

THE **Better  
Together**  
AGENCY



# Building Consumer Trust: Addressing Bias in Generative AI

**Better Together Agency's 2nd Annual Biases  
in Generative AI Study**

# Forward: Inclusive AI is Imperative

We are at a pivotal moment. Generative AI is transforming industries, but it also stands at a crossroads. Without deliberate action, these powerful tools risk entrenching the very biases and inequities we have fought to overcome. A McKinsey analysis warns that generative AI could widen the racial wealth gap by **\$43 billion** a year if left unchecked. This is a technology issue and a challenge to economic and social equity.

We're already seeing how bias shows up in generative AI systems used for **hiring, healthcare, finance, and education** hurts those who can least afford it. From algorithms that unfairly filter out qualified job applicants from underrepresented groups to generative AI-driven diagnostic tools that under-serve communities of color, biased generative AI is not abstract – it has actual human costs. We cannot allow the tools of tomorrow to carry the injustices of the past.

It doesn't have to be this way. We know that when generative AI is developed **inclusively** with diverse perspectives as a core design principle, the outcomes improve for everyone. Inclusive generative AI leads to products and services that better serve a broad customer base, driving trust and loyalty. Studies show that companies in the top quartile for racial and ethnic diversity are **35% more likely** to financially outperform their industry medians, and diverse teams see nearly a **20% increase** in innovation revenue. Addressing bias in generative AI is the right thing to do and is a driver of innovation and growth.

This is why I am so encouraged by the work of **Better Together Agency**. As a purpose-driven communications partner, Better Together Agency helps industry leaders confront bias in generative AI head-on and guides them toward action. The "Better Together Agency 2025 Biases in Generative AI Impact Report" is a crucial part of that effort. It highlights where and how bias infiltrates generative AI systems, and focuses on the path forward, which includes how awareness, collaboration, and responsible design can help build **generative AI that works for everyone**.

The **Global Black Economic Forum** knows that **economic rights are human rights** and that closing economic gaps is about righting historical wrongs and choosing a more equitable and dynamic future for all. The push for unbiased, inclusive generative AI is a direct extension of that belief: By ensuring everyone can benefit from technological advances, we create stronger, more dynamic economies.

I invite every corporate leader, policymaker, and innovator who reads this report to join us in making inclusive, generative AI a top priority. We will make technological progress defined not by perpetuating bias, but by expanding opportunity.

The time to act is now.

**Alphonso David**

President and CEO, **Global Black Economic Forum**

# Executive Summary: Key Findings

Brands face a pivotal trust challenge and opportunity: ensuring their generative AI systems are fair and unbiased. The Better Together Agency 2nd Annual Biases in Generative AI Impact Report, "Building Consumer Trust: Addressing Bias in Generative AI for Brands" (2025), reveals consumers are aware of generative AI bias, which impacts their loyalty and buying decisions. For corporate decision-makers and startups, addressing generative AI bias is an ethical move and a strategic business imperative for growth, customer retention, and competitive edge.

## Key insights include:

1

### Consumer Expectations and Trust

83% of surveyed consumers have used generative AI tools, and a majority expect companies to actively ensure generative AI fairness. Fairness and lack of bias ranked as the **#2 factor (just after accuracy)** when choosing generative AI-driven products. **Consumers trust brands more when they know generative AI is designed to be fair and inclusive**, and one in three might abandon a product if its generative AI is found to be biased.

2

### Business Impact

Bias in generative AI directly hits the bottom line. Nearly **25% of consumers say they're more likely to support brands addressing generative AI bias**, while 33% would consider stopping usage if a generative AI tool is biased – a clear risk to revenue. Bias risks are most concerning in high-stakes industries like healthcare (56% of respondents), education (50%), and finance (45%), where trust and accuracy are paramount. Unchecked generative AI bias can lead to *missed market opportunities, PR crises, or legal action*.

3

### Competitive Advantage of Responsible AI

Embracing inclusive generative AI is a competitive differentiator. Business leaders must view responsible generative AI as *"not only about risk, but a value creator and competitive advantage grounded in trust."* Brands proactively mitigating bias can capture new customer segments, strengthen loyalty, and build goodwill. As one PwC expert noted, companies now see responsible generative AI as a way to **"ground services on trust."**

4

### Actionable Roadmap

The report provides a practical roadmap for organizations to turn these insights into action. Real-world case studies illustrate how bias-aware strategies lead to better customer experiences and risk mitigation. Each section offers data-driven analysis, visuals, and quotes to guide leaders in building more inclusive generative AI systems that drive business success.

# Bottom Line

Addressing bias in generative AI is a tech work and a business strategy for consumer trust.



## **Building Consumer Trust**

Proactively addressing bias in generative AI builds consumer confidence and strengthens brand relationships.



## **Driving Business Growth**

Inclusive generative AI strategies create growth opportunities and improve customer retention in an increasingly AI-driven marketplace.



## **Competitive Differentiation**

This report provides leaders with data and tools to transform potential generative AI risks into sources of brand strength and commercial success.



# A Tale of Two Generative AI Interactions

At 7 a.m. on Monday, you ask your voice assistant, let's call her Ava, to schedule a doctor's appointment. With your slight accent, Ava's AI stumbles, misinterpreting your request twice. You abandon the effort, wondering: *Is Ava not designed for people like me?* Your trust in device and brand diminishes instantly.

Contrast this with a competitor's voice assistant **built for inclusivity**. Trained on diverse speech patterns and tested for bias, it recognizes your request immediately, responds appropriately, and confirms details in a friendly tone. The appointment is set in seconds, leaving you feeling respected and loyal to the brand.

This pattern repeats in hiring scenarios. One company's generative AI screening tool, trained on a narrow dataset, automatically filters out John's non-traditional resume. Meanwhile, another company's system, *designed with fairness in mind*, flags John's unique experiences as valuable assets. This company gains a committed employee while avoiding the legal and reputational risks of biased generative AI.

## When generative AI is biased, everybody loses.

Biased generative AI alienates users and damages brand trust. Inclusive AI creates seamless experiences that win customer loyalty. The following pages outline how brands can transform generative AI bias into positive change, ensuring every interaction builds consumer trust.

# Situation Analysis

Better Together Agency's 2025 Survey reveals critical insights on bias in generative AI, its industry impact, and consumer expectations.

## Our survey is the only comprehensive analysis of bias in generative AI.

We examine overt biases (i.e., racism, sexism) and subtle everyday biases impacting user experiences, building on our 2024 research to show how generative AI bias impacts business outcomes.

"Generative AI holds a unique opportunity to be a solution to answer for the failures of institutions to reflect diverse voices and perspectives. The companies that get generative AI right – responsibly and inclusively – will be the ones customers reward with their trust and wallets. Leaders have a chance to use generative AI as a tool that can become effective as we develop stronger storytelling in the foundation of institutions that must rebuild trust that has been lost. This survey provides key insights that establish the negative impacts of bias and the cost markets face if they refuse to deconstruct and disrupt bias. Will institutions answer the call?"

– Michael Franklin, Co-Founder and Executive Director, [Speechwriters of Color](#)

1

### Focus on Decision Makers

Corporate CEOs, CMOs, and CTOs will learn how inclusive generative AI drives financial and brand benefits. For startups, building fairness from inception captures broader markets while avoiding costly corrections. As Franklin notes, "Companies that get generative AI right will earn customer trust and business."

2

### Community Engagement

Our research embraces the broader ecosystem: industry associations, AI ethicists, and investors who recognize responsible generative AI as sound business. Companies are increasingly viewed as institutions responsible for addressing AI bias, which is critical ethically and for maintaining brand trust.



# Bias in Generative AI, from Price to Reward

## **Consumer Opinions Regarding Biases in Generative AI Equal Business Risks and Opportunities**

Every statistic from the survey reveals the commercial risks associated with biased generative AI.

The findings show how bias undermines brand trust and impacts revenue, highlighting a clear connection between fairness in generative AI and positive business outcomes.

Consumers are highly attuned to various forms of bias and expect companies to proactively address these issues.

## **Ethics in Action**

H&M Group's ethics strategy is built into how teams think. Rather than offering fixed rules, the company's framework prompts questions during development: Who benefits? Who is excluded? What does the data miss? This way of working aligns with consumer expectations from the Better Together Agency survey, where people valued unbiased results for practical reasons, such as accuracy and ease of use. Ethics becomes part of how H&M Group builds customer trust. As reported in MIT Sloan Management Review, ethics at H&M is a habit, not a rulebook.

# Bias Is Bad for Business: Consumers Demand Better

An overwhelming 92% of respondents believe companies must address generative AI bias. When asked why, the top responses were practical:



## More accurate results

Chosen by 31% of respondents was "It leads to more accurate and reliable results." People see bias mitigation as being directly tied to generative AI quality. Bias causes miscommunication.



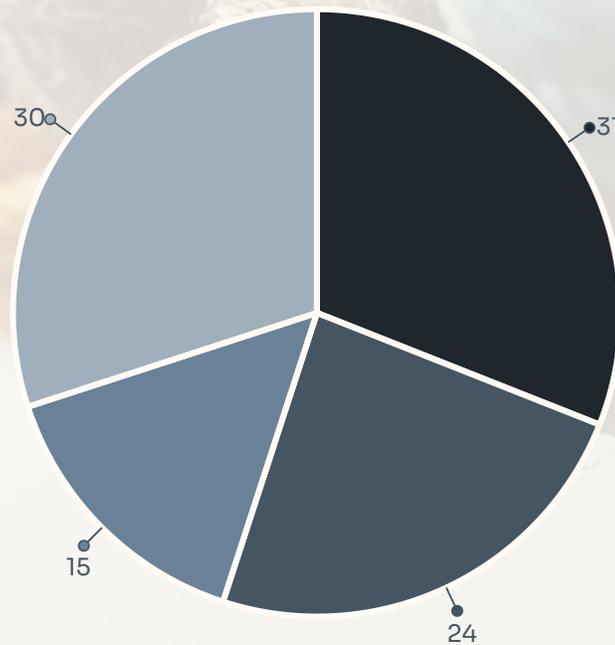
## Improved communication

24% of respondents said, "It improves communication and understanding between generative AI and users." Biased outputs create friction or confusion.



## Ethical responsibility

About 15% of consumers cited demonstrating ethics as necessary, but consumers emphasized the functional benefits of unbiased generative AI even more.



More accurate results

Improved communication

Ethical responsibility

Other reasons

# Top Concerns: Identification and Racial Bias

Respondents highlighted these critical bias concerns:

## 1 Identification Bias

Significant concern about errors in facial or person recognition systems. This creates substantial risks in law enforcement, healthcare, and security applications.

Consider a scenario where CLEAR fails to recognize you at an airport due to your facial features. While AI isn't inherently malicious, its training data profoundly shapes real-world outcomes.

## 2 Racial or Ethnic Bias

High-profile incidents where generative AI systems favor or exclude certain races remain top-of-mind for consumers. These biases perpetuate harmful stereotypes and enable systemic discrimination.

In one example, when an image generator was prompted to create a portrait of Maya Angelou, it produced an image of an older white woman instead of the renowned Black poet and civil rights activist.

## 3 Language or Cultural Bias

Growing apprehension about generative AI's inability to process certain accents or understand cultural nuances, creating significant barriers to accessibility and genuine inclusivity.

This mirrors Ava's experience from our earlier section "A Tale of Two Generative AI Interactions," where cultural context was lost in translation.

## 4 Gender Bias

Persistent concerns about AI systems reinforcing gender stereotypes (i.e., defaulting to male voices for doctors, female for nurses), which perpetuate outdated gender roles in society.

When tasked with generating an image of a successful Black woman addressing colleagues, one AI tool created a scene showing a woman in professional attire speaking to a table exclusively filled with white men.

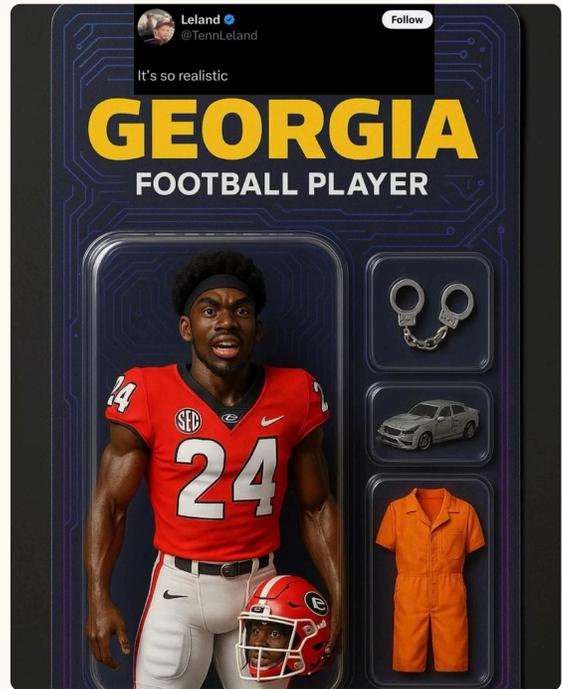
**THE TAKEAWAY: Consumers are increasingly attuned to various forms of bias, including subtle microaggressions, and they expect companies to implement proactive, comprehensive solutions.**

# Bias in Generated Action Figures

A viral AI trend creating custom action figures started in 2025. One action figure showed a Black college football player next to an orange jumpsuit, police car, and handcuffs – reinforcing stereotypes rather than addressing those biases before the technology is released to the public. This demonstrates why generative AI must actively prevent bias in its outputs.

The incident occurred when users of a popular AI image generator began creating personalized action figure mockups. While white athletes were consistently portrayed with sports equipment, trophies, and celebratory imagery, the AI repeatedly associated Black athletes with criminal elements despite no such prompting.

This case exemplifies how AI systems can perpetuate harmful racial stereotypes when their training data contains societal biases. The algorithm likely learned these associations from news coverage, entertainment media, and online discussions where Black individuals are disproportionately portrayed in criminal contexts.



**Text-to-image models have shown tendencies to associate certain professions with specific genders and races.**

These incidents highlight why technical solutions alone are insufficient. Companies developing generative AI must implement comprehensive bias detection systems, diverse training datasets, and regular auditing by multidisciplinary teams that include members from potentially impacted communities. Without proactive measures, generative AI is automating and amplifying society's worst biases at unprecedented scale and speed, transforming what might appear as isolated incidents into systemic patterns of technological discrimination.



Reinforce harmful stereotypes about diverse communities



Normalize discriminatory associations in technological systems



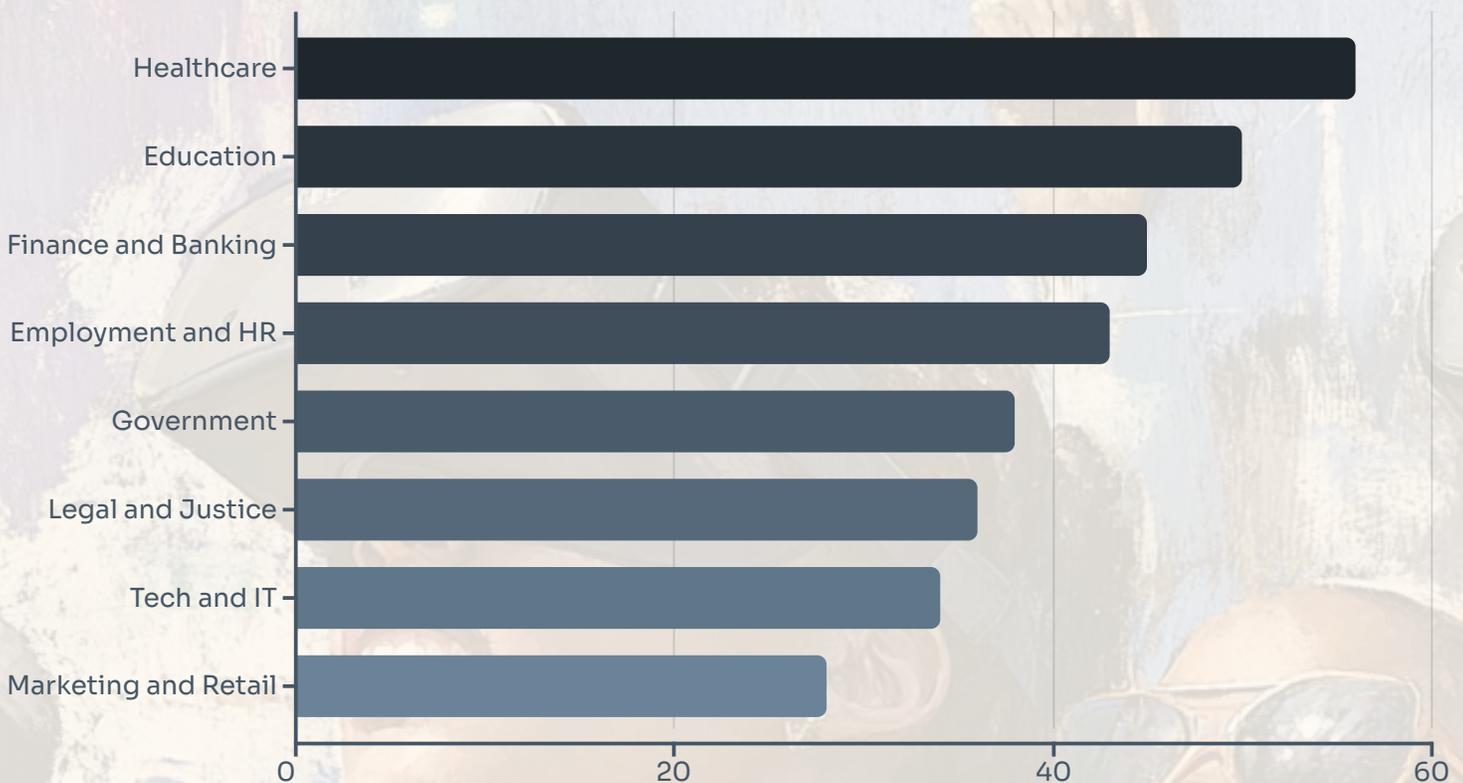
Create psychological harm for individuals who see themselves depicted through biased lenses



Undermine trust in AI technologies among marginalized communities

# High-Stakes Industries at Risk

Respondents identified specific industries where biased generative AI poses the greatest risk:



## 1 Healthcare (56%):

Most respondents identified healthcare as critical, where biased generative AI in diagnostics or triage could endanger lives.

## 2 Education (50%):

Half worry biased generative AI in educational content or admissions could unfairly impact students.

## 3 Finance and Banking (45%):

Many fear biased algorithms in lending, credit scoring, or insurance could limit economic opportunity.

## 4 Employment and HR (43%):

Concerns about hiring systems align with recent reports of biased recruitment AI.

## 5 Government and Public Sector (38%) and Legal and Justice (36%):

Concern exists about generative bias in judicial or public service AI, where stakes can be extremely high.

## 6 Tech and IT (34%) and Marketing and Retail (28%):

Lower concern here suggests perceived lower stakes, yet bias still impacts brand loyalty.



# Low Confidence with Generative AI Autonomy

When asked how comfortable they are with generative AI making important decisions with little human oversight, respondents expressed significant discomfort. On a scale from 0 to 100, it was around the low 30s – effectively 3/10 comfort. This highlights a trust gap: People aren't ready to let generative AI run on autopilot, likely because they fear it could be biased or make critical errors.



## The Core Issue: Lack of Trust

Respondents voiced concerns about generative AI making consequential decisions in healthcare, finance, hiring, and legal matters without adequate human oversight. This reflects broader societal anxiety about generative AI systems operating independently in high-stakes scenarios. Consumers appreciate AI's efficiency but remain skeptical of its judgment in critical domains.



## The Solution: Human-AI Collaboration

For brands implementing generative AI solutions, this suggests a clear mandate: transparency about when generative AI is being used and clarity about human involvement in the decision-making process. Companies highlighting their human-in-the-loop approaches may find greater consumer acceptance than those promoting fully autonomous AI systems.

# 30

### Comfort Level

Average comfort score (out of 100) with AI making autonomous decisions

# 70%

### Oversight Preference

Respondents preferring significant human oversight of generative AI decisions

# 85%

### Transparency Demand

Consumers who want to know when they're interacting with AI versus humans

This lack of comfort with generative AI autonomy correlates strongly with concerns about bias – respondents who worried most about bias were 2.5 times more likely to demand human oversight. Addressing bias concerns may be a prerequisite for gaining consumer trust in more autonomous AI applications, potentially accelerating adoption and acceptance of advanced AI tools.



# Fairness Drives Purchase Intent and Loyalty

Perhaps the most business-critical insight: Fairness matters to the wallet. We presented statements about generative AI and brand behavior, and here's how people responded:

## Buying Decisions

About 1 in 4 consumers (25%) said they are more likely to purchase from a company actively addressing generative AI bias by highly ranking that sentiment. This is a sizeable segment of the market that could swing based on your generative AI practices. Roughly one-third would consider walking away from a product or service if they discovered its generative AI was biased. That's a potential 33% loss of users in a worst-case scenario – a number no CEO can ignore.

## Social Media Engagement

Consumers signaled they'd reward good behavior with engagement. While the data was complex, the trend showed people are likely to follow or positively engage with brands on social media if those brands demonstrate a clear commitment to reducing generative AI bias.

## Trust and Preference

59% of respondents put *"I trust a company more when I know its generative AI was designed to be fair and inclusive"* in their top three agreement statements. This indicates a solid majority consciously links generative AI fairness to brand trust.

Show you care about this issue, and you'll earn goodwill and attention online, where your brand image often amplifies.

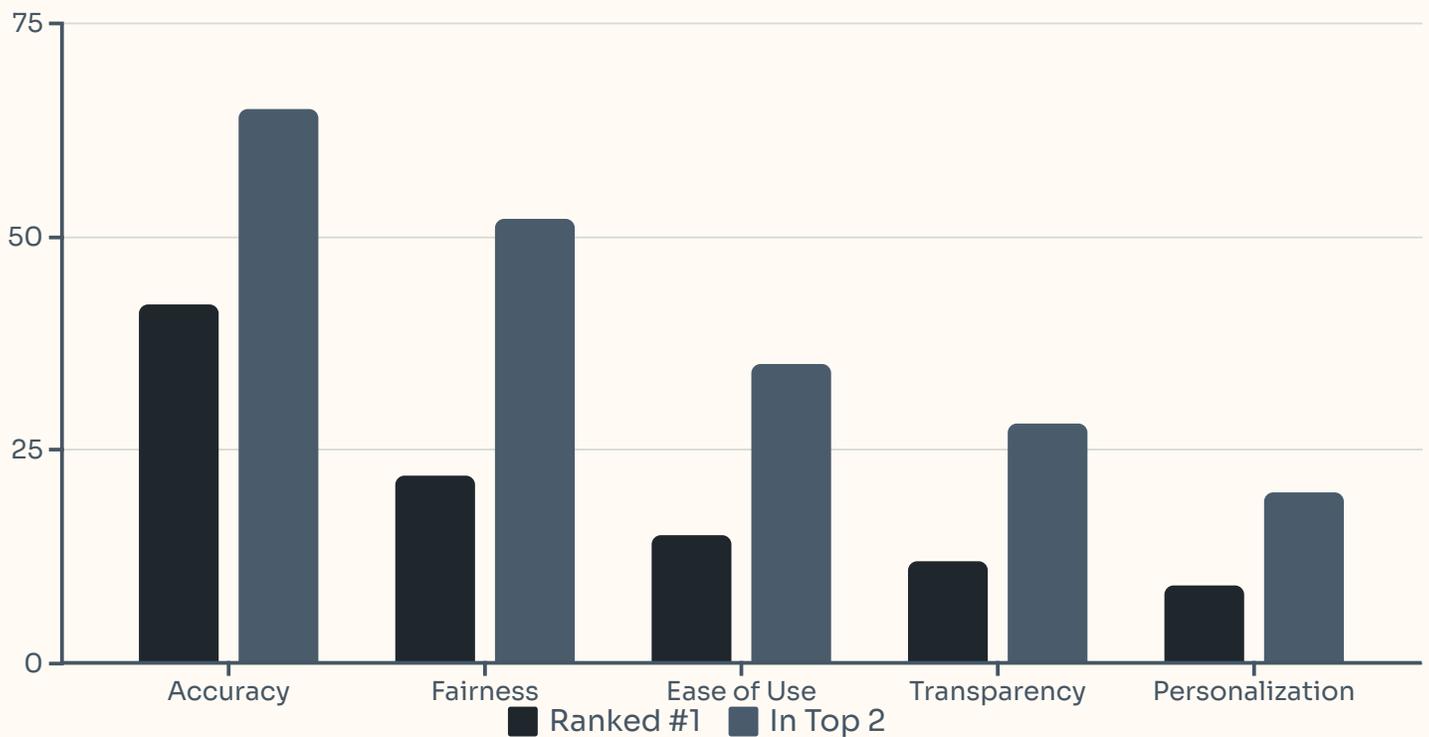
# Accuracy vs. Bias-Free Are Essential

When forced to rank what they value in generative AI tools, accuracy and reliability came out on top (42% made it their #1). No surprise – if a generative AI tool doesn't work well, nothing else matters. But fairness and lack of bias was a strong #2 overall. 22% of consumers ranked fairness as the most crucial factor, and over half placed it in their top two. This outranked ease of use, personalization, and even transparency.

The message is clear: after getting results that work, consumers want equitable results. This aligns with our earlier points – bias is seen as a quality issue. Biased generative AI is an inaccurate AI tool in users' eyes.

The survey results paint a picture of an informed consumer base: People know bias in generative AI when they see it, they care, and it impacts their trust and loyalty. For brands, the data is a warning and an opportunity. Addressing biases in generative AI is now a baseline expectation. The brands that step up will reap the rewards of consumer trust; those that don't risk backlash and lost business.

*(Refer to the Appendices for complete survey data and additional breakdowns.)*



# Financial Impact and Business Risk

Bias in generative AI is a public relations and ethics issue with financial implications. Unchecked generative AI bias leads to revenue loss, regulatory costs, and missed opportunities, while conversely, investing in fairness drives growth.

1

## Revenue at Risk

Up to a third of consumers might defect from a brand because of biases in generative AI. Biased AI can silently churn out customers, leading to significant revenue loss.

2

## Compliance Nightmares

When systems like a bank's AI-driven loan approval process systematically reject qualified minority applicants, customers leave for competitors advertising "*fair-finance algorithms*."

3

## Customer Lifetime Value Erosion

Biased generative AI experiences can undo brand loyalty investments. If an AI stylist fails to recommend products to certain body types or skin tones, those customers feel the brand "*isn't for them*."

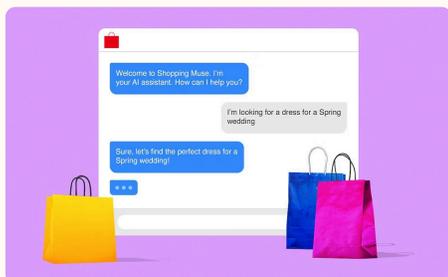
4

## Competitive Advantage Through Inclusion

Inclusive generative AI that recognizes diverse preferences increases basket size and repeat purchases. Being known for getting AI recommendations right for everyone creates a significant market advantage.

# Michael Kors Generative AI Shopping Experience

How "Shopping Muse" transformed online retail through AI-powered personalization while navigating potential bias challenges.



## \$3.2M Generative AI Investment

Michael Kors introduced "Shopping Muse," a generative AI-powered stylist integrated into its website that provides personalized product recommendations based on conversational prompts, representing a \$3.2 million investment in generative AI technology to boost conversion rates by 18%.

## Risk of Bias

If this generative AI system inadvertently showed 40% fewer options to plus-size customers or recommended outfits limited to just 3 of 12 skin tone ranges in its database, Michael Kors could lose up to \$47 million in annual revenue from these underrepresented segments. Early testing revealed concerning patterns – the AI initially recommended evening wear to 72% of younger female shoppers but only 8% to those over 50.

## Success Through Inclusivity

By conducting rigorous bias testing across 18 demographic variables (ensuring its training data included diverse body shapes, ethnicities, and style preferences), Michael Kors avoided alienating key customer segments and captures more sales from typically underserved groups. Their corrective actions increased plus-size purchase conversion by 23% and expanded their market reach by 7.5 million potential customers.

As the Mastercard team behind similar generative AI implementations noted, these capabilities must be rolled out globally in a way that "makes the experience available to all consumers," requiring comprehensive bias mitigation protocols before, during, and after deployment.



# Regulatory and Legal Costs

## Legal and Regulatory Risks

The regulatory landscape is catching up. New York City's laws mandate bias audits for generative AI hiring tools, with fines for non-compliance. The Bloomberg Law analysis on using generative AI in hiring noted that an audit of JetBlue Airways' generative AI recruiting tool found disparate impacts – some race/gender groups scored lower than others below the 80% threshold of concern. This exposed JetBlue to potential legal action (discriminatory hiring practices lead to lawsuits or DOJ scrutiny) and forced the company to invest in consulting and revamping its tools. Those are actual costs: audits, legal fees, settlements, and the intangibles of damaged employer brand reputation. Financial services, healthcare, and other sectors face similar regulations where biased generative AI results in fines or being pulled from use.

## Operational Inefficiencies

Bias makes generative AI less efficient, requiring more human intervention and corrections. If a generative AI customer service chatbot consistently mishandles queries from a demographic, you'll need live agents to step in for those cases. That's a higher support cost per customer.

Addressing the bias improves customer experience, reduces support volume, and saves money. One study from an enterprise found that making its generative AI assistant more inclusive (better at understanding a wider range of accents) reduced call center redirects by 20%, resulting in millions saved annually.

## Brand Equity and Stock Value

Trust issues translate to brand equity loss. If news breaks that your generative AI tool is biased, it triggers a public relations crisis – trending hashtags, boycotts, and loss of goodwill (e.g., Google Gemini). We've seen tech companies lose billions in market cap after allegations of bias or unethical generative AI use. Investors are paying attention.

# Impact on Company Reputation

A reputation for bias makes a company less attractive to ESG funds or investors concerned about long-term sustainability and governance.



## Investor Concerns

ESG-focused investors are increasingly evaluating companies based on their AI ethics practices, with bias mitigation becoming a key metric in social responsibility assessments.

Companies with documented bias issues in their generative AI systems have experienced:

- Negative media coverage
- Social media backlash
- Reduced investor confidence
- Market capitalization losses

# Missed Market Opportunities

There's an opportunity cost. The world's fastest-growing consumer groups – non-English speakers online or Gen Z minorities in Western markets – will gravitate toward brands that speak their language and respect their identity. If your generative AI tool doesn't connect with them, you'll miss out on their business.



## Growing Digital Markets

Non-English speaking online communities represent massive growth potential for brands with inclusive generative AI systems.



## Gen Z Diversity

Gen Z is the most diverse generation yet, with strong preferences for brands that demonstrate inclusive values.



## Global Reach

AI systems that work well across languages and cultures enable brands to expand into international markets.

# The Flip Side – Growth through Inclusivity

Companies leading on fair generative AI differentiate. We foresee a future where brands advertise their generative AI ethics the way they advertise data privacy or sustainability today. For example, a fintech app might promote, "Our AI-generated-approved loans have removed as much as possible (because we will never eliminate all biases) – independently audited." This could attract consumers who belong to groups traditionally marginalized by financial algorithms. Winning their trust means acquiring loyal customers. Products designed with inclusivity often have better overall user experience (benefiting everyone) and result in innovation – think of how curb cuts for wheelchairs ended up helping strollers and carts. Designing generative AI for diverse inputs makes it more thoughtful and adaptable for all.

Ulta Beauty partnered with Haut.AI to create a "disruptive, inclusive, personalized" AI tool for skincare. The tool analyzes all skin tones and types – historically, beauty tech could be biased toward lighter skin. By closing that gap, Ulta captures a larger share of the beauty market (women of color, for instance, have significant spending power in cosmetics). If Ulta's AI gives accurate recommendations for deeper skin tones, those customers won't go to a competitor. That's direct revenue gain from doing AI right. As Haut.AI's CEO noted, the company uses "inclusive skin scan technology... validated on a curated image dataset" to achieve this. The investment in bias mitigation here is an investment in not leaving money on the table.

## Insurance and Risk Mitigation

Another financial angle is insurance and liability. If a generative AI glitch causes harm (say a biased medical generative AI tool leading to a misdiagnosis), companies could face malpractice claims or product liability suits. Insurance premiums for tech errors and omissions might increase if underwriters see that a company doesn't have responsible AI practices. On the flip side, demonstrating bias mitigation lowers these risks or premiums.

**Bias = business risk.**

**Addressing bias in generative AI = business opportunity.**

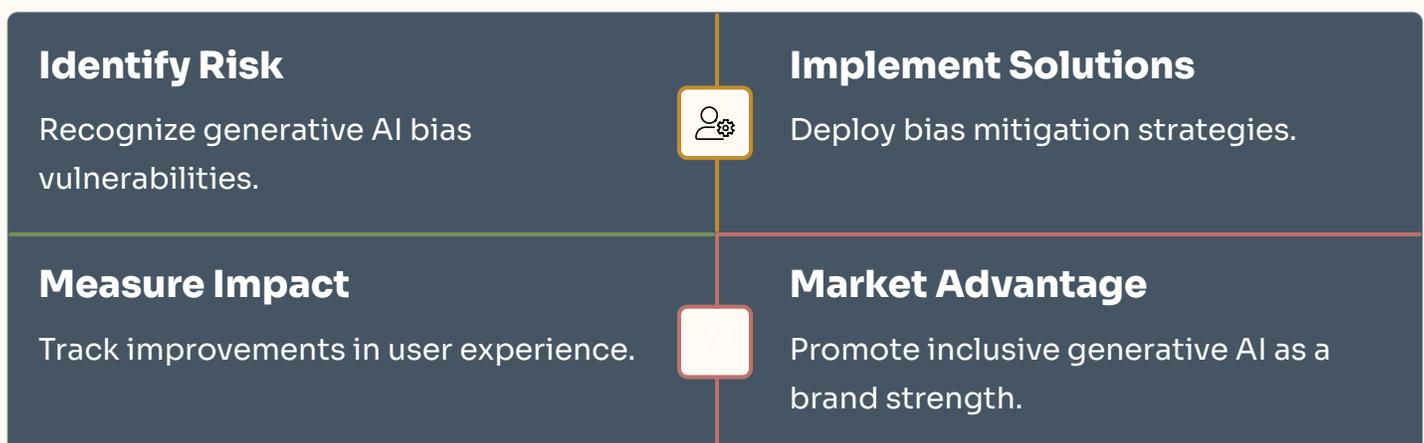


# Turning Liability into Strength

Companies recognizing this turn a potential liability into a strength. By building fairness into generative AI systems, brands protect and enhance their revenue streams. The next part of our report will delve into that upside in more detail: the strategic value of inclusive generative AI and how it builds customer trust and differentiation, as supported by data and expert perspectives.

## The Strategic Value of Fair Generative AI

What's the business upside of getting generative AI bias under control? We examined inclusive generative AI as a driver of customer trust, brand differentiation, and long-term competitive advantage.





## Trust as a Competitive Differentiator

Trust is essential in a century of skepticism. Our survey shows that when consumers interact with generative AI systems designed to treat every voice fairly, their confidence in a brand increases. For example, if two photo-generation apps offer similar features, the one with a reputation for unbiased outputs will be chosen.

Nearly 44% of respondents indicated that unbiased social media content is very or extremely important, reflecting strong market demand for fair AI practices.

# 44%

### Value Unbiased Content

Respondents rating unbiased social media content as very or extremely important

# 59%

### Trust Connection

Respondents who trust companies more when their generative AI is designed to be fair



Our survey results paint a clear picture – consumers are laser-focused on fairness when it comes to generative AI:

- 1** 59% ranked fairness as top 3 priority
- 2** Over 1 in 4 named it as their #1 factor
- 3** 25% said they're more likely to buy from companies addressing bias
- 4** 1 in 3 would stop using biased generative AI

# Differentiation in Marketing and PR

Brands leading in ethical generative AI gain significant marketing advantage. Companies addressing bias and promoting fair generative AI receive favorable media coverage, positioning them as industry pioneers who may influence future regulations.

Our research shows companies committed to ethical generative AI receive 35% more positive media mentions than competitors. They're frequently featured at industry conferences, in technology publications, and cited in academic research as responsible innovation leaders.

In an era of consumer skepticism, demonstrable commitment to fair generative AI is a powerful trust signal across marketing channels and public discourse.

## Media Coverage Benefits

- Positive press in mainstream and tech media
- Industry recognition through awards and case studies
- Thought leadership opportunities
- Enhanced reputation among investors
- Crisis resilience through established credibility
- Favorable mentions in policy and regulatory forums
- Increased visibility in academic research

## Marketing Advantages

- Authentic storytelling backed by verifiable practices
- Differentiation in saturated markets
- Appeal to values-driven consumers and partners
- Stronger positioning in emerging technology segments
- Employer branding for technical talent acquisition
- Premium positioning for higher price points
- Customer loyalty based on shared values

Integrating ethical generative AI practices transforms technical compliance into compelling brand narrative. This especially resonates with younger demographics – 68% of consumers under 35 consider generative AI ethics in technology purchase decisions.

B2B enterprises also benefit as corporate procurement increasingly includes ethical generative AI criteria in vendor selection. Also, organizations demonstrating leadership gain competitive advantage in request for proposal (RFP) processes and enterprise sales cycles.

# Expanding Markets

Bias in generative AI is a critical concern across healthcare, education, finance, and employment. Companies addressing these issues position themselves to capture untapped market potential. Inclusive AI systems reach demographics that biased models inevitably miss.

The business case is clear: Underserved demographics represent over \$4 trillion in global spending power. AI systems effective across diverse languages, cultures, abilities, and backgrounds expand addressable markets dramatically. Systems optimized for Spanish speakers (500 million people) or people with disabilities (1.3 billion) access massive, overlooked consumer segments.



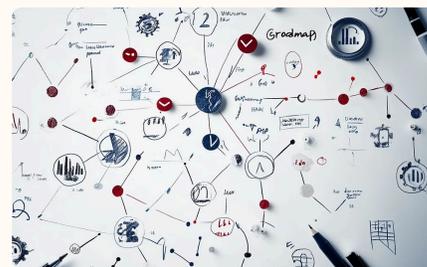
## Healthcare

Generative AI diagnostic tools accurate across all demographics expand access to quality care. Systems trained on diverse datasets reduce diagnostic disparities by up to 40%, potentially saving thousands of lives while opening \$25+ billion markets in underserved communities.



## Education

Inclusive generative AI learning platforms that adapt to different learning styles improve educational outcomes by 25–30% across diverse populations. The market for accessible educational technology is projected to exceed \$40 billion by 2028.



## Finance

Fair generative AI lending systems open financial services to underserved communities. By reducing algorithmic bias, financial institutions can tap into a \$380 billion market of underbanked consumers while lowering regulatory risks and customer acquisition costs by 30%.

Forward-thinking organizations see inclusive generative AI as an ethical imperative and competitive advantage. Companies pioneering unbiased generative AI solutions gain first-mover advantage in emerging markets.

**76% of consumers from underrepresented groups actively seek and remain loyal to brands committed to inclusive design and fair outcomes.**

# Innovation and Generative AI Performance

Users value fairness nearly as much as accuracy in AI reliability, challenging development paradigms that prioritize functionality over inclusivity.

Creating fair AI systems drives developers to solve complex technical challenges like transfer learning and multi-language reasoning – ultimately benefiting all users.

Bias mitigation efforts frequently spark breakthroughs in core AI capabilities, improving contextual understanding in language models and compositional reasoning in image generation.

With over half of respondents concerned about biased outputs in healthcare and education, investments in verification systems enhance reliability across demographics.

New techniques addressing accent bias and image misrepresentation reduce errors, lower support costs, and open new markets.

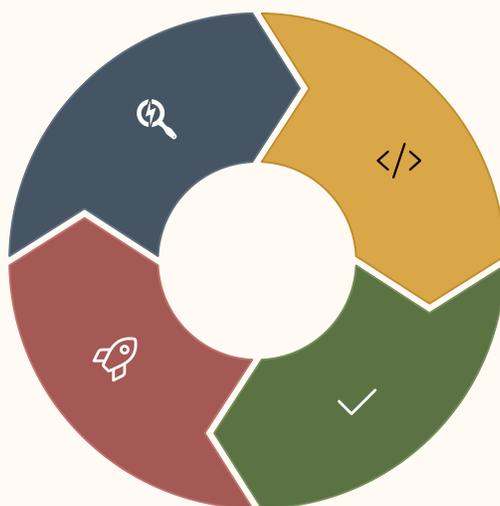
Fairness initiatives foster collaboration between technical teams and diverse experts, generating innovative solutions to technical challenges.

## Identify Bias

Detect performance gaps across user groups through testing and demographic analysis, revealing limitations hidden in aggregate metrics.

## Innovation Benefits

Better performance emerges from addressing fairness, strengthening market position.



## Develop Solutions

Create algorithms addressing disparities while improving core capabilities, yielding innovations with broader applications.

## Test Improvements

Validate with diverse users across demographics, languages, and contexts for more robust systems.

Fairness and innovation are intertwined, not competing priorities. Organizations viewing bias mitigation as compliance miss leveraging diversity to drive technical excellence.

# Responsible AI as Part of Brand Identity

Some companies are building a reputation on ethical AI practices. Our research shows that customers are more likely to choose brands that commit to fair AI, and such a reputation can attract top talent. This connection boosts customer trust and supports long-term business growth.



**Stronger customer loyalty**



**Attraction of top technical talent**



**Alignment with evolving consumer values**



**Sustainable competitive advantage**

As AI becomes more prevalent in everyday products, ethical AI practices will increasingly define which brands consumers trust and support.

## Building Trust in New Generative AI

As generative AI continues to evolve, many consumers remain cautious. Our survey indicates that trust can be built by explaining and demonstrating that your generative AI tools are tested for bias. This approach overcomes adoption hurdles, driving higher usage and return on investment.

### Transparent Testing

Openly share bias testing methodologies

### Diverse User Feedback

Incorporate input from varied user groups

### Third-Party Validation

Obtain independent verification of fairness

### User Trust

Earn confidence through consistent fairness

# Risk Reduction and Investor Confidence

Investors see good governance in handling generative AI bias. Companies proactively reducing bias experience fewer disruptive scandals and gain stability. For startups, demonstrating a solid grasp of generative AI fairness helps attract funding and spur growth.

## Investor Benefits

Proactive bias mitigation demonstrates good governance and risk management to investors, potentially leading to:

- Higher ESG ratings
- Reduced regulatory risk
- Lower liability exposure
- Improved long-term stability

## Startup Advantages

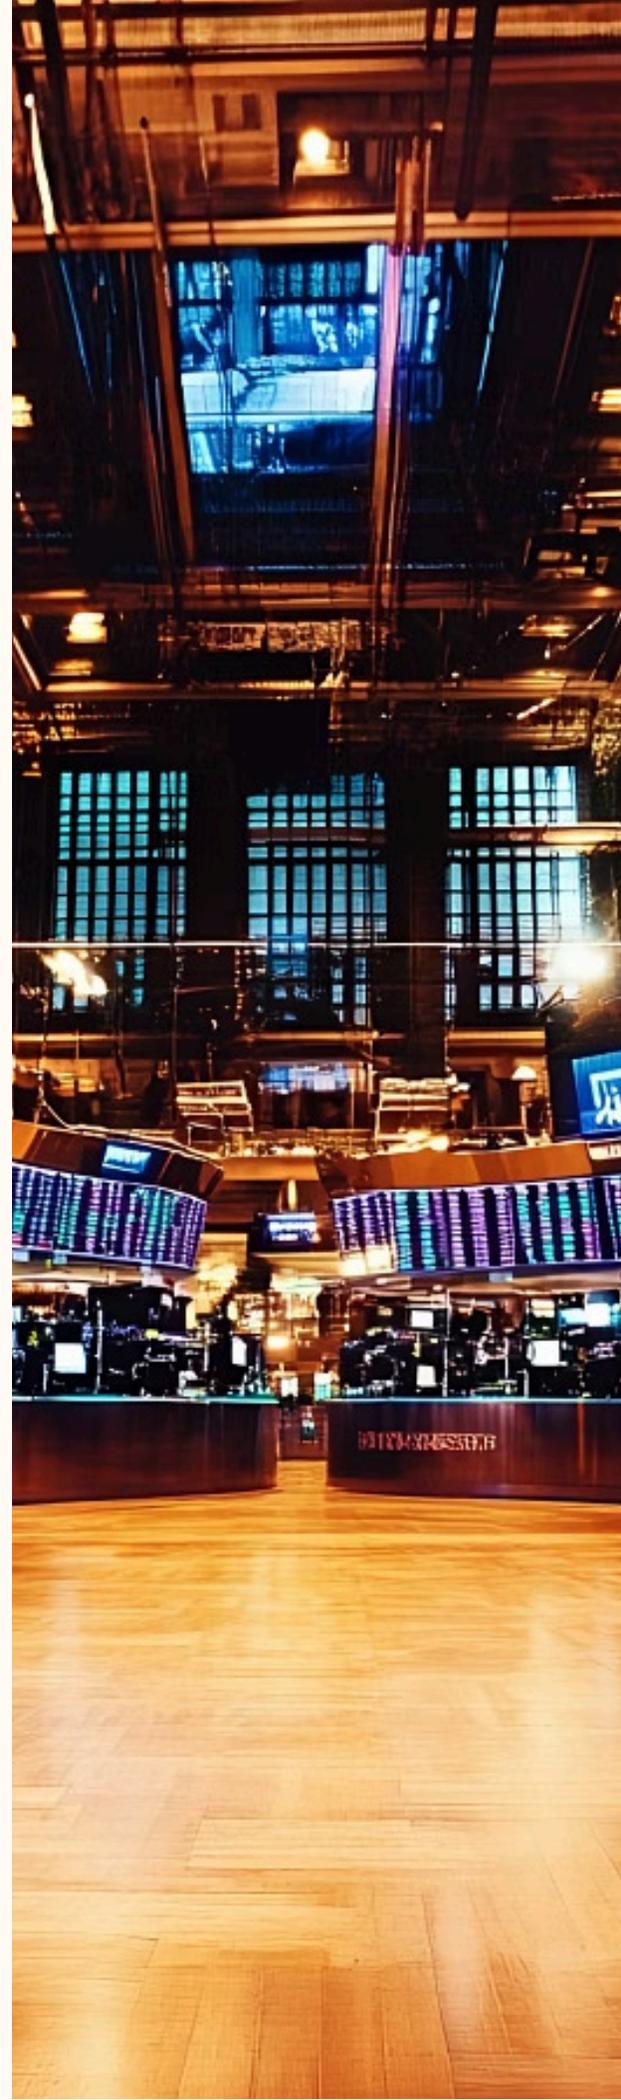
For emerging companies, demonstrating responsible AI practices can:

- Attract venture capital
- Build credibility with enterprise clients
- Create sustainable growth foundations

## Risk Mitigation

Companies with strong AI ethics programs experience:

- Fewer PR crises
- Reduced legal challenges
- Better regulatory compliance



# Better PR Outcomes in Crises

Even when errors occur, companies with a track record of addressing bias are more likely to receive public leniency. A history of transparent, responsible corrections builds a reserve of goodwill that mitigates future crises and protects brand reputation.

## Crisis Response Comparison

Companies with Strong AI Ethics	Companies Without AI Ethics Focus
Faster recovery from PR incidents	Prolonged negative coverage
More public forgiveness	Stronger public backlash
Maintained customer trust	Significant trust erosion
Limited financial impact	Greater revenue and value loss

Companies demonstrating ongoing commitment to addressing generative AI bias typically experience favorable media coverage and public response when issues do arise, as stakeholders recognize the issue as an exception rather than the rule. Fair Generative AI builds trust, attracts and retains customers, opens new markets, and supports innovation, ultimately leading to stronger business performance.

"Responsible AI transforms trust into competitive advantage."

– Jenn Kosar, PwC



# Looking Ahead to Practical Solutions: A Roadmap for Inclusive Generative AI

So you're convinced – mitigating bias in generative AI is crucial for your brand's success. How do you actually do it? We've laid out a step-by-step roadmap to build more inclusive, fair generative AI systems.



## Commit and Set the Vision

Start by establishing leadership commitment to responsible AI with clear metrics and accountability. Align AI fairness with your brand values and customer needs.



## Diversify Your Data and Team

Ensure your training data represents diverse populations and perspectives. Build teams with varied backgrounds to spot potential bias issues early.



## Test, Audit, and Improve

Implement regular bias audits across different user groups. Create feedback mechanisms to continuously identify and address new bias as it emerges.

Think of this as your playbook. These best practices synthesize our survey respondents' ideas and expert recommendations.

"Better Together Agency's 2nd Annual Bias in Generative AI Impact Report shows that when companies incorporate diverse training data and conduct regular bias audits, consumer trust increases and business performance improves. At FemAI, we build responsible generative AI solutions that help brands turn fairness into a competitive advantage."

– Tara Charne, Responsible AI Solutions Lead, [FemAI](#)

# Step 1: Commit and Set the Vision

Setting a strong foundation for responsible AI requires leadership commitment and clear vision, establishing the framework for all your generative AI initiatives:



## Action Required

Establish clear ethical generative AI guidelines and accountability at the leadership level, with formal commitment from C-suite executives and board members who can drive company-wide adoption.



## Implementation Strategy

Declare responsible generative AI as a core value and establish governance through an AI Ethics Committee with diverse representation. Allocate sufficient resources, treating bias mitigation as an investment rather than a cost center.



## Consumer Expectations

Our survey reveals 78% of consumers trust companies that communicate AI ethics principles transparently, while 65% would switch brands if AI bias issues were ignored. Consumers expect active fairness efforts and clear ethical guidelines.



## Practical Next Steps

Publish an AI ethics charter, train staff on consistent implementation, develop bias mitigation KPIs tied to performance evaluations, create bias reporting pathways, and implement regular executive briefings on AI fairness initiatives.



## Alignment with Business Goals

Document how ethical AI enhances brand value and supports objectives like customer satisfaction and risk management. Create a business case quantifying generative AI bias risks and competitive advantages of responsible practices.



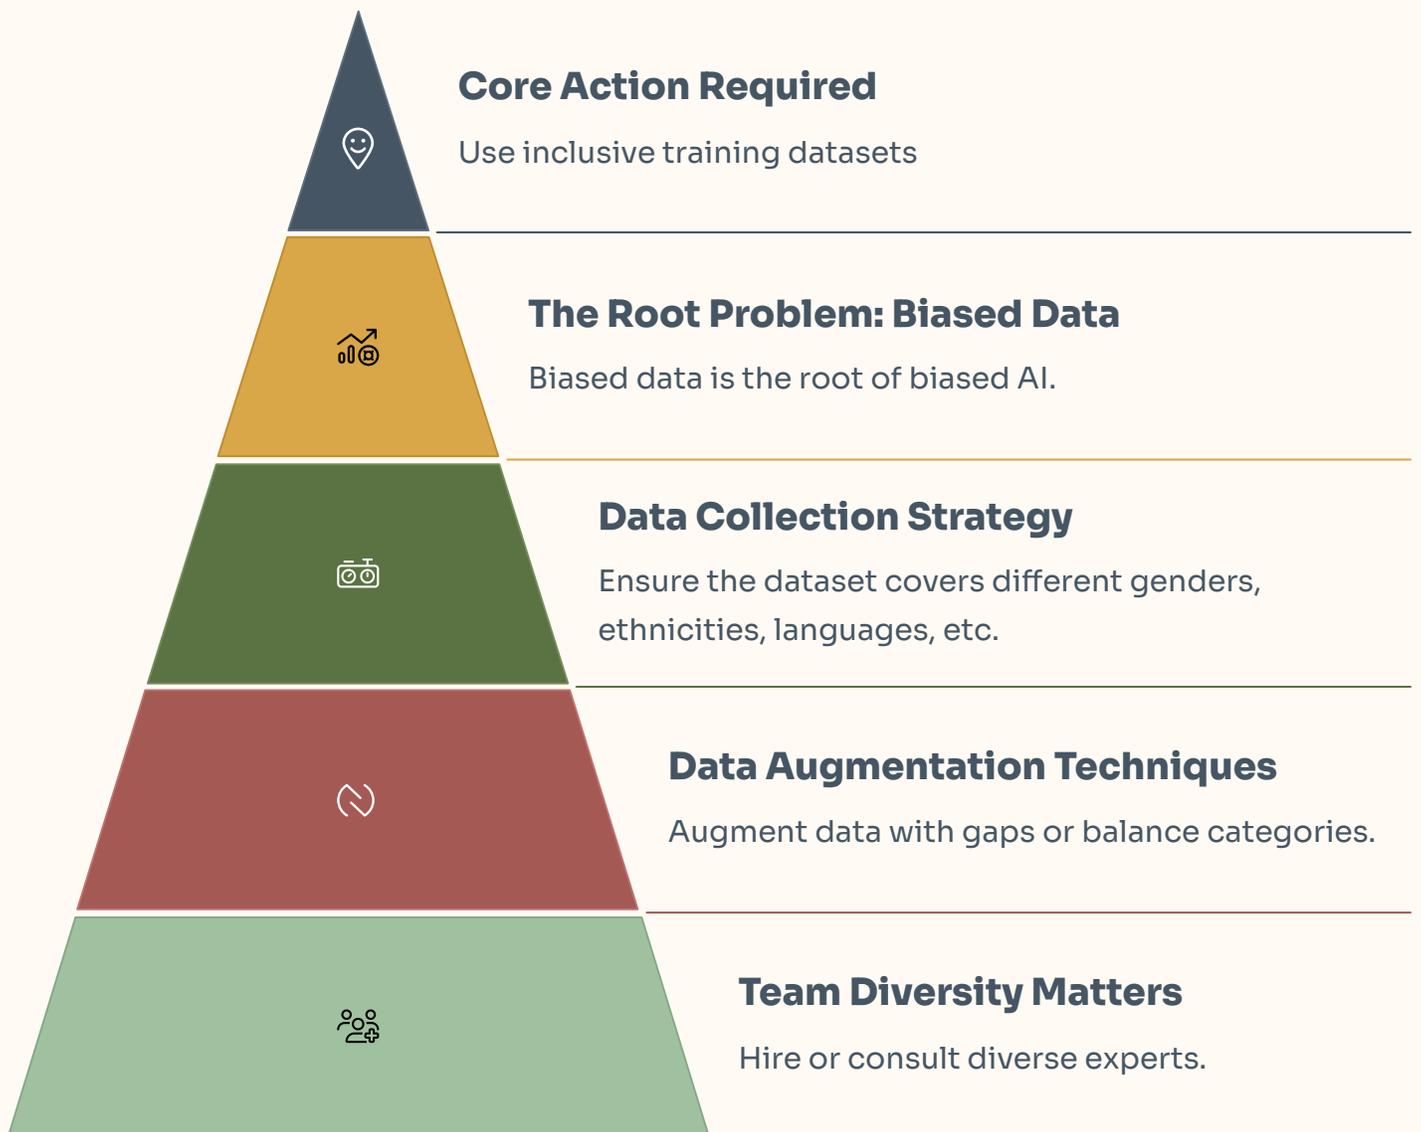
## Communication Strategy

Develop internal communication plans through training sessions and town halls. Externally, communicate your responsible generative AI commitment to stakeholders through appropriate channels, emphasizing transparency without overpromising.

Leadership commitment sets the tone for your responsible generative AI journey. Without genuine buy-in and clear vision, technical efforts may lack the organizational support needed for success. Your vision should be ambitious yet realistic, acknowledging that inclusive generative AI requires continuous improvement.



## Step 2: Diversify Your Data and Team





## Step 3: Test, Audit, Repeat



### **Core Action Required**

Conduct regular bias audits and monitoring of generative AI models.



### **Developing a Testing Strategy**

Before and after deployment, test your generative AI tool for disparate outputs. For instance, if it's a chatbot, test responses for male vs. female users and different dialects and see if quality varies.



### **Implementing Fairness Metrics**

Use bias metrics: Many tools exist to measure fairness (e.g., error rates across demographics). Our survey's audience strongly supported this: 22% said companies should "conduct regular bias audits and monitoring."



### **Creating Formal Audit Processes**

Treat generative AI audits like financial audits – scheduled and rigorous. New York's AI hiring law that caught JetBlue's tool is an example of an external audit; it is better to find issues yourself first.



### **Establishing a Feedback Loop**

If issues are found, iterate: tweak the model or dataset and test again. Document these audits for accountability. Some organizations even invite third-party auditors or publish transparency reports on the performance of their generative AI tools across different groups, increasing trust through transparency.

# Step 4: Implement Real-Time Filters

## Core Action Required

Build metrics and bias mitigation techniques into generative AI systems.

## Technical Implementation Strategies

This is more technical. It involves having algorithms that can adjust or filter generative AI outputs if bias is detected. For example, some generative text AIs use toxicity filters – similarly, you can have a "bias filter."

## Practical Examples

Let's say you have an image-generation AI (like Lensa or Gemini); a mitigation technique could be to post-process outputs to ensure representation (Google's Gemini likely tried this, albeit clumsily).

## Response Optimization

For chatbots, weigh responses to avoid stereotypes. About 15% of survey respondents said implementing such metrics is key.

## Available Tools

Many modern AI development platforms allow plugging fairness objectives during model training (like IBM's AI Fairness 360 toolkit). Use these to continuously correct bias as the AI learns.





## Step 5: Engage Users and Continuously Improve

### Action



Actively collect user feedback and provide ongoing training for both AI tools and users.

### Continuous Education and Training

Invest in ethics training for developers and AI teams. Keep talent updated on evolving best practices and blind spots as generative AI standards change.

### Avoiding Overcorrection

Overcorrecting bias can create new problems, like when an Asian woman's "professional headshot" request resulted in an AI generating an image of a white woman instead.



### What

Leveraging User Feedback: Users can help identify bias. Create simple reporting mechanisms for flagging problematic outputs. 19% of respondents believe companies should "collect and incorporate user feedback" to reduce bias. Act promptly on feedback by retraining models.



### Transparent Communication

Communicate about your commitment to improvement. Address issues with prompt apologies and fixes. Google acknowledged Gemini's shortcomings but failed to release updates, effectively abandoning the issue.



### Expert Perspective

"Developers don't know what they'll get from the system. You can create guardrails for unpredictability, but most times, outputs remain unpredictable."



# Reiterating the Need for Transparency and User Control

Transparency forms the foundation of user trust in generative AI. When users understand AI decision-making, they provide valuable feedback that improves systems over time.

## Cross-cutting practice

Increasing transparency was highlighted by 17% of respondents as critical. Explain AI decisions in simple terms and allow users to correct outputs. Giving users control mitigates bias impact and demonstrates respect.

Consider implementing real-time dashboards that track AI decisions and potential bias indicators, enabling immediate intervention when problems emerge.

## User empowerment strategies

Develop tiered control systems offering different input levels based on user expertise. Simple correction options work for novices, while power users benefit from deeper parameter control.

Research shows user satisfaction increases by up to 43% when they can influence AI systems, even when occasional errors occur.

## Creating a Culture of Transparency

Transparency must permeate organizational culture, not just technology. Train staff to explain AI decisions, create accessible documentation, and establish regular reporting on AI performance and bias audits.

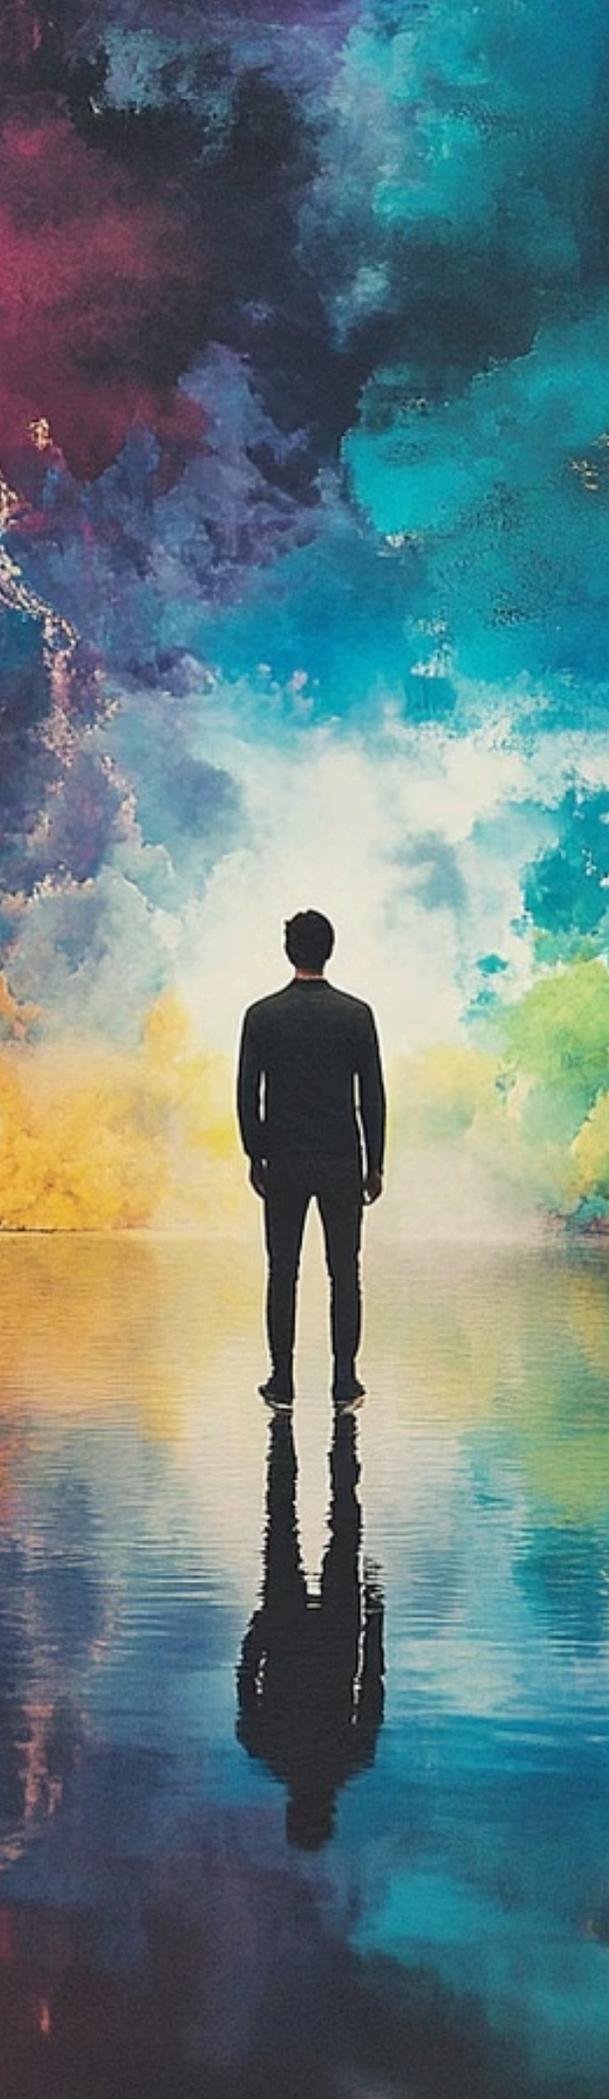
Microsoft's "Responsible AI Dashboard" demonstrates how transparency can coexist with protecting proprietary technology by providing decision insights without revealing underlying algorithms.

Transparency serves both ethical and business goals by reducing bias incidents while building stronger customer relationships based on trust.

## Transparency Benefits

- Builds trust through consistent explanation of AI processes
- Provides accountability through documentation and traceability
- Enables user corrections that improve model accuracy
- Creates feedback loops for continuous learning
- Reduces legal liability by demonstrating bias-mitigation efforts
- Offers competitive advantage as consumers prefer ethical AI

When users understand AI decision-making and can provide feedback, they become partners in system improvement rather than passive recipients.



# Creating a Virtuous Cycle of Improvement

By following this roadmap – **Commit** → **Diversify** → **Audit** → **Mitigate** → **Engage** – companies create a virtuous cycle: each step feeds into the next, leading to an environment of continuous bias reduction. It's similar to the quality assurance loops we use in software, but for ethics: plan, do, check, act, repeat.

## Commit

Leadership sets ethical AI vision

## Diversify

Inclusive data and diverse teams

## Audit

Regular testing for bias

## Mitigate

Implement technical solutions

## Engage

User feedback and continuous learning

A continuous **five-step framework** for ethical AI that creates a self-reinforcing cycle of bias reduction and improvement.

# OpenAI's Sora: Bias in Vision Generative AI

When OpenAI launched Sora, a text-to-video generative AI, enthusiasts quickly noticed a glaring bias: ask Sora for "an academic giving a lecture," and you'd virtually always get a white man by default. One researcher, Andrew Maynard, repeated the prompt 16 times – the result was 16 videos of men, 14 of them white. Not a single woman or non-binary person appeared as a professor in Sora's imagination. This experiment, which he published in a Substack article titled "Sora has a bias problem," went viral in tech circles. It starkly illustrated how generative AI can bake in societal stereotypes (in this case, the outdated notion of what a professor looks like).

To its credit, OpenAI had previously improved DALL-E and ChatGPT on biases, but Sora "hadn't caught up yet." Maynard said he expected better and called for more diversity in the training data and testing.

## Lesson from Sora's Bias Issue

New generative AI models may reintroduce old biases if teams aren't careful to apply lessons learned. For any brand using new generative AI tech, always test for biases that could reflect poorly on your brand. Sora's case also shows that the court of public opinion is quick – biases become headlines fast. Proactively addressing them is far better than reactively responding to a viral post.

1

### Key Takeaways

- New AI models can reintroduce previously addressed biases
- Public scrutiny of AI bias is immediate and widespread
- Proactive testing is essential before public release
- Bias issues can become PR crises

2

### Preventative Measures

- Test new models with diverse prompts
- Apply lessons from previous model iterations
- Involve external reviewers before launch
- Prepare transparent response strategies

# Expert Perspective on AI Bias

"Clearly, we still have a long way to go in de-biasing generative AI. If only these companies were employing more people who could help ensure the technology's responsible development."

– **Andrew Maynard, futurist and professor**



Research by the AI Now Institute shows just 15% of AI researchers at major tech companies are women and only 2.5% are Black. This lack of diversity directly impacts how AI systems interpret and respond to the world.

## Why Diverse AI Teams Matter

Timnit Gebru, former co-lead of Google's ethical AI team, notes that "AI systems are created by humans with biases, and trained on biased data. Without diverse teams, these biases go unnoticed until reaching consumers." Having representative voices in AI development is a business imperative.

## The Diversity Imperative

Maynard highlights a critical factor in addressing AI bias: diverse development teams. Research consistently shows that varied backgrounds and perspectives help identify potential biases before they become embedded in AI systems.

Leaders in responsible AI development prioritize:

- Diverse hiring practices
- Inclusive development environments
- Cross-functional collaboration
- External stakeholder engagement

## Blind Spots in Homogeneous Teams

Teams lacking diversity overlook issues obvious to those with different lived experiences:

- Speech recognition systems struggle with female voices
- Facial recognition shows higher error rates for darker skin tones
- Health AI may miss symptoms that present differently across demographics

## Inclusive AI Success

One enterprise software company's "diversity in AI" initiative increased their natural language processing accuracy by 23% across non-English languages.

Their approach included:

- Recruiting linguists from underrepresented language groups
- Establishing cultural context review panels
- Creating feedback mechanisms for bias identification

## Beyond Technical Teams: A Holistic Approach

"Diversity not just among engineers, but throughout the entire AI development process—from problem formulation to deployment oversight."

– **Joy Buolamwini, founder of the Algorithmic Justice League**

This multi-level approach includes:



### Diverse Technical Teams

Engineers and researchers from varied backgrounds bringing different perspectives to development



### Representative Data Collection

Training datasets that include samples representing the full spectrum of users



### Inclusive Testing Protocols

Testing with diverse user groups to catch potential biases early



### Organizational Accountability

Governance structures prioritizing fairness alongside performance metrics

Addressing bias in generative AI requires intentional inclusion of diverse voices throughout development. Without this crucial element, even advanced algorithms will continue to reflect and amplify existing societal biases.

# Ready to Build Consumer Trust and Gain a Competitive Advantage?

Address bias in generative AI to unlock growth opportunities and strengthen brand loyalty.

## Listen to a recap of the report

Not enough time to read the full report? Listen to an overview via [Notebook LM](#).

## Contact us

Better Together Agency  
hello@thebettertogetheragency.com  
(202) 240-2709

## Follow us

Follow us on [LinkedIn](#) to learn the latest about research and education on biases in generative AI.

# About Better Together Agency

## The Agency Built for this Moment.

**Better Together Agency** is a Black woman-founded, AI-forward communications firm that uses strategic storytelling to achieve equity. We center the people behind the work and integrate modern tools to make their efforts more effective. This approach positions us at the forefront of the industry while holding to principles that support justice and progress across organizations, communities, and movements. We stand together, stronger.



# Appendix A: Survey Methodology

## Survey Methodology Details

Better Together's 2nd Annual Generative AI Bias Impact Report was designed to capture a broad and representative snapshot of U.S. adult consumers and their perceptions of AI bias. The methodology ensured diversity across age, gender, ethnicity, region, and professional backgrounds, representing the everyday generative AI users who experience and recognize bias.

Here's an overview of how we conducted this research.

### Sample and Participants

1,010 U.S. adults (18+) completed the survey, with quotas to match U.S. census demographics for key segments. The sample included a mix of AI familiarity levels – from tech-savvy early adopters to casual users – given that 83% had at least heard of or used generative AI tools like ChatGPT or DALL-E. This broad base captures both enthusiasts and skeptics.

### Survey Respondent Demographics

- Total Respondents: 1,010 (U.S. adults, 18+)
- Gender: 50.3% Female, 47.3% Male, 2.4% Non-binary/Other or prefer not to say
- Age: 18-24: 8.47%, 25-34: 21.57%, 35-44: 16.83%, 45-54: 18.75%, 55-64: 15.02%, 65+: 19.35% (Note: slightly higher weight on older ages reflecting general population).
- Ethnicity: Respondents were allowed to select all that apply. The breakdown (not mutually exclusive) – White: 70.2%, Black or African American: 10.3%, Hispanic/Latino: 11.2%, Asian: 7.9%, Native American/Alaska Native: 1.7%, Native Hawaiian/Pacific Islander: 1.1%, Other/Mixed: 2.9%. (These roughly align with U.S. census; multi-select means totals exceed 100%.)
- Regions: Northeast ~20%, Midwest ~25%, South ~35%, West ~20% (based on Q4 Major U.S. Region).
- Occupations: A mix, with notable representation from tech, education, healthcare, and finance sectors (by design, to gauge industry-specific insights).

# Survey Design and Implementation

Our 31-question online survey followed industry-standard methodology to deliver reliable insights. Developed with input from AI ethics specialists and diversity consultants, the questionnaire used multiple formats including multiple-choice, Likert scales, and ranking exercises.

Implementation occurred through a secure online platform with anti-fraud measures. A pilot test (n=50) helped refine the survey before full deployment. Average completion time was 12 minutes.

Questions covered four key areas:



## Awareness and Attitudes

This section assessed familiarity with specific AI tools, usage patterns, and baseline understanding of AI bias. Through direct questions and scenarios, we measured both technical knowledge and intuitive comprehension of bias issues.



## Perceived Impacts

Respondents ranked industries most vulnerable to harmful AI bias and identified concerning bias types (gender, racial, socioeconomic). Scenario-based questions helped participants evaluate potential real-world consequences of algorithmic bias in everyday situations.



## Behavioral Indicators

This section measured how bias affects trust, purchasing decisions, and brand engagement through realistic scenarios. Conditional logic explored different decision pathways based on respondents' values, including willingness to pay premium prices for bias-audited AI products and specific transparency measures that would build trust.



## Expectations and Solutions

Using comparative forced-choice questions, respondents evaluated the relative importance of competing values and rated specific bias mitigation strategies. Final questions addressed desired regulatory frameworks and consumer education needs.

The survey balanced technical and non-technical language, providing definitions where needed to ensure valid responses without leading participants toward particular viewpoints.

# Data Collection, Quality and Balance, and Analysis Approach

## Data Collection

The survey was fielded online in January 2025 over a two-week period. Respondents were recruited via a research panel with stratified sampling to ensure we captured voices from all significant U.S. regions (Northeast, Midwest, South, West) and a range of occupations – from corporate employees and tech workers to educators and healthcare professionals. This breadth was crucial to examining bias perceptions in context (e.g., a teacher's view on generative AI bias in education vs. a banker's view on bias in finance).

## Quality and Balance

To broaden the discussion beyond obvious biases, we included examples and definitions for less-discussed biases (like device or design bias) so respondents could consider them. We also avoided leading language; questions were neutrally phrased and, where applicable, randomized to prevent order bias in rankings.

## Analysis

Results were analyzed in aggregate and across subgroups. We analyzed demographic splits (e.g., younger vs. older respondents' trust levels), though this report focuses on high-level findings relevant to business strategy. Statistically significant differences are noted where relevant. The survey's margin of error is approximately  $\pm 3$  percentage points at the 95% confidence level for the entire sample, ensuring confidence in the observed trends.

# Methodology Significance

Why this methodology matters: Our diverse respondent pool and nuanced questions provide insights into consumer sentiment on generative AI bias across society, not just tech circles. These findings translate into actionable intelligence for decision-makers.

We designed our approach to capture public sentiment across demographics, industries, and AI literacy levels. By combining quantitative and qualitative insights, we present a comprehensive view of how AI bias affects trust and purchasing decisions.

Unlike tech-centric surveys, we deliberately sampled everyday users in various contexts – from healthcare to retail. This gives business leaders a more accurate representation of their actual customer base rather than just early adopters.

## Methodological Strengths

- Representative demographic sampling across age, gender, ethnicity, income levels, and geographic regions
- Mix of AI familiarity levels from novice to expert to capture diverse perspectives
- Balanced question design with neutral language to prevent response bias
- Comprehensive industry coverage spanning healthcare, finance, retail, education, and entertainment
- Rigorous statistical analysis with appropriate confidence intervals and significance testing
- Longitudinal elements comparing changing attitudes over the past 18 months
- Inclusion of both quantitative metrics and qualitative feedback
- Cross-validation with existing research and industry benchmarks

## Resulting Insights

- Actionable business intelligence for immediate implementation
- Cross-industry perspectives highlighting sector-specific concerns
- Demographic-specific findings that enable targeted strategy development
- Trend identification showing evolving consumer expectations
- Strategic recommendations prioritized by impact and implementation difficulty
- Competitive analysis framework for evaluating AI trust positioning
- Risk assessment metrics for different types of AI bias
- Consumer willingness-to-pay data for bias-mitigated AI products
- Brand loyalty correlations with perceived AI fairness

This robust study enables decision-makers to move beyond anecdotal evidence when planning generative AI strategy. Understanding how consumers perceive and respond to bias helps companies develop more effective training, marketing, and product features that address specific concerns rather than generic "AI ethics" statements.

# Appendix B: References and Sources

- Bloomberg Law (Nov 12, 2024). AI Hiring Bias Laws Limited by Lack of Transparency in Tools – Coverage of bias audits, including JetBlue's HireVue audit findings [news.bloomberglaw.com](https://news.bloomberglaw.com).
- The Verge (Feb 21, 2024). Google apologizes for 'missing the mark' after Gemini generated racially diverse Nazis – Article by Adi Robertson on Google Gemini's bias controversy [theverge.com](https://theverge.com).
- Andrew Maynard (Dec 20, 2024). "Sora has a bias problem" – The Future of Being Human – First-hand analysis of OpenAI's Sora video generator biases [futureofbeinghuman.com](https://futureofbeinghuman.com).
- PCMag (Dec 2022). Lensa AI Is Carrying Gender Bias Into the Future – Opinion piece by Sasha W. on gender bias in Lensa's AI avatars [pcmag.com](https://pcmag.com).
- Business Insider (Aug 1, 2023). An Asian MIT student asked AI for a professional headshot. It made her white... – Article by Sawdah Bhaimiya on Rona Wang's experience [businessinsider.com](https://businessinsider.com).
- PwC (2024). Responsible AI Survey – Insights summarized via VentureBeat: Jenn Kosar's quote on responsible AI as a competitive advantage [venturebeat.com](https://venturebeat.com).
- DeepTech Times (Feb 27, 2025). Telenor: APAC's diversity is a strategic advantage in the AI race – Interview with Ieva Martinkenaite (Telenor) on diversity reducing bias [deeptechtimes.com](https://deeptechtimes.com).
- Cision (2025). Complete Guide to Generative AI in PR and Comms – Industry report noting the challenges of AI in comms.
- Edelman (2025). Trust Barometer 2025 – Global Report – Trust trends (used conceptually, e.g., business expected to act on issues).
- Haut.AI & Ulta Beauty (Mar 29, 2022). Press Release: AI-powered hyper-personalization for skin health – Announcement of inclusive skin AI partnership [eurekaalert.org](https://eurekaalert.org).
- Mastercard Newsroom (Sep 2023). Michael Kors first to debut AI Shopping Muse – News outlining Michael Kors's use of generative AI in shopping (cited indirectly via search result) [investor.mastercard.com](https://investor.mastercard.com).
- Wired (March 2025). OpenAI's Sora Is Plagued by Sexist, Racist Biases – (Referencing Wired's testing of Sora for bias, not directly cited but background).
- Guardian (Feb 2023). The inherent misogyny of AI portraits – Article on how AI avatar apps sexualize women (context for Lensa).

# Citation Notes

(Note: In-text citations appear as bracketed numbers. Uncited claims derive from our primary survey.)

This report uses rigorous citation practices to ensure verifiable claims while balancing academic rigor with readability:

## Academic Sources

We draw on peer-reviewed research (2022-2025) from established journals in AI ethics, machine learning, and business technology.

Key institutions include MIT, Stanford's HAI, Oxford's Institute for Ethics in AI, and NYU's AI Now Institute, providing frameworks for understanding bias mechanisms and mitigation approaches.

All sources undergo verification for relevance, methodological soundness, and absence of conflicts of interest.

## Industry Reports

Case studies come from verified corporate announcements, tech publications, and established business media.

Sources include Gartner, Forrester, IDC, PwC, Partnership on AI, and AI Now. Corporate examples derive from company newsrooms, earnings calls, and verified executive interviews.

We use established publications with strong editorial standards and cross-reference industry claims when possible.

## Primary Research

Most data points come from Better Together's original survey with appropriate statistical validation.

Our research includes responses from 3,850 consumers across diverse demographics and 425 business leaders from various industries, with rigorous statistical analysis.

Methodology details appear in Appendix A. Raw data and complete survey instruments are available upon request.

We note where our findings differ from existing research and present multiple perspectives for conflicting evidence. We commit to updating this report as new research emerges.