The AI Series

Governance and guardrails for regulated industries embracing AI-driven marketing

ACXIOM

Governance and guardrails for regulated industries embracing AI-driven marketing

From perceived impossibility to growing opportunity

Data privacy laws lead many brands in regulated industries like healthcare and financial services to believe AI-driven marketing is impossible. Brands may view complying with many laws and regulations while feeding sensitive data into an AI engine as a tricky challenge. For example, the Health Insurance Portability and Accountability Act (HIPAA) covers "protected health information," while the Gramm-Leach-Bliley Act (GLBA) protects non-public personal financial information, and the Fair Credit Reporting Act (FCRA) regulates the use of "consumer reports."

But it is possible for brands in regulated industries to use AI in an ethical, privacy-conscious way to create more personal, engaging experiences. It's just about putting the right governance guardrails in place.

Let's start by looking at four common challenges brands face when they first start exploring AI.

Challenge 1:
Overcoming AI misconceptions

Before regulated brands even consider data privacy and governance concerns, several myths dissuade them from embracing AI-driven marketing's transformative opportunities.

"AI is inaccurate": This misconception can stem from companies sometimes employing an AI model for issues the technology is not well-suited to solve. But fundamentally, the accuracy of AI comes down to how it's trained. If AI is trained on incomplete or inaccurate data, it will produce inaccurate outputs. AI also needs to understand business and market context to be able to accurately interpret the nuances of economic trends.

"All AI is the same": AI inaccuracies and hallucinations can also occur when brands use one AI solution for all their marketing problems. Brands need a diverse AI ecosystem so they can pick the right solution for the right situation, such as using machine learning for audience segmentation and generative AI for customer communications. Another example is when brands build ads on platforms like Google and Meta, they need an AI model trained on proprietary data for each platform.

"AI is just for tech companies and engineers": Many AI marketing tools are specifically designed with user-friendly interfaces. This empowers non-technical people to perform technical tasks with AI, meaning brands are at a disadvantage if they're not upskilling employees in AI technology.

Once regulated brands get past these initial misconceptions, they can then start focusing on how to implement an effective data privacy and governance strategy for AI-driven marketing.

Challenge 2:
Addressing consumer consent

When it comes to data privacy risks, AI supercharges concerns that already exist around consent and data misuse.

With AI-powered marketing, even if customers are happy to express affirmative consent at every turn, it's difficult to provide total transparency regarding all the AI use cases. Brands should prioritize ethical, responsible use of customer data and be able to demonstrate their responsible use, if necessary.

> "Hyper-personalization demands a higher standard of data ethics. Instead of relying on consumers to navigate complex consent requirements, brands should proactively ensure responsible data use and transparent practices."
>
> — **Jordan Abbott**, Chief Privacy Officer, Acxiom

© 2025 Acxiom LLC. All rights reserved. Acxiom is a registered trademark of Acxiom LLC.

Challenge 3:
Guarding against bias

Data privacy isn't the only risk with AI; there's also the much-talked-about issue of bias.

If some bias didn't exist, companies wouldn't be able to segment audiences. But discrimination based on protected classes runs afoul of Fair Lending, Reg B, and other laws. The challenge with marketing is that if AI is trained with biased data, that will carry through into the output, which could have a disparate impact on underrepresented or disadvantaged populations. This is particularly important for regulated industries, for example, when issuing credit or determining the specifics of insurance policies and interest rates.

So what can regulated brands do to guard against bias? There are several things.

Model monitoring.

Many enterprise models have built-in safety settings that allow companies to control the freedom of the model to interact on controversial topics like religion and sexuality.

Companies can also use toxicity monitoring, which semantically searches inputs and outputs for whatever the company defines as toxic. Users can then revisit toxic outputs, explore what led to them, and restrict the model from generating unsafe outputs. The safety settings don't just restrict the model from toxic outputs; they also prevent unsafe inputs by telling users the AI can't talk about their suggested topic.

Drift metrics.

Drift metrics allow companies to measure divergence from true representation when building audiences by selecting attributes. When building audiences for advertising, brands might, for example, create a segment of people of a certain income group and select the relevant attributes. But attributes, like location, could be lost as a result of that selection.

Challenge 4:
Navigating evolving risks

Regulated brands are particularly vulnerable to AI security concerns, such as cyber attacks, because of the highly sensitive nature of people's data. AI models can also be manipulated through malicious inputs. This makes it critical to have a robust information security framework to protect customer data by anticipating and responding to AI cybersecurity threats and vulnerabilities as they emerge.

At the same time, the regulatory landscape for AI is also rapidly evolving, often in unpredictable patterns. This is a challenge for providing oversight of boards of directors in financial institutions, where new frameworks, such as the EU AI Act, are introducing mandates for transparency and human supervision that require continuous adaptation.

Another reason why brands need to pay close attention to data regulations is to minimize concerns about leaks of proprietary data. That's why regulated brands are often cautious about using third-party AI solutions, like chatbots. To comply with stringent regulations like GDPR and U.S. state privacy laws like the California Consumer Privacy Act (CCPA), brands must carefully assess how AI models ingest and retain personal data.

Best practices and governance advice for regulated brands considering AI-driven marketing

Whether brands use AI to communicate directly with customers and prospects or to generate content that represents the brand, in effect, they're deploying the technology to speak on their behalf. With AI acting as an extension or ambassador of the brand, and a wide range of complex regulations to comply with, data governance must be airtight.

Here are 10 AI governance measures brands can implement:

1. Be transparent with people.

Companies need to clearly articulate their intentions regarding AI, execute those intentions, and demonstrate they've done so. In practical terms, brands must be as clear as possible about how data will be used and what AI will be used for, such as making predictions and inferences or determining which products to offer. And they should give people as much control as possible, including opting out of their personal data being used for training future iterations of the model.

2. Extend current compliance programs and obligations.

Brands should embrace privacy and security by design. In other words, they should adopt ethical AI use at every phase of the AI management lifecycle. Marketing teams must collaborate closely with legal, compliance, and risk teams from the initial design phase of AI applications.

Once brands have determined whether AI accurately addresses the problems they're trying to solve, they should establish the data they need to directly answer their questions. If brands can't achieve business objectives without using personal data, then they must properly classify it so they are always clear how they're using it.

Brands' AI governance frameworks should have clearly defined roles, reinforced by training programs to close employees' skills gaps in using AI responsibly.

3. Practice data minimization.

The same principles apply to AI that have held for years in maintaining data privacy. Limiting data's purpose – using the least amount of data possible to achieve business objectives – is key. Brands should only collect what they need and use it only for the purposes they disclose.

4. Maintain independent oversight.

Brands should always keep a human in the loop who can pull the plug on AI outputs if necessary. On the front end, brands need AI conformity to know the model is doing what it was built to do. Then, on the back end, brands need bias and fairness validation and training, as well as continuous monitoring with ongoing audits to ensure they understand and fix any unexpected outputs.

5. Check AI providers' policies.

It's important to ask AI providers about the data they used to train the models and whether the company provides warranties against the unauthorized use of the brand's outputs. This allows brands to maintain ownership over their licensed AI, so no outputs leave the model or are used to train future iterations.

The knowledge storage layer is equally important. When brands permit AI to train models using their data, the provider collects and stores it in a database. Brands should ask any company they license AI from how that works with their solution and if it's possible to delete all prompts and interactions from the database, if needed.

6. Facilitate data removal.

If there's one thing brands have learned about privacy, it's that the ability to audit a system and remove people's data is imperative. In AI's case, this often means removing conversations and personally identifiable information (PII) from interactions.

People are more likely to enter PII into an AI interface than into other types of systems, as they feel like they're having a conversation with someone. But it's becoming common to use an AI filter agent that strips this information out of prompts before passing inputs on to the model.

7. Enable reason tracing.

The "black-box" complexity of an AI model makes it difficult to understand the decisions it makes to arrive at an outcome. But thanks to recent AI developments, companies now use AI systems that combine multiple AI agents with specific purposes or tasks. This makes AI explainability theoretically easier, as companies are better able to trace the decisions, parameters, and features behind a given output to particular AI agents.

8. Balance generalization and specificity.

As AI adoption increases, brands are starting to use more specific models. On the one hand, these models have lower inference costs compared to using massive, generalized models for specialist tasks. On the other hand, brands won't be able to answer every question effectively if they use a smaller, more specifically trained model. They must be mindful of the balancing act between specialization and generalization.

9. Invest in data quality.

By prioritizing data quality management, brands ensure AI solutions are supported by clean, validated, and standardized customer data. For companies like banks, this usually requires addressing the data fragmentation and complexity that's common in internal systems. Establishing processes for regular audits and updates ensures data is consistently accurate.

10. Source data ethically.

Regulated brands must build ethical rigor into data sourcing, model training, and deployment. Diverse and inclusive data prevents the amplification of inequitable social bias. Brands should use only properly

© 2025 Acxiom LLC. All rights reserved. Acxiom is a registered trademark of Acxiom LLC.

credentialed and licensed datasets and avoid publicly scraped sources to mitigate the risks of AI outputs incorporating copyrighted materials.

Where regulated brands should start with AI

Internal operations is a great place to find low-risk use cases to start testing AI, demonstrate immediate value, and incentivize investment. Think of the relatively trivial tasks people first used AI for, such as writing emails. Large language models (LLMs) are highly proficient in these tasks, while users can still feel secure that they have to hit send before the output is released.

Anything that requires writing text, drafting documents, and turning unstructured enterprise data into something more meaningful is low-hanging fruit. What's more, the data and documentation used in these initial AI use cases are already stored using privacy-based rules, regulations, and governance. And these use cases integrate seamlessly in legacy infrastructure and existing workflows.

So before brands dive into higher-risk use cases like dynamic pricing, credit-related decision-making, or fraud and regulatory change management, here are some recommended low-risk, high-impact use cases tailored to regulated industries:

Product and service recommendations: AI can give brands 360-degree views of customers, helping them recommend hyper-personalized products and services. Brands can also use these detailed customer profiles to predict future needs, such as recommending investment strategies based on AI analyzing investment holdings, risk tolerance, and market trends. In the healthcare sector, examples may include recommending a preventive health strategy based on risk markers.

Customer support and query resolution: Brands can drive unparalleled customer service efficiency with intelligent chatbots and virtual assistants providing 24/7 support. This is particularly important in regulated industries where people need prompt responses about issues related to their health and financial situation. In addition to handling customer interactions around the clock, AI can analyze enquiries to derive insights that lead to truly personalized resolutions.

Proactive customer monitoring: Predictive analytics have long been a mainstay for financial services, but they will only get more effective with AI, enabling more complex pattern recognition to detect early signs of potential issues, including customer churn. Brands can then improve customer loyalty and retention by offering incentives and advice for managing credit or healthcare payment plans.

Automating data entry, reconciliation, and reporting: AI can unlock substantial time savings and reduce human error by automating repetitive tasks like creating financial reports, triggering internal notifications regarding suspicious behavior, and collating patient records.

> "Hyper-personalization is perhaps the most transformative application of AI in the financial services sector. It means that financial institutions can move from more generic forms of marketing and deliver individualized experiences that we know will foster deeper customer relationships and cultivate more enduring loyalty."
>
> — **Graham Wilkinson**, EVP, Chief Innovation Officer, Global Head of Artificial Intelligence, Acxiom

How we help

Acxiom has championed ethical data use and led the industry in data privacy since our company's founding in 1969. We are trusted by leading financial institutions and insurance providers.

We appointed the world's first chief privacy officer and have had a dedicated privacy team ever since. We've long adhered to the ethical data use framework championed by the Information Accountability Foundation. A member of our privacy team was part of the first International Association of Privacy Professionals' AI Governance Professional class, with certifications to boot. This expertise empowers regulated brands to keep pace with the ever-expanding compliance landscape.

As part of our privacy program, we offered data subject rights to consumers nationwide, well before the law required it, providing data access and deletion in the U.S. since 2022. Our Snowflake partnership is one of the latest examples of our commitment to supporting people's data sovereignty. Our clients can now run AI models in the same place they're storing data by making their data sovereign to their Snowflake solution.

Our data sources are also robustly credentialed to ensure we understand how the data is collected and whether it can be used for our clients' intended purposes. In fact, we perform more than 500 privacy impact assessments a year to identify potential harm from data use and implement controls to eliminate or mitigate risk.

We aspire to go beyond the minimum legal requirements and set a higher bar for ethical data use and demonstrable accountability. Acxiom empowers brands to leverage AI responsibly, maintaining trust and upholding ethical standards. We unlock the full potential of AI-driven marketing. Specializing in data orchestration for the AI era, we unify, connect, and prepare data, accelerating brands' success with innovative solutions from our expertly curated AI ecosystem.

About Acxiom

Acxiom puts data to work. We solve complex challenges for the world's leading brands and agencies by unifying, connecting, and preparing data for AI-driven marketing and decision-making, maximizing technology investments. As leaders in data ethics and governance, Acxiom brings a privacy-first approach to serving clients globally, with locations in the U.S., UK, Germany, China, Poland, and Mexico.

Connect with Acxiom on LinkedIn and discover more at Acxiom.com.

About IPG's Interact

An AI-powered, end-to-end marketing platform that enables brands to deliver hyper-personalization at scale across every touchpoint, rooted in Acxiom's connected identity and data foundation, activated wherever it is needed.

© 2025 Acxiom LLC. All rights reserved. Acxiom is a registered trademark of Acxiom LLC.

ACXIOM

For more information, please
contact Acxiom at **acxiom.com**
or email **info@acxiom.com**.