

Ethics and Responsible Practice in Digital Marketing



- GDPR fines now exceed €7.1 billion cumulative as of January 2026



- FTC Operation AI Comply: 14 actions, roughly \$51 million recovered



- Scope: data and consent, targeting fairness, content disclosure, AI governance



- The cost is operational — penalties, platform actions, lost trust, lost revenue



- Objectives: spot the four legal surfaces, apply one decision framework, build a 30/60/90 plan



Why Trust Is Now a Growth Metric



52% will pay ~7% more to brands that handle data transparently



47% took a revenue-consequential action over data concerns in six months



Consumers trusting AI less than humans with data:
48% (2025) → **52%** (2026)



60% are uncomfortable with their data training AI models at all




46% still don't understand how their data is collected and used



WILLINGNESS TO PAY A PREMIUM




CONFIDENCE IN AI HANDLING DATA



UNDERSTANDING OF HOW DATA IS USED



The 2026 Regulatory and Platform Landscape



-  • EU AI Act Article 50 transparency duties live August 2, 2026; machine-readable marking follows December 2, 2026
-  • Penalties up to €15M or 3% of worldwide turnover; Article 5 bans up to €35M or 7%
-  • High-risk Annex III deferred to December 2027 by Regulation (EU) 2026/1744, but manipulation bans apply now
-  • US state enforcement leads: CPPA and California AG settlements — Honda, Disney, Ford, Healthline, Sling TV, LocateSmarter
-  • Platform and industry policy hardening: Meta's multimodal HEC classifier; IAB, NAI, ICC, ANA frameworks



Core Ethical Principles in Practice



Transparency: disclose data use, sponsorship, AI involvement, and material connections — targeted, not universal labeling



Honesty and accuracy: no misleading claims, fabricated social proof, or unsourced AI statistics



Respect for autonomy: meaningful consent, symmetric accept/reject friction, easy withdrawal



Fairness and non-discrimination: no exclusionary targeting, sensitive-category proxies, or exploitation of vulnerable audiences



Accountability: named owner, documented review decision, and a record you could hand to a regulator

Dark Patterns, Deceptive Design and Consent Validity



DSA Article 25 bans designs that deceive, manipulate, or impair free choice



CPPA: pre-ticked boxes and deceptive interfaces do not equal valid consent



Four audit targets: asymmetric accept/reject, pre-ticked boxes, excess fields, buried disclosures



Honda settlement: Accept All vs. Manage Preferences split is a cited dark pattern



SpiceJet fined for pre-ticked loyalty enrollment and misleading consent wording



Fix by design spec: equal buttons, no pre-checks, one-click withdrawal



DECEPTIVE DESIGN

Pre-ticked and Asymmetric

We use your data to personalize content and ads.



I agree to the use of my data for the purposes above.

By clicking “Accept All”, you agree to our Data Policy and Terms.

Accept All

Manage Preferences

You can withdraw consent through account settings.



CORRECTED DESIGN

Balanced and Transparent

We use your data to personalize content and ads.



I agree to the use of my data for the purposes above.

By clicking “Accept” or “Decline”, you agree to our Data Policy and Terms.

Accept

Decline




You can withdraw consent at any time with one click.

[Withdraw Consent](#)

Data, Consent and Tracking That Holds Up

- Decide the lawful basis per purpose, channel and recipient first
- Consent must travel to tag managers, SDKs, ad platforms, CRM, warehouse
- Failure pattern: tags firing before consent, opt-outs scoped to one device
- Server-side consent tied to a first-party identifier
- Per-purpose granularity plus timestamped logs with policy version

THE CONSENT SIGNAL JOURNEY



SERVER-SIDE CONSENT TIED TO A FIRST-PARTY IDENTIFIER

Captured once, stored server-side, and applied consistently across all systems and devices.

CONSENT INTERFACE CLOSE-UP: CHOICES AND GRANULARITY

Banner Choice

We use your data to improve your experience and for relevant advertising.

Decline All

Accept All

Granular Controls (Per Purpose)

Essential



Analytics



Marketing



Always Recorded

- ✓ Timestamped log
- ✓ Policy version
- ✓ Purpose-level decisions
- ✓ Linked to first-party identifier



Targeting, Personalization and Fairness



- **Discriminatory delivery can be unintentional:** metro-wide ad sets can optimize toward wealthier neighborhoods — a Fair Housing Act risk even without intent.



- **Regulated verticals:** housing, employment, financial products. Meta requires a Special Ad Category; Google strips gender, age, parental status, marital status, ZIP.



- **Meta's 2026 multimodal classifier** reads image, copy and audio — declaration gaps can become Evasion integrity violations.



- **Controls:** paired delivery tests, documented audience definitions, and a 'never target / never infer' list.




- **Inferring financial distress or emotion** to time offers is an Article 5 red flag, not an optimization win.

Declared Audience
Metro-wide footprint



Delivered Footprint
Optimized toward wealthier neighborhoods



Regulated Verticals

- ✓ Housing
- ✓ Employment
- ✓ Financial products

Meta requires a Special Ad Category; Google strips gender, age, parental status, marital status, ZIP.



STOP

Do not proceed.



FAIRNESS CHECKPOINT

Go / No-Go Threshold

PROCEED

Document, monitor, and review.

Content Claims, Influencer Disclosure and Affiliate Integrity



- FTC standard: clear, conspicuous, in the same medium and language



- Platform tags supplement but never replace plain-words disclosure: "ad," "sponsored"



- Material connection goes beyond cash: gifts, affiliate commissions, trips, employment



- September 2026 FTC sweep: 14 publishers, \$28M+ in penalties



- Undisclosed AI-generated "expert review" drew a \$4.1M penalty




- Phia cookie-stuffing case led to an Impact.com suspension




AI in Marketing: Four Surfaces That Actually Apply

- Most marketing AI is limited-risk, not Annex III high-risk under the AI Act
- Four surfaces bite: Article 5 manipulation, Article 50 transparency, GDPR profiling, GPAI duties
- Chatbots must disclose they are AI at first interaction — not buried in terms
- Label deepfakes and public-interest text; routine ad copy usually exempt
- Cox Media Group: \$930K penalty and 20-year order for AI capability claims
- Stacked regimes: GDPR Article 22, DSA minors rules, FTC Section 5 AI-washing



Governance, Review Workflows and Team Accountability

- Only 18% of marketing teams formally review AI-generated content
- Three-tier risk model: low-risk logs, standard reviews, high-risk legal sign-off
- Operating cadence plus a 24–48 hour rapid-review lane
- Audit trail: draft, AI findings, reviewer decision, published asset
- Marketing owns the brand layer; agentic AI fails on undocumented brand rules
- Tier 3 covers personal data, health/financial claims, synthetic depictions
- No formal path means teams route around governance under deadline pressure



AUDIT TRAIL



DRAFT

Initial content or asset draft.



AI FINDINGS

AI insights, flags, and notes.



REVIEWER DECISION

Human reviewer outcome and rationale.



PUBLISHED ASSET

Final approved asset published.

Case Scenarios: Where Marketing Teams Actually Get It Wrong



- **Health pixels:**
FTC v. Hims & Hers
(July 2026) —
condition data sent
to ad platforms



- **AI "slop" endorsements:**
fake doctors pushing
recalled supplements
on TikTok Shop



- **Agentic affiliate interference:**
Phia overrode rival
referral codes,
suspended by
Impact.com




- **Regulator on autopilot:**
\$930,000 fine plus
20-year order for
unproven AI claims



Discussion: do your pixels match your privacy promise?




A Decision Framework Under Performance Pressure





Action Plan: Building Responsible Practice Into Daily Work



PersonWise.

PersonWise • AI digital-human course platform



Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/41def7b2/public