Deloitte.
Together makes progress

When innovation outpaces protection: The new face of UDAAP risk

May 2026

Center *for*
Regulatory Strategy
US

Product innovation, digitization, and artificial intelligence (AI) have fundamentally reshaped banking. Customer financial journeys are increasingly digital, and banks are competing for customer attention through more complex products, highly personalized offers, and AI-driven marketing, lending decisions, and customer services. These innovations and advances have transformed traditional banking and increased value for customers. These advancements have also introduced new avenues for customer harm that traditional Unfair, Deceptive, or Abusive Acts or Practices (UDAAP) risk frameworks may not catch and, if left unaddressed, quietly erode customer trust. When innovation outpaces governance, the very features designed to improve customer experience can become the sources of UDAAP risk. **As banks race to modernize and launch new offerings, an important question emerges: Are the bank's processes, controls, and customer experience designs keeping up?**

Meeting this challenge requires a fundamental shift in how institutions approach UDAAP risk—from a compliance exercise focused on avoiding regulatory scrutiny to a proactive, AI-enabled discipline centered on sustaining customer trust. Customers expect their bank to be transparent, treat them fairly, and deliver experiences that work in real life, not ones that simply "check the box" for compliance. This shift requires institutions to take a proactive approach to consumer protection, monitoring trends, deploying AI to detect complaint patterns before they escalate, and supporting customers during key moments in their financial journeys.

While these efforts may be viewed by some as a compliance exercise, getting products and customer journeys right can create significant business value. A more customer-centered experience can reduce complaints and remediation costs, but more importantly, it strengthens loyalty and protects brand reputation—outcomes that ultimately support long-term growth and profitability.

The pages that follow explore how UDAAP risk is evolving, from the shifting nature of today's risk environment to the organizational blind spots that allow it to persist, to the governance gap that AI has introduced. The following customer journey walks through these principles in plain terms, with a deeper analysis and implications to follow.

UDAAP at a glance

Consumer protection remains a core expectation for US financial institutions, requiring banks to proactively identify and prevent potential customer harm.

Unfair, Deceptive, or Abusive Acts or Practices (UDAAP)[1] encompasses a set of US consumer protection regulations that require financial services companies to market, sell, and service products in ways that are truthful, clear, and do not create harm to consumers. UDAAP is enforced by federal and state regulators, including the Consumer Financial Protection Bureau (CFPB), Office of the Comptroller of the Currency (OCC), and the Federal Reserve.

In practice, UDAAP extends beyond regulatory compliance. Regulators expect institutions to embed consumer protection principles throughout the product life cycle—from product design and marketing through servicing and account closure. Specifically, banks should avoid practices that:

- Cause harm that customers cannot reasonably avoid (unfair);
- Mislead customers through inaccurate or incomplete information (deceptive); and
- Take advantage of a customer's lack of understanding or bargaining power (abusive).

At its core, UDAAP asks a simple question: Are institutions protecting customers from harm?

To understand where UDAAP risk hides, consider a typical customer journey and how seemingly small experience gaps can create UDAAP risk. Meet Catie Client, a longtime customer of $Best-Bank who is shopping for a new credit card. As we follow Catie's journey, we will highlight the key moments where $Best-Bank should pay close attention to spotting and mitigating potential unfair, deceptive, or abusive practices.

Customer journey scenario	UDAAP risk	$Best-Bank
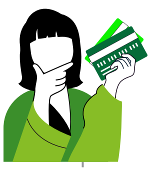 **Product design: The comparison trap** Catie is comparing credit card options at $Best-Bank. She wants a premium card that provides competitive benefits and rewards her highest spend category, dining.	$Best-Bank offers multiple cards with complex reward tiers, eligibility requirements, and restrictions, making it difficult for Catie to compare card options. This complexity is further compounded when AI-driven recommendation engines are trained in business profitability rather than customer suitability. In practice, this means Catie may be steered toward a higher-margin card that does not actually reward her top spend category.	$Best-Bank should make the cards easy to compare by offering a standard comparison framework with the same fields in the same sequence for all products. For example: annual fee, earn rate, caps, and exclusions. AI-driven product recommendation tools should also be designed and monitored to ensure they surface genuinely suitable options rather than defaulting to higher-margin products.
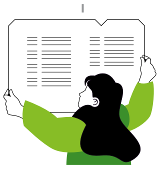 **Marketing: The definition dilemma** Catie is drawn to the Diamond Dining card, which advertises 5% cash back on dining transactions. She visits $Best-Bank's website, where the claim appears in large, bold, animated text. The page does not clearly define what qualifies as dining.[2] Catie reasonably assumes all types of dining are included, like ordering through third-party food delivery apps.	$Best-Bank's digital marketing is designed to capture attention and build interest quickly. If AI is used to personalize benefit claims, such as prominently displaying Catie's dining offer because her browsing history flagged her as a frequent diner, this becomes a targeted, misleading experience. If that 5% cash back applies to only in-person restaurants and excludes delivery apps, Catie may choose a product she otherwise would not have chosen if provided clear and complete information.	$Best-Bank's material claims and limitations should appear clearly and prominently wherever the benefit claim appears, and not be relegated to fine print or hover text. If "dining" has a specific definition that excludes common customer behaviors, that limitation deserves equal prominence to the claim itself.[3] This standard applies equally to AI-generated content because personalization does not reduce the obligation for accurate disclosure; it raises it.
 **Application: The on-the-go experience** Catie now understands that food delivery will not earn 5% cash back. She is slightly disappointed, but she still plans to apply for the Diamond Dining card since she recently launched a food blog reviewing local restaurants and expects her in-person dining spend to increase. Catie decides to apply through the $Best-Bank app during an afternoon walk.	The mobile application includes optional add-ons that may be presented in a way that results in inadvertent enrollment, such as pre-checked boxes, confusing toggles, or affirmative language framed as default sections. As Catie quickly scrolls through, she does not realize she is being enrolled in a free 30-day credit monitoring trial that will automatically charge her card after the trial period. In this case, what feels like a seamless experience to Catie actually obscures meaningful consent.	$Best-Bank should not automatically enroll customers and should present optional add-ons as truly optional. When a customer actively selects an add-on, $Best-Bank should send immediate confirmation summarizing what was chosen, the post-trial cost, and cancellation instructions. $Best-Bank could also add a brief "confirm you understand" screen so customers can pause, review key points, and proceed with confidence.
 **Underwriting and credit decisioning: The preapproval assumption** Catie is optimistic because her preapproval indicated a 20% annual percentage rate (APR) and a $20,000 credit limit, terms she feels confident managing. When she reaches the decision page, large text flashes "You Are Approved!"	If the approval screen does not immediately and clearly display the actual credit terms, Catie has no indication that a change occurred, or why, and may assume she has received her preapproved offer. Accepting less favorable terms (also referred to as an "adverse action"), like a lower credit limit or higher APR, based on that misunderstanding could be considered misleading.	$Best-Bank should clearly state when the final offer differs from the expected terms. Where AI or automated models drive underwriting decisions, customers should receive a plain-language explanation of why their terms have changed. The decision screen should state clearly "Approved with changes," display the credit limit and APR prominently on that same screen, and include a simple side-by-side comparison of preapproved versus final terms.[4] Transparency is non-negotiable.

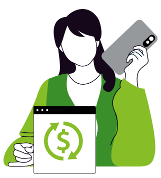

Customer journey scenario	UDAAP risk	$Best-Bank
Account onboarding: The autopay clarity gap Catie accepts her modified offer and is prompted to enroll in one of the recommended autopay options: Minimum, Statement balance, or Fixed amount.	Account onboarding flows are often optimized using AI to achieve completion rates, which can inadvertently steer customers toward lower-friction options rather than the option most aligned with their financial interest. If autopay definitions are hidden behind an icon, and the "Minimum" option is surfaced as the default option, Catie may select it assuming it will cover her monthly obligation and prevent late payments. However, if the definitions had been displayed clearly on the screen, she might have chosen the "Statement Balance" option to pay in full each month and avoid interest charges entirely.	Before confirming the selection, $Best-Bank should display a final review screen that clearly shows the selected autopay option with a plain language definition, the linked funding account, the pull date, and the date when autopay takes effect. AI-driven onboarding flows should be tested to ensure that default selections align with customer financial outcomes, not just optimize completion metrics.
Servicing: The messaging disconnect Catie dines at a new local restaurant and pays $62.79 with her new Diamond Dining card, but her statement shows $162.79. She calls $Best-Bank to dispute the charge, and the call center agent assures her they will investigate and issue temporary credit. Two weeks pass and the credit never appears, so Catie follows up through the in-app chat feature. The chat agent tells her $Best-Bank cannot act until she contacts the merchant first and, if that fails, she should submit an online dispute form. Regardless, she remains responsible for any interest and fees tied to the charge in the meantime.	Inconsistent guidance or mixed messaging across channels (e.g., phone, chat), including those powered by AI chatbots, could mislead Catie about the correct next steps or timelines, delaying resolution. In multichannel servicing environments, each handoff is a potential point of failure. As her payment due date approaches, she may feel compelled to pay the full amount to avoid interest and fees, causing financial harm.	$Best-Bank should train AI chatbots on consistent, accurate servicing and dispute procedures and include clear human escalation protocols for unresolved or complex servicing issues. After any customer interaction, whether human or AI-assisted, $Best-Bank should send an automated confirmation that documents: (1) the action taken, (2) the temporary credit amount and posting date, (3) any customer next steps, and (4) the amount due by payment date to avoid interest and fees. Introducing automation can help standardize messaging and ensure customer interactions are documented and managed according to the organization's policy.
Account closure: The failed closure After some time, Catie realizes that the card's benefits no longer align with her lifestyle and decides to close her credit card. She calls $Best-Bank, and the call center agent agrees to process the account closure. Believing the account is closed, Catie stops monitoring it, only to receive an alert a year later when an annual fee has been charged.	The agent may process the closure incorrectly, leaving Catie's account active without her knowledge. Incurring unexpected fees on an account Catie believes is closed may cause financial harm and a breach of trust that is difficult to recover from.	$Best-Bank should require agents to remain on the phone to confirm that the account closure has been successfully completed, not just submitted. $Best-Bank should also send an automated account closure confirmation to the customer, providing a clear record that the request was processed and the account is closed.

Across each stage of Catie's journey, $Best-Bank had the opportunity to build trust or erode it. Each touchpoint represents a decision point where small design choices and communication gaps created potential customer harm and significant operational, financial, and even regulatory exposure for the bank. The lesson is not about a single failure, but that in a world of digital channels, complex products, and algorithmic personalization, these gaps multiply faster than oversight can catch them.

Today's UDAAP risk is more dynamic

UDAAP risk has always existed in banking. What has changed is the speed, scale, and complexity at which risk emerges. Traditional frameworks were built to evaluate static artifacts such as disclosures, fee schedules, or advertisements. **Today's UDAAP risk lives in dynamic systems: an AI model adjusting credit offers in real time, a chatbot providing inconsistent guidance, and dynamic applications nudging customers toward certain products**.

Digital channels compress the time between a flawed design and widespread customer harm. Personalization means a single misleading claim can be algorithmically targeted to millions of customers simultaneously, each receiving a tailored version designed to maximize response. Periodic review cycles designed for a slower, more linear process may struggle to keep pace. **Managing UDAAP risk today involves continuous monitoring, embedded governance, and systems designed for the complexity banks are delivering**.

UDAAP risk hides in the handoffs

UDAAP risk is rarely owned by a single team; it emerges within handoffs across functions where intent can get lost and customer impact may go unnoticed. Consider what it took for Catie to receive her personalized offer: Product defines the value proposition. Marketing shapes the messaging. Technology builds the experience. Credit decisioning sets the final terms. At the same time, AI quietly shapes decisions across functions, at each step. When these teams operate independently, customers are caught in the middle, facing mixed messages and abrupt handoffs at every turn.

Cross-functional collaboration is not just an operational preference; it is how risk is managed in practice. Without clear ownership and accountability, fixing problems can turn into a game of "hot potato" where everyone agrees something should change, but responsibility gets passed around. Progress can stall, and customers may end up paying the price through unexpected fees, missed deadlines, and longer resolution times. When broken journeys persist, confidence in organizations often erodes, and with it, the loyalty and brand equity that is not easily restored.

The answer is not more layers of review, but smarter coverage across the full customer journey, with accountability and governance structures that match the complexity of modern banking.

AI introduces UDAAP risks that require additional governance[5]

AI is already reshaping how banks design products, make credit decisions, personalize marketing, and serve customers, transforming UDAAP risk in ways that traditional compliance frameworks were not built to manage. When an algorithm drives a lending decision or a personalization engine determines which offer a customer sees, the institution is still fully accountable for the outcome. Model drift can cause a system that was fair at deployment to produce systematically different outcomes over time, without visible warning signs. Training data can embed historical disparities that disadvantage certain customer segments in ways standard monitoring will not catch. When an institution cannot articulate in plain language why a specific customer received a specific outcome, it cannot defend that outcome to the customer or regulators.

What makes this particularly challenging is that AI risk does not live in one place. It is distributed across the organization, embedded in models built by technology, campaigns designed by marketing, decisioning logic owned by credit, and service interactions managed by operations. Yet, in most institutions, compliance is brought in after the fact, reviewing what has already been built rather than shaping what is being built. Closing that gap involves more than a process change. Compliance should have a seat at the table from the start, not as a final checkpoint, but as an active design partner.

That expectation should extend beyond the compliance function itself. Everyone in the organization, from product managers and data scientists to marketing leads and frontline agents, should understand that AI outputs carry the same consumer protection obligations as any other human-driven output. **The speed and scale of AI do not reduce accountability, they amplify it**.

Used well, AI can also be part of the solution, monitoring customer journeys at a scale no manual process can match, surfacing complaint patterns before they escalate, and testing whether automated decisions produce fair and consistent outcomes across customer segments. Realizing that potential involves treating AI governance and consumer compliance as the same discipline, not parallel efforts managed by different teams with competing priorities.

From awareness to action

Banks across the industry are navigating accelerating product complexity, expanding digital channels, and AI capabilities that are evolving faster than the governance frameworks designed to oversee them. Keeping pace requires more than incremental adjustments to existing programs. Institutions should consider fundamentally rethinking how consumer protection is designed, owned, and measured. **It should be embedded into the way products are built, AI is deployed, and customers are served, rather than layering it on at the end**. Institutions that get this right protect more than their compliance posture; they protect their customer relationships and brand reputation in an increasingly competitive market.

The gap between where most institutions are and where they need to be is rarely a question of intent. It is a question of design and whether the structures, processes, and practices governing customer experience are built for the environment banks now operate in. In a landscape where a single flawed AI model can reach millions of customers overnight, that gap is no longer just a compliance concern. It is a business risk. The good news is that designs can be changed.

The questions below are a practical starting point, not a scorecard, for the cross-functional dialogue that meaningful change requires.

Operating model

- Do both your business teams and compliance function maintain documented policies and procedures that reflect how your products, channels, and AI tools actually operate today?
- Does your organization provide role-specific UDAAP training that helps product, technology, marketing, and operations teams understand how consumer protection obligations apply to their day-to-day work, or is training limited to a general annual compliance requirement?

Risk assessment

- Are your UDAAP risk assessments keeping pace with your product and technology roadmap, or are they primarily reviewing what has already been built?
- Have you mapped your highest-risk customer journeys end to end, across every channel, to identify where experience gaps or harm are most likely to occur?

AI governance

- Can you explain, in plain language, why a specific customer received or did not receive a specific offer, decision, or servicing outcome in a way that would satisfy a regulator or the customer themselves?
- Do you have model risk management and consumer compliance working together, with shared visibility into where AI is driving customer-facing decisions?

Monitoring

- Are you monitoring your AI-driven outputs for patterns of disparate impact or fairness concerns on an ongoing basis, or primarily at the point of model deployment?
- Are you monitoring complaint trends, servicing escalations, and customer feedback across channels for early signals of UDAAP risk, or primarily responding after issues escalate?

The design choices, communication gaps, and process failures that create customer harm are often the result of systems and structures that were not built to keep pace with the complexity of institutional growth, customer demand, and competitive pressures. That is a core challenge facing banks today, and it is one that no single team, tool, or policy can solve in isolation.

The institutions that are best positioned to navigate this era successfully are those that stop treating consumer protection as a compliance deliverable and start treating it as an operating principle with UDAAP considerations embedded in how products are designed, how AI is governed, how teams collaborate, and how performance is measured. In a world where digital channels, complex products, and algorithmic personalization can move faster than traditional oversight can follow, that shift should not be viewed as optional. It is the foundation of sustainable, trustworthy banking.

When banks get this right, the benefits extend beyond compliance. Clearer experiences reduce complaints and remediation costs. Stronger trust deepens loyalty. A well-protected reputation becomes a competitive differentiator—the kind that takes decades to build and only one bad cycle to damage. Consumer protection becomes what it should be, not a constraint on innovation, but an accelerator for sustainable success.

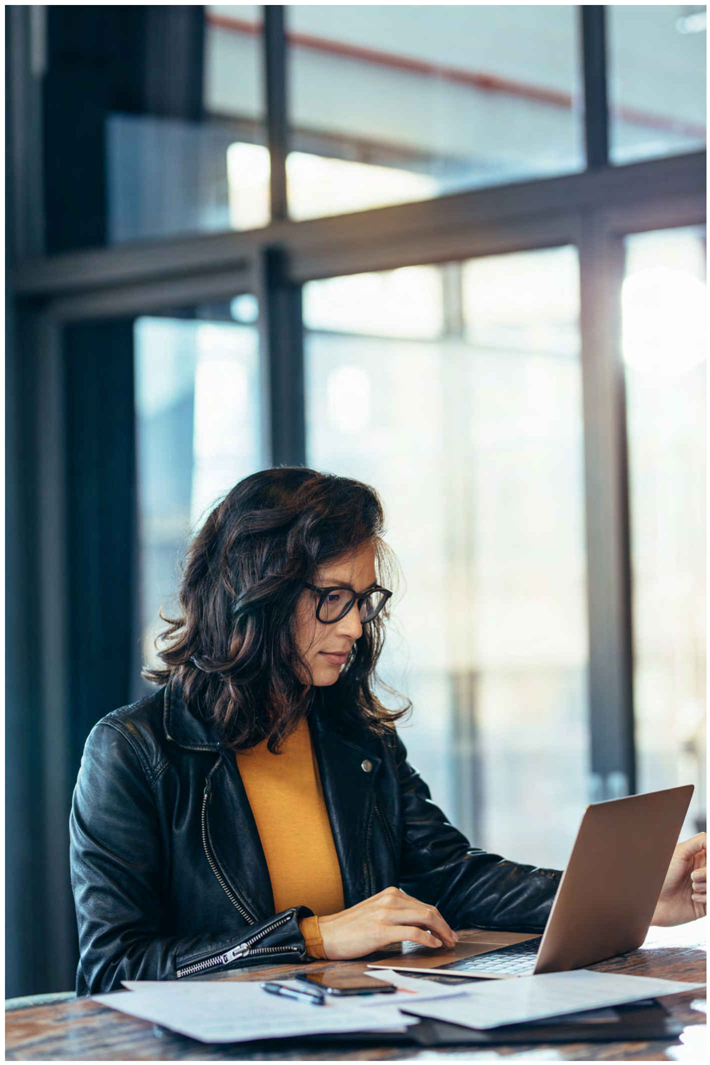

Regulatory insights to action

The Deloitte Center for Regulatory Strategy, US

DCRS helps financial services firms anticipate regulatory change and respond with confidence. We focus on four sectors: banking, capital markets, investment management, and insurance.

As global regulators adapt to shifting economic, geopolitical, and technological forces, they continue to reshape business models and operating frameworks. These shifts create both risks and opportunities for firms.

Our team combines private and public sector experience, supported by Deloitte's access to a global network and regional hubs in Asia Pacific and Europe, the Middle East, and Africa. With former regulators, industry specialists, and business advisers, we deliver practical insights to help organizations navigate complexity and stay ahead.

Additional insights can be found on our hub page here.

How can we help you?

Register for our monthly newsletter

Contacts

Mom Nguon
Managing Director
Deloitte & Touche LLP
mnguon@deloitte.com

Jonathan Valenti
Principal
Deloitte Consulting LLP
jvalenti@deloitte.com

Nick Ingham
Principal
Deloitte Consulting LLP
niingham@deloitte.com

Roxane Li
Senior Manager
Deloitte & Touche LLP
roxli@deloitte.com

Mackenzie Donnell
Senior Consultant
Deloitte & Touche LLP
mamott@deloitte.com

Shaun Nabil
Managing Director
Deloitte & Touche LLP
snabil@deloitte.com

Tom Nicolosi
Principal
Deloitte & Touche LLP
tnicolosi@deloitte.com

Paul Sanford
Independent Senior Advisor to
Deloitte & Touche LLP
pasanford@deloitte.com

Heather Decker
Manager
Deloitte & Touche LLP
hdecker@deloitte.com

Charlize Marie Melendez
Senior Consultant
Deloitte & Touche LLP
chmelendez@deloitte.com

Deloitte Center for Regulatory Strategy, US

Irena Gecas-McCarthy
FSI Director, Deloitte Center for Regulatory Strategy, US
Principal
Deloitte & Touche LLP
igecasmccarthy@deloitte.com

Lananh Nguyen
Co-Director, Deloitte Center for Regulatory Strategy, US
Managing Director
Deloitte Services LP
lananhnguyen@deloitte.com

Aaron Salerno
Manager
Deloitte Services LP
asalerno@deloitte.com

Kyle Cooke
Manager
Deloitte Services LP
kycooke@deloitte.com

End notes

1. 12 USC § 5531.
2. Consumer Financial Protection Bureau (CFPB), "Consumer Financial Protection Circular 2024-07," December 18, 2024.
3. 12 USC § 5532.
4. 12 USC § 5532(b)(2).
5. Department of the Treasury, *Artificial intelligence in financial services: Uses, opportunities, and risks of artificial intelligence in financial services*, December 2024.

Center *for*
Regulatory Strategy
US

About the Center
The Deloitte Center for Regulatory Strategy provides valuable insight to help organizations in the financial services industry keep abreast of emerging regulatory and compliance requirements, regulatory implementation leading practices, and other regulatory trends. Home to a team of experienced executives, former regulators, and Deloitte professionals with extensive experience solving complex regulatory issues, the Center exists to bring relevant information and specialized perspectives to our clients through a range of media, including thought leadership, research, forums, webcasts, and events.