markup ai

Content Guardian Agents™

Sounds Like Everyone Else:

The 2026 Content Credibility Report

What 1,000 U.S. adults say makes branded content worth trusting, and what makes them doubt it.

A MARKUP AI REPORT

markup ai
Content Guardian Agents™

Executive summary

AI disclosure isn't a substitute for content credibility.

Brands are publishing more content with AI than ever. Many are asking whether to tell audiences. Readers are asking something more basic: Does this sound like it came from a brand that knows what it's talking about?

Markup AI surveyed 1,000 U.S. adults about what makes them doubt branded content, what earns their trust, and what they've noticed as more companies use AI. The answer is clear: Readers judge content by how it reads and whether it holds up, not just by how it was made.

The credibility problem in numbers:

73.3%

have noticed brands' content sounding more similar as companies use AI (33.0% definitely, 40.3% somewhat)

59.2%

name writing that sounds fake, robotic, or overly polished as a reason for skepticism — the most common reason selected

80.6%

rate accuracy as very or extremely important when deciding whether to trust branded content

Disclosure still matters to many readers. But it works best alongside content that is distinct, resonant, accurate, and supported: the four tests this report walks through.

markup ai
Content Guardian Agents™

Readers believe AI content is part of their daily routine

Most readers believe they're already reading AI-created content every day.

In this survey, 31.4% of respondents think they encounter AI-created content multiple times a day. Another 34.8% think they encounter it daily. That's 66.2% who see it as a daily presence. A further 20.2% say weekly, and just 7.3% say never.

These are perceptions, not measurements. But perception is what readers bring to every page a brand publishes.

66.2%

believe they encounter AI-created content daily or more often.

Trust is divided, not collapsing

We asked respondents to compare their trust in branded content today with two years ago:



This is respondents' recollection, not a comparison against an earlier survey.

The takeaway: About three in 10 readers say their trust has slipped, and fewer than one in four say it has grown. Brands can't assume trust. They have to earn it with every piece.

markup ai
Content Guardian Agents™

Disclosure matters, but it can't carry the weight alone

Readers are split on what earns trust. No single answer wins a majority.

We asked what would do more to earn trust. The answers divided four ways:

32.9%
say a label and high-quality content matter equally

28.8%
say relevant, authentic, accurate content matters more, regardless of how it was created

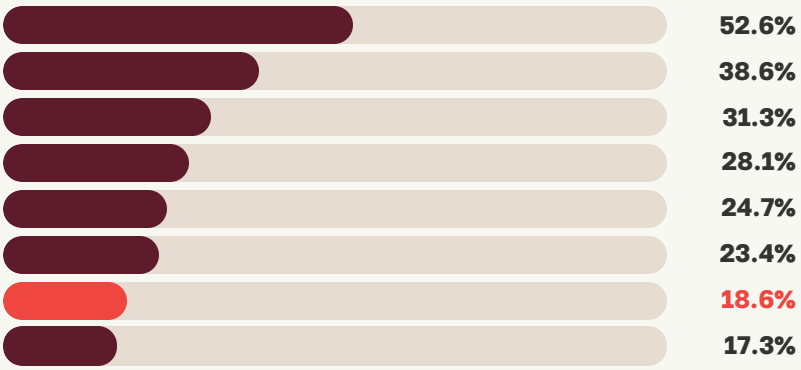
25.0%
say labeling alone matters more

13.3%
say neither

Nearly one in three readers want transparency and quality together. Only one in four put a label first.

What builds trust in an unfamiliar brand

When respondents could select everything that would make them more likely to trust an unfamiliar brand, they chose:



The AI-use label ranked seventh of the eight named options. Respondents could pick as many as they wanted, so this isn't a head-to-head test. It shows what readers reach for first.

Does it matter that AI wrote it?

When content is already relevant, accurate, and authentic, does it matter that AI created it? Readers split almost evenly: 52.8% say it matters little or not at all, while 47.2% say it matters somewhat or a lot.

And the label itself isn't neutral for everyone. 25.1% included disclosed AI use among their reasons for skepticism.

The takeaway: Transparency explains how content was made. Quality gives readers a reason to trust it. Brands need both.

markup ai
Content Guardian Agents™

Distinct: the risk is sounding like everyone else

The risk isn't that content sounds like AI. It's that it sounds like everyone else.

Robotic or overly polished writing was the most common reason respondents gave for doubting branded content. **59.2%** selected it, ahead of unsupported claims (40.4%). Respondents could choose up to three reasons.

73.3%

say they've noticed brands' content sounding more similar as companies use AI



59.2%

cite writing that sounds fake, robotic, or overly polished

23.2%

cite content that sounds generic enough to come from any brand

One in three respondents (33.0%) say they've *definitely* noticed it. Another 40.3% say they've *somewhat* noticed it.

These are readers' perceptions. The survey doesn't measure whether brand content has changed over time.

The takeaway: A distinct voice is a credibility signal. When content sounds interchangeable, readers notice, and many say it makes them skeptical.

markup ai
Content Guardian Agents™

Resonant: relevance and authenticity count

After accuracy, relevance is the trust builder readers select most.

Asked what would make them more likely to trust an unfamiliar brand, 38.6% chose content that's relevant to their needs, and 31.3% chose content that sounds authentic. Only clear, accurate information ranked higher (52.6%).

When respondents rated how important each quality is to trusting branded content, both landed well above the halfway mark:

63.8%

rate relevance very or extremely important (38.0% very, 25.8% extremely)

63.1%

rate authentic-sounding writing very or extremely important (40.0% very, 23.1% extremely)

The flip side shows up in skepticism. 13.4% say content that isn't catered to them gives them a reason to doubt a brand.

The takeaway: Content can be accurate and still miss. Relevance and an authentic voice are part of the credibility test, not extras on top of it.

markup ai
Content Guardian Agents™

Accurate: the baseline for trust

Accuracy is the entry requirement.

80.6%

rate accuracy very or extremely important when deciding whether to trust branded content.



That's a higher share than relevance (63.8%) or authentic-sounding writing (63.1%).

Readers also name inaccuracy as a reason to doubt a brand:

40.4%
cite unsupported claims

23.7%
cite outdated information

And for some, the cost has been personal. **22.0%** say they've made a decision based on AI-generated information they later learned was inaccurate. Another 33.3% say definitely not. These are self-reports, so they show experience, not how often AI content is wrong.

The takeaway: Accuracy is the cost of entry, not what sets a brand apart. Content that's wrong, unsupported, or out of date gives readers a reason to doubt it.

markup ai
Content Guardian Agents™

Supported: readers want to see the evidence

Readers want to see what a claim rests on.

Unsupported claims were the second most common reason for skepticism, selected by **40.4%**. Only robotic writing ranked higher. Another 20.6% pointed to missing sources or citations.

The same theme shows up in what builds trust:

WHAT RAISES DOUBT

40.4%

unsupported claims

20.6%

missing sources or citations

WHAT BUILDS TRUST

28.1%

say clear citations would make them more likely to trust an unfamiliar brand

24.7%

say third-party data or research would

These numbers come from two separate questions, so they shouldn't be added together. Taken one at a time, each shows that a meaningful share of readers look for evidence.

The takeaway: Readers set a bar for evidence, and many notice when content doesn't meet it. Claims that show their support meet that bar.

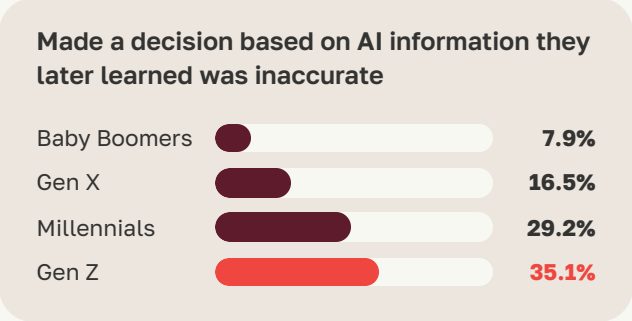
markup ai
Content Guardian Agents™

Younger readers report more friction with AI content

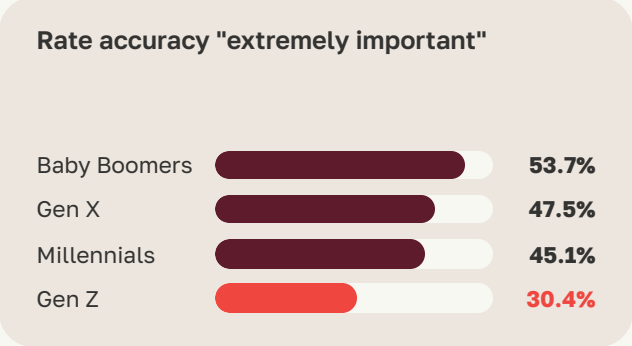
Gen Z reports noticing sameness and inaccuracy most. Treat it as a signal, not a verdict.

On three of the four measures below, results climb steadily from Baby Boomers to Gen Z.









Sample sizes: Baby Boomers (n=190), Gen X (n=284), Millennials (n=339), Gen Z (n=168). Bars share one scale.

That last measure runs the other way, and it needs careful reading. It captures how strongly respondents rate accuracy, not whether they value it. Gen Z is simply less likely to pick the top of the scale.

All four are self-reported. They don't show that younger readers are exposed to more AI content, or that older readers care less about quality.

The takeaway: Don't assume one content strategy fits every audience. If your readers skew younger, expect more of them to say they notice sameness.

markup ai
Content Guardian Agents℠

Five questions to ask before you publish

If readers judge content by how it reads and whether it holds up, the review that matters happens before you hit publish.

The survey measured reader perceptions. The questions below are Markup AI's recommendations, not practices the survey tested.

1. **Is it distinct?** Would a reader know this came from your brand, and only your brand?
2. **Is it resonant?** Does it answer what this reader came to find out, in a voice that sounds human?
3. **Is it accurate?** Is every claim correct and current?
4. **Is it supported?** Can the reader see what each claim rests on?
5. **If you disclose AI use, does the content earn the trust the label asks for?** Transparency and quality work together.

Checking all five by hand slows AI down. Content Guardian Agents℠ close that gap. They review content against your brand, accuracy, terminology, and compliance standards before it publishes.

> AI creates the draft. Markup AI makes it sound like you.

markup ai

Content Guardian Agents℠

Key takeaways

- Disclosure matters to many readers, but only 25.0% say a label alone does more to earn trust than quality.
- Robotic or overly polished writing is the top reason for skepticism (59.2%), and 73.3% have noticed brands sounding more alike.
- Accuracy is the baseline: 80.6% rate it very or extremely important.
- Relevance, authenticity, and visible support for claims all shape what readers trust.
- Younger respondents report noticing sameness and acting on inaccurate AI information more often.

Distinct **Resonant** **Accurate** **Supported**

About Markup AI

Markup AI's Content Guardian Agents℠ are the industry's first always-on AI content guardrail and editor. Markup AI is the content control layer that checks every piece — human- or AI-written — against your brand voice and style guidelines, so everything you publish is on-brand, high-quality, and easily citable.

markup.ai

markup ai

Content Guardian Agents™

About this research

Markup AI commissioned Dynata to run this survey online in September and October 2026. It collected 1,000 responses from U.S. adults age 18 and older.

Unless otherwise noted, all statistics reflect the full sample of 1,000 respondents. Percentages may not sum to 100% due to rounding or multi-select responses.

How to read the data

- Questions on reasons for skepticism let respondents select up to three options. Questions on what builds trust let them select all that apply. Options within those questions aren't additive.
- Generational results reflect these sample sizes: Baby Boomers (n=190), Gen X (n=284), Millennials (n=339), and Gen Z (n=168).
- Comparisons across questions are comparisons of separately asked questions, not choices respondents made side by side.

What the survey measures

The survey measures what respondents perceive and what they report about their own behavior. It doesn't show whether brand content has changed over time, because there is no earlier wave to compare against. It also didn't test any product or whether any practice increases trust.