State of Digital Trust 2026

The AI Transparency Report

>>>>> What 11,000 Consumers Told Us About Trust, Data, and Paying More

USERCENTRICS × Cookiebot by Usercentrics

Executive summary

Consumers will **now pay 7% more for the brands that have earned their trust on AI**. They've already walked away from the ones that don't.

Two-thirds did so in the past six months.
One in four canceled a subscription.
One in five switched to a competitor.

The AI product you used this morning probably asked permission to access something. Your calendar. Your inbox. Your bank account. Your customers. A year ago, that question barely existed. Today, agentic AI is the default mode of the tools enterprises run on, and consumers are quietly forming opinions about how every brand they touch is handling it.

2026 is the year that opinion started showing up in revenue.

> **"Consumers are making purchasing decisions also based on how brands handle their data, and over half are willing to pay more to the ones that get it right. The brands that move first won't just earn the premium. They'll earn a category position that's almost impossible to compete against once it's established."**

Tilman Harmeling
Strategy & Market Intelligence, Usercentrics

In last year's State of Digital Trust, we told you consumers were worried about how brands use their data. This year they stopped talking and started acting.

The year-on-year data confirms it. The share of consumers who trust AI less than humans with their personal data has risen from 48% to 52%, the biggest single movement in the entire survey. Cookie acceptance is declining, the numbers aren't spiking and resetting. **They're compounding.**

24% of consumers, almost one in four, canceled a subscription or stopped purchasing from a brand altogether in the past six months because of concerns about how their data was being used in AI. One in five walked to a competitor they trusted more. **Almost half (47%) took at least one action with direct revenue consequence, canceling, switching, or reducing spend.** On the other side of the ledger, over half (52%) of consumers are willing to pay 7% more to the brands that get it right.

The United States (US) makes the commercial case most clearly. Half of American consumers will pay more for a brand that is transparent about AI, in line with the global average. What sets the market apart is the institutional vacuum behind that number: only 39% of US consumers trust government services with their data, the lowest of any market in the study. Where institutional trust is at an all-time low, brands can step in and earn it.

And here is the structural finding that connects 2025 to 2026: 46% of consumers still don't have a good understanding of how their data is collected and used, identical to last year. Two years of privacy headlines have not moved it. That's the gap brands can close, and it's the one that pays.

For the second year running, Usercentrics commissioned Sapio Research to survey 11,000 consumers across seven markets, the United Kingdom, United States, Germany, Netherlands, Sweden, Spain, and Italy, on data privacy, consent, and AI.

This report shows you what consumers expect from the brands they reward, where the commercial opportunity lies, and how to capture it. The **T.R.U.S.T. Framework** at the end of this report is how to build your trust infrastructure for the next era of AI demands.

Global Data Highlights

24% of consumers canceled a subscription or stopped purchasing from a brand altogether in the past six months because of concerns about how their data was being used in AI

71% are concerned that AI-driven personalization feels intrusive

47% took at least one action with direct revenue consequence, canceling, switching, or reducing spend

52% trust AI less than humans with their personal data, up from 48% in 2025, the biggest single YoY movement in the dataset

52% would pay more for a brand that is transparent about how it uses AI with their data

48% click "accept all" on cookies less often than three years ago, up from 46% in 2025

7% average price premium consumers are willing to pay for AI-transparent brands, rising to 73% willingness in Germany

CHAPTER 1

The trust economy

For years, trust has been the kind of word that sits on the values slide. Important, intangible, generally agreed upon, rarely measured.

Those days are over. In 2026, trust shows up in churn rates, switching behavior, and willingness to pay more. This report is the first full read on how much that's worth.

Last year's report captured consumers at a tipping point. This year shows the tipping point has passed. Consumers stopped describing how they felt about AI and started acting on it: two-thirds have already changed a purchasing decision over it in the past six months.

Three structural forces collided in 2025–26 to turn consumer concern into purchasing decisions.

Agentic AI went mainstream.

For most of 2023 and 2024, AI was something consumers interacted with. By late 2025, it was acting on their behalf: booking meetings, accessing inboxes, and connecting to financial services. That changed the consent question entirely. It stopped being about what a user chooses to share and became about what an AI agent accesses before making decisions.

Regulation fragmented.

The EU AI Act moved from phased implementation into active enforcement. Over twenty US states now have comprehensive privacy laws in effect, with no federal standard unifying them. On 31 March 2026, four UK regulators jointly published a foresight paper mapping agentic AI governance requirements for the advertising and marketing sector. Consumers are reacting to an environment where the rules are visibly, publicly catching up to the practice.

Evidence kept accumulating.

The "accept all" decline isn't a single event but a compounded result of high-profile data breaches, AI training corpus controversies, and intensified cookie banner enforcement across Europe through 2025. Each incident moved a different segment of consumers one step further from passive acceptance toward active decision-making.

These three forces are what the data in this report is measuring the consumer response to.

What consumers will pay for

Over half (52%) of consumers globally will pay more for a brand that is transparent about how it uses AI with their data, with an average premium of 7%. Of those, 28% would stretch to 10% above the regular price, and another 16% would pay 11–20% more.

52% would pay more

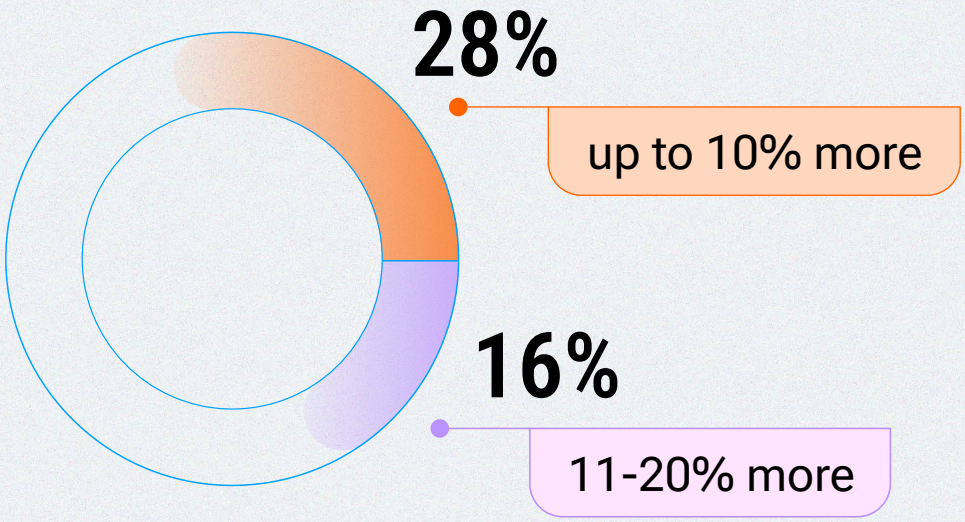

The willingness runs deepest among younger consumers. Among 18- to 29-year-olds, 67% will pay more. Among consumers aged 60 and above, only 23% say the same. The generation most willing to pay is the one brands most need to win long-term.

Market spread: *Germany leads at 73% (9% average premium). The Netherlands is lowest at 35%. Range: 38 points, the widest spread on any commercial metric in the study.*

In Germany, willingness peaks at 73%, the highest of any market in the study. In the US, half of consumers are willing to pay more for AI transparency, against a backdrop where only 39% trust government services with their data at all. Two very different markets. Same commercial signal: consumers are ready to reward the brands that get this right.

The trust trinity

Consumers are clear about what would make them pay it. **Three drivers come up again and again:**

44%

clear explanations of how their data is used

42%

strong security guarantees

41%

ability to control what they share

Together, they describe one consistent demand: **transparency, security, and control.**

These three drivers have not changed in two years. The formula is known. The gap is in execution.

What that gap looks like in practice: while nearly half of consumers want clear explanations, **46% still don't have a good understanding of how their data is collected, identical to 2025. The knowledge gap is structural.**

> **Market spread:** *The understanding problem is most acute in markets often assumed to be the most privacy-mature: Sweden (56%) and Germany (53%) don't have a good understanding of how their data is collected and therefore have the largest comprehension gaps. The US and UK are the most informed markets at 41% each.*

The commercial cost shows up in the privacy-aware split: consumers who understand what's happening with their data are nearly three times more likely to be comfortable with personalization online, and significantly more willing to pay the AI transparency premium.

Consumers have told brands exactly what trust is worth and what it takes to earn it. The brands that listen will command the premium.

Take action — Build the trust infrastructure

Consent is where trust is earned or lost. Cookiebot by Usercentrics is trusted by over 2.4 million websites to handle consent across web, app, and CTV — with automatic compliance across GDPR, CCPA, and 20+ US state laws. For enterprise deployments, Usercentrics CMP offers deeper customization across your full stack. Privacy Policy Generator produces plain-language policies consumers actually read.

EXPLORE CMP

CHAPTER 2

AI has changed the game

AI has shifted from something you interact with to something that acts within your systems on your behalf, booking meetings, drafting responses, shaping search results, and connecting to the services you use every day. The technology moved into daily life faster than the trust infrastructure built to govern it.

Consumers noticed.

31%

have warned friends and family or complained publicly

35%

have taken two or more actions - canceled, switched to a competitor or reduced spend

- Another 24% avoided trying a new product from that brand
- 20% switched to a competitor they believed handled AI data more responsibly
- 20% reduced their spending

Market spread: *Spain leads at 76%. The Netherlands trails at 53%. Range: 23 points. The markets most likely to act are not always the markets most likely to feel concerned.*

These actions don't carry equal weight. Cancellation and competitor-switching are the actions brands feel first in their numbers. Reduced spend and avoided trial are slower drains that often show up as unexplained softness in the quarterly review. Public complaints damage acquisition for consumers a brand hasn't even met yet, which is harder to quantify, even harder to mitigate.

Why 2026 is a different kind of year

The sentiment behind this year's behavior change has a longer history than the AI conversation. Consumers have been getting more cautious about data sharing, more selective about consent, and more skeptical of institutions for some time. AI didn't create that skepticism, but it gave consumers something specific enough to act on.

Over half (52%) of consumers now trust AI less than humans with their personal data, up from 48% in 2025. That four-point shift is **the biggest single year-on-year movement in the entire dataset**. And 60% are uncomfortable with their personal data being used to train AI models at all.

The shift is not uniform. Men are more willing than women to pay a premium for AI transparency, 58% versus 47%. **That 11-point gap cuts against the assumption that AI wariness is a single, undifferentiated consumer mood.** People process AI risk differently, and one-size communications will not land with most of them.

Personalization isn't the problem

The same awareness that makes consumers more cautious about data-sharing also makes them more receptive to personalization done right.

Seven in ten (71%) consumers find AI-driven personalization intrusive. Read in isolation, that sounds like a mandate to pull back. But among privacy-aware consumers, the ones reading cookie banners, who understand their data rights, and make active consent decisions, 53% are comfortable with companies using their data to personalize their experience. Among privacy-unaware consumers, that figure drops to 19%.

Both numbers are real. They describe different points on the personalization spectrum.

Market spread: *The Netherlands leads concern at 77%. Sweden (60%) and Germany (66%) are lowest, see Chapter 5 for why the German figure looks counterintuitive.*

The constraint on personalization isn't technology or regulation. It's communication. **Brands that explain what they're doing to the right audience unlock nearly three times the consent.**

Nowhere is this clearer than in Germany with 86% of German consumers aged 18 to 29 having stopped using a company's services due to data privacy concerns, versus 44% aged 60 and above. This is a 42-point gap, the largest generational divide on this question in the entire dataset. Young German consumers are not just the most privacy-conscious audience globally. They're also the most commercially motivated, with 73% willing to pay more for a brand that gets AI transparency right.

The consumers most likely to leave are also the most likely to stay, and pay more, for the brands that earn it.

And these aren't isolated reactions. 35% of consumers, more than one in three globally, took two or more actions against the same brand over AI data concerns in the past six months. Cancelation and warning friends and family. Switching and reducing spend. Public complaint and competitor migration. The consumer response to AI data concerns is compounding across multiple behaviors at once, which is the clearest indicator that this is a structural shift in how consumers relate to brand trust, not a passing reaction to a news cycle.

Take action — From AI anxiety to AI advantage

The consumers most concerned about AI data use are also the most willing to reward brands that handle it well. MCP Manager by Usercentrics enables AI connectivity and access governance — so your teams can move fast without losing oversight.

EXPLORE MCP MANAGER BY USERCENTRICS

CHAPTER 3

The consent shift

Giving consent used to be a passive transaction. A banner appears, the user clicks through, a signal is logged. For most brands, that was good enough.

That can't be the case anymore, and the reason isn't that consumers have suddenly become more anxious. Two distinct shifts are happening at once, and they pull in opposite directions. The first is **consent fatigue**: users worn down by years of cookie banners, clicking through to reach the content. The second is quieter but more consequential, a **privacy awakening**, where consumers click into the "more information" link more often, read what's actually being done with their data, and choose more deliberately. The decline of "accept all" in this year's data is the visible outcome of both forces combining.

Almost half (48%) of consumers click "accept all" less often than they did three years ago. Only 23% click it more often. That two-to-one ratio is a structural behavioral shift playing out across every market in the study. **The campaign signals brands have built their measurement infrastructure around are degrading at the source.**

Market spread: *The Netherlands leads the decline at 57%. Sweden and the US are the slowest movers at 42% each. Every market is shifting in the same direction; only the pace varies.*

What's replacing passive acceptance is active consideration. Over half of consumers now selectively manage their cookie preferences, with 39% accepting only necessary cookies and 16% customizing their settings. The user who clicks "accept all" without thinking is becoming the minority. The user who reads, interprets, and decides is becoming the norm.

For brands, distinguishing between the two groups, the fatigued and the awakened, is the difference between optimizing the experience and being penalized by it. Both groups are rejecting "accept all" at higher rates. **Only one of them can be won back with better design.**

The 31-point gap

Not all consumers are moving at the same pace, and the distance between the fast movers and the rest is worth paying closer attention to.

Among privacy-unaware consumers, 57% still accept all cookies by default. Among privacy-aware consumers, that figure drops to 26%. A 31-point gap, on a single behavior, driven entirely by understanding.

As awareness spreads, driven by AI headlines, data breach coverage, and an increasingly educated digital public, the gap is likely to narrow. The 57% represents the consent rate brands are still relying on. It's probably not the consent rate they should be banking on.

The understanding gap is a measurement problem

The constraint underneath all of it: 46% of consumers still don't have a good understanding of how their data is collected and used. That figure has not changed in two years.

The systems built to explain data use to consumers, banners, policies, notices, aren't doing the job they were designed to do. According to NordVPN, a typical internet user would need a full working week each month to read the privacy policies of every site they visit. Forrester's privacy research reaches the same conclusion from a different angle: confronted with a wall of choices in the first seconds of arriving at a site, the rational response is to skip through. **This failure isn't on the user. It's a design failure, and it compounds every year it goes unfixed.**

The implication for marketers is direct. If nearly half of your audience doesn't understand what they're consenting to, they'll make the fastest choice, which trends toward rejection as awareness grows. Suppressed consent rates are a direct reflection of communication quality.

Better consent design is where privacy compliance becomes performance. Shorter explanations, plain language, interfaces that give users a genuine sense of control rather than the appearance of it, these are more than UX nice-to-haves. Comprehension goes up. Opt-in follows. So does the quality of the data your campaigns run on.

CHAPTER 4

The AI access frontier

The capabilities of agentic AI go far beyond answering questions to acting on your behalf. It can access your calendar and books your meetings, reviews your contracts, monitors your bank accounts, and makes decisions inside your preferred tools and services.

This isn't a future scenario. The infrastructure connecting AI assistants to third-party apps and services is being built and adopted right now, by the platforms consumers already use every day. The question for every brand and developer in that ecosystem is the same one running through this report: **are you building the trust infrastructure ahead of the access, or after the fact?**

The consumer data gives a precise answer about what that infrastructure needs to be.

Where consumers draw the line

Comfort with AI access is not uniform. It falls on a spectrum from low-stakes to high-stakes choices.

49% are comfortable with AI assistants accessing work tools - the highest of any category tested

37% are comfortable with access to financial accounts - the lowest

23% will allow it, but only if they can approve each individual request

17% are uncomfortable but would allow access anyway – known as resigned consent

Only 8% are fully comfortable with AI access without any conditions

It's worth pausing on resigned consent. **It's the most unstable form of permission a brand can have. It reflects friction, not trust.** A user who allows access because opting out feels harder than going along is one bad experience away from becoming a very public reason not to trust your platform.

> **Market spread:** *Swedish and Dutch consumers are the least comfortable with AI assistants accessing personal data. Comfort increases with younger demographics in every market.*

The results:

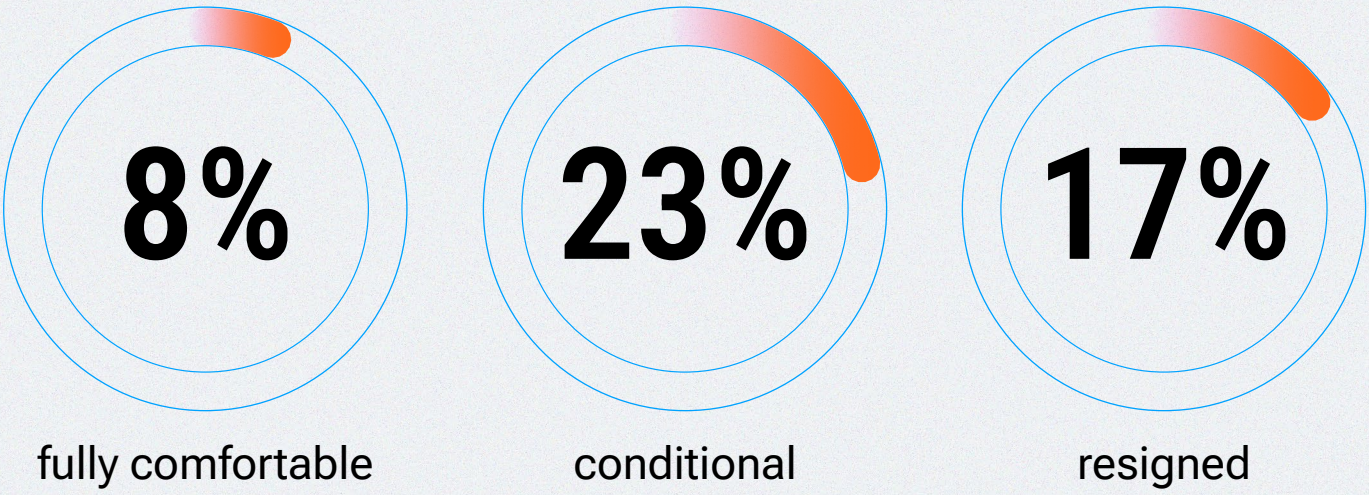

That leaves just over half of consumers (52%) who haven't yet formed a view. **That's where the commercial opportunity is.**

Human preference is a trust gap, not a ceiling

Two-thirds (65%) of consumers prefer humans for healthcare interactions, 61% for financial advice, and 61% for handling complaints.

One way to read these numbers is as a call for hard limits on AI in high-stakes contexts. Another is to read them as a map of where the trust infrastructure needs to be strongest before agentic AI can be brought in.

The preference for human interaction in sensitive situations isn't about consumers doubting what AI can do. It's about trust, control, and the reassurance they need before handing over access to the things that matter most: clarity on what's being accessed, confidence that they can revoke consent at any point, and visible evidence

that the organization on the other side has built permission into the product rather than bolted it on afterward. That's exactly what agentic AI infrastructures built on explicit, revocable, granular consent make possible.

The contexts where human preference is highest are the ones where getting AI trust right pays off the most. Building the permission layer now, before scale forces the question, builds trust with consumers and unlocks the access that matters.

Take action — Build the compliance layer before you need it

MCP Manager by Usercentrics is built for exactly this moment. As AI agents connect to a new class of data, MCP Manager brings AI connectivity and access governance to the systems AI agents are starting to access, so the same oversight that protects your website and app extends to what AI can reach and action. Brands that deploy it now are building the foundation that consumers are already asking for.

EXPLORE MCP MANAGER BY USERCENTRICS

CHAPTER 5

Where is trust in the world?

Trust is not uniform. Across the seven markets in this year's study, consumer attitudes to data privacy, AI, and brand accountability vary significantly – and so does the commercial opportunity. The concerns driving a German consumer in their late twenties are not the same ones shaping a Dutch consumer in their fifties. The institutional vacuum in the US has no equivalent in Sweden.

The table below gives the comparative view at a glance. The profiles that follow give the details.

(Sweden joined the study for the first time in 2026; year-on-year comparisons exclude Sweden throughout.)

MARKET	Will pay more for AI transparency	Acted over AI data concerns	Find AI personalization intrusive	Don't understand data collection	Would stop using service over misuse	Trust banking with data
Germany	**73%** (~9% avg)	75%	66%	**53%**	66%	55%
United States	50%	67%	75%	41%	78%	63%
United Kingdom	50%	61%	75%	**41%**	**80%**	65%
Spain	49% (~6% avg)	**76%**	70%	42%	72%	52%
Italy	42% (~5% avg)	63%	73%	44%	73%	51%
Sweden	42%	62%	**60%**	**56%**	79%	**69%**
Netherlands	**35%**	**53%**	**77%**	48%	74%	61%
Global average	**52%** (7%)	**67%**	**71%**	**46%**	**74%**	**60%**

Bold = highest or lowest in the study.
Premium percentages in parentheses for markets where willingness-to-pay is highest.

What holds across all seven markets is the foundation. Consumers everywhere want the same three things: clear explanations of how their data is used, strong security guarantees, and the ability to control what they share. The ask hasn't changed in two years. What changes by market is the urgency, the trigger, and the audience most likely to act on it first.

Three patterns emerge:

Where commercial concern has matured into commercial action — Spain and Germany lead, with action rates of 76% and 75%, both well above global.

Where institutional trust has collapsed and the trust vacuum is widest — the US, where government trust sits at 39% but commercial willingness-to-pay holds at 50%.

Where high concern has not yet translated into behavior — the Netherlands, the only market that ranks lowest on three separate measures: paying more, acting, and feeling comfortable with personalization.

The UK is moving fastest along the awareness curve. Sweden joins the study for the first time, appearing in the spotlight section at the end of this chapter.

UNITED STATES

THE TRUST VACUUM NO INSTITUTION IS FILLING

The US is the market where the commercial case for trust is most urgent.

Only **39% of American consumers trust government services with their data**, the lowest figure of any market in the study, in a political context that shows no sign of reversing: federal rollback of privacy protections, sustained scrutiny of Big Tech, and a generation that has watched institutions fail them on data with reliable consistency. Almost three-quarters (74%) of US consumers are cautious about sharing data with US companies. Americans are among the most wary of businesses from their own country.

The gap that institutions are vacating is real, measurable, and open. **Half of US consumers will pay more for a brand that is transparent about how it uses AI with their data.** 78% would stop using a service over data misuse, above the global average. AI anxiety runs high too: 75% find AI-driven personalization intrusive, tied with the UK for the highest figure of any market, and the US has the largest "very concerned" segment of any market at 39%.

What makes the US distinctive is not the level of concern (several markets match it), but the absence of any institutional counterweight. In Europe, GDPR gives consumers a framework to stand on, even if awareness of it is imperfect. In the US, the regulatory floor is lower and the public trust in those who set it is lower still. Every point of trust a brand earns here, it earns on its own terms.

39% of US consumers trust government services with their data, the lowest of any market. Half will pay more for a brand that earns what institutions no longer provide.

UNITED KINGDOM

THE MARKET THAT'S CATCHING UP

The UK is not the loudest market in this dataset, but it's one of the most instructive.

Rights awareness improved more in the UK than in any other market this year: **43% of UK consumers are now unaware of their data privacy rights, down from 50% in 2025.** That seven-point shift is the largest single-market rights awareness movement in the study. It matters because awareness is the variable that drives everything else: more informed consumers reject more cookies, take more action, and make more deliberate decisions about the brands they stay with. The UK is moving along that curve faster than anywhere else.

The behavior is already reflecting it with **80% of UK consumers saying they would stop using a service if their data was misused**, the highest of any market in the study, six points above the global average. Six in ten (61%) are uncomfortable with their data being used to train AI models. And UK consumers are also among the most cautious about sharing data with companies from China (84%) and the US (79%), both above global averages.

What makes the UK distinctive is the gap between attitude and action. UK consumers have the lowest rate of stopping use of a company over privacy concerns in the dataset (56% vs. 65% global), but the highest threshold for what would make them stop. The bar is high, but once it's crossed, there's no return.

A 7-point drop in rights unawareness in a single year, the largest improvement of any market. The UK consumer is becoming more informed, more deliberate, and less forgiving of brands that assume otherwise.

GERMANY

THE MARKET WHERE ACTION HAS ALREADY OUTPACED SENTIMENT

Germany is the market where the data stops being theoretical.

Almost three-quarters (73%) are willing to pay more for AI transparency, the highest globally and by a wide margin, at an average premium of 9%. Sentiment alone doesn't drive that number; behavior does. **Three-quarters of German consumers have taken action against a brand due to AI data concerns**, second only to Spain (76%). These are not attitudinal signals. They are commercial ones, from a consumer base that has spent decades in a culture shaped by some of the strictest data privacy norms in the world.

The comprehension picture complicates the read. Over half **(53%) of German consumers don't have a good understanding of how their data is collected and used**, the second-highest comprehension gap in the study after Sweden, and seven points above the global average. The market most willing to pay for transparency is also one of the markets least clear on what's actually happening. That's a paradox brands can use: in Germany, explaining clearly is rewarded more than anywhere else.

One counterintuitive finding worth noting: only 66% of German consumers are concerned about AI personalization feeling intrusive, the second-lowest figure in the study after Sweden. The apparent contradiction resolves when you look at the action data: German consumers aren't more relaxed about AI. They've simply moved past concern into response.

73% of German consumers will pay more for AI transparency, the highest of any market, at a 9% average premium. Concern here long ago became behavior.

NETHERLANDS

THE SKEPTIC MARKET WITH A GENERATIONAL FAULT LINE

The Netherlands sits in a category of its own: high concern, low commercial activation, and a generational gap that makes the overall numbers misleading.

77% of Dutch consumers find AI-driven personalization intrusive — the highest of any market. Only 25% are comfortable with companies using their data to personalize their experience, the lowest comfort level in the study. And yet only 53% have taken AI-related action against a brand, the lowest of any market, and just 35% are willing to pay more for AI transparency, also the lowest globally.

The Dutch consumer, in aggregate, is deeply uncomfortable and largely unmoved. That pattern, high anxiety, low activation, is a signal in itself. It suggests a population that has not yet found a framework for translating concern into action, or brands that have not yet given them a compelling reason to. Dutch consumers are also among the most cautious about sharing data with both US (77%) and Chinese (81%) businesses, the suspicion is generalized.

The generational data complicates that read. Among Dutch consumers aged 18 to 29, the willingness-to-pay figure jumps materially, well above the 35% overall figure and closer to the global average. The skepticism is concentrated in older cohorts. What looks like a low-activation market at the aggregate level contains a younger segment behaving like a different market entirely.

35% overall willingness to pay for AI transparency — the lowest of any market. But concern (77%) is the highest. The brand that solves "what to do about this" wins a market currently holding all its commercial energy in reserve.

SPAIN

THE MARKET THAT ACTS FIRST AND ASKS QUESTIONS LATER

Spain has the highest activation rate of any market in the study.

92% of Spanish consumers have taken steps to protect their personal data in the past six months, five points above the global average and the highest of any market. And **76% have taken action specifically against a brand over AI data concerns, also the highest of any market**, fractionally ahead of Germany. Spain is a market where concern converts into behavior at a higher rate than anywhere else in the study.

Spanish consumers also pull above their weight on commercial action. 73% have stopped using a company over privacy concerns (tied with Germany for the highest of any market), and they're more likely than the global average to have switched to a competitor (25% vs 20%) or reduced spend (23% vs 20%) over AI data concerns specifically.

The commercial opportunity sits in the gap between behavior and premium. Almost half **(49%) of Spanish consumers are willing to pay more for AI transparency, at an average premium of 6%**. That's broadly in line with global figures, but considering the unusually high behavioral engagement, it represents a consumer base that is already primed to reward the brands that meet them where they are. In Spain, the question isn't whether consumers are activated. It's whether brands are ready to capitalize on it.

76% of Spanish consumers have taken action against a brand over AI data concerns, the highest of any market. The activation is there. The pricing power follows the brand that earns it.

ITALY

THE MID-MARKET WITH MARGIN PRESSURE

Italy sits close to global averages on most dimensions, and the one place it diverges most sharply is where it matters most commercially.

The majority (86%) of Italian consumers have taken data protection action in the past six months: 73% find AI personalization intrusive and 63% have taken AI-related action against a brand. On each of these measures, Italy tracks within a few points of the global figure, a market that reflects the broad consumer mood without amplifying it.

The outlier is the AI transparency premium: **42% of Italian consumers are willing to pay more for a brand that is transparent about AI data use. But the average premium they'd pay is just 5%, the lowest of any market in the study.** Italy is a market where the willingness exists but the commercial intensity behind it is thin. Trust matters here, but it competes with price sensitivity in a way that doesn't show up as clearly in other markets.

For brands operating in Italy, the implication is less about commanding a premium and more about protecting against churn: 73% would stop using a service over data misuse, consistent with global norms. The downside risk of getting trust wrong is real. The upside of getting it right is more modest. In Italy, trust is a defensive asset, not an offensive one.

5% average AI transparency premium, the lowest of any market. Italian consumers value trust. They're just less willing to pay extra for it. The commercial case here is retention, not premium capture.

NEW MARKET SPOTLIGHT

SWEDEN

Sweden joined the State of Digital Trust study for the first time in 2026. With no prior-year data, what follows is a baseline rather than a trend, a first read of a market we'll be tracking from here on.

Sweden enters the dataset with some of the strongest trust fundamentals of any market, and some of the sharpest comprehension gaps.

Three numbers that stand out in Sweden's first year:

- 69% trust banking with their data, the highest banking trust figure of any market in the study.
- 56% don't have a good understanding of how their data is collected and used, the highest comprehension gap of any market, ten points above the global average.
- 60% find AI personalization intrusive, the lowest concern figure in the study, eleven points below global.

The pattern is unusual. Swedish consumers show strong protective behavior, 85% have taken data protection action in the past six months and 79% would stop using a service over data misuse, but operate with less systemic understanding of what's actually happening with their data than any other market. They're cautious about AI access (Swedish and Dutch consumers are the least trusting of AI assistants having access to personal data and documents), but less likely than other Europeans to describe AI personalization as intrusive, possibly because their default position is to assume the worst already.

What this means commercially: Sweden is a market where the comprehension gap is the single clearest opportunity. Brands willing to explain, clearly, in plain language, in a way that closes the gap rather than performing transparency, will likely find Swedish consumers more receptive than the surface data suggests.

First year in the study. Strong instincts, limited framework. Over half (56%) don't understand how their data is collected, the highest knowledge gap of any market. The brand that closes that gap wins a market that takes trust seriously and acts on it.

Sweden will have year-on-year tracking from 2027 onwards.

CHAPTER 6

From insight to action

The data tells you where consumer trust is breaking down. This chapter tells you what to do about it.

Start with the diagnostic below. It will tell you where your organization sits today — and which part of the framework that follows is most relevant to you.

Where do you sit today?

Most organizations are in one of three places.

Tier 1

You have a consent problem you haven't measured. You have a banner. You have a policy. You're broadly compliant. But you don't know your current consent rate, you haven't benchmarked it, and you don't know whether your banner is driving acceptance or rejection. The 2026 data point that should concern you: 48% of consumers click "accept all" less often than three years ago. If you're not measuring, your consent rate is almost certainly degrading without your visibility.

Tier 2

Your consent experience isn't earning trust. You measure your consent rate. You've fixed the obvious dark patterns. But the broader experience — preferences, AI disclosures, data subject requests — feels disconnected. The 2026 data point that should concern you: privacy-aware consumers are nearly three times more comfortable with personalization online than privacy-unaware ones. The audience most willing to engage commercially is the one most attuned to inconsistency.

Tier 3

Your consent is solved. Your trust strategy isn't built yet. Your consent experience is strong. Your data flows are governed. But you don't yet have a strategy for the next category of risk: AI agents acting on your customers' behalf, requesting access to email, calendars, financial data, and customer records. The 2026 data point that should concern you: only 8% of consumers are fully comfortable with AI assistants accessing their data without conditions.

The T.R.U.S.T. Framework

Five steps, in the order that works. Each builds on the one before it. Most organizations get stuck by jumping to a later step before the earlier ones are solid — or by treating them as parallel workstreams when the sequence is the point.

1. Translate — get the consent moment right

Start at the banner. The most effective consent experiences feel like an extension of the brand: clear, considered, and written for the person reading it. Until the consent moment itself works, nothing further in this framework will. Tier 1 starts here.

2. Remove — strip out what's undermining the choice

Audit what's getting in the way of an honest answer. Give equal weight to accept, decline, and customize. Make controls reachable in two clicks. Dark patterns produce opt-ins, not consent — and the cost shows up in churn, deletion requests, and legal exposure rather than your consent rate. Tier 1 completes here before moving on.

3. Unify — extend the standard everywhere consent matters

The banner, the preference center, the DSAR tool, the AI disclosure — they all need to feel like the same brand making the same promise. Inconsistency signals privacy as an afterthought. Consistency signals intentionality. Tier 2 starts here.

4. Secure — extend governance into AI and third-party data flows

Map where consent signals go. Ensure AI tools don't become shadow data processors outside the visibility of users or internal governance teams. With agentic AI now acting on users' behalf without a visible consent moment, governance infrastructure must be in place before deployment, not retrofitted afterward. Tier 2 completes here. Tier 3 starts here.

5. Track — measure trust as a commercial signal, not just an opt-in metric

Opt-in rate alone tells you nothing meaningful. The numbers that matter are what happens after: retention, engagement, DSAR volume, churn, complaint rates. Clean, consented data is more than a compliance output — it's a performance input. Coerced opt-ins degrade the ad platform signals that your campaigns depend on. Tier 3 focus.

The infrastructure behind the framework

Each step of T.R.U.S.T. requires infrastructure. Below is how the Usercentrics platform maps to each.

STEP	WHAT IT REQUIRES	USERCENTRICS
1. Translate	Consent experiences across web, app, and CTV, designed to communicate clearly	Usercentrics CMP
2. Remove	Dark pattern detection and consent rate benchmarking	Cookiebot by Usercentrics
3. Unify	A single branded experience for consent, preferences, and DSAR	Usercentrics Preference Manager
4. Secure	AI connectivity and access governance	MCP Manager by Usercentrics
5. Track	Clean consented signals flowing to ad platforms and attribution	Server-side tagging — Usercentrics sGTM and Meta Signals Gateway

Where to start

Five questions.
Answer honestly — they'll confirm which tier you're in and where to focus first.

1. **Do you know your current consent rate and how it compares to your industry median?**
2. **When did someone last audit your consent banner for dark patterns?**
3. **Do your consent banner, preference centre, and DSAR flow feel like the same brand?**
4. **Do you have a governance plan for how AI agents access customer data?**
5. **Do you measure how consent quality affects retention, churn, or ad performance?**

Mostly no, you're in Tier 1. Mixed, Tier 2. Mostly yes, you're in Tier 3 and the next frontier is AI agent governance.

Wherever you sit, the next move is the same: a clear-eyed look at the gap between where you are today and what the consumer data in this report says will earn the premium.

We'd be glad to help you do that.

› **Talk to a Usercentrics specialist** about where you sit today and what closing the gap looks like for your business.

› **See how leading brands are doing it**, explore case studies from organizations turning privacy infrastructure into measurable performance.

The brands consumers will pay more to trust in 2027 are deciding what kind of brand they are right now.

About the research:
11,000 consumers. 7 markets. Conducted by Sapio Research in March 2026. Accurate to ±0.9% at 95% confidence.
Sweden is a new market this year; YoY comparisons exclude Sweden.

About Usercentrics:

Usercentrics is a leading data privacy technology company that helps businesses collect, manage, and activate consented data with confidence. Trusted by 2.4 million websites and apps across 195 countries, the company processes more than 8.8 billion user consents every month. Through its platform, spanning consent management, server-side tagging, and AI data governance, Usercentrics gives businesses the compliance infrastructure to grow, innovate, and operate responsibly in an AI-first world.

Learn more at usercentrics.com.

USERCENTRICS × Cookiebot by Usercentrics