Patient Survey Report

Coalition for Health AI

January 2026

Introduction

Artificial intelligence in healthcare is advancing faster than federal policy can keep pace.

Nearly every state is introducing legislation to fill the gap, but approaches vary widely. Without a shared evidence base, policymakers risk creating a fragmented patchwork of rules that confuse providers, burden innovators, and leave patients unprotected.

By convening stakeholders and surfacing community perspectives, we can ensure that transparency requirements are grounded in real-world practice, preference and behaviors, aligned across jurisdictions, and effective in building trust.

Acknowledgments

This patient survey was made possible by the California Health Care Foundation, NORC at the University of Chicago and their AmeriSpeak® Panel, CHAI's policy workgroup, and the 150+ focus group attendees across patients, clinicians and vendors.

Thank you to all of you. We are looking forward to working alongside each of you to translate this into meaningfully higher trust in AI in healthcare.

California Health Care Foundation NORC at the University of Chicago CHAI

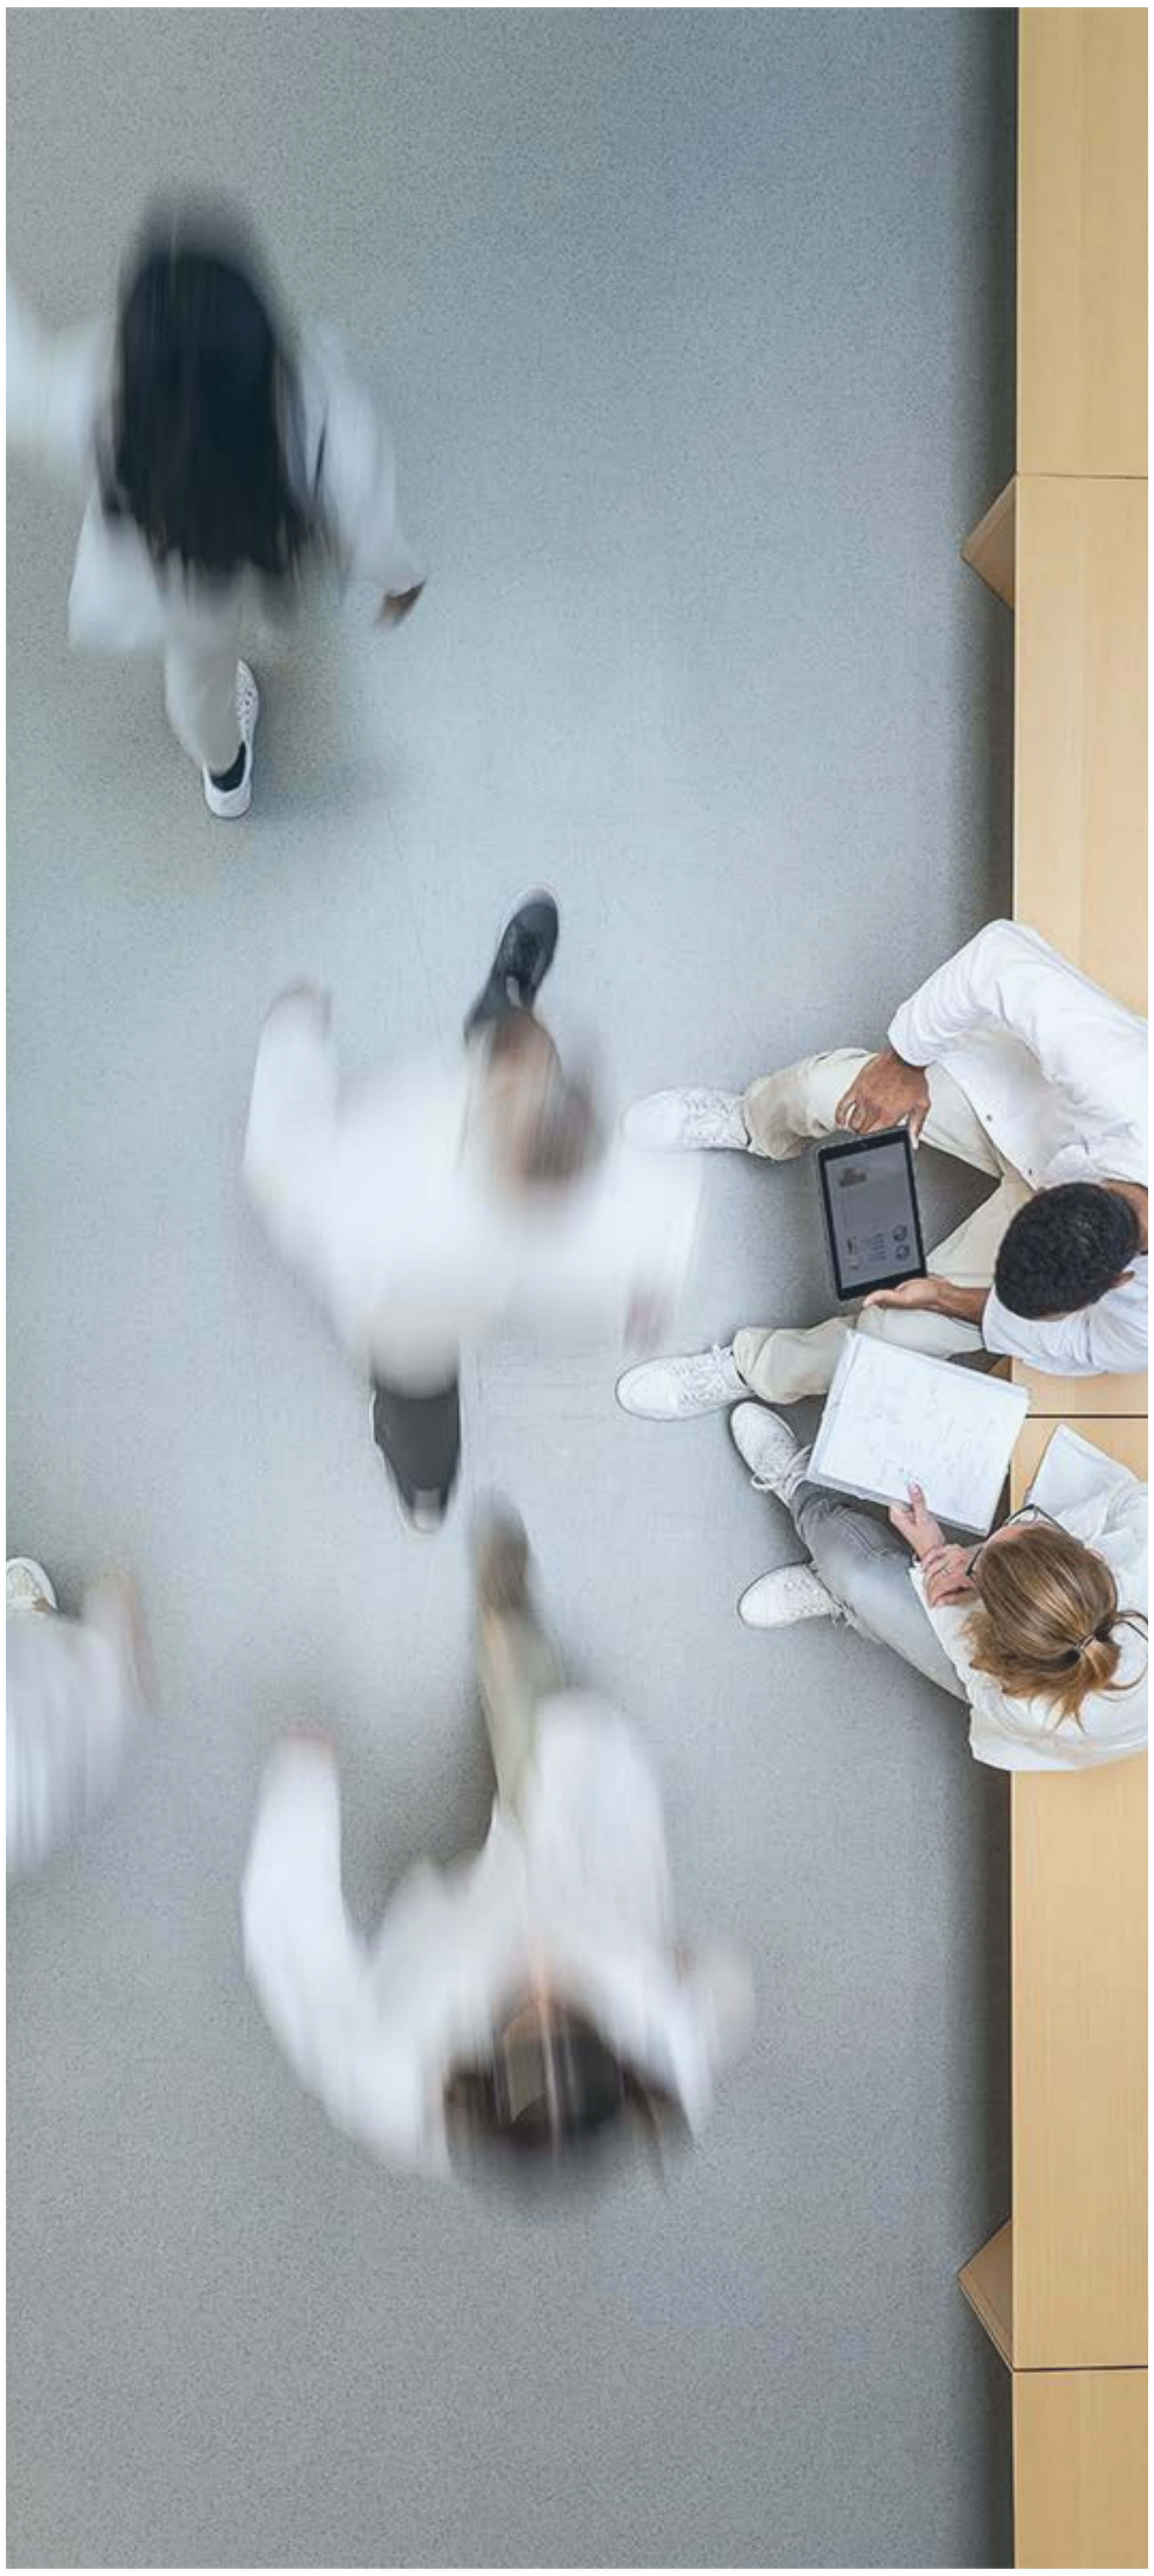

sanjeri via Getty Images

What is Transparency?

In this context, transparency is the extent to which information about an AI solution (e.g., capabilities, limitations, and purpose) and its output is available to all relevant stakeholders.

Purpose and Goals

Why we developed this

This research was developed to understand how the public experiences and evaluates the growing use of AI in healthcare. As AI becomes embedded into clinical and administrative workflows, public trust, consent, and accountability mechanisms have not kept pace.

The goal of this survey is to provide policymakers and health system leaders with data-driven evidence of where trust breaks down, what reassures the public, and which governance levers meaningfully affect legitimacy, so that emerging AI policies align with real public expectations and preferences.

Key Benefits

Why it matters

As we enter a new state legislative session, and as health systems rally to implement their own AI Governance, this work offers a foundation for shaping health AI governance that is grounded in public behaviors and preferences. We find that public trust is driven by oversight, accountability, patient control over data, and understanding data commercialization.

By identifying where AI use reduces trust, where human oversight restores confidence, and which accountability signals resonate most across populations, the findings contribute to informed policymaking on governance mechanisms that actually improve public confidence and protect patients.

Target Audience

Who it's for

This research is designed for policymakers at the federal and state level, as well as deployers of Health AI and those responsible for AI adoption and governance.

It is also relevant to standards-setting bodies, oversight organizations, and patient advocacy groups seeking evidence to inform accountability frameworks, consent models, and data governance approaches.

For developers and vendors, the findings provide insight into the conditions under which AI earns or loses public trust.

1456 patients surveyed

♀ 51% ♂ 49%

Age

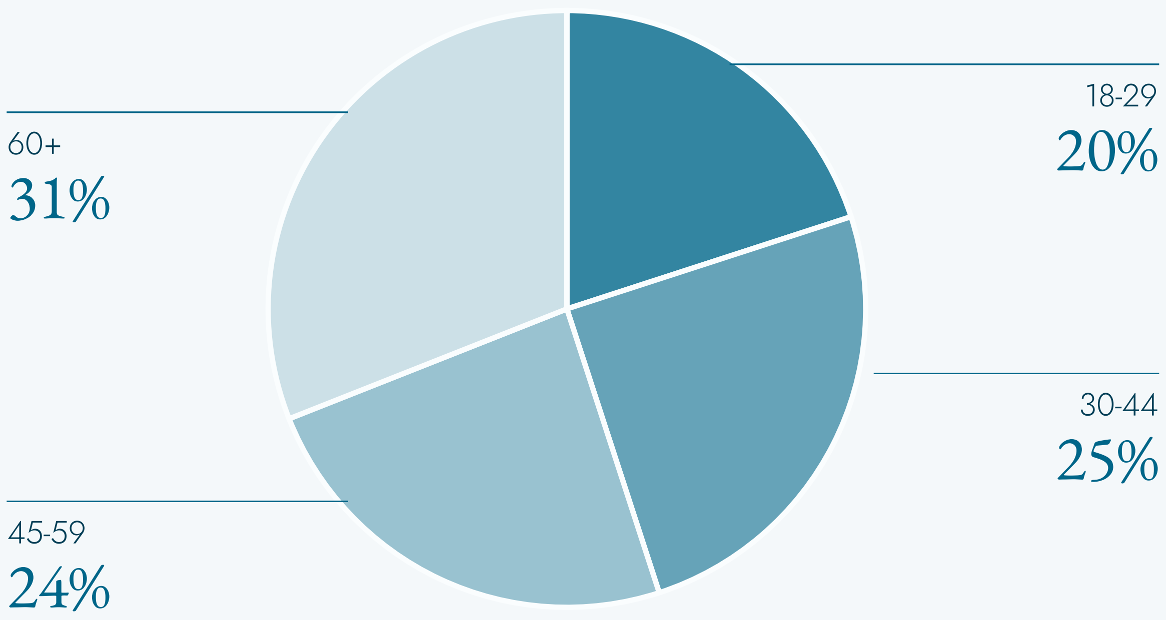

Education Level

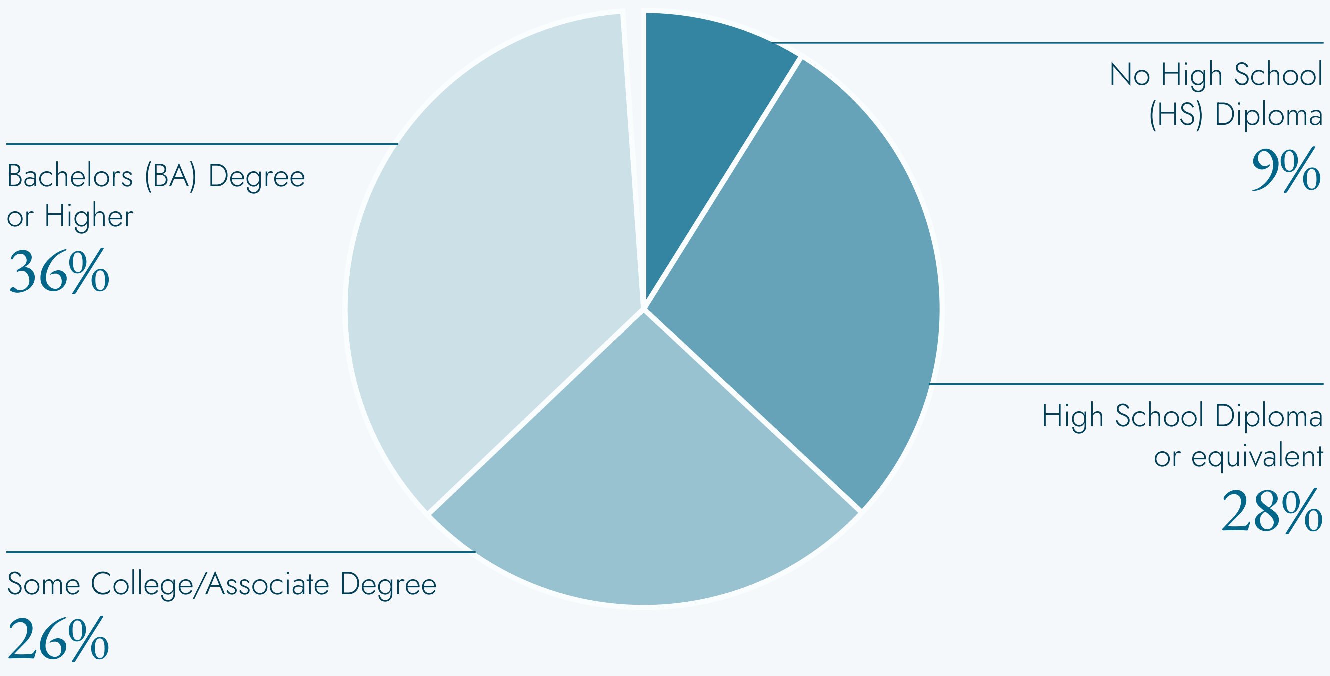

Annual Income

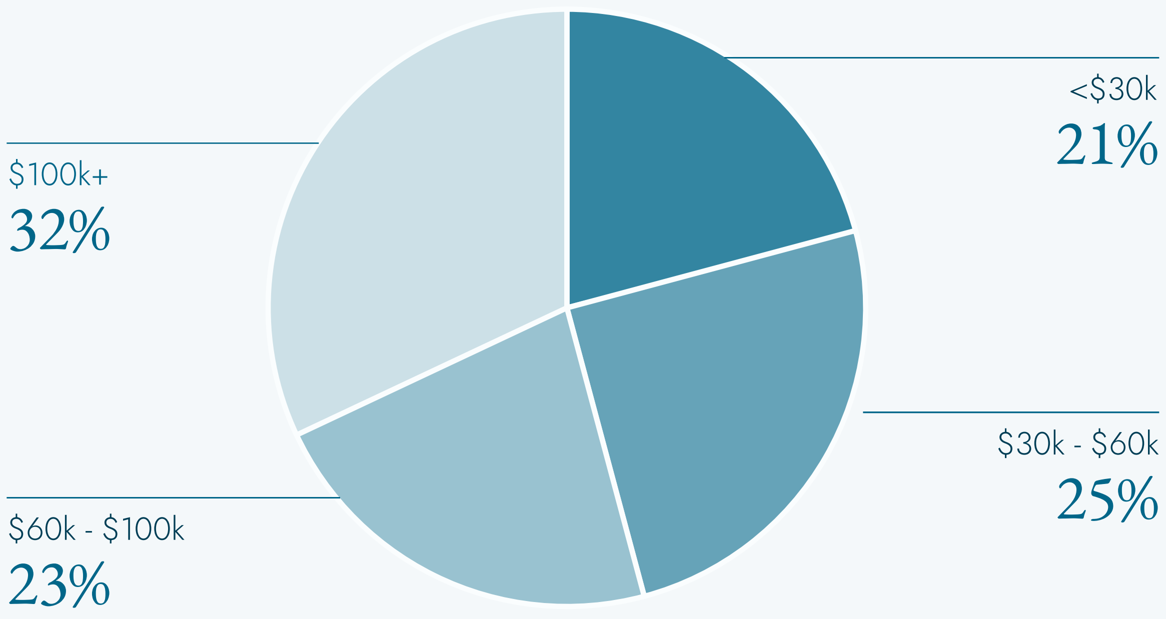

Race Or Ethnicity

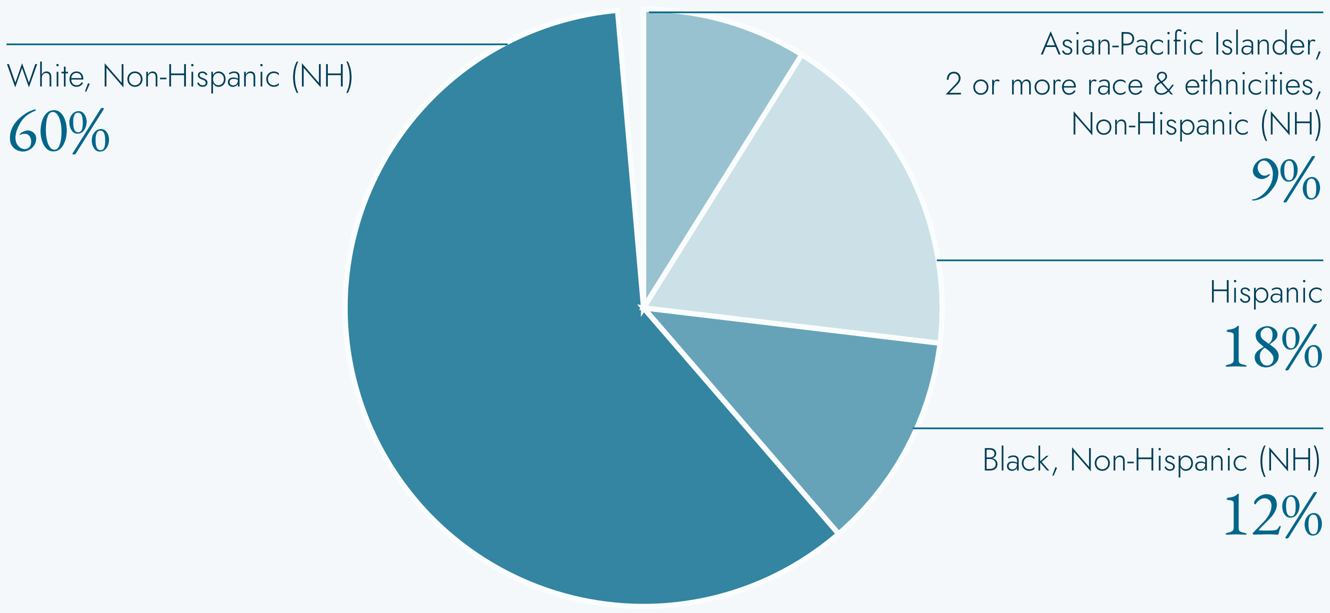

Executive Summary

People are using AI, but they don't trust it

Three-quarters of respondents report using AI, intentionally or unintentionally, yet only 13% feel very comfortable with it. Use, comfort, and trust are only modestly linked and highly context-dependent: people can be frequent users of AI while trusting it only in limited scenarios.

Concerns about data commercialization outweigh concerns about algorithmic bias; and notably 12% report never having considered AI bias at all. Most respondents know their data is being used to train AI and consistently want to retain primary control over it.

Transparency and disclosure alone will not increase trust

The public wants to know when AI is used in healthcare, but disclosure alone does not build trust. In many cases, awareness that AI was involved in care actually reduces trust. Trust is more closely linked to understanding safeguards, accountability, and the role AI plays in clinical decisions. Despite operational challenges, respondents strongly favor informed consent and opt-in approaches, particularly for emotionally or financially consequential decisions. Simple disclaimers are widely rejected.

Human oversight, multi-layered governance, and accountability measures will shift the dial

Human-in-the-loop oversight is critical, especially in high-stakes contexts affecting care pathways and reimbursement. Trust consistently increases when clinicians review AI outputs, with the largest gains seen in insurance-related use cases (despite overall trust remaining low).

No single institution is viewed as a trusted overseer of health AI. Instead, respondents favor multi-layered governance involving independent non-profits, health systems and provider networks, and federal regulators. Independent evaluation, bias testing, performance transparency across populations, and data transparency all significantly increase confidence.

So what

Trust in health AI will not be achieved through transparency or education alone. Public legitimacy depends on enforceable guardrails that preserve clinical judgment, enable meaningful consent, protect patient data, and establish clear accountability for outcomes. Policymakers should prioritize proportional, context-sensitive governance frameworks that reflect AI's growing invisibility in care delivery while reaffirming human oversight, patient authority, and shared responsibility across the health system.

This research provides an in-depth study into how the public currently experiences, understands the use of AI in healthcare today. Across demographics, we find that concern is less about AI itself, but more about how it is governed, restricted and accountable in ways that protect patients, their data, and improves their lives and experiences.

93%

of respondents report at least one concern about AI in healthcare

---

More than

4x

as many respondents say AI use makes them trust healthcare less (51%) than more (12%), on average across use cases and demographics

---

More than

80%

say their trust would increase if clear accountability measures were in place.

Contents

Introduction 1

Executive Summary 3

Usage of AI 5

Comfort with AI 6

Trust in Health AI 7

Trust by use case 8

Trusted bodies 9

Disclosure 10

Health Data 12

Concerns 17

Expectation 18

Excitement 19

What Matters Most 20

Age 20

Education Level 21

Race and Ethnicity 22

Household Income 23

Gender 24

About 25

Usage of AI

We found that the same number of respondents report intentionally using AI vs those who use it unintentionally (37%).

This presents a new challenge: AI is becoming embedded in everyday products faster than public awareness or consent mechanisms are keeping pace.

While intentional AI use is concentrated among younger, more advantaged populations, large segments of the population, particularly older adults, women, and white and multiracial respondents, are interacting with AI passively.

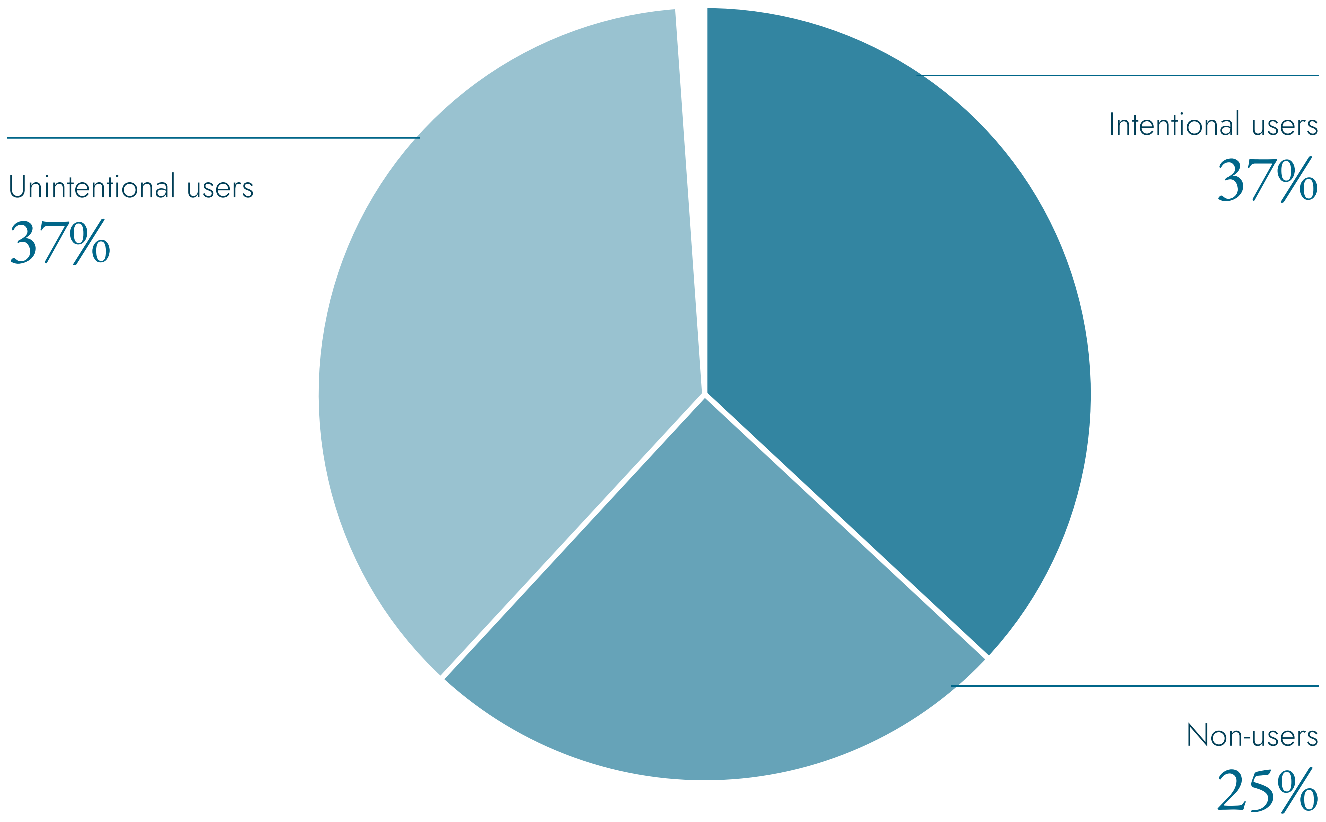

Intentional avoidance of AI is not confined to a single demographic group.

These findings underscore the need for meaningful transparency measures across all populations, particularly as AI becomes an invisible layer of decision-making in everyday use.

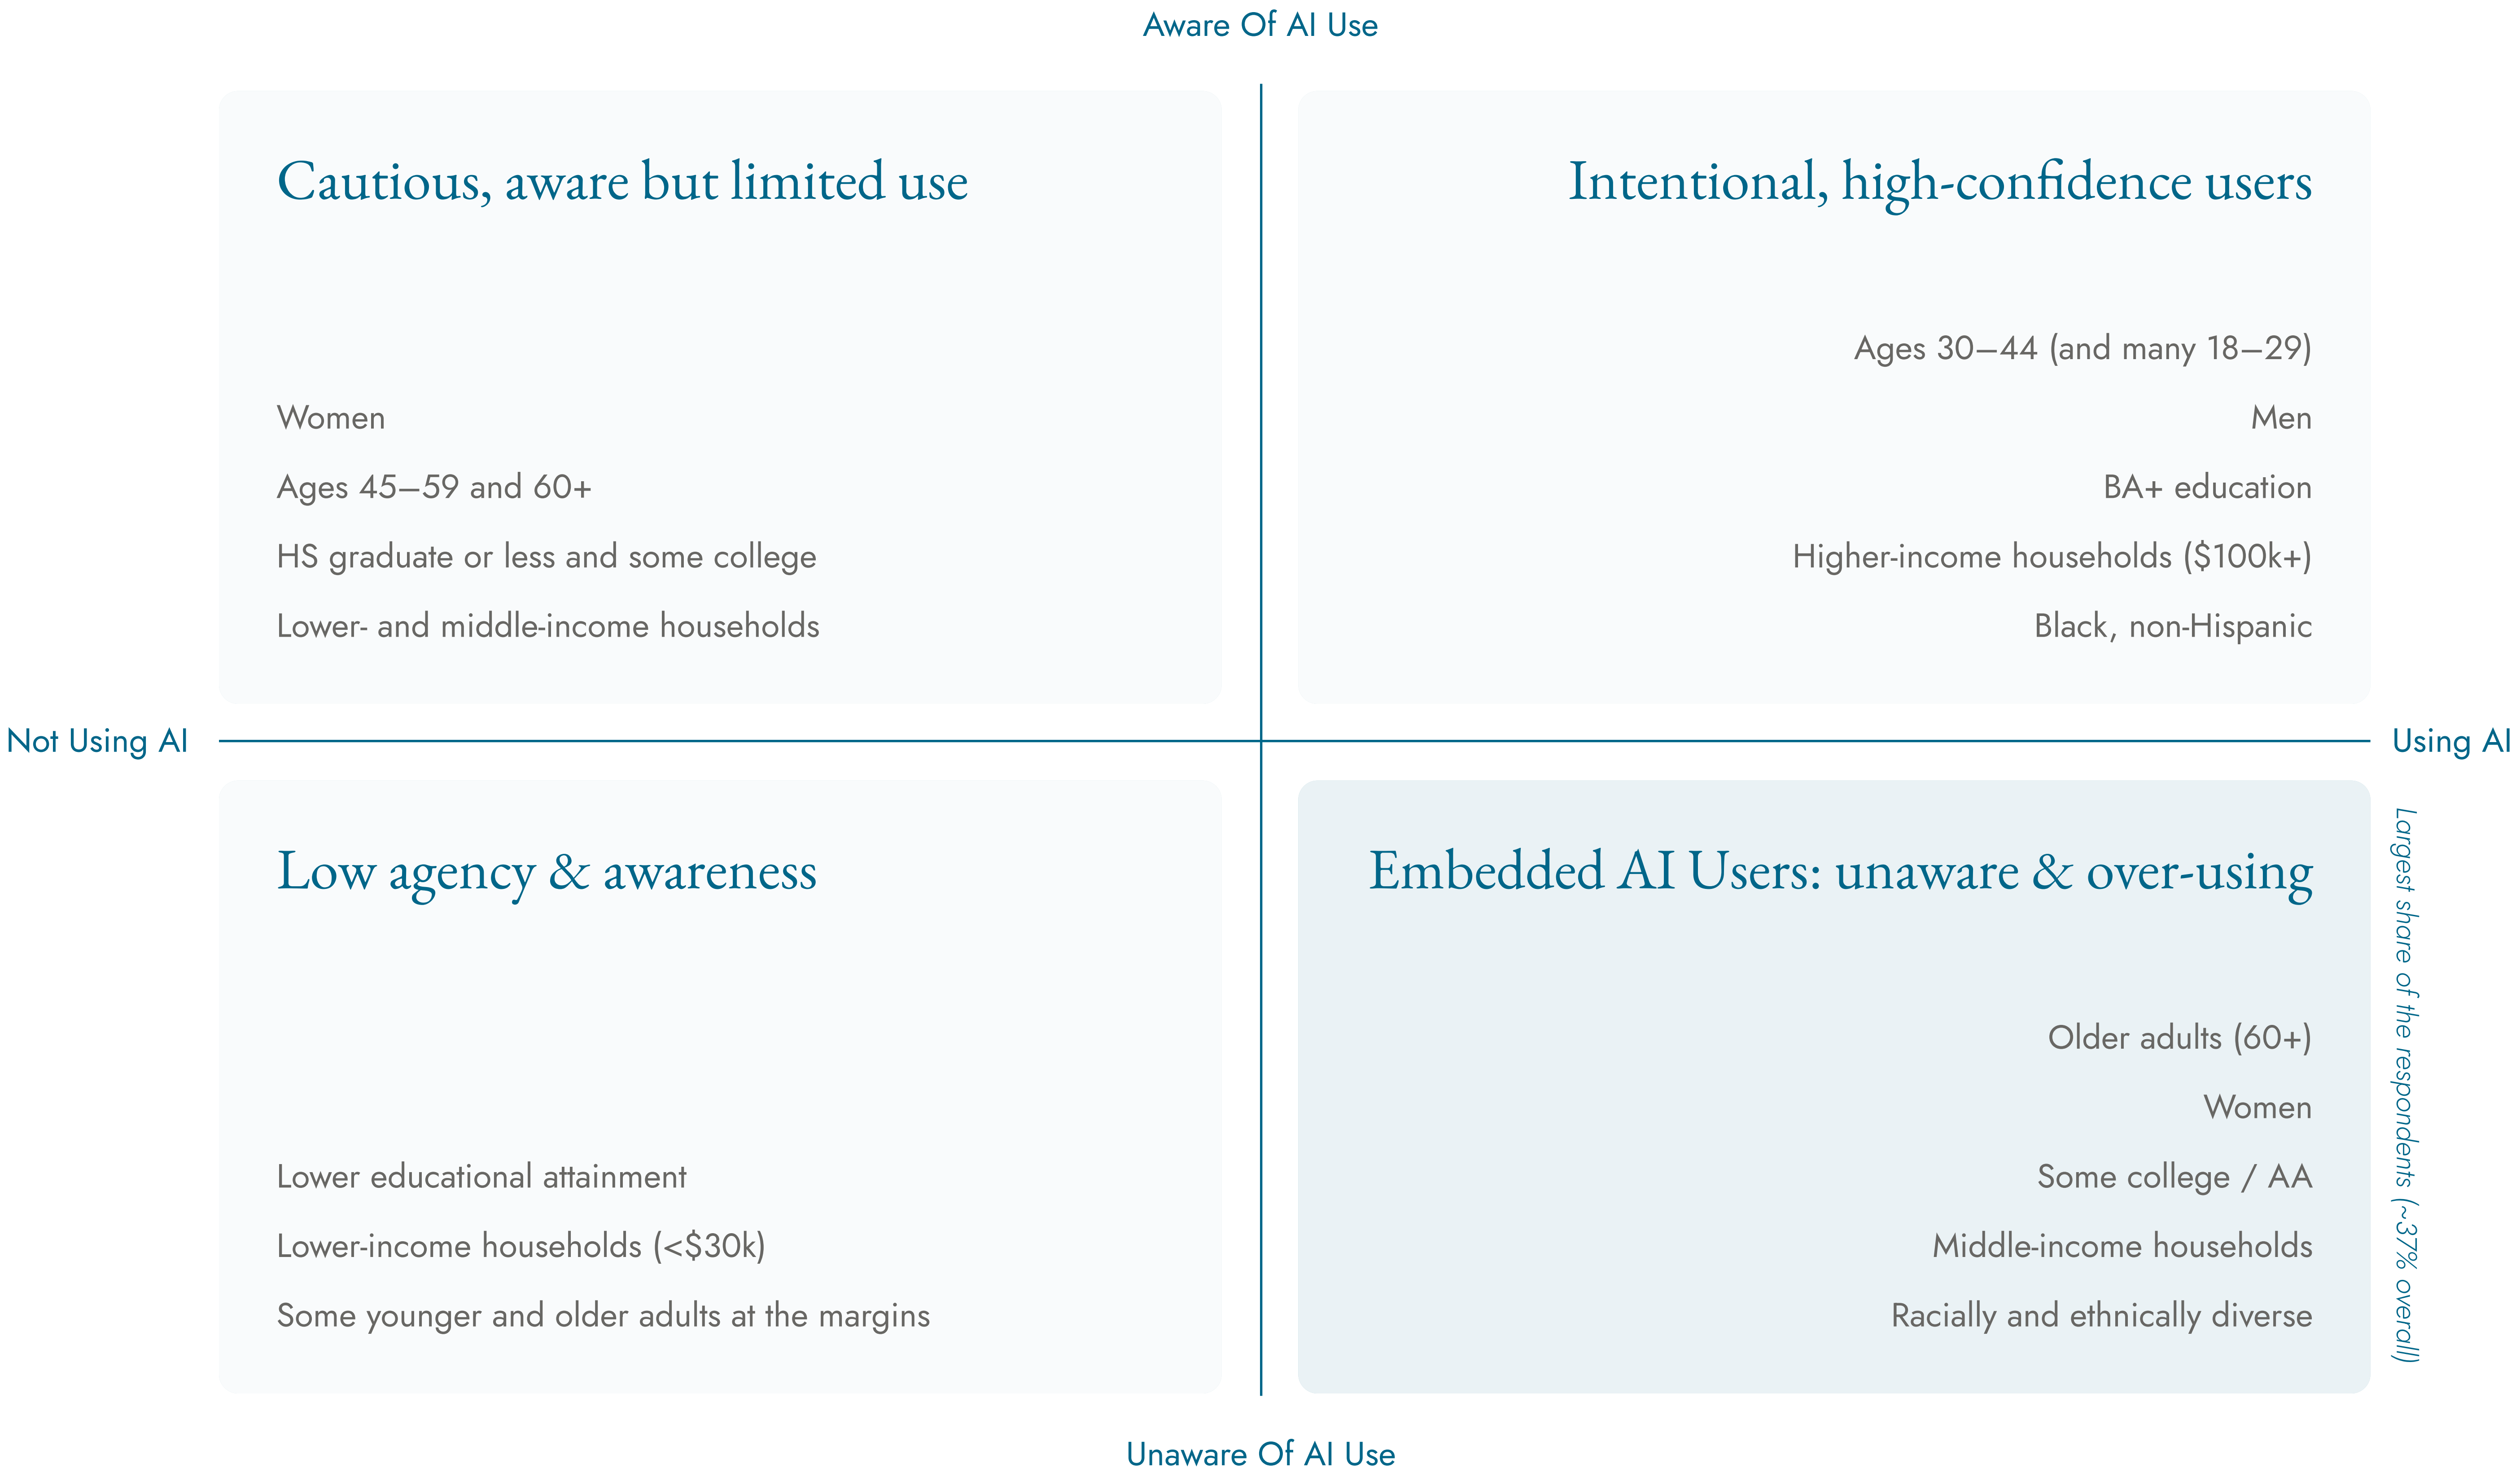

Cautious, aware but limited use

- Women
- Ages 45–59 and 60+
- HS graduate or less and some college
- Lower- and middle-income households

Intentional, high-confidence users

- Ages 30–44 (and many 18–29)
- Men
- BA+ education
- Higher-income households ($100k+)
- Black, non-Hispanic

Low agency & awareness

- Lower educational attainment
- Lower-income households (<$30k)
- Some younger and older adults at the margins

Embedded AI Users: unaware & over-using

- Older adults (60+)
- Women
- Some college / AA
- Middle-income households
- Racially and ethnically diverse

Largest share of the respondents (~37% overall)

Comfort with AI

On average, only 13% of respondents say they felt very comfortable with AI.

Analysis by demographic reveals unequal distributions of comfort across population groups:

- Higher levels of comfort are concentrated among adults aged 30—44, men, higher-income households, and Black respondents
- Notably, Black respondents report the highest comfort levels overall, challenging common assumptions that marginalized groups are uniformly more skeptical of emerging technologies.
- Older adults, women, lower-income households, and White non-Hispanic respondents are less comfortable with AI.

Collectively, these patterns indicate that public comfort with AI is not merely a function of awareness or technical literacy, and **understanding of its concerns or benefits are not well understood or widespread.**

Policy interventions should therefore move beyond education alone and focus on transparency, safeguards, and demonstrable benefit.

How comfortable are you with AI in general?

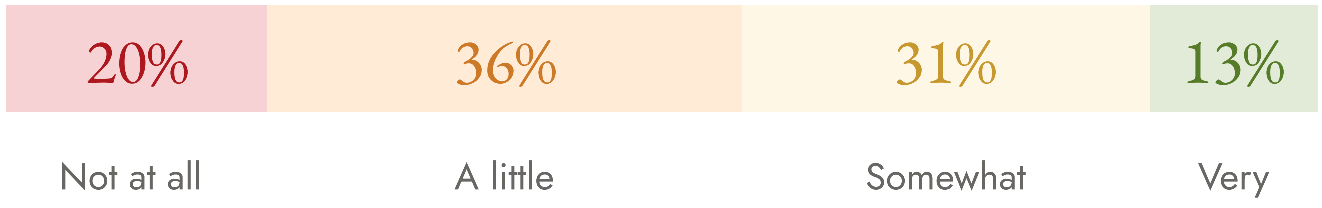

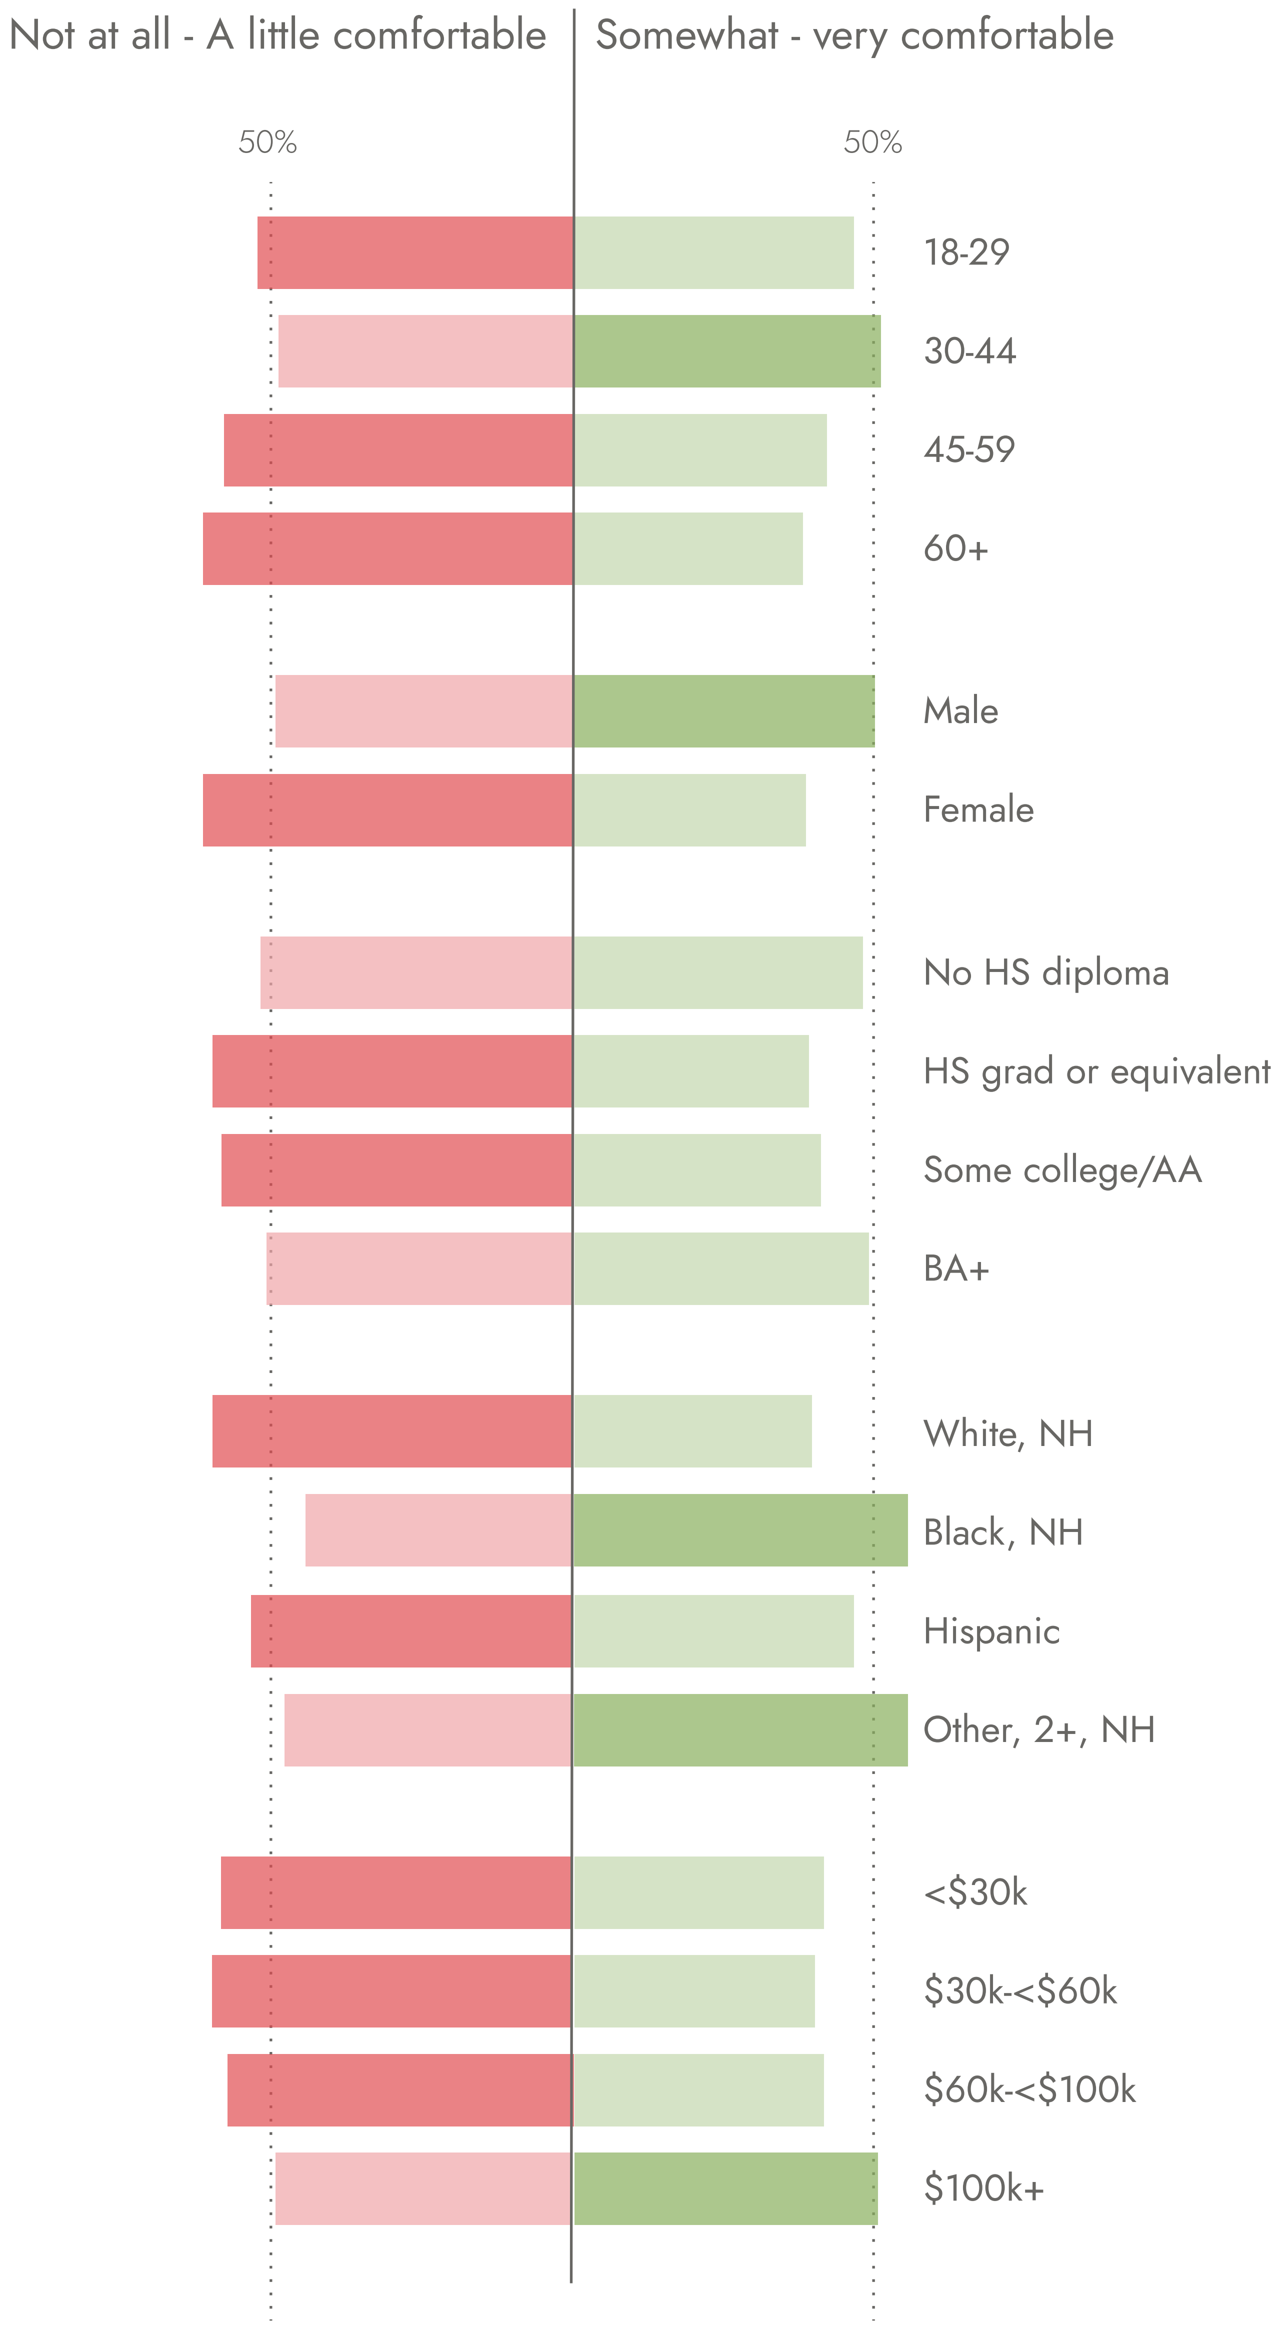

Trust in Health AI

Trust in healthcare AI appears to be modestly associated with intentional use and overall comfort.

Active AI users report slightly higher trust than passive users, and higher trust is more common among individuals with higher income or a college degree.

While comfort with AI is generally correlated with trust, the relationship is not uniform: some groups report high comfort but lower use, while others report higher use and trust despite lower comfort.

Across the data, there is no clear evidence of a group that regularly uses AI in healthcare while simultaneously expressing strong distrust.

Relationship Between Comfort and Trust in AI

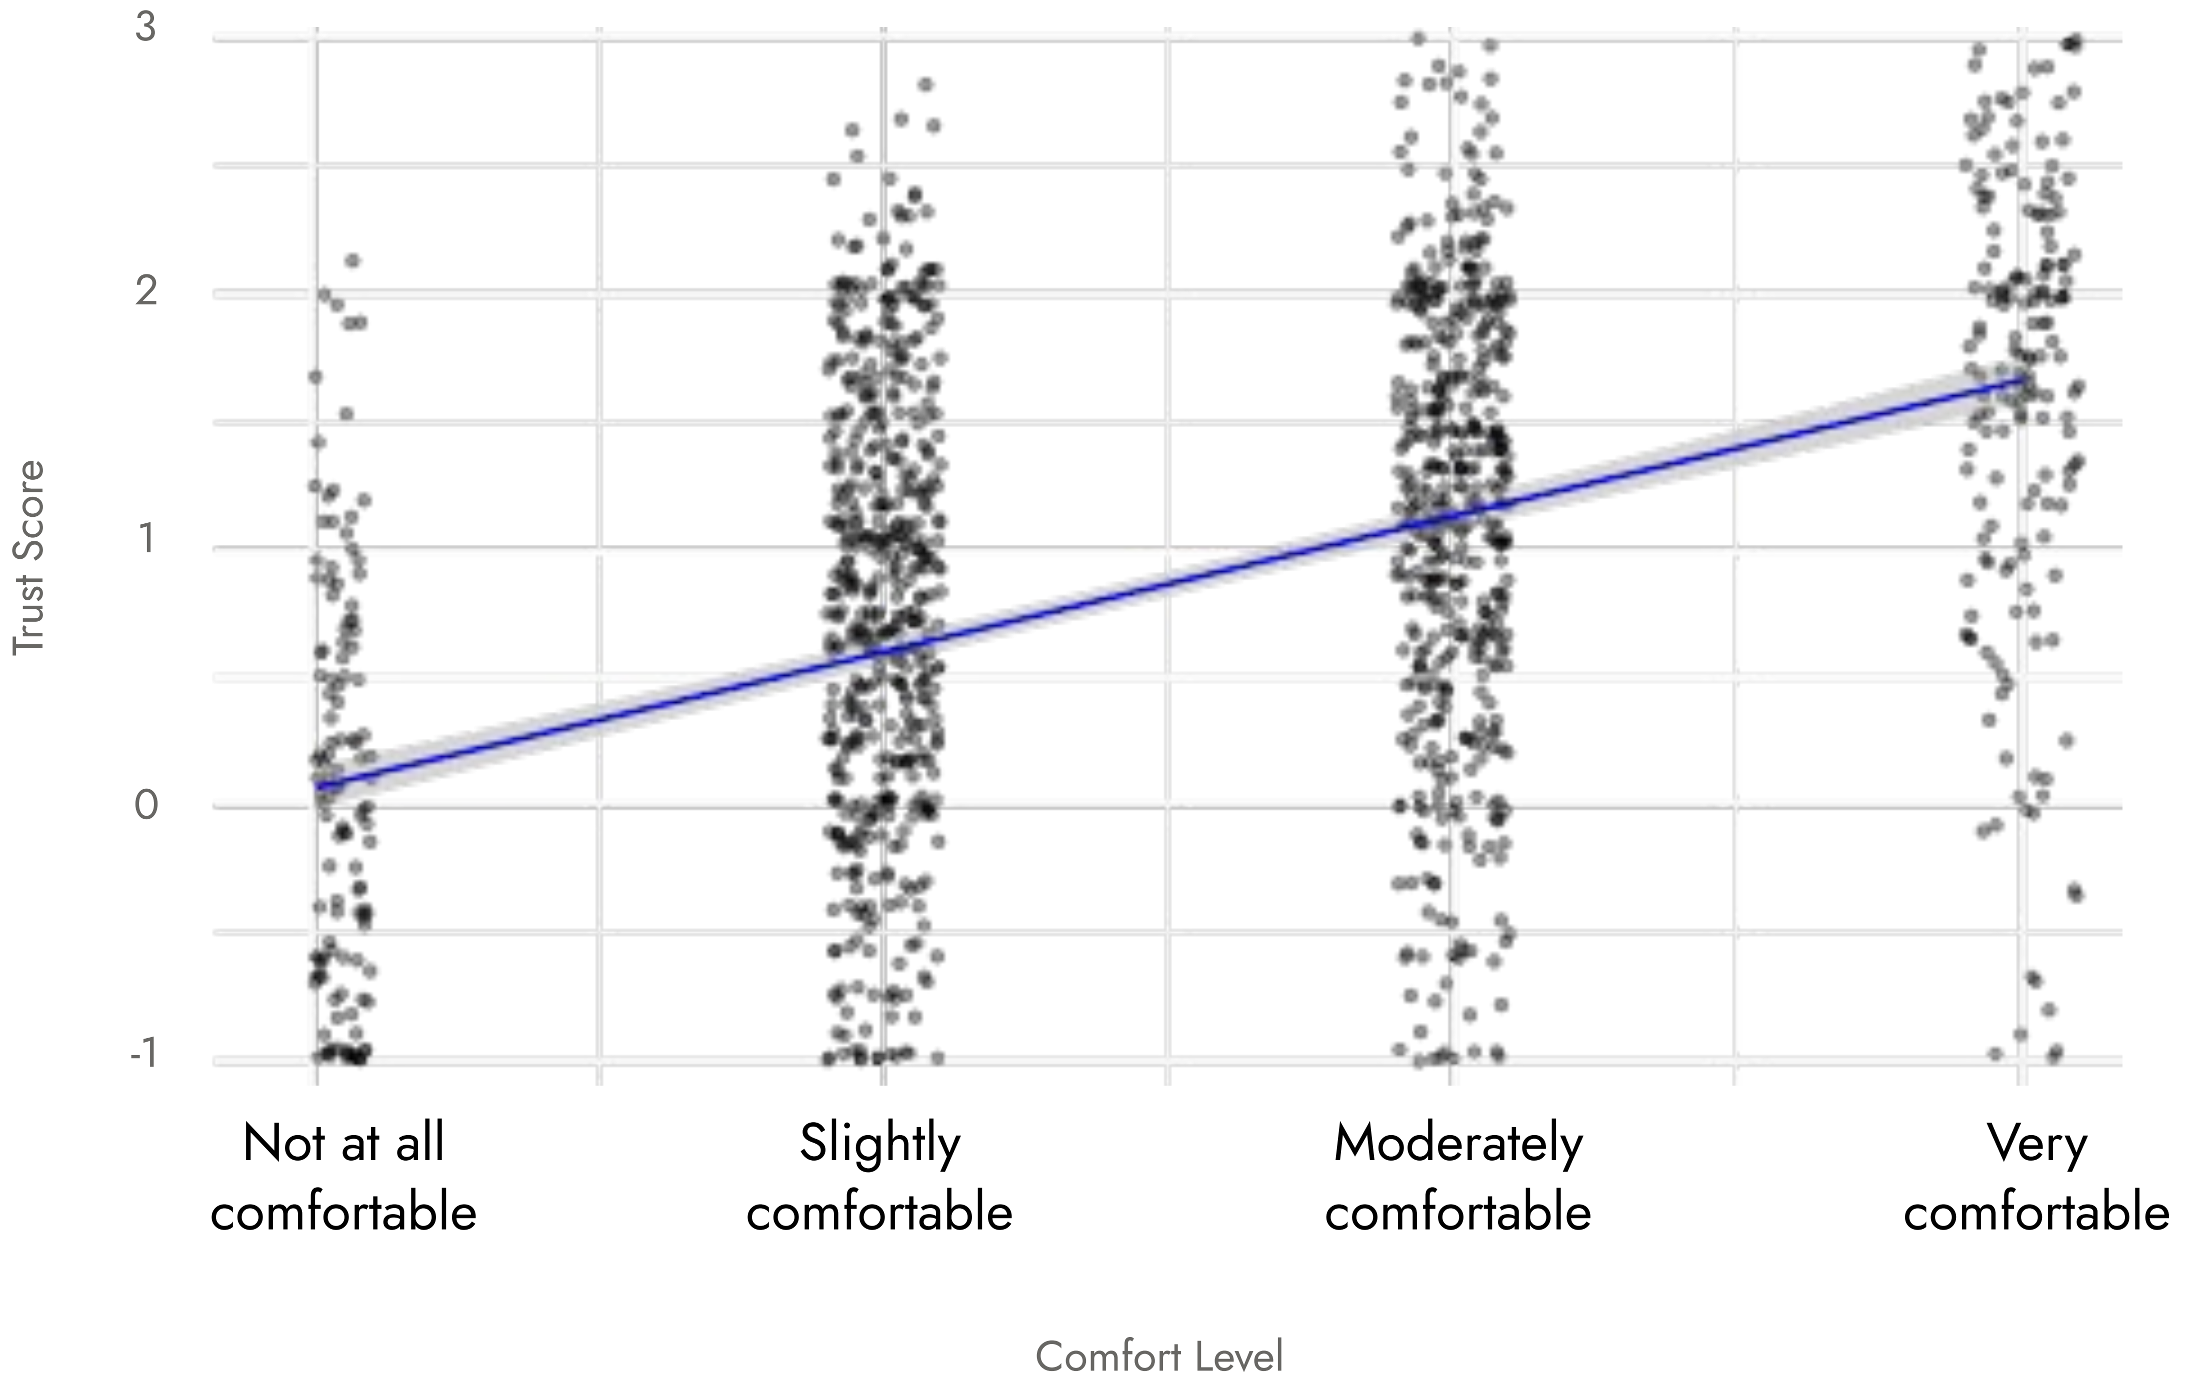

Trust by use case

On average, trust in AI-enabled use cases increases by roughly 7% when outputs are reviewed by a doctor or nurse.

This effect is strongest in clinical contexts, such as summarizing visit notes or assisting with diagnosis and treatment, where trust levels reach their highest (around 50% with oversight).

In contrast, insurance-related applications remain **substantially less trusted** overall, even when clinician review is introduced, with trust remaining at around 28% even with oversight. Notably, however, insurance pre-approval exhibits the largest marginal gain from oversight (+9%).

The largest gains from oversight are concentrated among respondents with higher education (+14% for BA+), higher income (+13% for $100k+ households), younger adults (+12% for ages 18–29), and men (+12%).

Taken together, these findings indicate that **oversight meaningfully reassures the public**, especially for low-trust use cases.

For policymakers, this underscores the importance of clearly defined human-in-the-loop requirements. However, oversight alone cannot resolve deep-seated trust issues in insurance and utilization management without broader transparency and accountability reforms.

How much do you trust using AI to help with your health care in each of the following ways?
How much would you trust AI in these same situations if a doctor or nurse reviewed and approved the results first?

Trusted bodies

While an independent non-profit is the most trusted oversight option, trust is not concentrated in any single institution.

With the top four options each receiving roughly 40–50% support, the public appears to favor shared responsibility over a single guarantor of AI safety and performance.

Trust in AI and oversight varies by demographic group: older, higher-income, and more highly educated respondents favor independent, clinical, and federal oversight, while younger adults, lower-income groups, and different racial and ethnic communities distribute trust across developers, insurers, and state or public health authorities.

No single oversight model emerges and effective AI governance will therefore require layered oversight structures.

Who should make sure the AI is safe and works well?

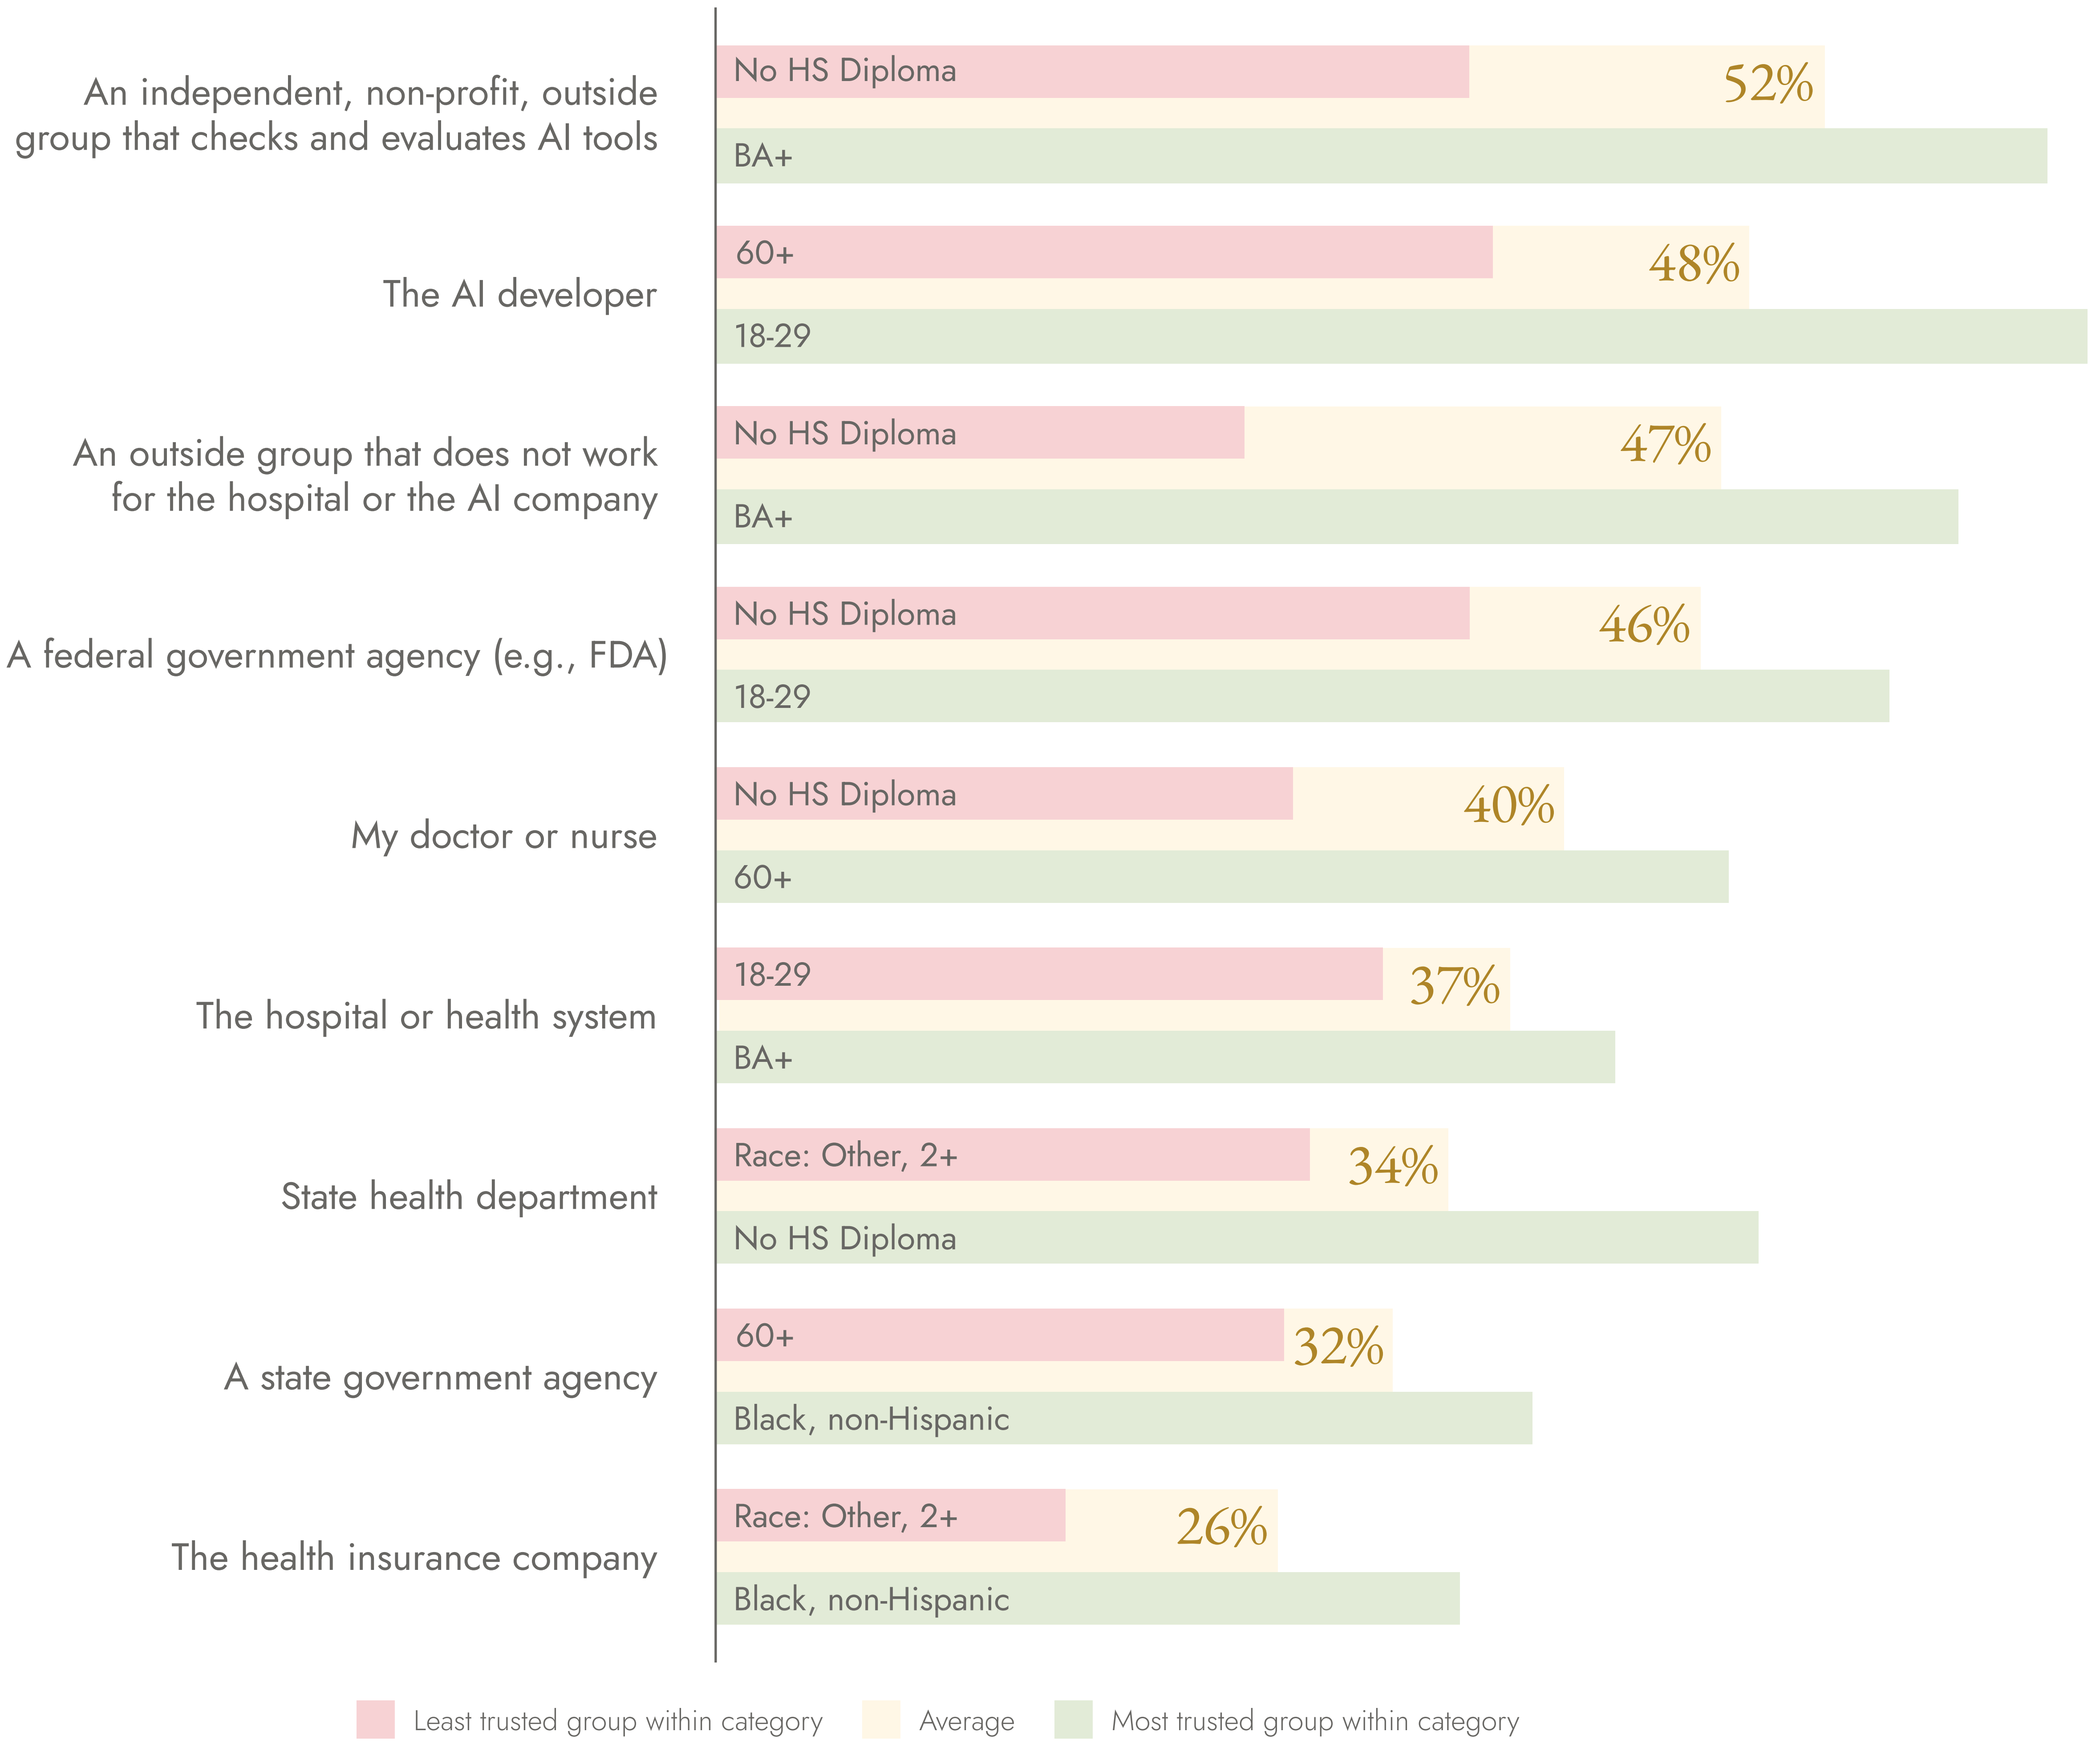

Disclosure

Though unsurprising, **the findings show a consistent call for disclosure across all use cases,** with an average of ~83% calling for some level of communication when AI is being used.

The expectation is highest for perceived "high-stakes", consequential use cases, such as diagnosis, treatment, and insurance decisions.

Across all healthcare AI use cases, the public strongly favors informed consent or "opt-in" and rejects AI use with disclaimers almost entirely.

Support for opt-in consent is highest in high-stakes uses like mental health, insurance, and diagnosis. Preference for opt-in is strongest among younger, lower-income, and less-educated respondents, and among Black and Hispanic respondents (~78%), while White respondents are more open to opt-out, with about 24% preferring it.

These patterns suggest that consent preferences reflect differing experiences of healthcare systems. Policymakers should prioritize proportional, context-sensitive consent frameworks that preserve user agency while accommodating varying expectations across populations.

Should your doctor or nurse tell you when AI is used when...

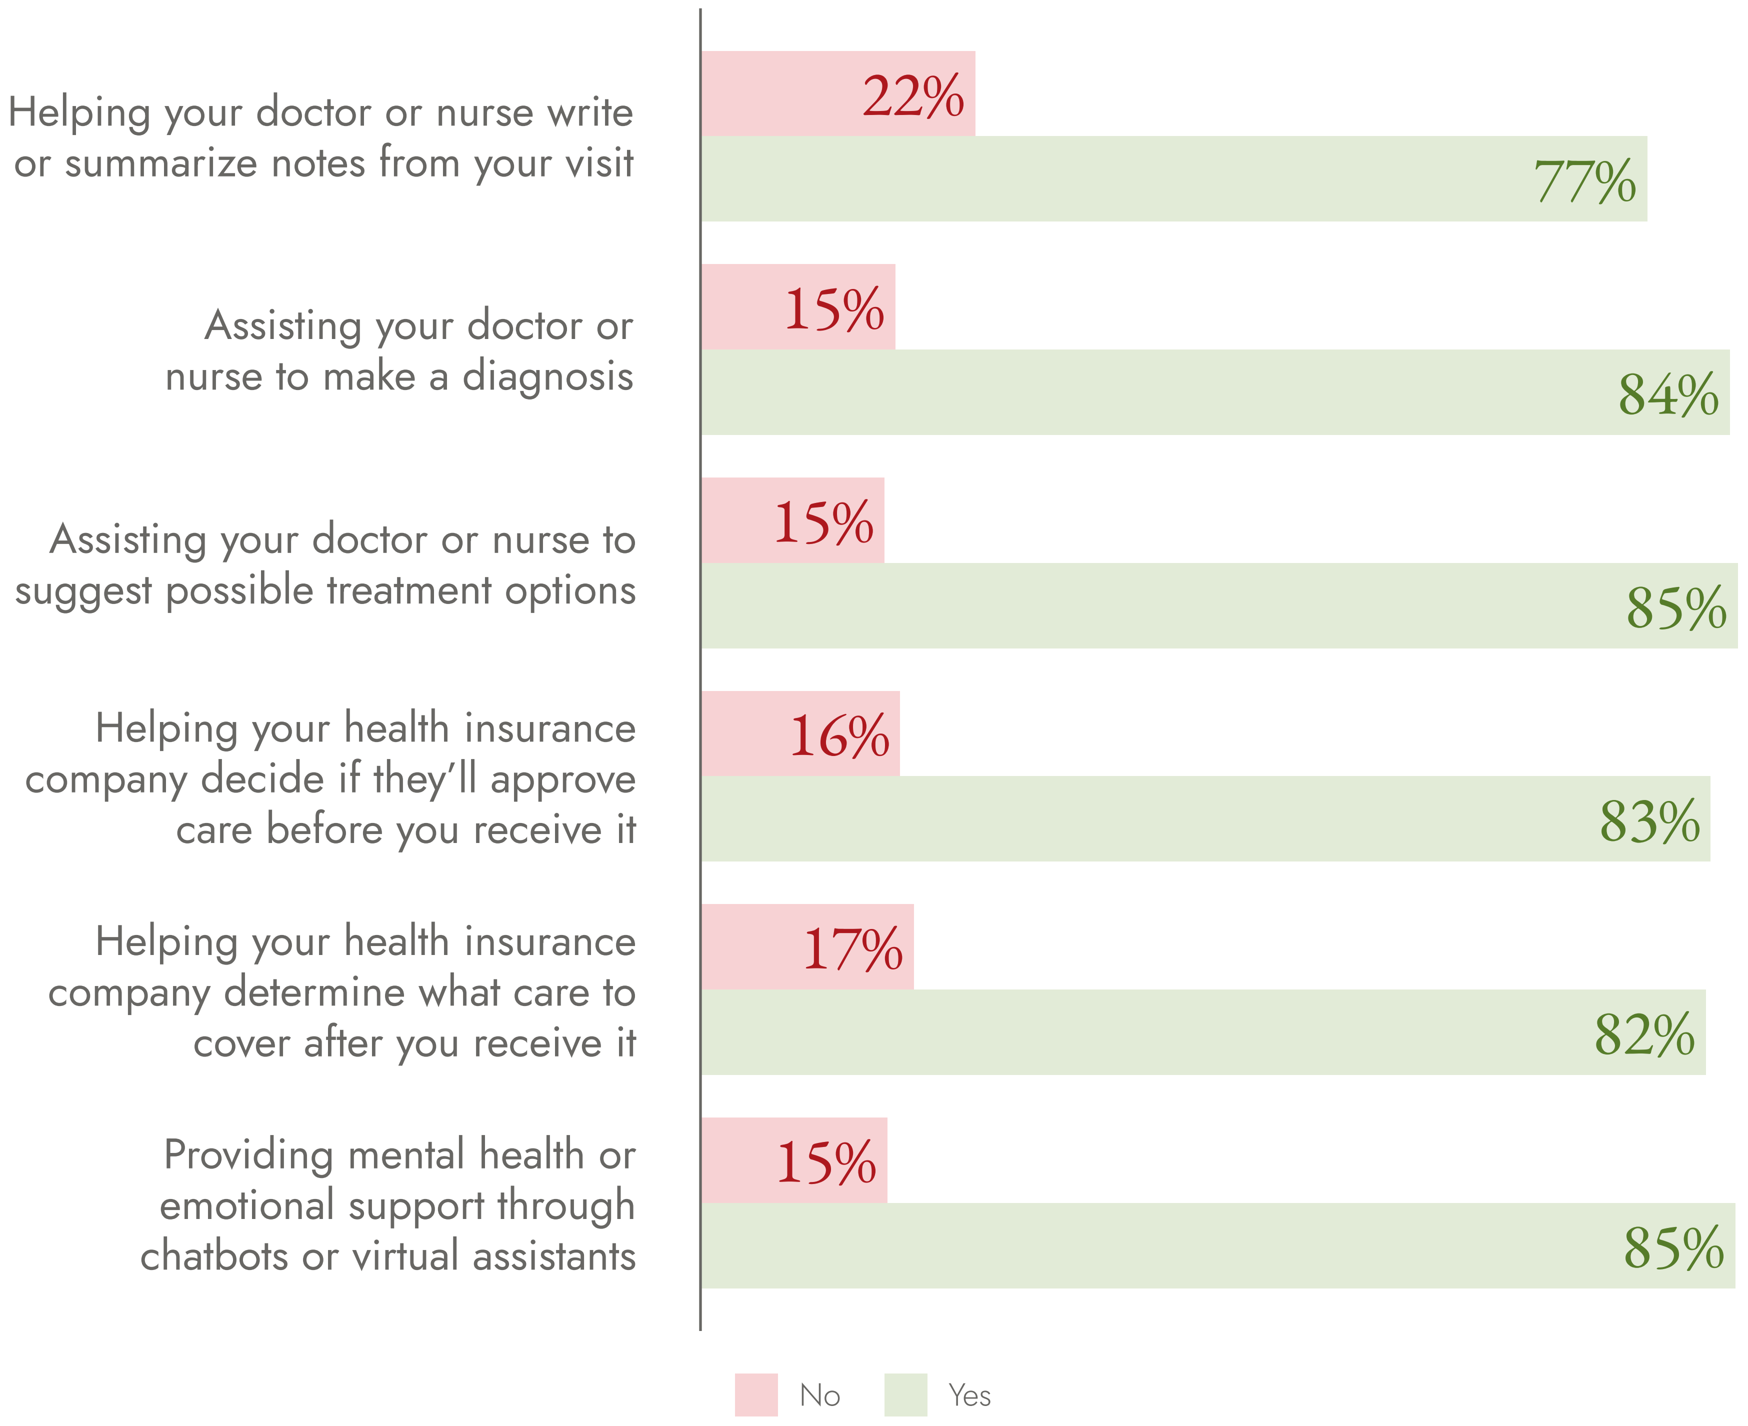

Should your doctor or nurse ask for your permission before using AI when...

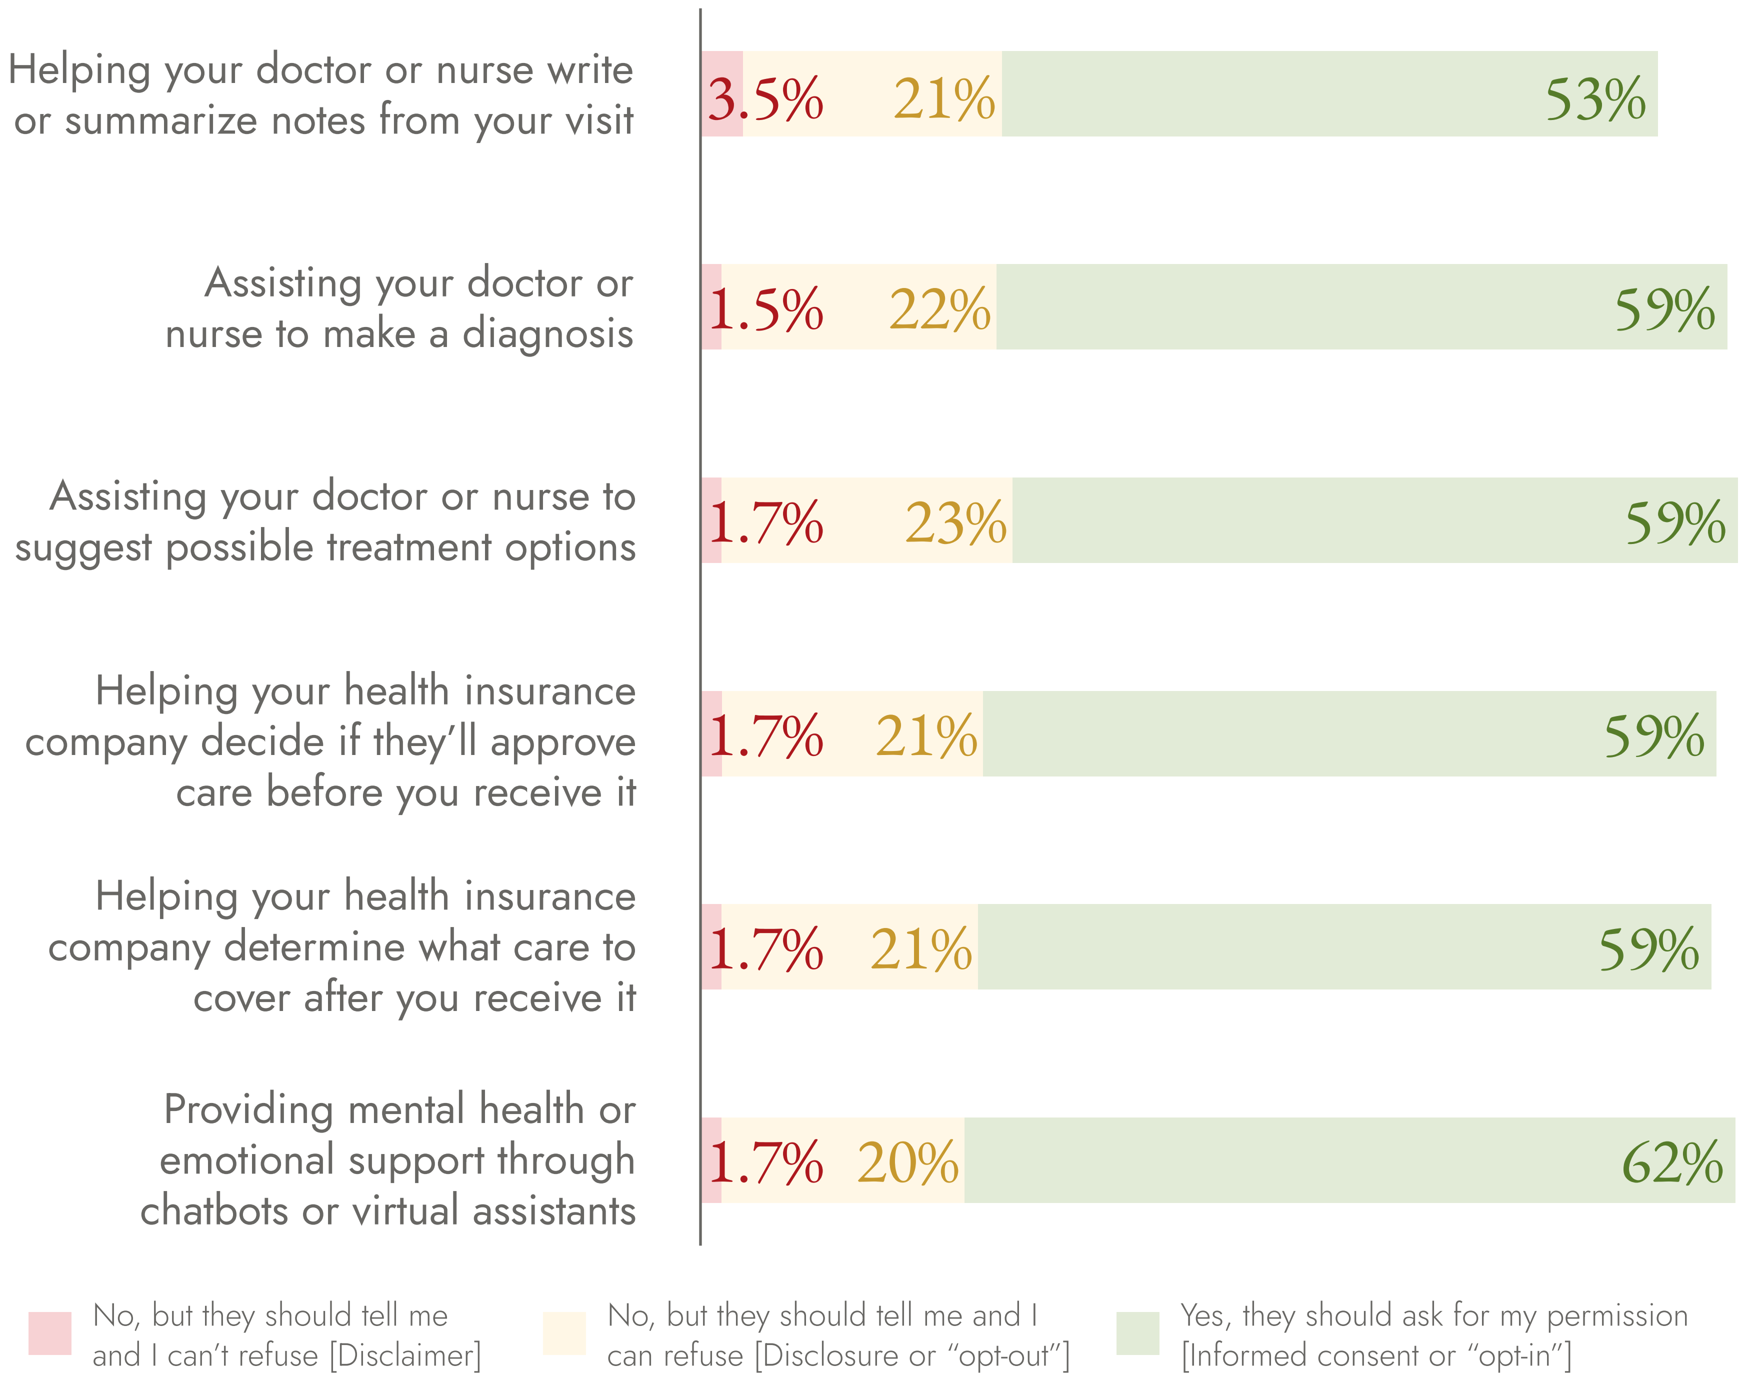

Use of AI, even with disclosure, does not increase trust in healthcare, and in many cases, it reduces it.

Across all use cases, more respondents report decreased rather than increased trust when told AI was involved. Loss in trust is smallest for lower-risk, supportive functions such as summarizing notes and matching patients to clinical trials, and largest for insurance decisions and mental health support, where more than half of respondents say AI use would make them trust the care less.

Importantly, disclosure does not affect all groups equally: older adults, men, higher-income respondents, college graduates, and Black respondents experience smaller trust declines, while younger adults, women, lower-income groups, and those with some college education show sharper drops in trust.

This pattern suggests that the use of AI in healthcare, even if you disclose it, is not trusted. Disclosure alone is insufficient in trust-building.

If your doctor or nurse told you that AI was used in your care, how would that affect your trust in the care you received?

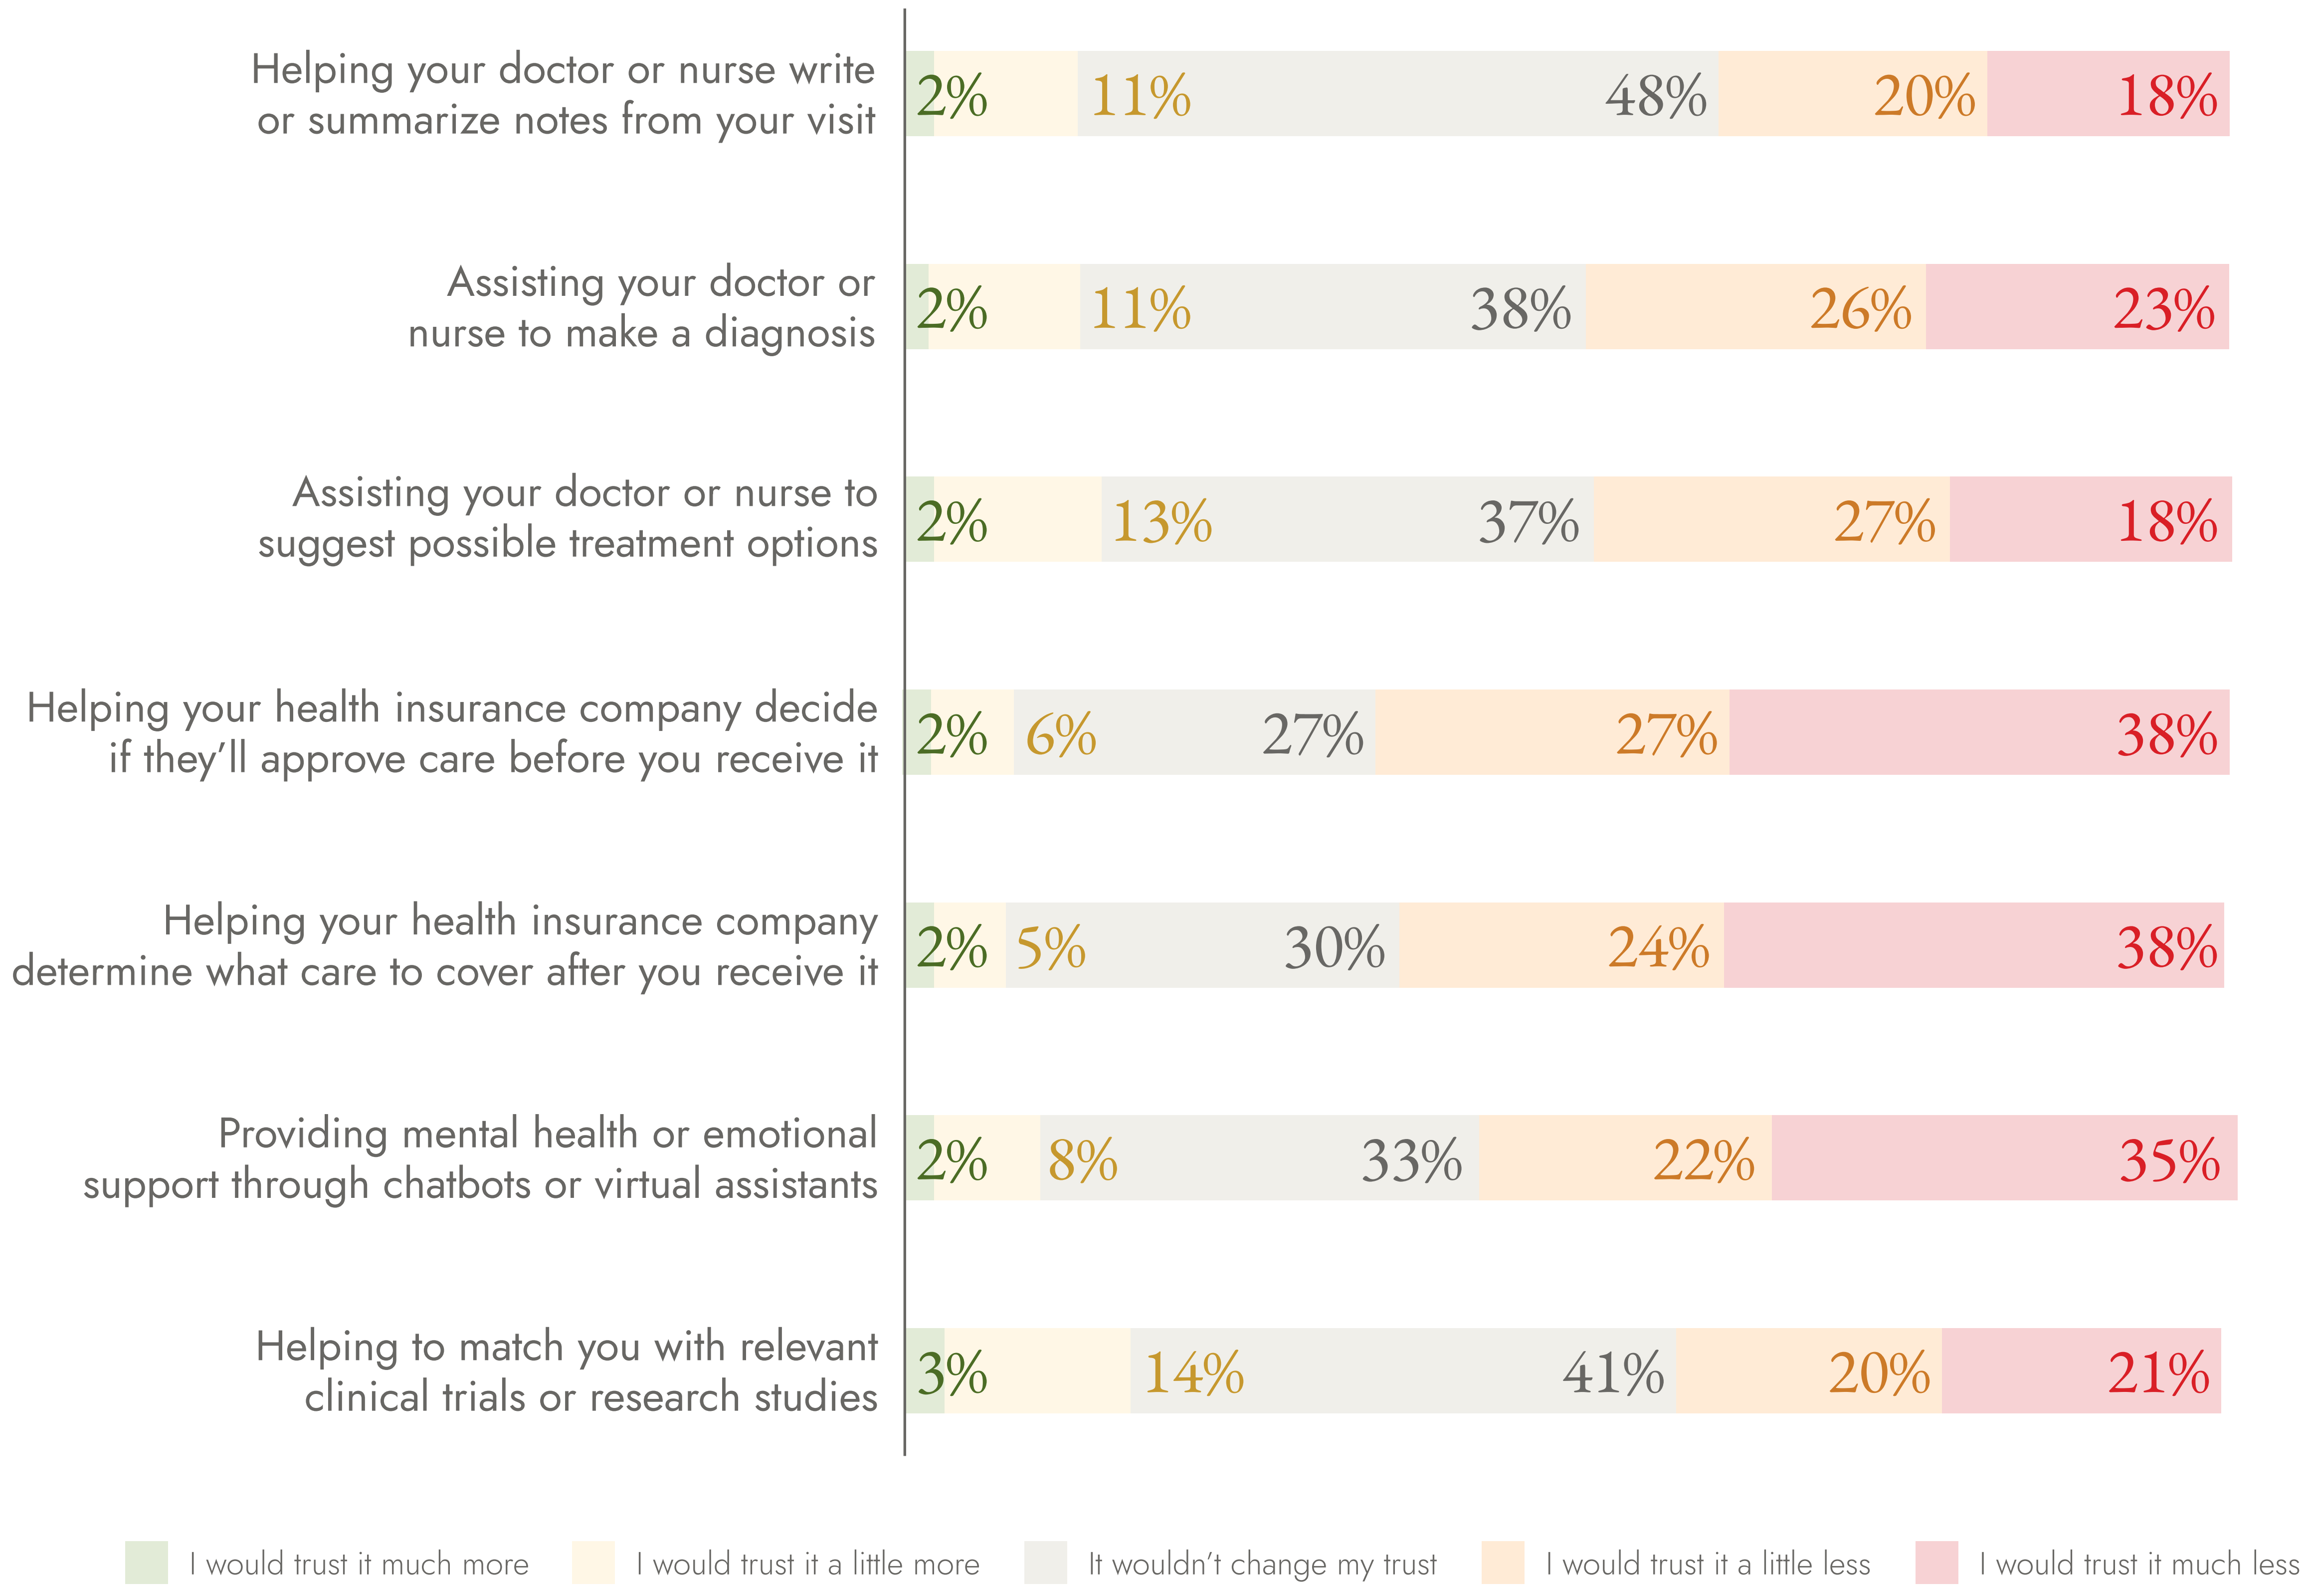

Health Data

Awareness of health data use for AI training is consistently high, with an average of 70% across all populations.

There is notably lower awareness among adults 60+ (59%) and higher awareness among higher-income households ($100k+: 78%).

Interestingly, higher awareness of data use does not equal more distrust in AI. For instance, women report lower awareness than men (66% vs 74%) yet markedly higher distrust (56% trust less vs 46% for men). The highest-income group is most aware (78%) but least likely to say they trust healthcare less (46%) if AI is used.

Do you think your personal health data is currently being used to train AI?

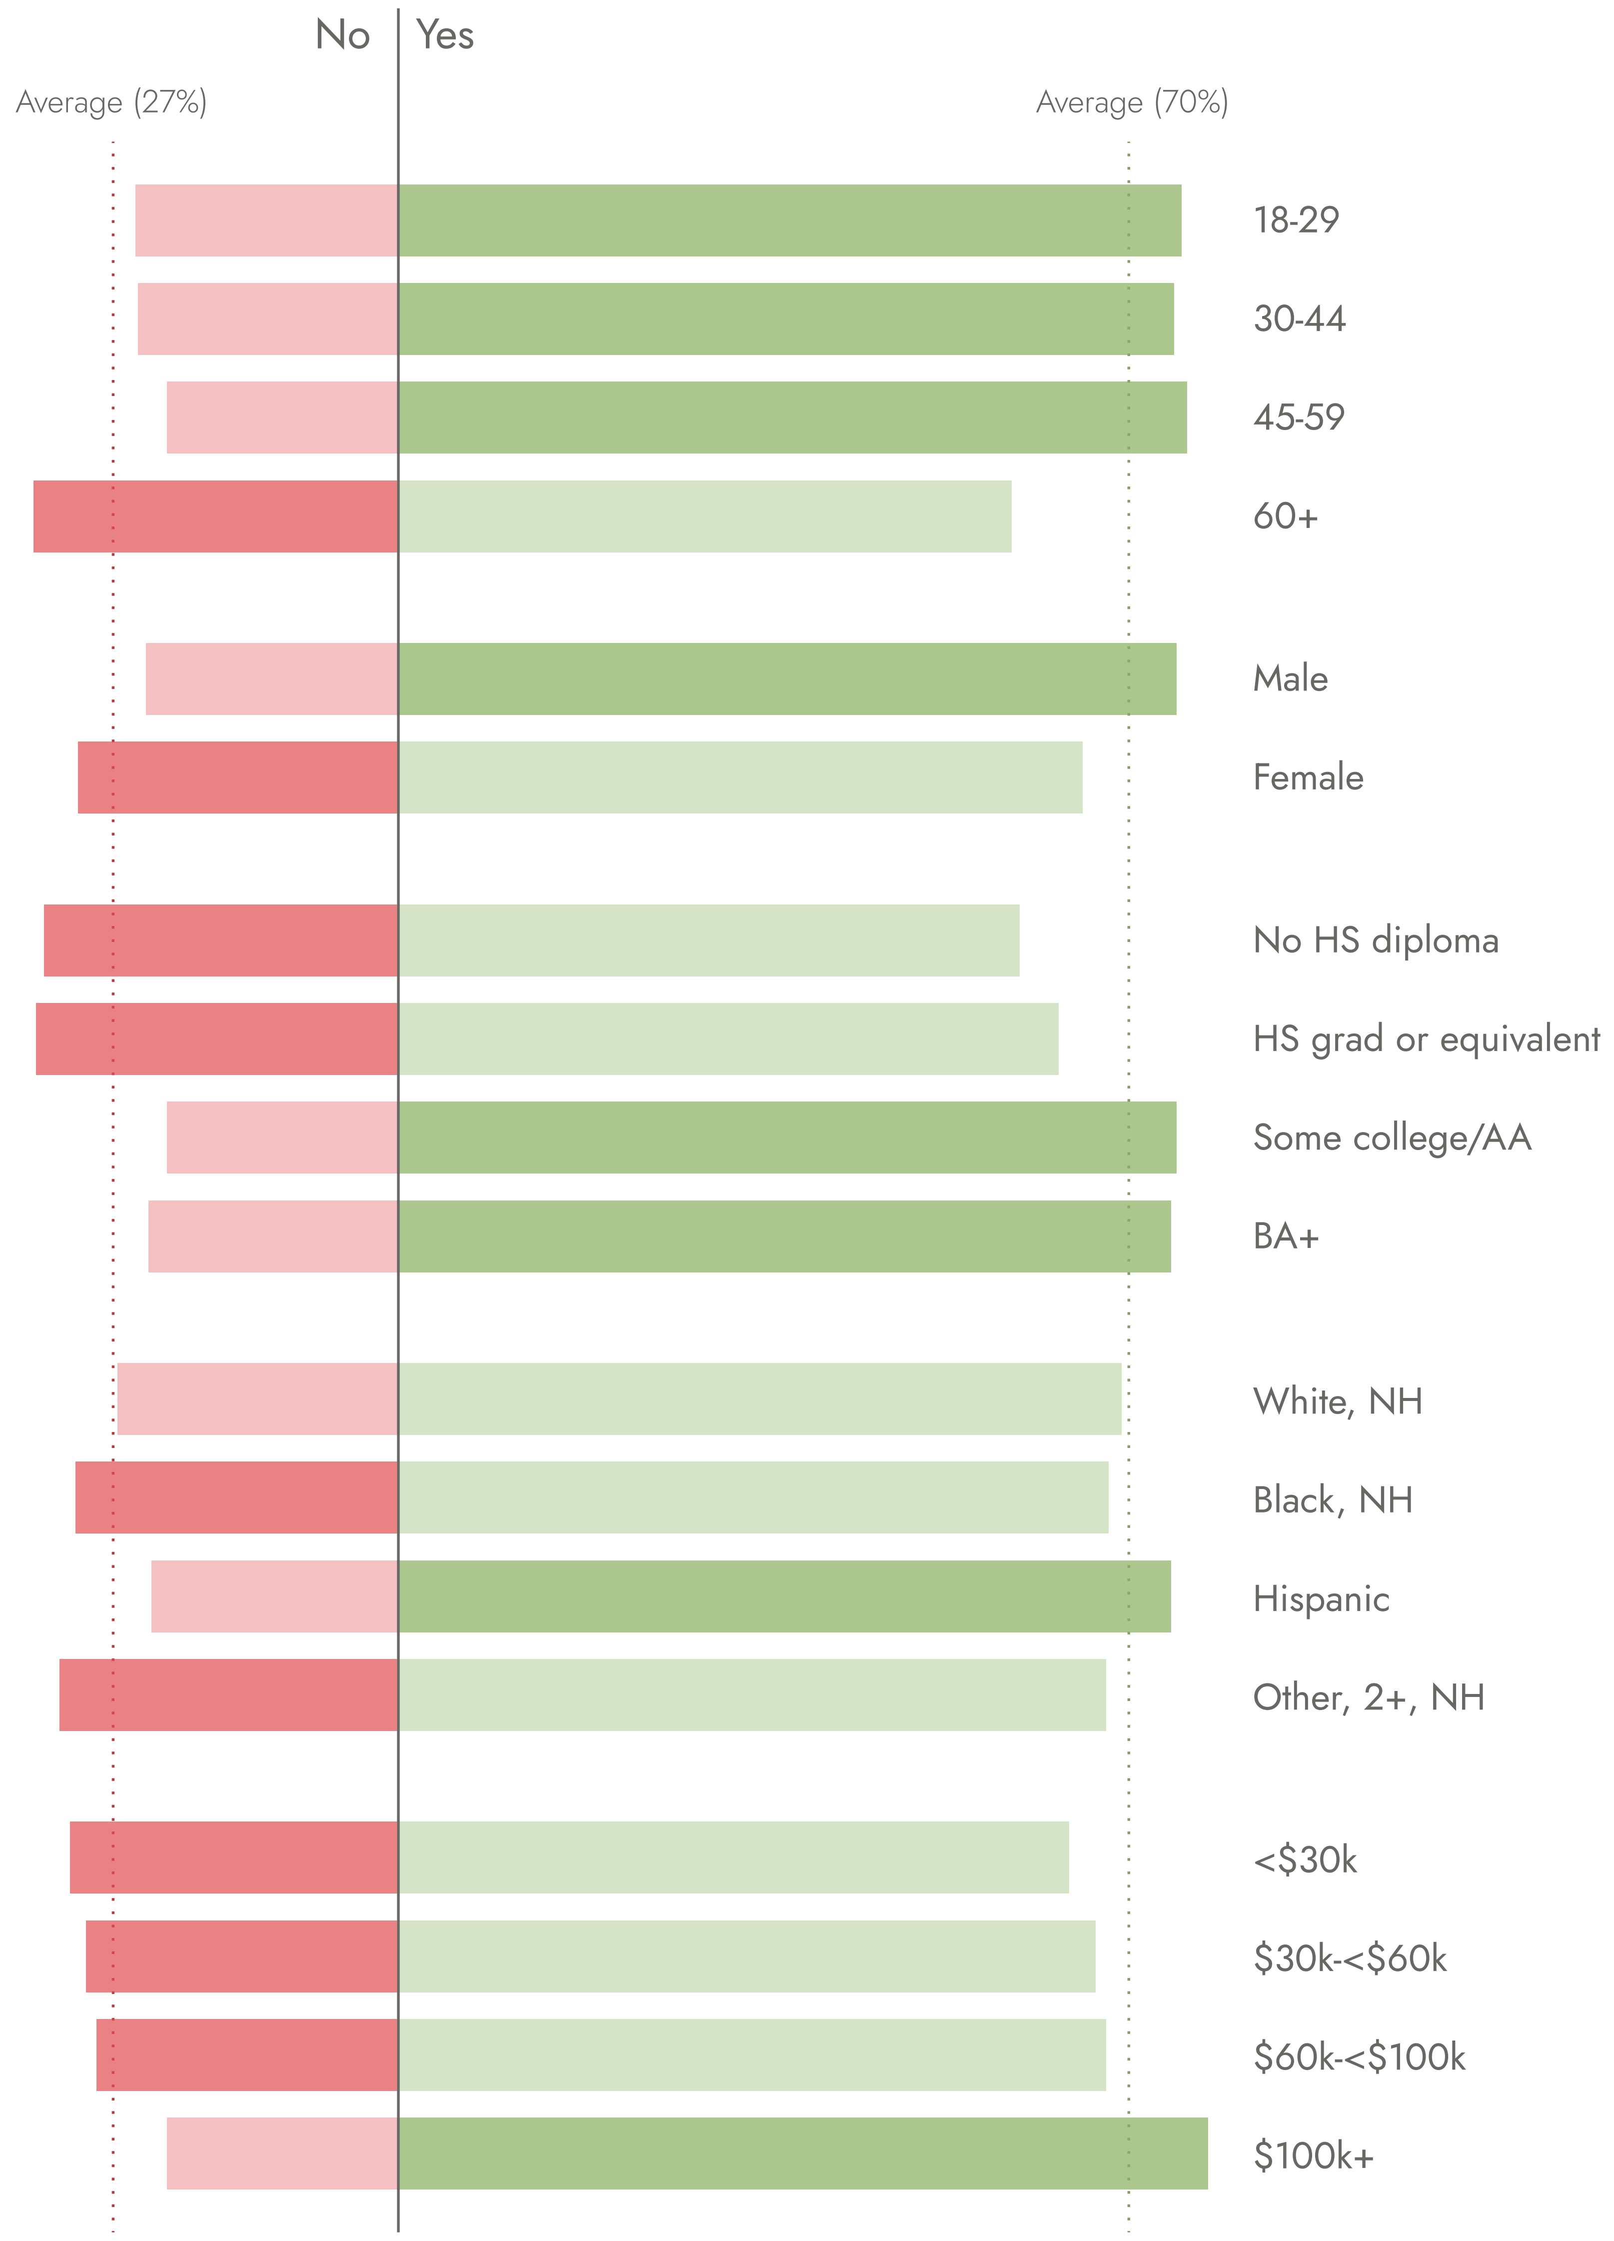

Public support for using personal health data to train AI is fairly high, but conditional.

While nearly two-thirds of respondents (64%) are open to personal health data use to improve healthcare, a substantial minority (35%) reject it outright, regardless of safeguards.

Older adults are more supportive overall but place greater emphasis on anonymization, Black respondents show a stronger preference for explicit opt-in permission, and lower-income and less-educated groups are more likely to refuse data use entirely. Even among higher-income and more educated respondents (who show the highest overall support), opt-in consent remains far more accepted than unconditional support.

These findings underscore that legitimacy for health-data use in AI will depend on layered protections that combine meaningful choice, strong de-identification standards, and clear limits on use, rather than relying on a single consent or anonymization mechanism to secure public trust.

Would you want your personal health data to be used to train AI models designed to improve health care?

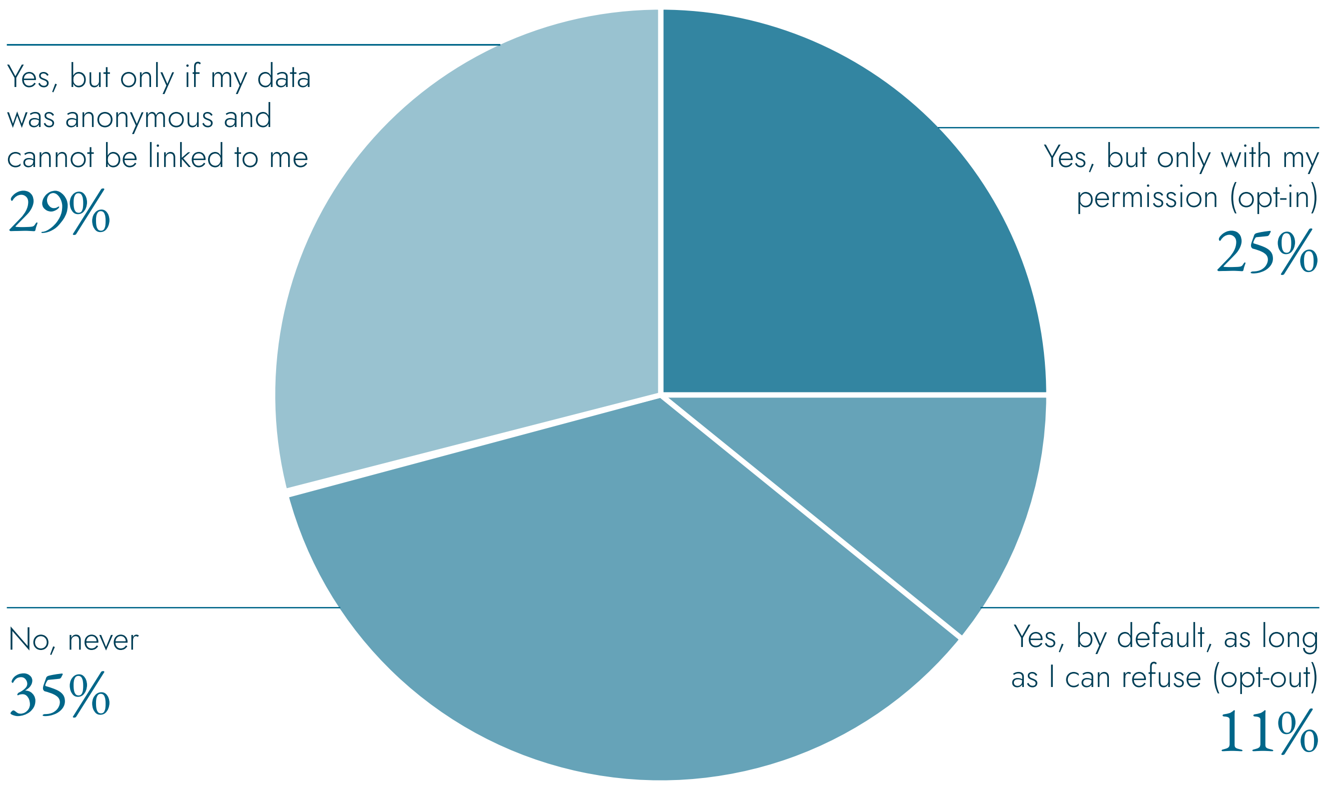

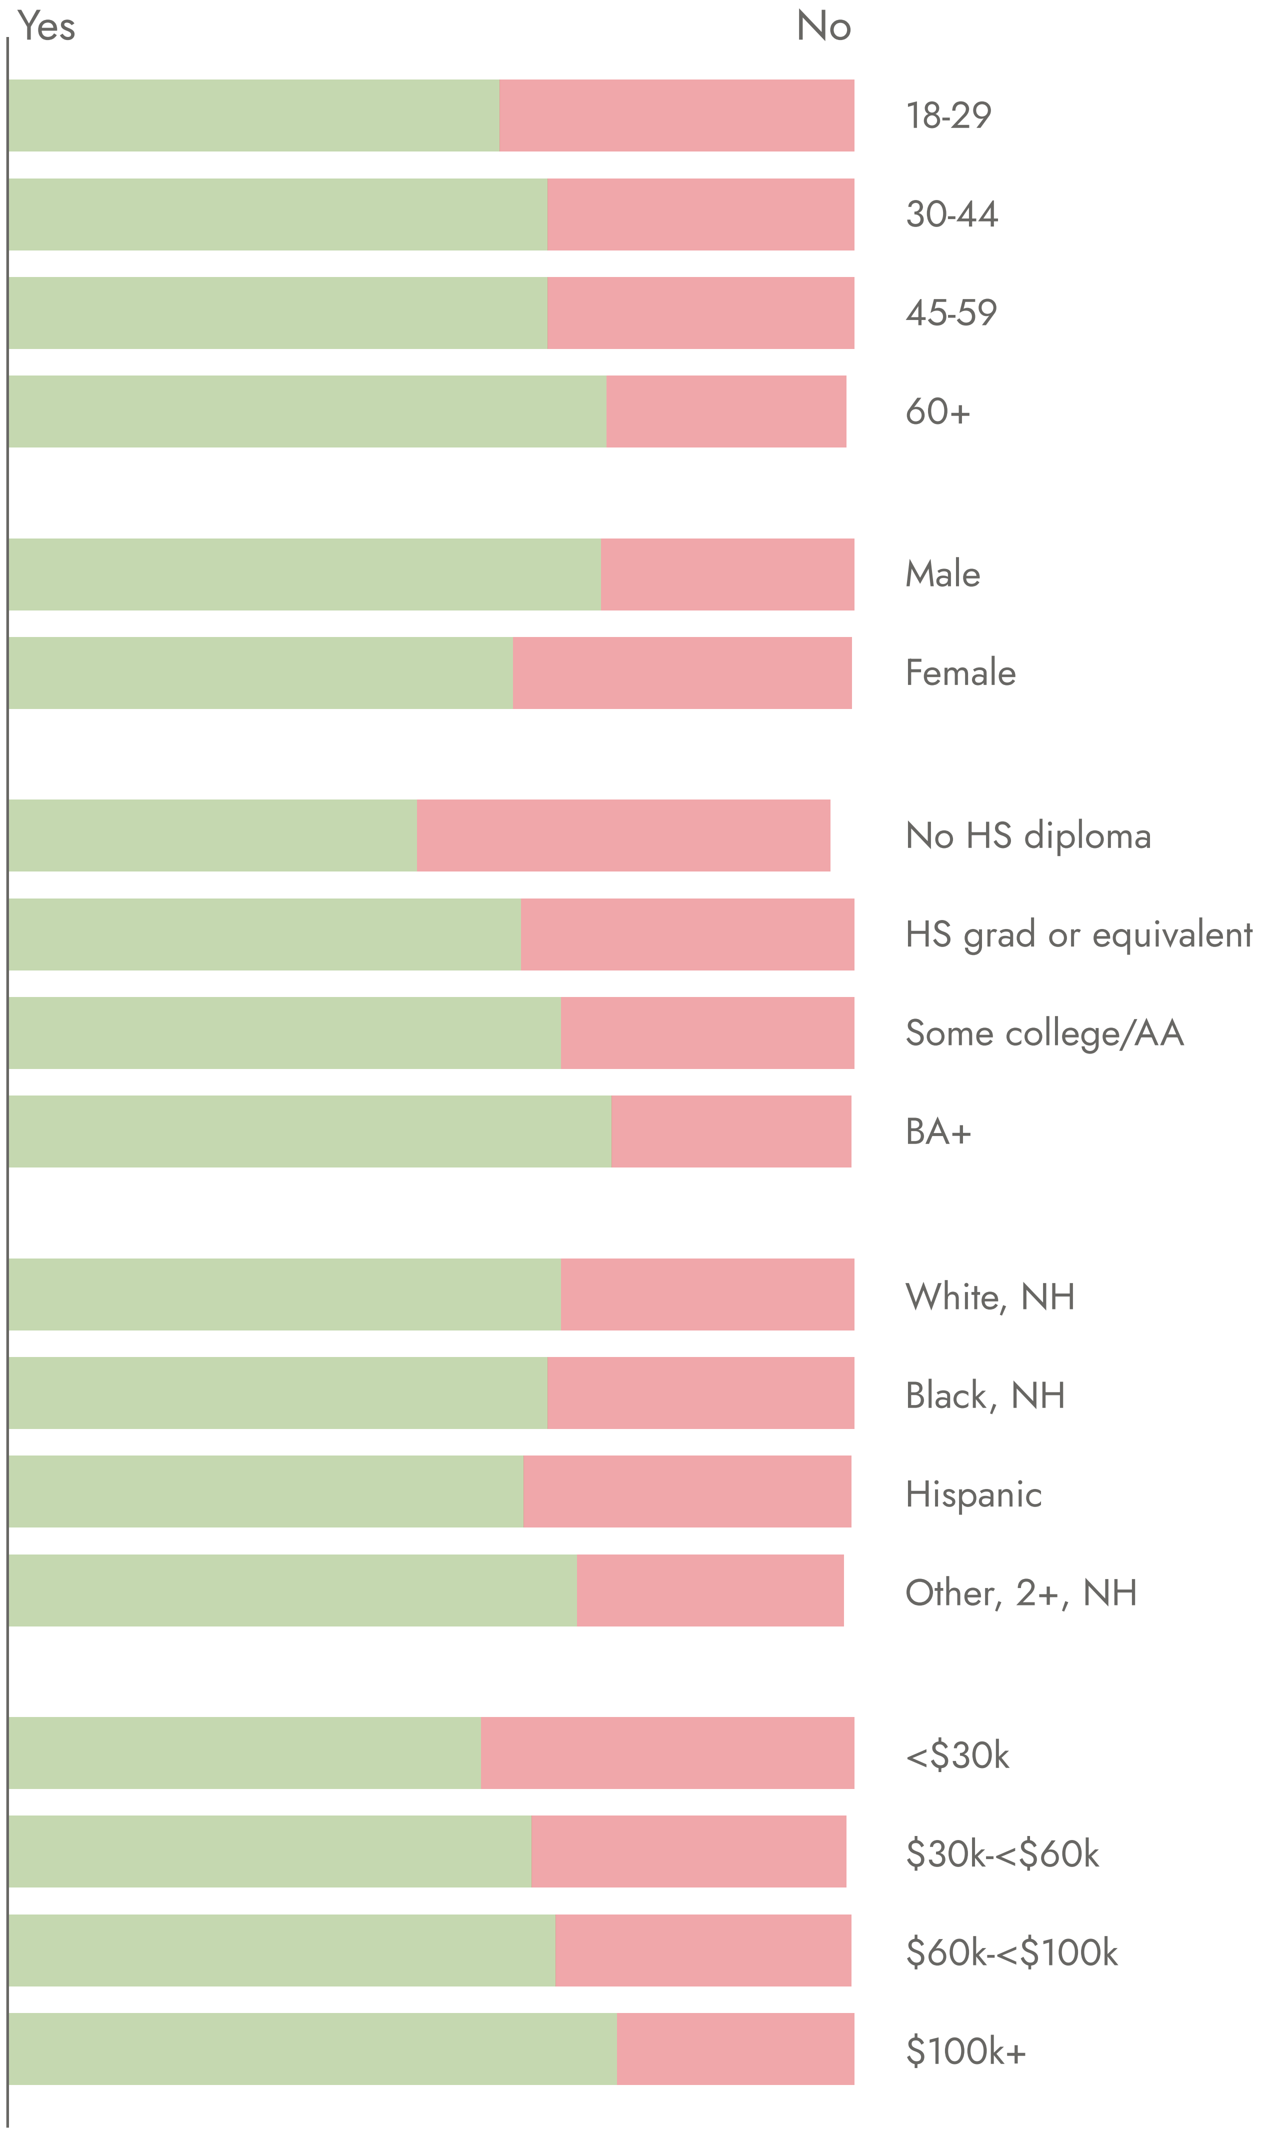

The majority of the public (~70%) think that the patient alone should own and control their health data.

Only 15% believe that this control should sit wholly apart from patients, with 14% preferring shared ownership between patients, providers, and AI developers.

Support for patient control is highest among higher-income and more educated respondents, while lower-income, less-educated, older adults, and Black respondents are more likely to favor shared ownership.

The data indicates a strong public expectation that patients remain the primary decision-makers over their health data when AI is used.

While a meaningful minority support shared stewardship models, there is very low support for exclusive control by providers or AI vendors. Governance approaches need to center patient authority and explicitly define any shared control arrangements.

Who do people believe should control health data when AI is used?

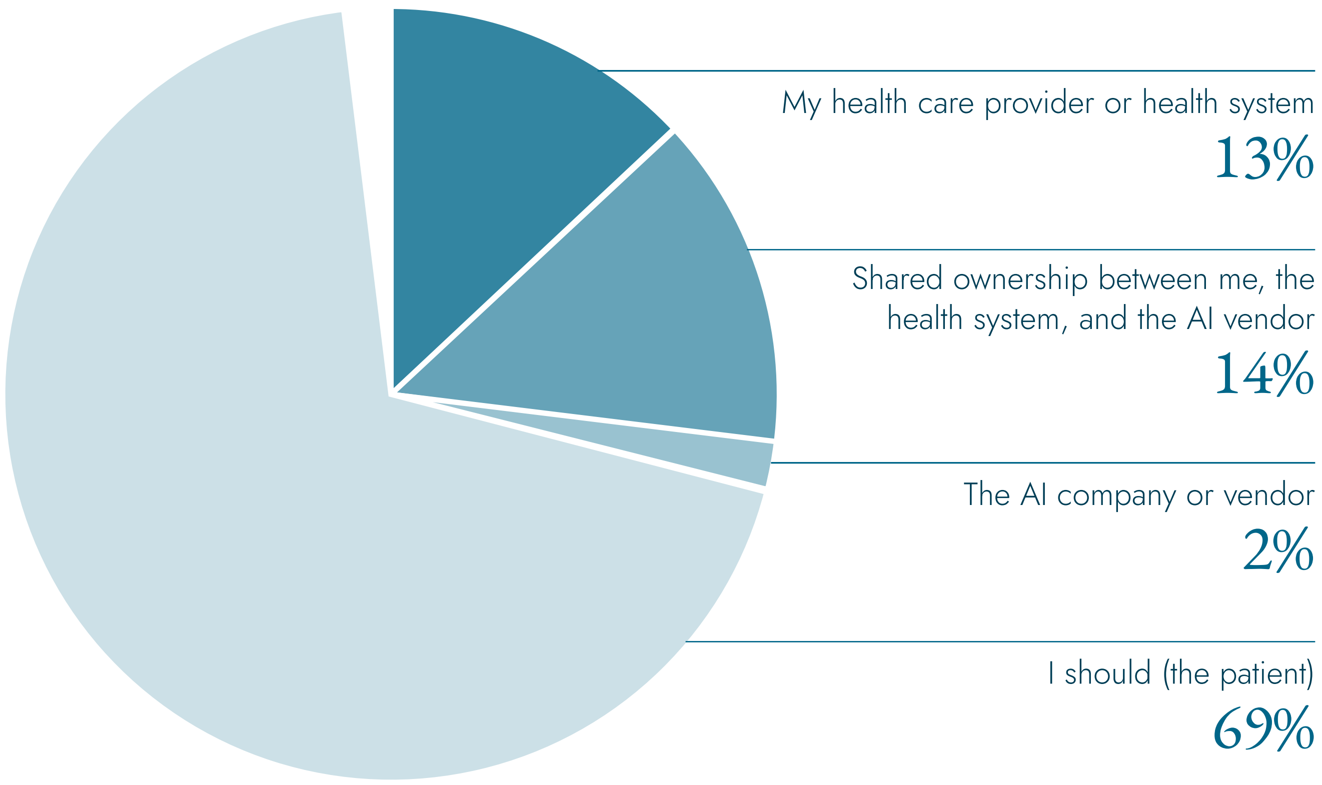

Concern about the commercialization of health data is widespread and applies to all demographic groups.

Nearly two-thirds of respondents (63%) say they are somewhat or very concerned about their health data being sold or shared for profit. Concern is highest among adults aged 45–59, individuals with some college education, Black and Hispanic respondents, and middle-income households. It is not limited to traditionally marginalized groups or to those with low digital literacy. Interestingly, concern is lowest amongst high-income respondents ($100k+).

How concerned are you about your health care data being sold or shared to make money?

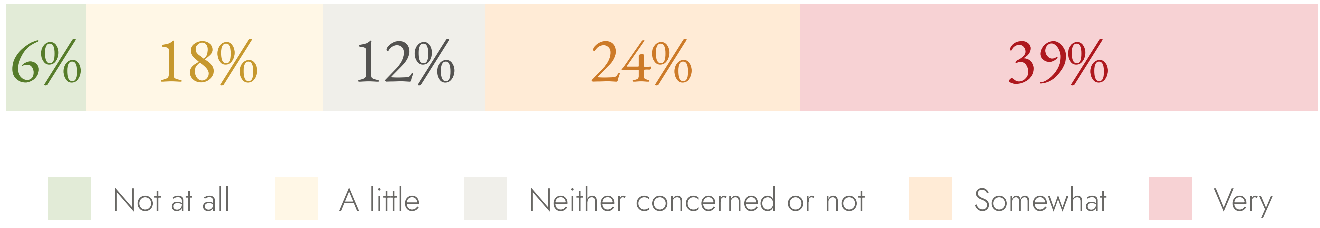

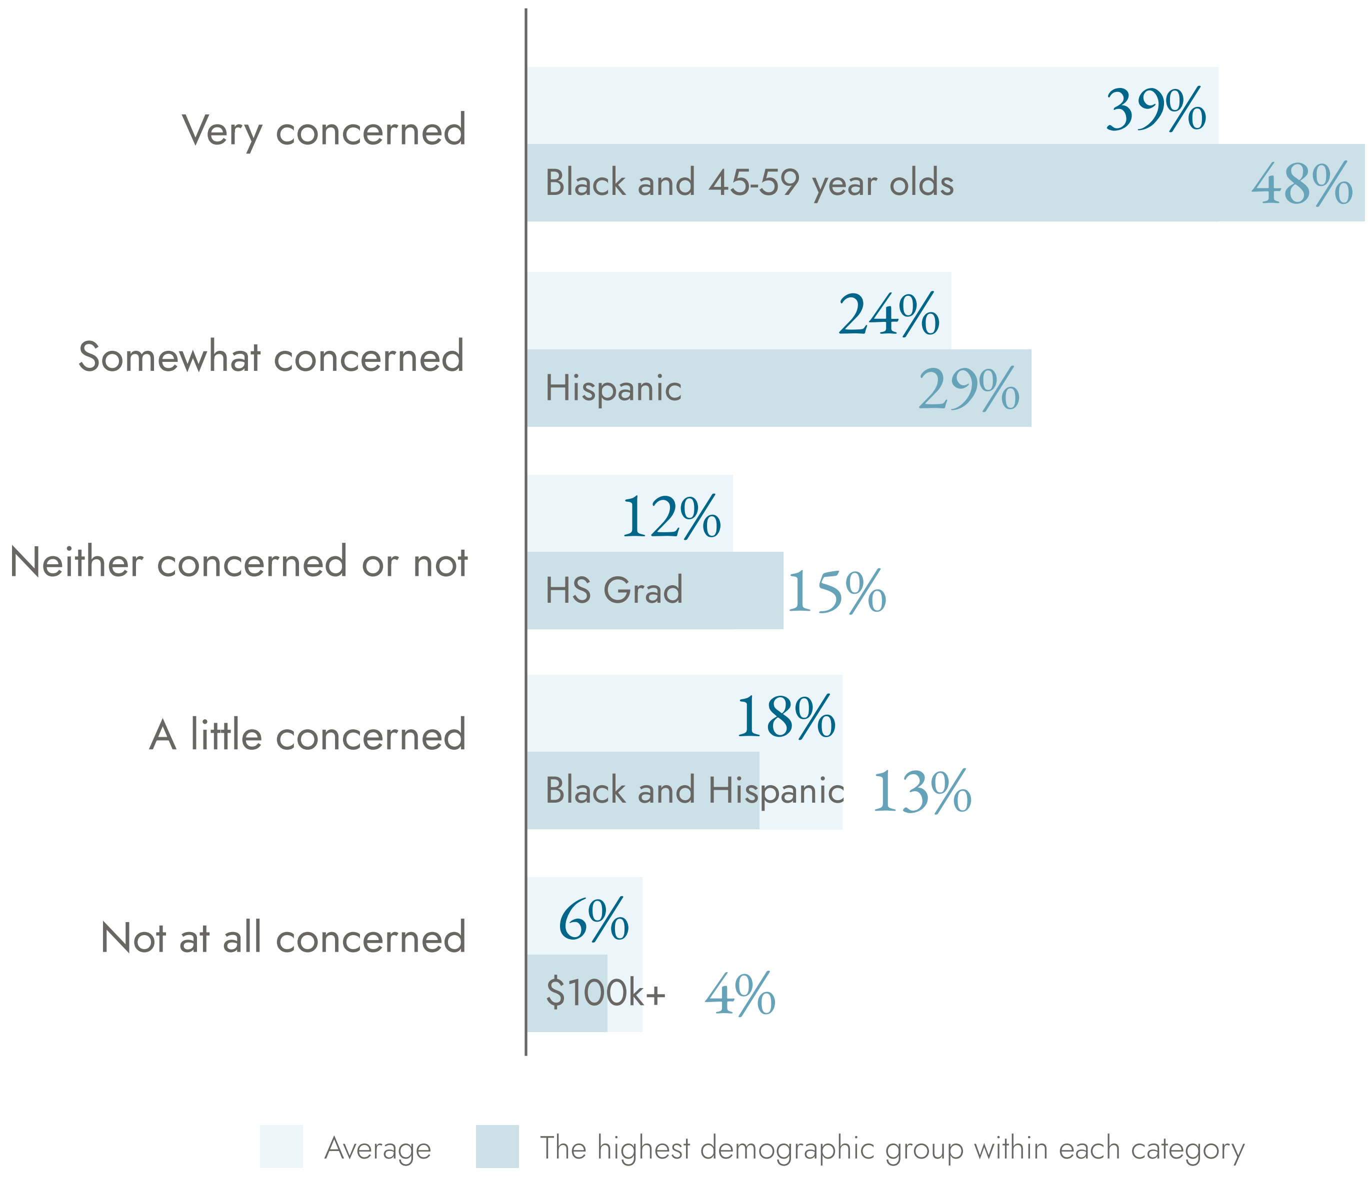

Concern about unfair or biased treatment by health AI is substantial, but notably lower than concern about data commercialization.

While a majority of respondents (55%) say they are somewhat or very concerned that AI systems may treat some groups unfairly, this trails the level of concern over health data being sold or used for profit (63%). This points to perhaps an awareness gap, reinforced by 12% never having considered AI bias at all.

Concern about bias is relatively consistent across demographic groups, with modestly higher levels among adults aged 45–59, college-educated, and Black respondents.

This presents an opportunity to educate the public on the real consequences of fairness in AI performance being evenly distributed. But at present, the economic use and misuse of health data currently poses a more salient trust risk to patients than algorithmic fairness.

How concerned are you that AI systems in health care may treat some groups of people unfairly?

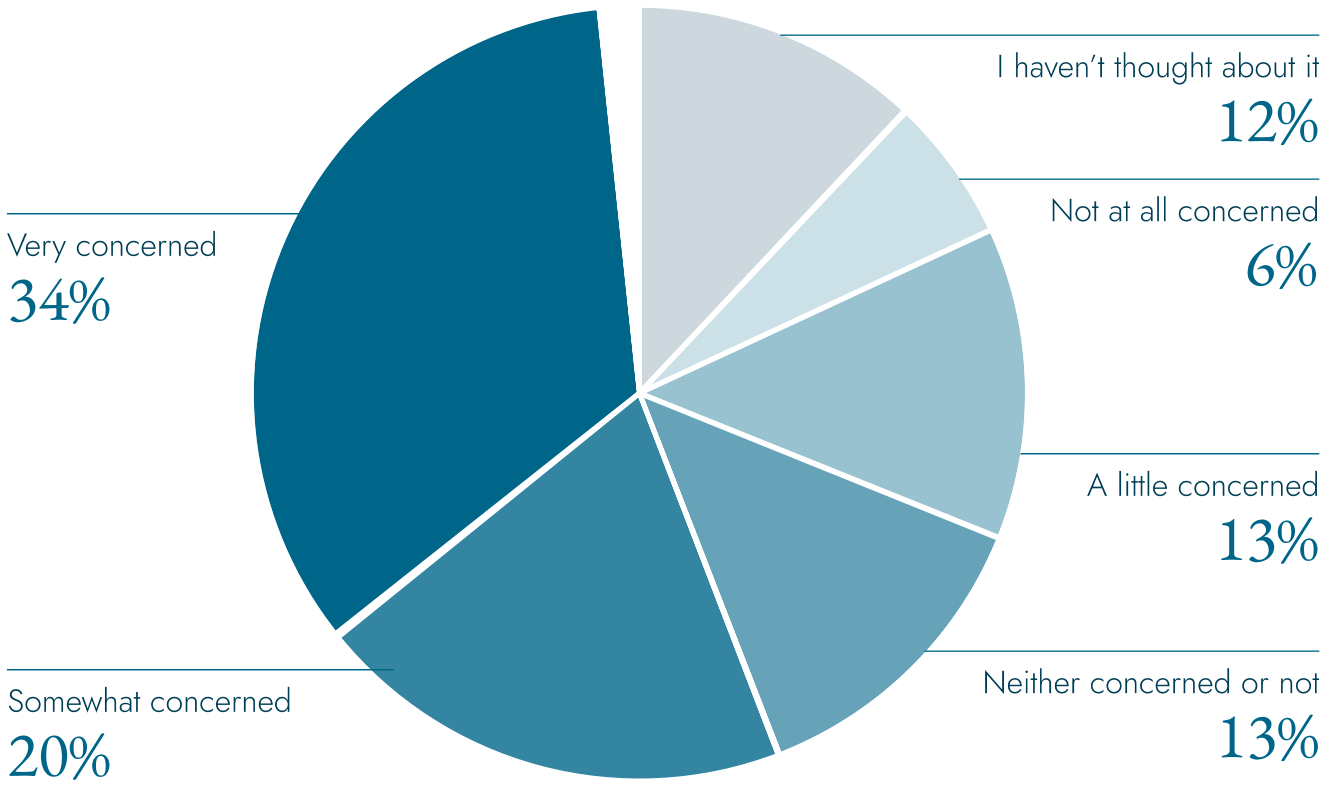

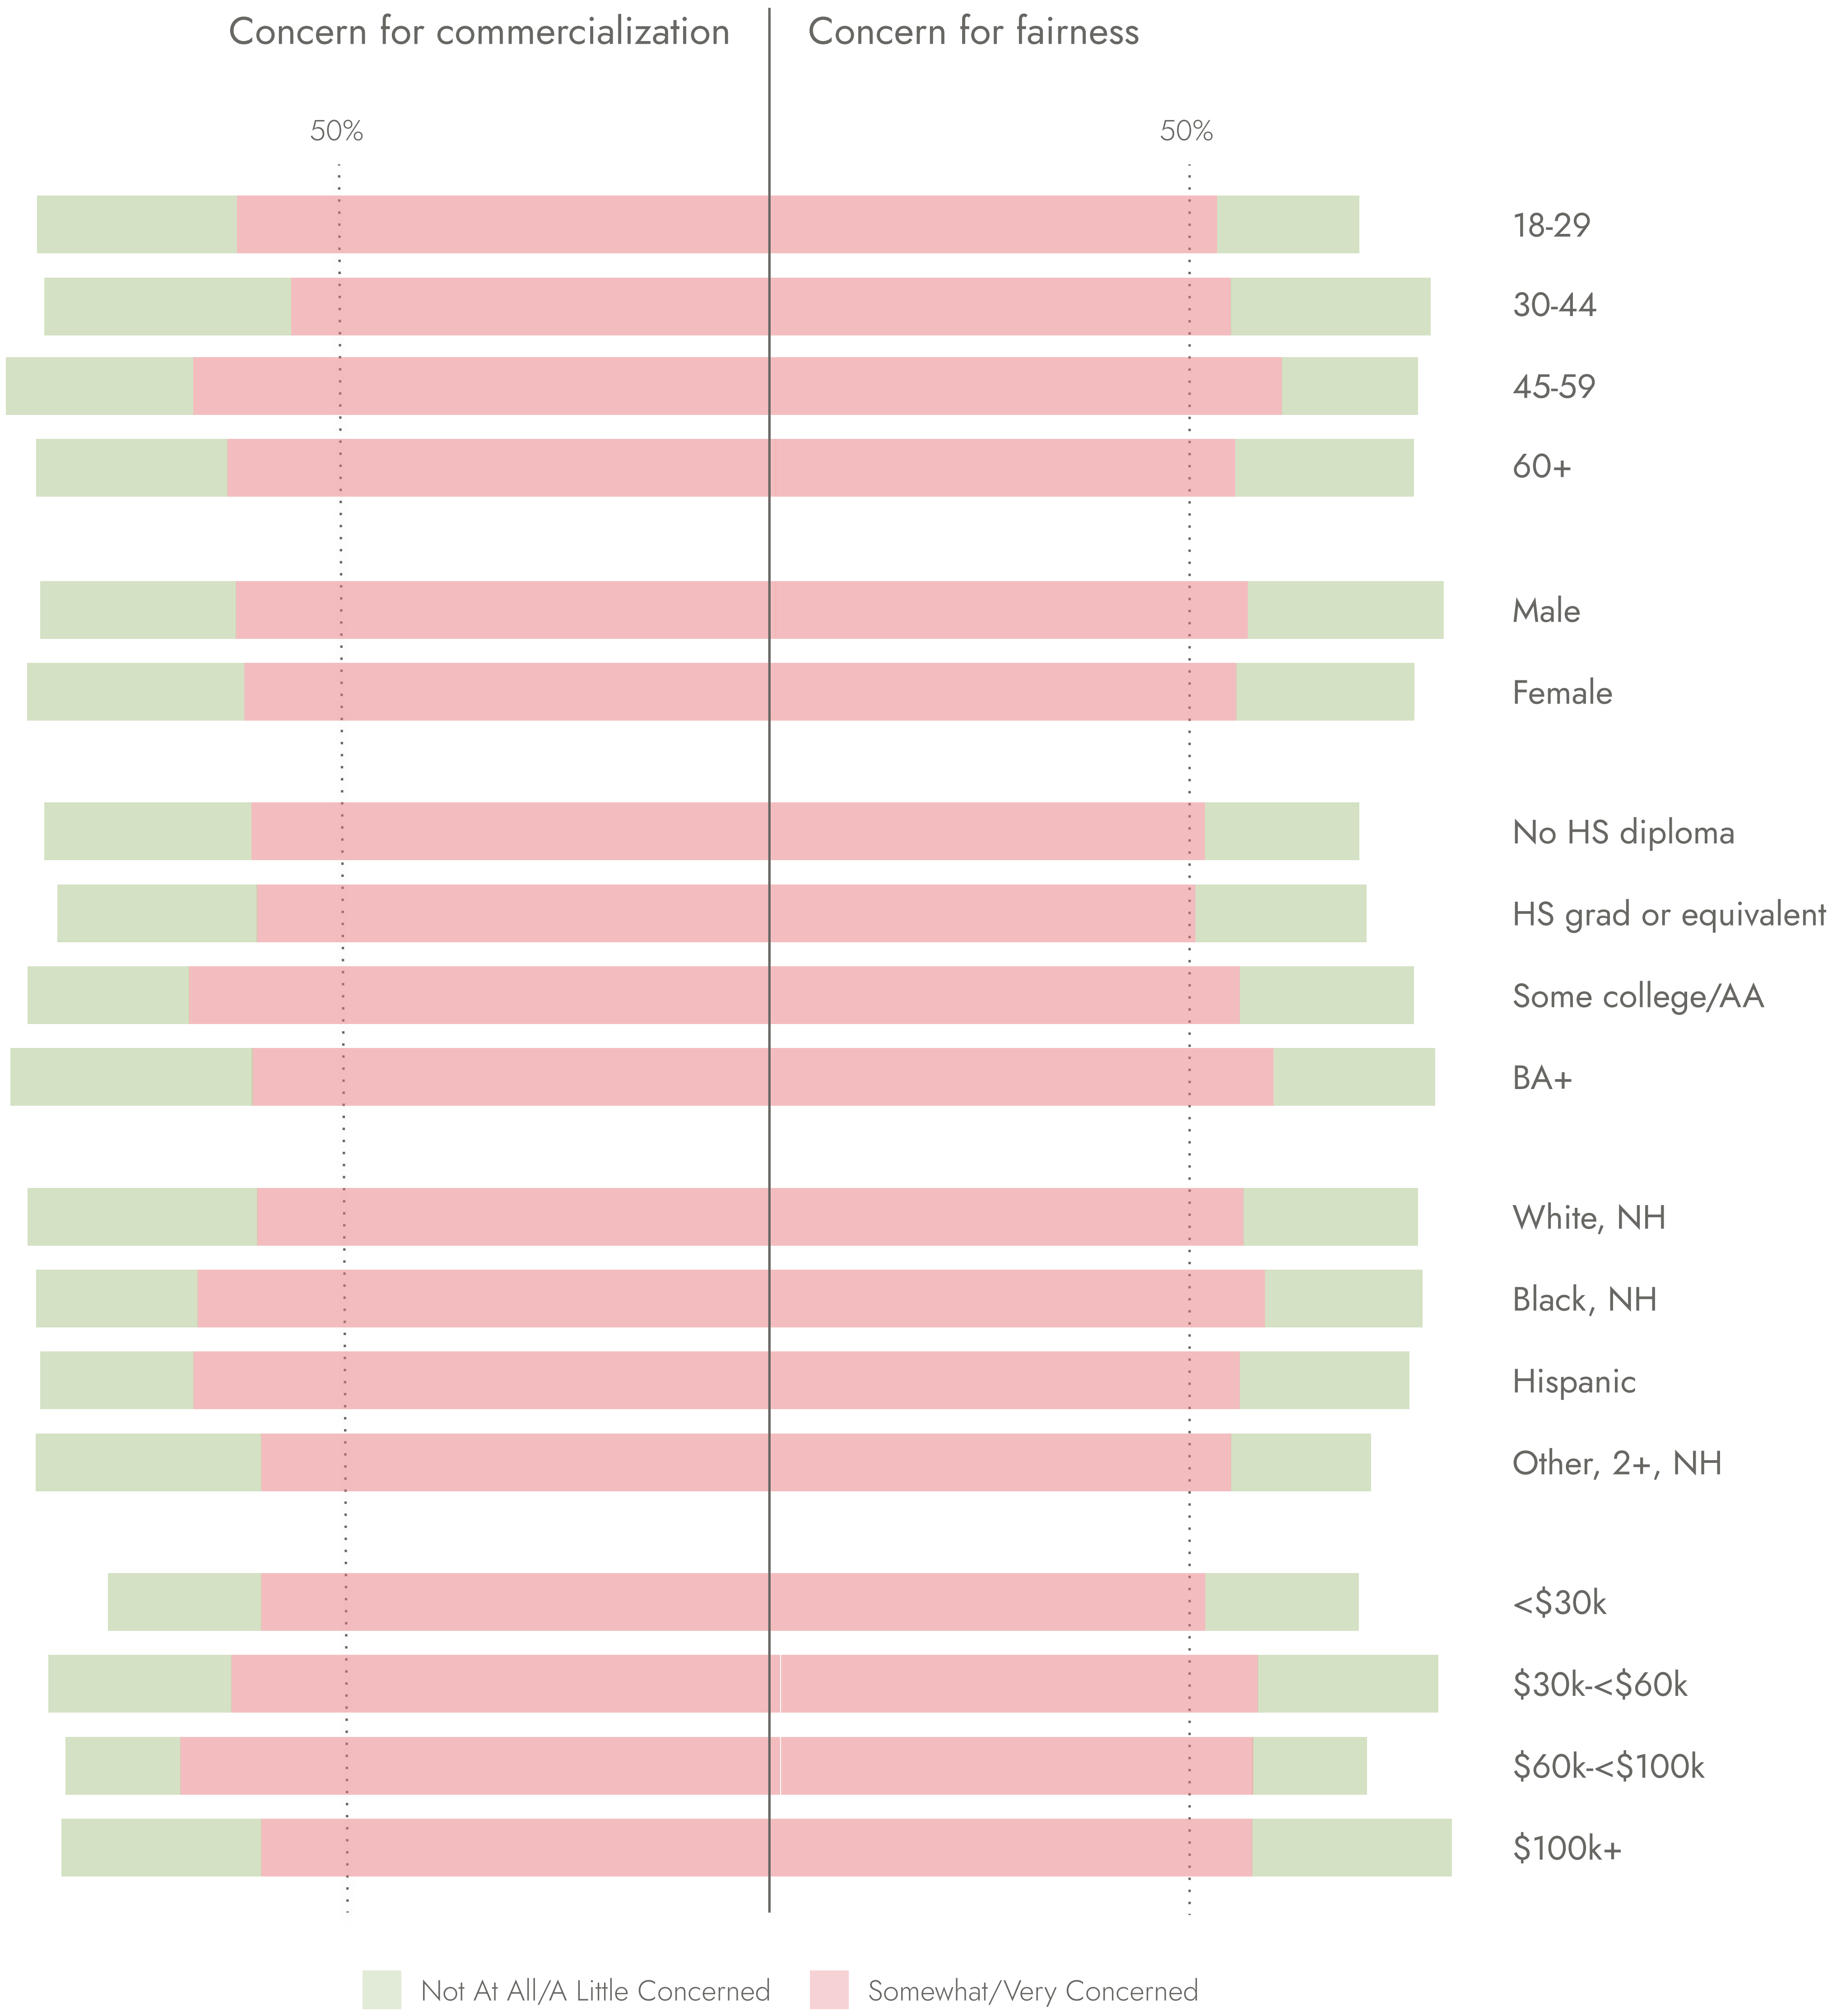

Clear accountability signals are what most increase public confidence in health AI being fair and unbiased.

More than four in five respondents (82%) say their trust would increase if at least one accountability measure were in place. The top four responses all attract support from about half of respondents; a multi-pronged combination of independent review, bias testing, performance transparency, and data transparency should be considered rather than reliance on any one measure alone.

These preferences are broadly shared across demographic groups, but intensify with age, education, and income: older adults, college-educated respondents, and higher-income households place especially high value on independent review and rigorous bias testing. Patient involvement in AI design boosts trust for a smaller but meaningful minority, particularly among Black, Hispanic, and lower-income respondents.

Taken together, trust in health AI is most effectively built through a range of oversight and accountability measures: independent evaluation, and demonstrable performance across populations, not through transparency alone.

What would make you more confident that AI is fair and unbiased in health care?

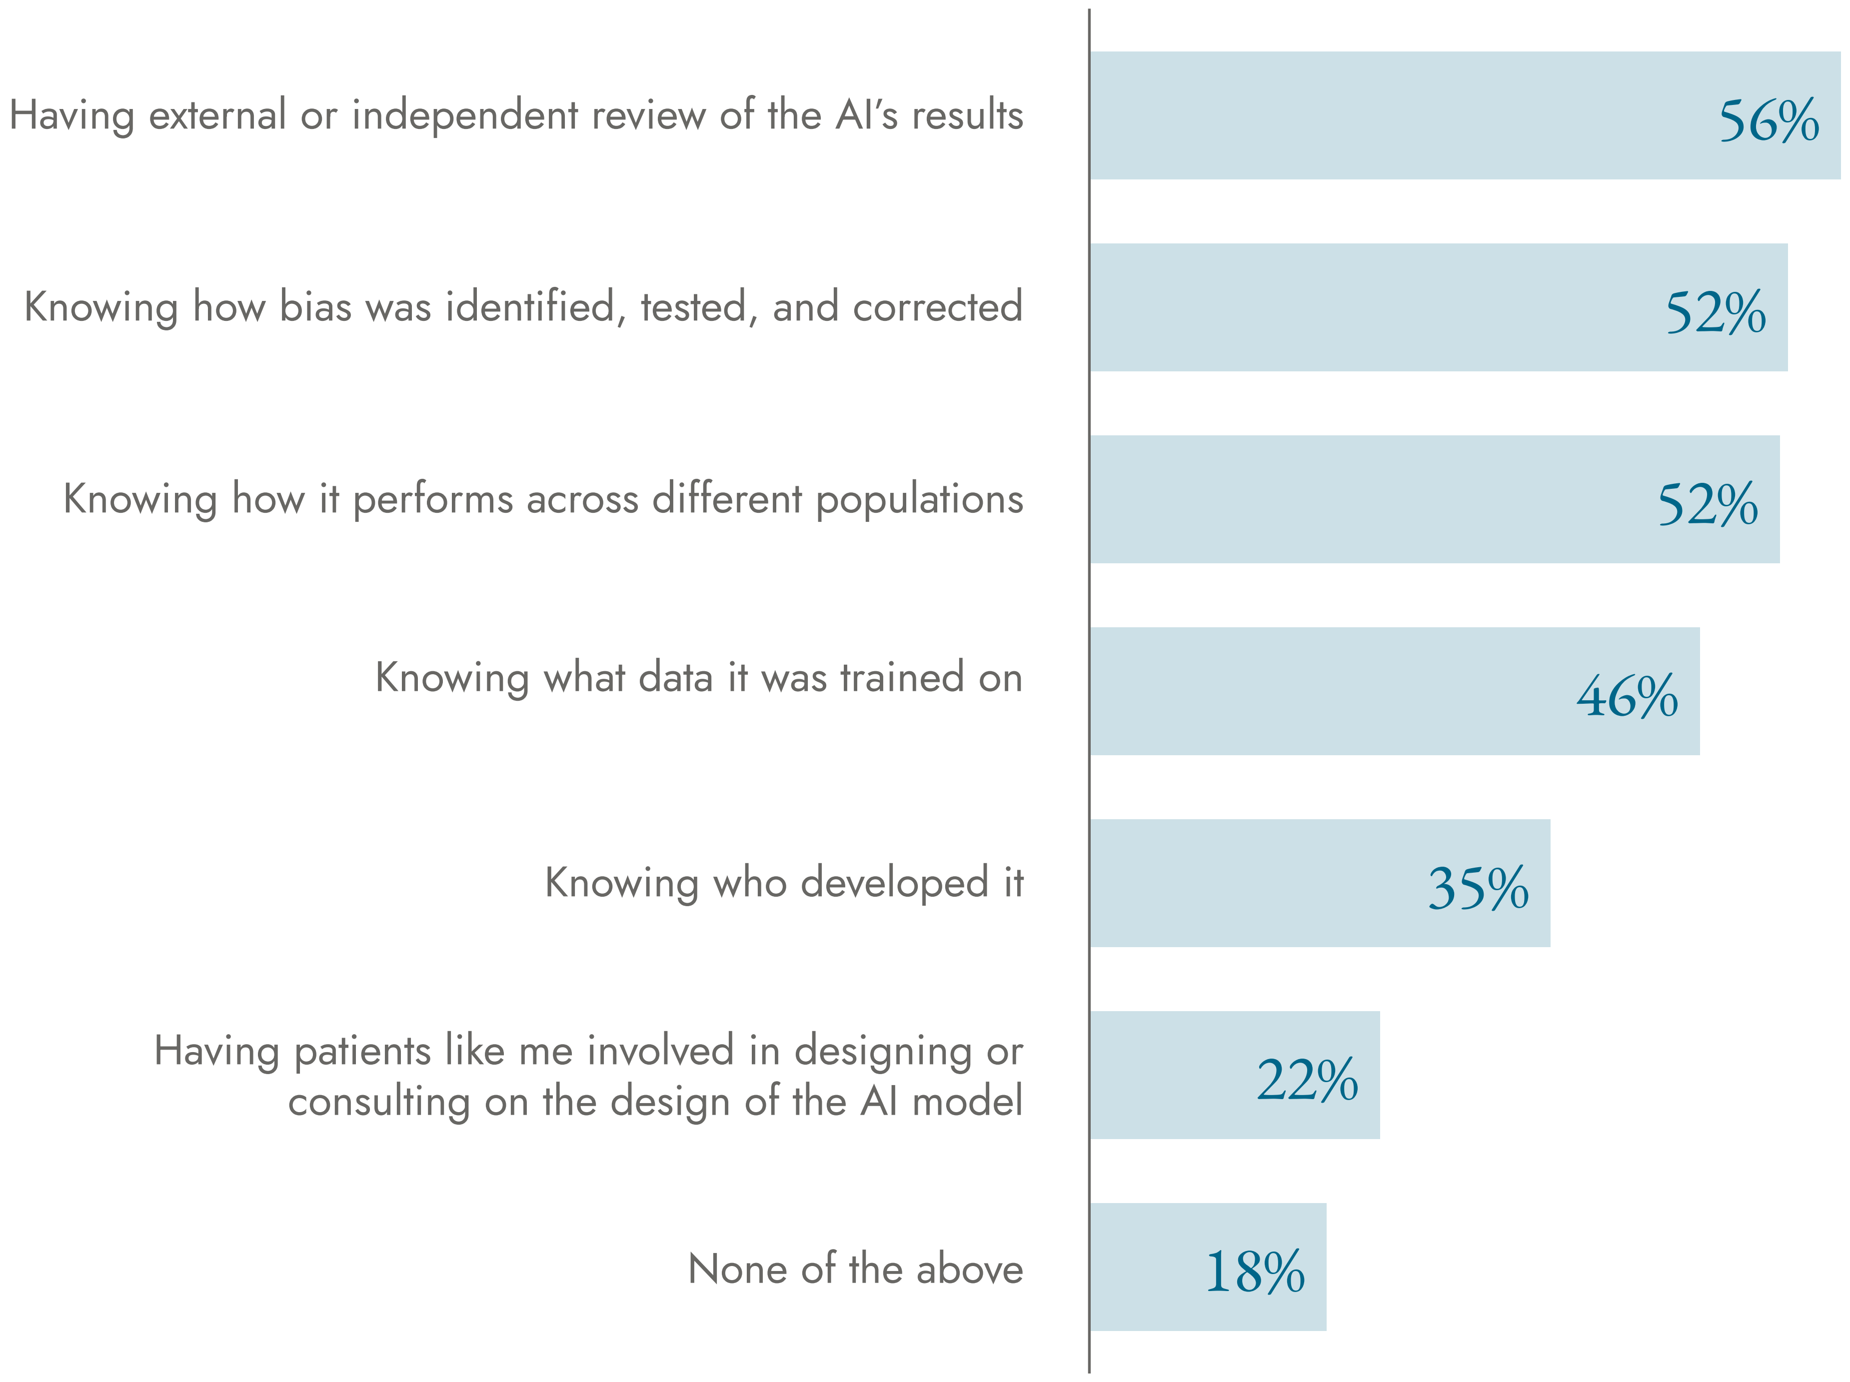

Concerns

Over 93% of respondents reported concerns about AI in health care.

Public concern is less about AI replacing clinicians or limiting access, and more about whether adequate oversight exists to ensure humans remain firmly in control of its use: whether clinicians defer too readily to algorithmic outputs, whether meaningful human oversight remains in place, and whether errors or miscommunication are introduced rather than prevented.

This signals that trust in health AI will be driven by guardrails that preserve clinical accountability, limit automation bias, and protect patient data within everyday care delivery.

What worries you about using AI in health care?

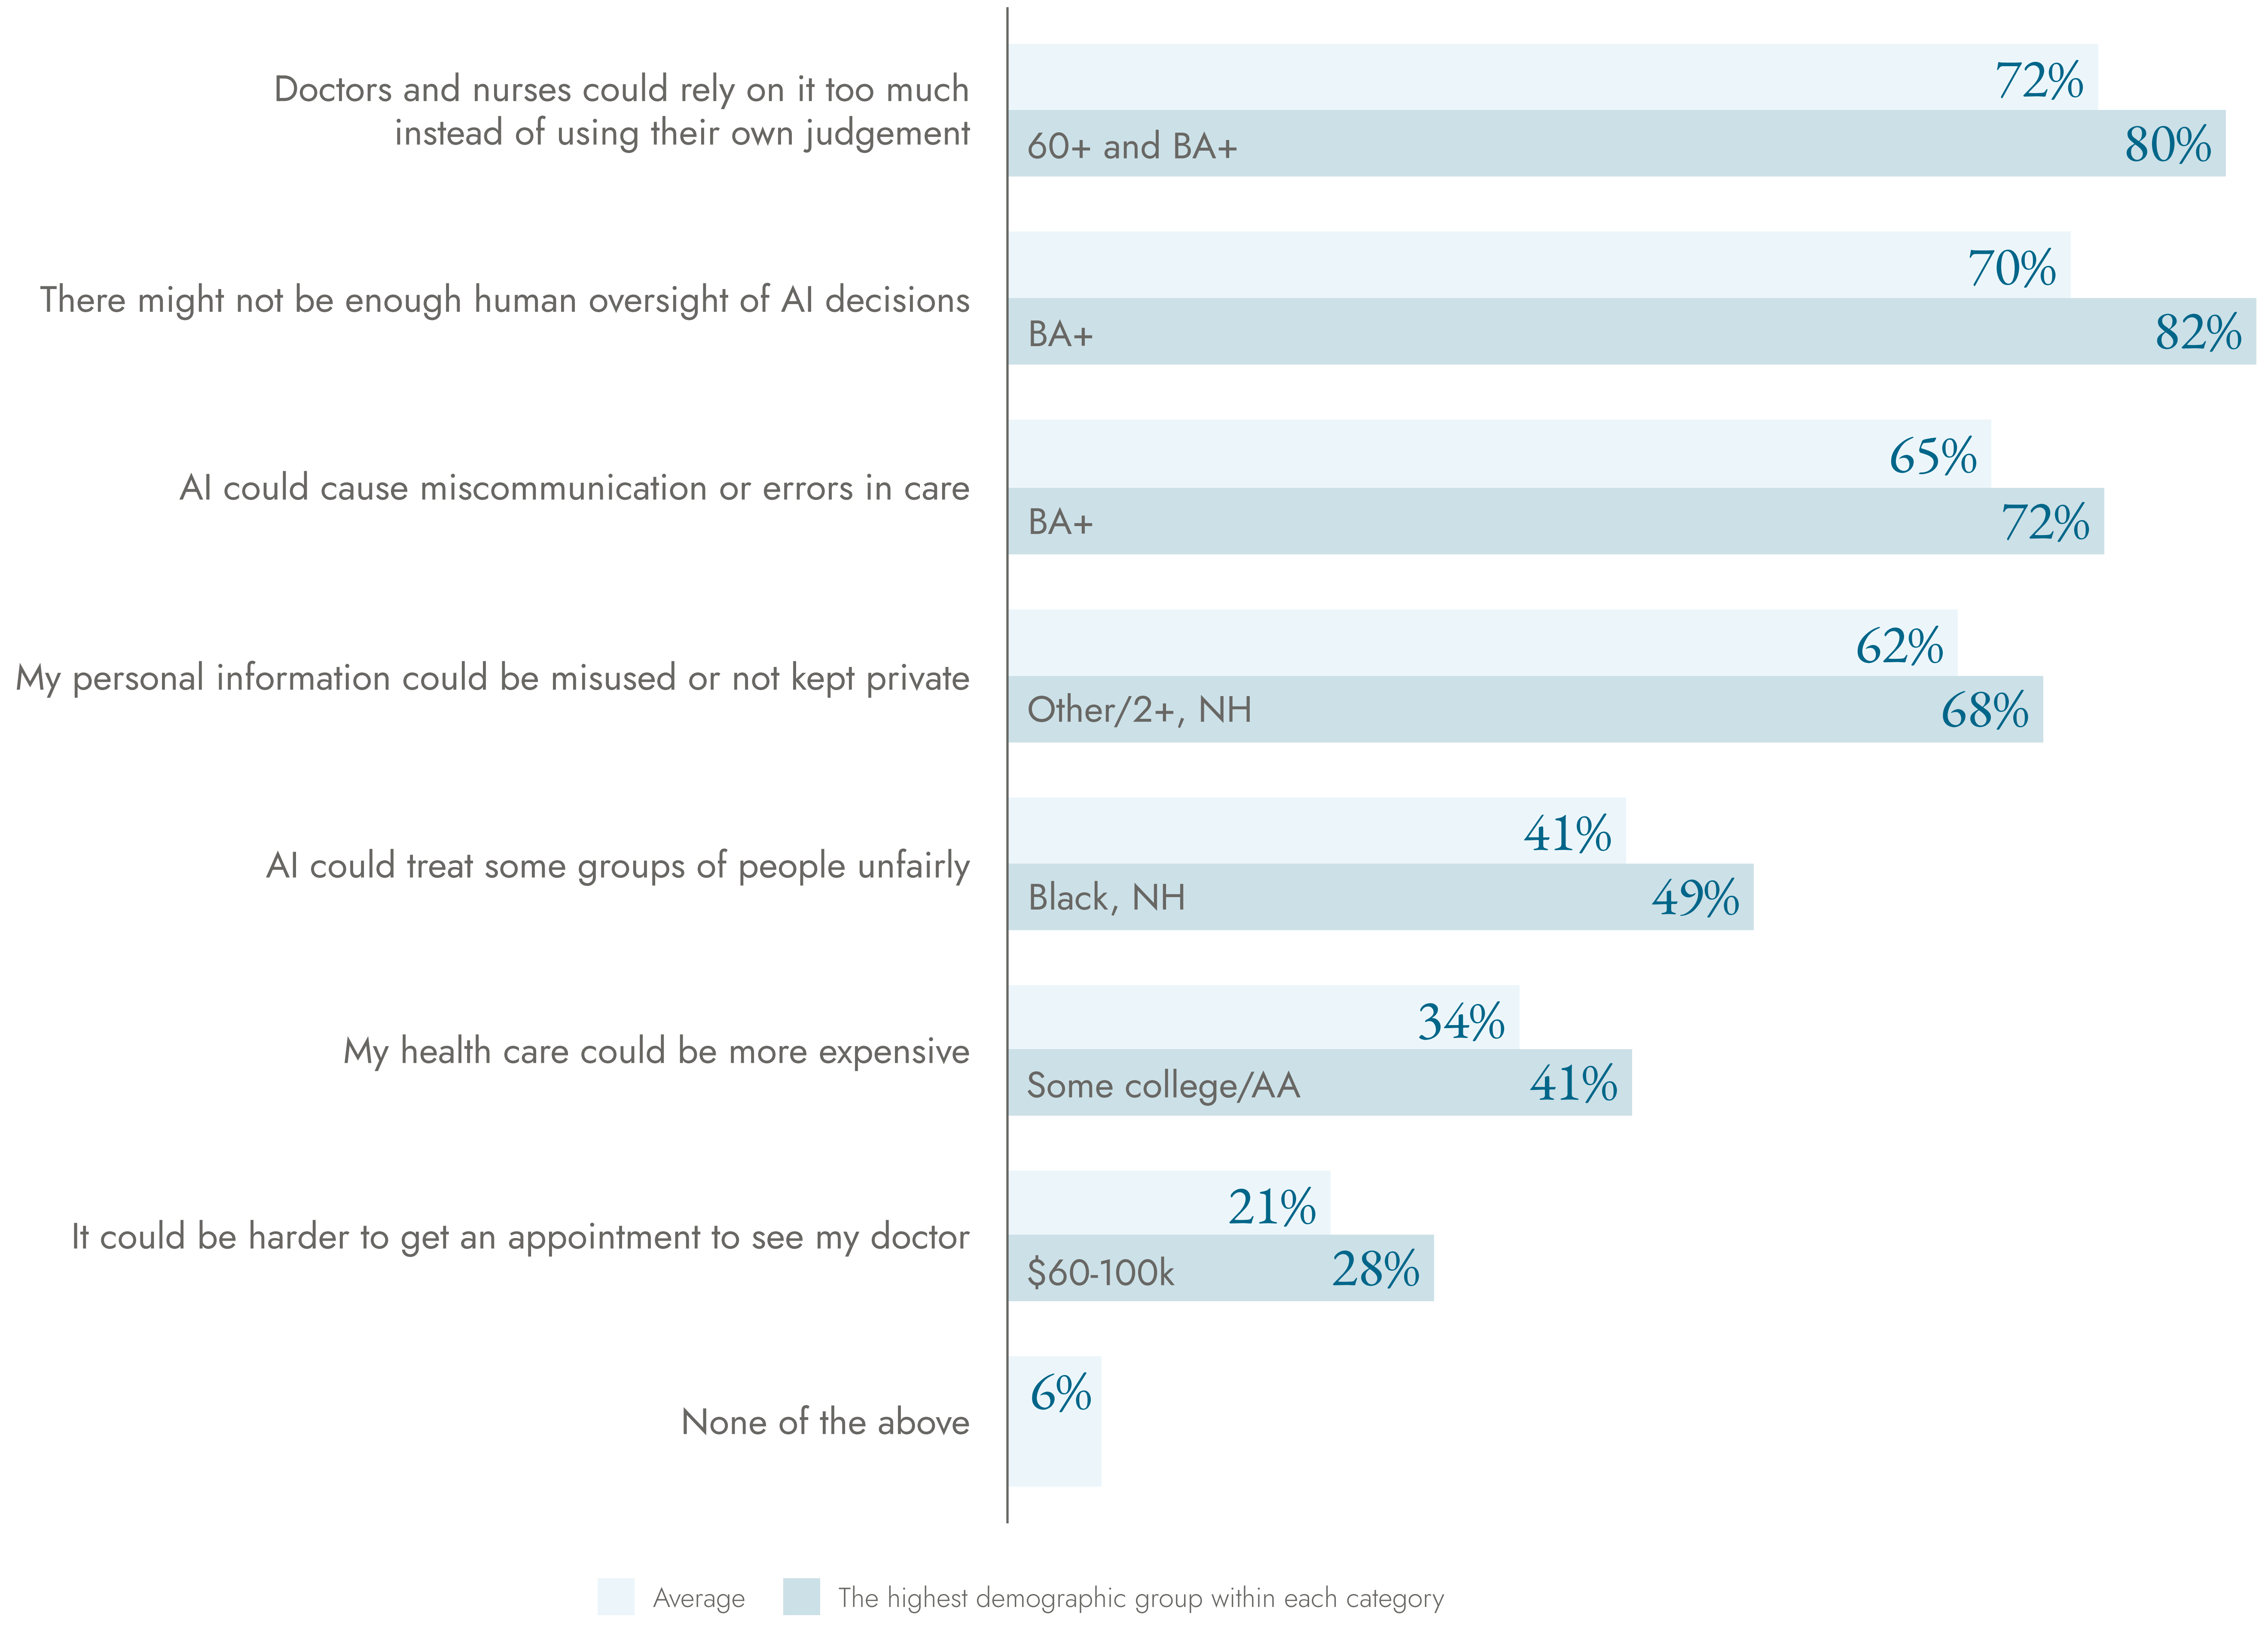

Expectation

Support for the idea that clinicians have a responsibility to use AI is closely tied to baseline trust in AI-enabled care.

Support is strongest among groups that already express higher trust in AI - older adults, men, higher-income, and college-educated respondents - while younger adults, women, and lower-income groups are more skeptical and more polarized.

This divide suggests that the public distinguishes between permission to use AI and a duty to do so: many are open to AI as a tool, but resist the idea that clinicians should be compelled to rely on it.

As policy emerges, we need to consider that mandates or expectations for providers to use AI may outpace public confidence. Governance frameworks are more likely to gain acceptance if they emphasize clinician discretion, patient benefit, and accountability over any presumption that AI adoption is inherently required.

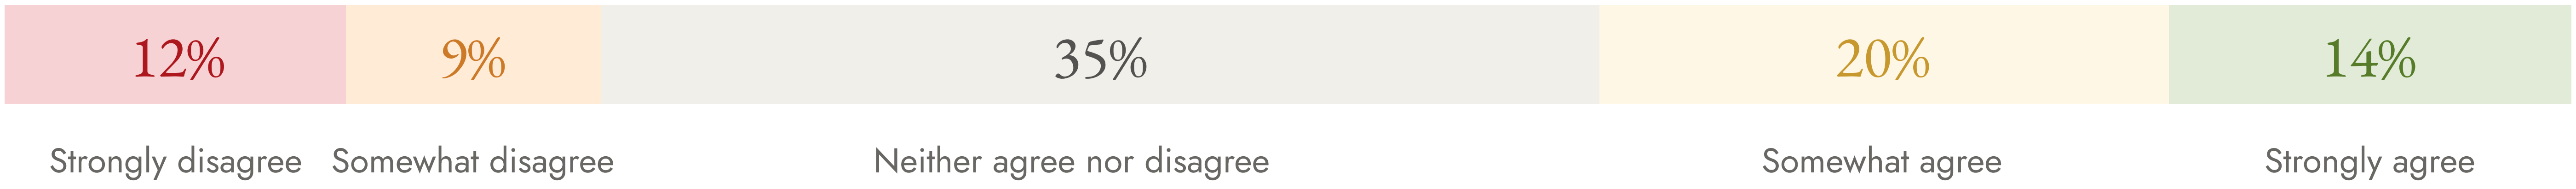

Do you agree that health care providers have a responsibility to use AI tools if they will improve patient care?

Excitement

High-level themes (what excites people most)

Respondents were offered an optional open-text answer to "What excites you most about using AI in health care?"

Improved accuracy & earlier diagnosis

"Possibly easier and quicker and/or early diagnosis of illnesses that doctors may miss or overlook."

"AI could help doctors analyze more information than a human can and identify patterns that lead to more accurate diagnoses."

Efficiency, speed, and time savings

"Making things quicker and more efficient so patients don't have to wait as long for answers or care."

"Speeding up processes that currently take too long, like reviewing test results or records."

Reducing clinician burden & burnout

"Giving healthcare providers more tools and options so they can focus more on patients instead of paperwork."

"As a former healthcare professional, I think AI could really help reduce burnout by handling repetitive tasks."

Better use of data & decision support

"It incorporates technological advancements in medicine that can analyze large amounts of data more effectively than humans."

"Using data in smarter ways to support decisions instead of replacing doctors."

Improved access and system-level benefits

"If it helps people in rural or underserved areas get better care or faster answers, that would be a big improvement."

"Lowering costs and making healthcare more accessible to everyone."

Important Nuance

Excitement is often cautious

Even among "positive" responses, excitement is frequently conditional—tied to oversight, accuracy, and human involvement.

"I think it could make doctors be able to handle more patients effectively, as long as it's monitored and not fully relied on."

"This is hard because I know AI makes mistakes, but if used carefully it could be very helpful."

What Matters Most

Age

18-29

This group is simultaneously over-represented among intentional AI users and among the most volatile in trust responses.

Transparency mandates alone are insufficient and may be counterproductive without accompanying agency-preserving mechanisms.

Policies that emphasize choice architecture (clear opt-in, refusal without penalty) will matter more for legitimacy than disclosure language alone.

Youth familiarity with AI does not translate into tolerance for institutional AI use in healthcare.

43.3% are intentional AI users (over-represented vs 37.3% overall).

37.0% are unintentional AI users.

13.4% say they avoid using AI.

51% report being comfortable with AI (highest comfort group).

Highest or second-highest "Highly trust" rates across most healthcare AI use cases.

Also highest or second-highest "Not at all trust" rates across use cases.

Trust increases by 12 percentage points with clinician oversight for insurance pre-approval AI.

75% believe their health data is used to train AI.

56.6% say trust decreases when aware of data use.

58.8% say permission should be required for data use.

42% say they would never want their data used to train AI.

Highest opt-in consent rates across use cases.

Largest trust decrease after AI disclosure across use cases.

30-44

High intentional AI use, high awareness of data use, and consistent trust declines after disclosure.

This group represents a mainstream stress test for health AI policy: engaged, informed, but not deferential.

Regulatory approaches that emphasize independent oversight and clinician accountability are more likely to resonate than developer assurances.

Policy credibility depends on showing that AI use is governed, not assumed.

46.5% are intentional AI users (strongly over-represented).

37.0% are unintentional AI users.

74% believe their health data is used to train AI.

51.6% say trust decreases when aware of data use.

57.9% say permission should be required.

37% say they would never allow data use.

Highest comfort after 18—29 (51%).

Trust decreases after AI disclosure across all use cases.

Clinician oversight increases trust across clinical use cases.

45-59

High awareness of data use, high concern about trust loss, and strong expectations of institutional responsibility

More trust placed in hospitals and health systems than in vendors.

Shared governance models that include provider accountability are likely to be seen as appropriate.

38.7% are intentional AI users (near threshold).

39.0% are unintentional AI users (above threshold).

76% believe their health data is used to train AI.

51.7% say trust decreases when aware of data use.

60.3% say permission should be required.

37% say they would never allow data use.

Highest trust in hospital oversight among age groups.

Trust decreases after AI disclosure across diagnosis, treatment, insurance, and mental health use cases.

60+

This group is comparatively receptive to AI if safeguards are structurally in place

Policies emphasizing data de-identification, independent evaluation, and clinician oversight may be more effective than complex consent regimes.

Seniors are not anti-AI, but expect risk reduction to be built into systems rather than managed individually.

24.7% are intentional AI users (strongly under-represented).

40.0% are unintentional AI users (above threshold).

10.2% say they avoid using AI.

27.6% report being least comfortable with AI (lowest comfort).

59% believe their health data is used to train AI.

47.0% say trust decreases when aware of data use.

57.3% say permission should be required.

28% say they would never allow data use (lowest "never" rate by age).

Highest preference for anonymization as a condition for data use (38%).

Smallest trust decrease after AI disclosure across age groups.

What Matters Most

Education Level

No High-School Diploma

Polarization: high trust in some healthcare AI uses alongside the highest refusal of data use for training.

Strong emphasis on patient ownership and permission.

Policies should avoid ambiguity around commercialization and ownership.

Plain-language consent and strong default protections will matter more than layered governance frameworks.

33.1% are intentional AI users (under-represented).

48.1% report being comfortable with AI.

59% believe their data is used to train AI.

50.3% say trust decreases when aware of data use.

61.5% say permission should be required.

48% support data use under conditions.

49% say they would never allow data use (highest).

Highest "Highly trust" rates across most healthcare AI use cases.

High School Graduate or Equivalent

Moderate awareness and distrust, steady opt-in preference, and persistent concern about data use.

Policy design should emphasize clarity and predictability.

Clear rules about when AI is used, who controls data, and how refusal works will matter more than nuanced technical distinctions.

30.8% are intentional AI users (under-represented).

63% believe their data is used to train AI.

52.7% say trust decreases when aware of data use.

60.0% say permission should be required.

60% support data use under conditions.

40% say they would never allow data use.

Some College / Associate Degree

High unintentional AI exposure combined with some of the strongest trust declines after disclosure.

Elevated distrust despite familiarity.

This group illustrates a legitimacy gap: exposure without perceived agency.

Policies should prioritize visibility of control—who decides, who can stop AI use, and how oversight operates.

Failure to address this group risks widening skepticism among the "middle" of the population.

32.9% are intentional AI users (under-represented).

39.3% are unintentional AI users.

13.9% avoid AI.

75% believe their data is used to train AI.

54.0% say trust decreases when aware of data use.

58.3% say permission should be required.

66% support data use under conditions.

34% say they would never allow data use.

Most or second-most "Not at all trust" group across use cases.

BA+

High intentional AI use, high awareness, lower distrust, and greater acceptance of conditional data use.

Preference for anonymization over repeated consent.

This group is receptive to risk-based, proportional regulation.

Policies that differentiate between low-risk and high-risk AI uses, and emphasize evaluation and monitoring, will align with expectations.

Overly rigid consent mandates may be seen as inefficient rather than protective.

46.3% are intentional AI users (strongly over-represented).

38.2% are unintentional AI users.

74% believe their data is used to train AI.

48.1% say trust decreases when aware of data use.

56.4% say permission should be required.

71% support data use under conditions.

29% say they would never allow data use.

Highest anonymization preference (38%).

Second-highest "Highly trust" rates in some clinical use cases.

Among most "Not at all trust" groups for insurance and mental health use cases.

What Matters Most

Race and Ethnicity

White, Non-Hispanic

Moderate awareness and distrust, strong preference for patient ownership, and consistent trust declines after disclosure.

Legitimacy hinges on patient control and transparency.

Policies should clearly articulate ownership rights and refusal mechanisms.

Trust is sensitive to insurer-led AI uses, suggesting the need for tighter governance in coverage decisions.

35.2% are intentional AI users (slightly under-represented).

37.8% are unintentional AI users.

11.0% avoid AI.

69% believe their data is used to train AI.

52.6% say trust decreases when aware of data use.

58.4% say permission should be required.

65% support data use under conditions.

34% say they would never allow data use.

Black, Non-Hispanic

High AI use, high trust in healthcare AI, but the strongest insistence on explicit opt-in for data use.

Lower trust loss after disclosure compared to other groups.

Consent is central to legitimacy, even where baseline trust is higher.

Policies should avoid default data use without permission and clearly separate care improvement from commercialization.

Opt-in frameworks may be especially important for equitable governance.

44.6% are intentional AI users (over-represented).

68% believe their data is used to train AI.

46.5% say trust decreases when aware of data use.

58.4% say permission should be required.

63% support data use under conditions.

37% say they would never allow data use.

Highest explicit opt-in share (31%).

Highest “Highly trust” rates across healthcare AI use cases.

Hispanic

High AI awareness, moderate trust loss, and strong concern about data commercialization.

Preference for consent and oversight.

Policies should explicitly address how data is monetized and shared.

Transparency around secondary uses of data will be essential.

Oversight mechanisms must be visible and understandable, not implicit.

38.8% are intentional AI users (slightly over-represented).

74% believe their data is used to train AI.

50.5% say trust decreases when aware of data use.

59.6% say permission should be required.

61% support data use under conditions.

38% say they would never allow data use.

Frequently second-highest “Highly trust” group across use cases.

Asian–Pacific Islander, 2+ Races & Ethnicities, Non-Hispanic (NH)

High overall AI exposure (both intentional and unintentional), moderate awareness of health data use, moderate trust declines after disclosure, and comparatively low trust in institutional actors, including hospitals, state agencies, and insurers.

Legitimacy hinges on clear accountability and visibly independent oversight, given comparatively lower trust in traditional healthcare and government institutions.

Policies should emphasize external validation, transparency, and independent governance structures, rather than relying primarily on hospital- or state-led oversight models.

37.9% are intentional AI users (slightly over-represented).

40.7% are unintentional AI users (over-represented).

Not over-represented among those who avoid AI.

68% believe their data is used to train AI.

49.7% say trust decreases when aware of data use.

56.8% say permission should be required.

67% support data use under conditions.

31% say they would never allow data use.

What Matters Most

Household Income

<$30k

Lower AI use, higher avoidance, high refusal of data use, and strong demand for ownership and consent.

Economic vulnerability amplifies sensitivity to perceived exploitation.

Policies must clearly separate public benefit from private profit.

Strong default protections and limits on commercial data use will be critical for legitimacy.

31.5% are intentional AI users (under-represented).

12.6% avoid AI.

64% believe their data is used to train AI.

52.0% say trust decreases when aware of data use.

60.9% say permission should be required.

56% support data use under conditions.

44% say they would never allow data use.

Among highest "Highly trust" groups for healthcare AI use cases.

$30k - <$60k

High distrust associated with data awareness and strong consent demands.

Clear consent, refusal, and accountability mechanisms will matter more than technical assurances.

30.8% are intentional AI users (under-represented).

12.3% avoid AI.

67% believe their data is used to train AI.

57.7% say trust decreases when aware of data use.

63.1% say permission should be required.

61% support data use under conditions.

38% say they would never allow data use.

$60k–<$100k

Moderate trust declines and growing acceptance of conditional data use.

Risk-tiered policies with clear oversight and anonymization standards are likely to be acceptable.

Transparency must be paired with demonstrated governance.

36.0% are intentional AI users (near threshold).

38.1% are unintentional AI users.

11.3% avoid AI.

68% believe their data is used to train AI.

51.0% say trust decreases when aware of data use.

59.0% say permission should be required.

65% support data use under conditions.

35% say they would never allow data use

$100k+

Highest AI use and awareness, lowest refusal rates, strong preference for anonymization.

Policies emphasizing evaluation, auditing, and post-deployment monitoring will align well.

Consent should remain available but not overly burdensome.

47.0% are intentional AI users (strongly over-represented).

37.8% are unintentional AI users.

78% believe their data is used to train AI.

46.0% say trust decreases when aware of data use.

53.1% say permission should be required.

72% support data use under conditions.

28% say they would never allow data use.

Highest anonymization preference (37%).

What Matters Most

Gender

Men

Higher comfort, higher trust after disclosure, and lower opt-in demands.

Greater tolerance for developer and FDA involvement.

Men are more likely to accept institution-led governance models.

Regulatory clarity, technical standards, and federal oversight signal legitimacy.

Policy emphasis on performance, validation, and safety standards may resonate more than consent mechanics alone.

30.8% are intentional AI users (under-represented).

12.3% avoid AI.

67% believe their data is used to train AI.

57.7% say trust decreases when aware of data use.

63.1% say permission should be required.

61% support data use under conditions.

38% say they would never allow data use.

Women

Lower comfort, larger trust losses after disclosure, and higher opt-in demand.

Less trust in vendors and regulators relative to clinicians and independent bodies.

Consent and disclosure requirements will be judged not just on existence but on how meaningful they feel.

Policies must ensure consent is proactive

Trust-building requires visibly independent oversight rather than reliance on developer compliance.

34.4% are intentional AI users (under-represented).

39.1% are unintentional AI users (above threshold).

10.3% avoid AI.

66% believe their health data is used to train AI.

56.0% say trust decreases when aware of data use.

60.2% say permission should be required.

59% support data use under conditions.

40% say they would never allow data use.

24% report being least comfortable with AI.

Larger trust decreases after AI disclosure than males.

About CHAI

The Coalition for Health AI (CHAI) is a non-profit, industry-led public-private partnership dedicated to advancing responsible AI in healthcare. Representing over 6,000 members from over 3,000 organizations—including more than 200 health systems, leading academic centers, patient advocacy groups, startups, and technology innovators—CHAI brings together diverse stakeholders to create consensus-driven best practices, assurance frameworks, and practical tools.

Our mission is to foster trustworthy, transparent, and equitable AI adoption that improves care quality, safety, and outcomes. Through initiatives like the Blueprint for Trustworthy AI, the Responsible AI Guide, and the widely adopted Applied Model Card, CHAI translates high-level principles into actionable standards that support developers, health systems, policymakers, and patients alike.

Get in touch at
admin@chai.org ✉

Visit our website ↗

Connect on
LinkedIn ↗

CHAI

Coalition for Health AI