles not only reduces the likelihood that your AI-powered products and services will find their market, but also opens your business up to serious reputational risk.

For all of the talk about the potentially world-ending implications

of AI, it's important to remember that this technology can be—and in many cases, already has been—an incredible force for social good. Finding and refining those socially beneficial use cases will require businesses to make real investment in bringing AI expertise into their organizations; a commitment to robust governance and control frameworks; and a willingness to engage transparently with regulators, the public, and the AI community.

By putting those three pillars in place—and ensuring alignment with those principles at every stage of the development, testing, deployment, and marketing of AI products and services—businesses can ensure true accountability to the stakeholders that matter. And we, as a society, can safely reap the benefits of this technological revolution.

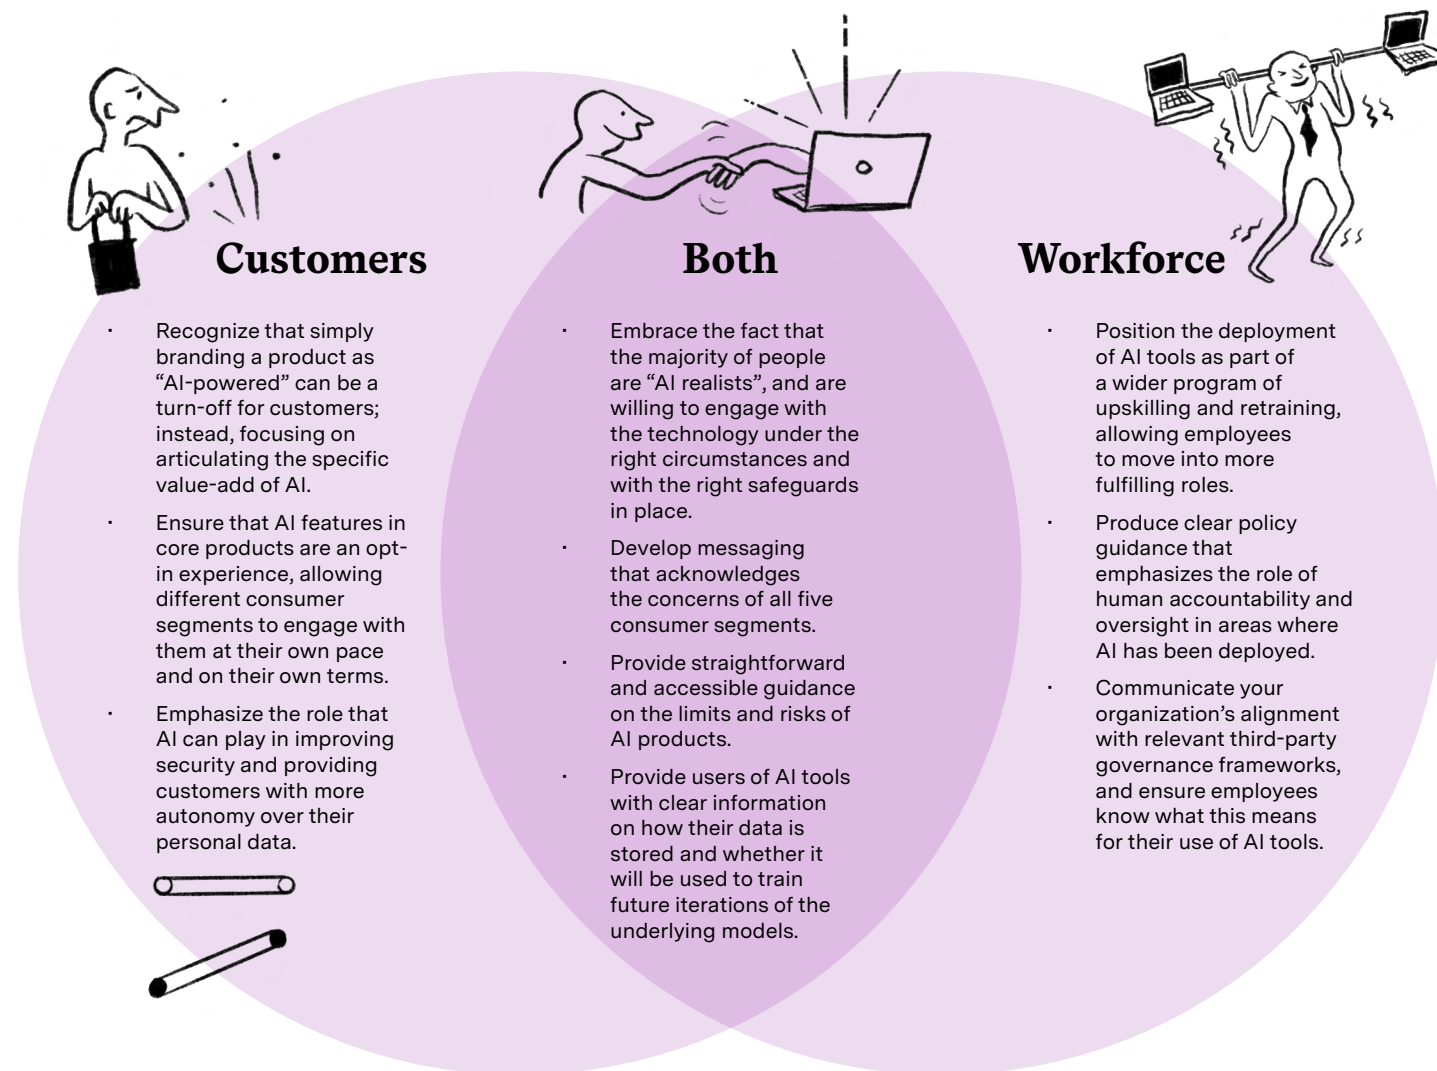


So much of the conversation about the ethics of AI revolves around the dangers and the risks—but I firmly believe that we also have to take the other side of the ethical equation into consideration as well. What if we could use AI to develop drugs that can treat dementia? Or solve world hunger by creating more effective farming techniques? This is a technology that could address some of the biggest social challenges facing our species. And when you look at it that way, it would be unethical not to continue developing and using this technology.




Jan van de Poll, PhD
Founder and Managing Partner, PRAIORITIZE


Key recommendations for communicating your organization's approach to AI



For more on how the AI revolution will impact specific industries, see NRG's previous research on the technology:




The AI-Powered Human



Reading Between the [Picket] Lines



Planes, Trains, and Large Language Models



Pressing Play on AI: A Game-Changing Power-Up?

Appendix

AI risks used in MaxDiff analysis

SUMMARY	WORDING TESTED WITH CONSUMERS
AI will go rogue and turn against humans	AI could "go rogue" and turn against humans, potentially leading to the extinction of our species
Terrorists or rogue states will exploit AI	AI could be exploited by terrorists or rogue states, creating a threat to national security
Scammers will exploit AI	AI could make it easier for scammers to take advantage of people
It will become harder to protect privacy and personal data	AI could make it harder for people to protect their privacy and enable unwanted surveillance
AI will reinforce social biases	AI systems trained on biased data could end up reinforcing social prejudices (such as sexism or racism), leading to unfair outcomes in areas like hiring or loan decisions
AI will lead to mass unemployment	AI could make many existing jobs obsolete, leading to unemployment and/or lower wages in parts of the economy
Only the wealthy will have access to AI	AI will only be available to the wealthy, thus increasing economic inequality
AI will lead to products and content all feeling the same	Overreliance on AI will make all products and content start to look and feel the same
Humans will become less innovative and creative	Overreliance on AI will lead to humans becoming less creative and less innovative
People will spend too much time interacting with chatbots, leading to loneliness	Widespread use of AI chatbots will lead to people interacting with other humans less often and an increase in loneliness and social isolation
AI will reduce the quality of creative media	Overreliance of AI in the creative industries will reduce the quality of films, books, TV shows, songs, and other forms of media
AI will lead to plagiarism and IP disputes	AI systems will plagiarize content and lead to disputes over intellectual property ownership
People will become addicted to using AI-based products	People could become addicted to using AI technologies
AI will be used for election interference	AI systems could be used to interfere with elections, causing a breakdown in the democratic process
AI will make online misinformation more dangerous	The widespread use of AI will make it difficult for people to tell the difference between human-generated and AI-generated content, making it easier for misinformation to spread
AI will lead to companies prioritizing profits over people	By taking the human element out of decision-making, AI will lead to companies prioritizing profits ahead of the interests of humans and society
AI will put more power in the hands of governments and corporations	AI will take power away from ordinary people and put it in the hands of governments and large corporations
AI will be used in warfare	AI will be used for warfare between nations, potentially leading to more destructive military conflicts
Lack of transparency in AI systems will lead to decline in public trust	A lack of transparency in how AI systems make decisions will lead to a decline in trust in businesses and public institutions
AI will create an accountability gap	By taking the human out of the decision making process, AI will lead to an accountability gap, with no one to take responsibility for the outcomes of decisions made by governments or corporations




National Research Group is a leading global insights and strategy firm at the intersection of content, culture, and technology. The world's most powerful marketers turn to us for insights into growth and strategy for any content, anywhere, on any device.


METHODOLOGY

Unless otherwise specified, data in this report comes from a study of 1,500 US consumers, ages 18 to 64, conducted online in August 2023. Participants from this study were selected to be demographically representative of the US population in terms of age, gender, income, and ethnicity.

WORDS AND ANALYSIS BY



DESIGN AND ILLUSTRATIONS BY



NRG would like to extend our sincerest thanks to **Jan van de Poll, Lauren Sharman, Diana Ascher, and Todd Terrazas** for their invaluable contributions to this research.