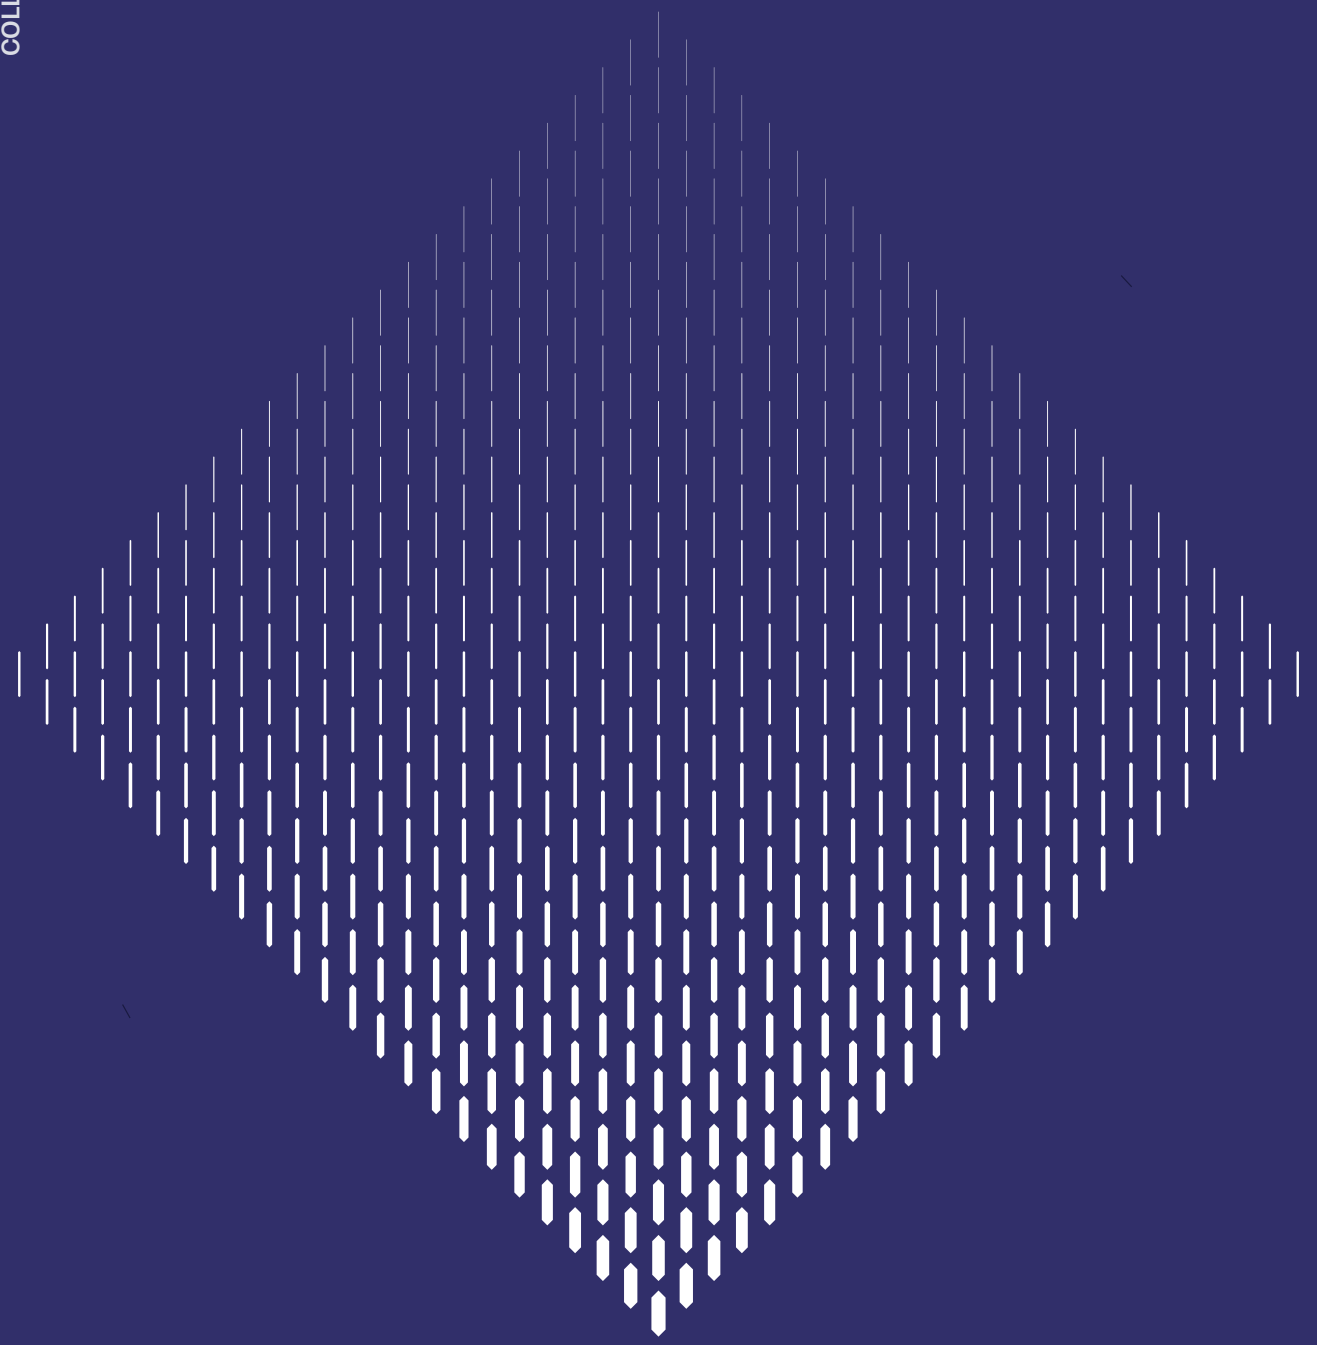
zango
COLLECTED ESSAYS

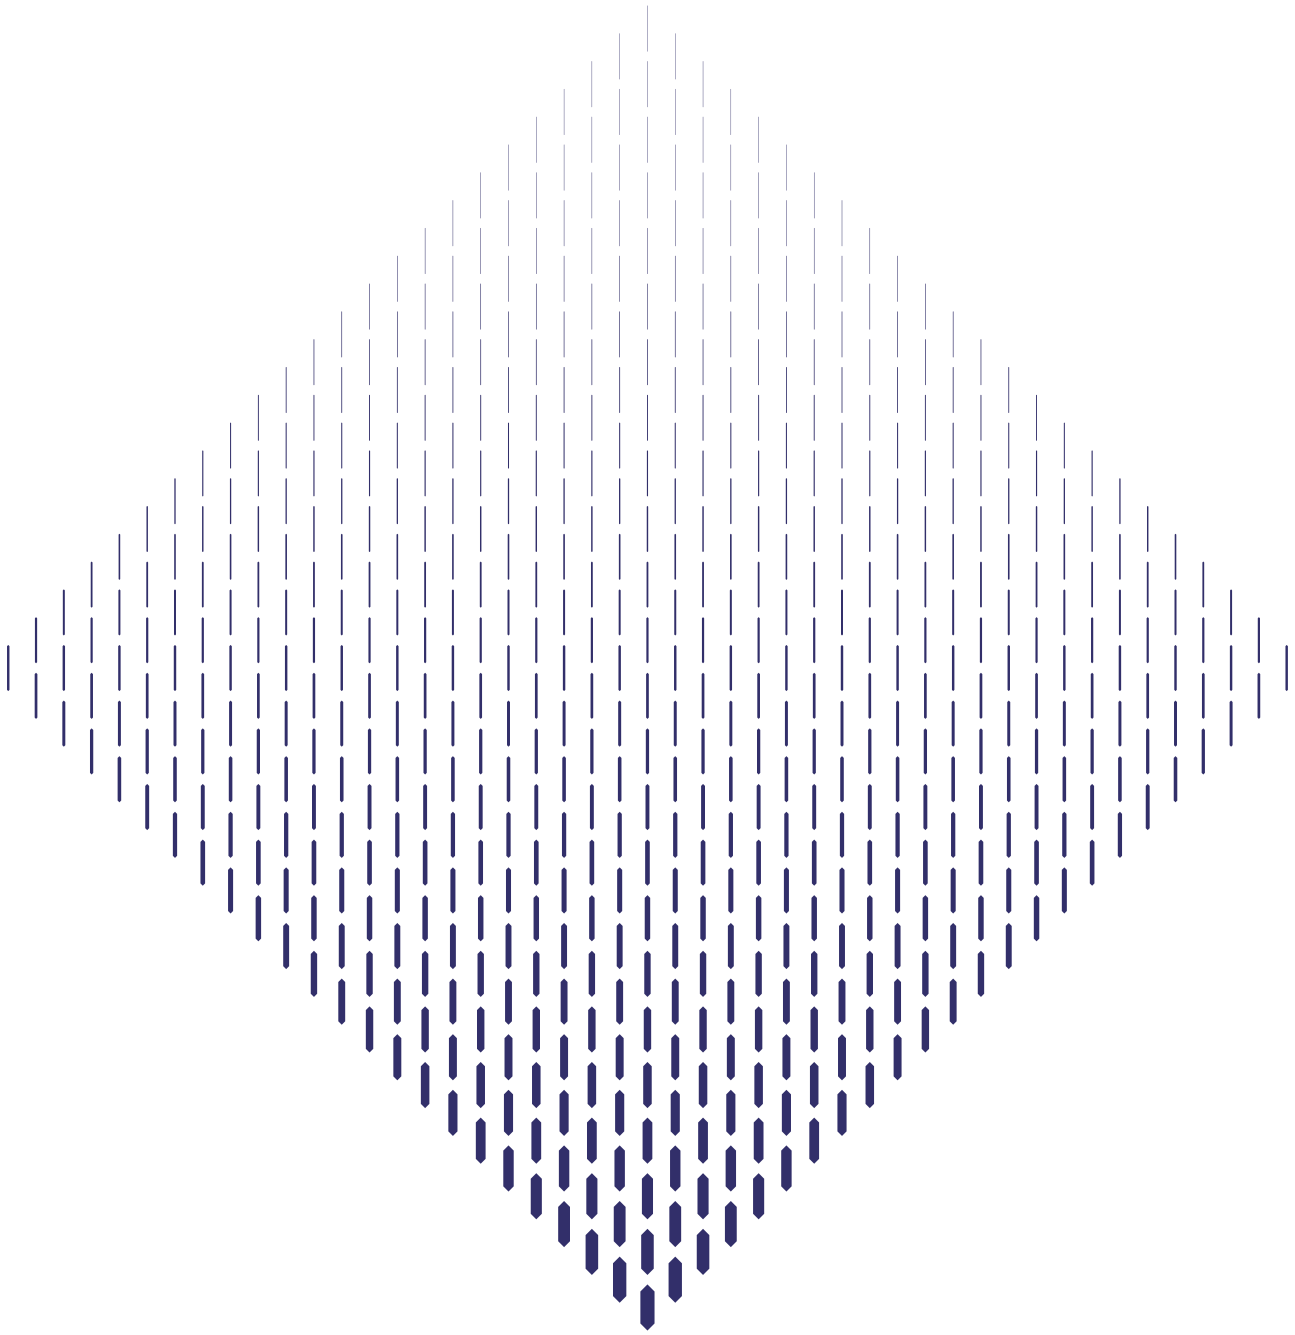
2026



THE NEW ERA OF MARKETING IN FINANCIAL SERVICES

How AI is reshaping communications, compliance and customer trust

2026



THE NEW ERA OF MARKETING IN FINANCIAL SERVICES

COLLECTED ESSAYS

zango

Contents

Introduction p.4
Ritesh Singhania, CEO and Co-Founder, Zango

1. The big picture: AI, trust and the customer

Trust as a foundation of marketing in the age of AI p.8
Miguel Maya, CEO, Millennium bcp

Do customers understand what they are buying? p.12
Rogério Campos Henriques, CEO, Fidelidade

AI in asset management: democratising the ability to compete p.16
Emanuel Silva, CEO, IMGA

AI in insurance: better risk assessment requires better supervision p.20
Manuel Caldeira Cabral, former Member of the Board, ASF

2. Communicating compliantly: conduct and the law

How existing financial marketing rules co-exist with AI-generated content p.26
Inês Palma Ramalho, Partner, Pérez-Llorca

Marketing compliance in the age of AI p.30
Raquel Vila Verde, Group Head of Compliance, Caixa Geral de Depósitos

When AI talks to AI: conduct risk and customer communications p.34
Francisco Luís Alves, Head of Legal, Insurance, Fidelidade Group

3. Keeping pace: governance, capability and culture

Rethinking the Three Lines of Defence p.40
Pedro Oliveira Machado, Country Senior Director, Grupo Ageas Portugal

Ahead of the curve: how compliance functions get ready for AI p.44
André Prazeres Henriques, Chief Compliance Officer, BPI

Why AI is opening a marketing compliance skills gap in banking p.48
João Pedro Tavares, Programme Director, Católica Lisbon

Building a culture that adopts AI responsibly while oversight keeps pace p.52
Inês Caria, Chief Compliance Officer, Banco Santander Totta

Conclusion p.56
Sam Green, Senior Policy & Partnerships Lead, Zango

Contributors



Miguel Maya
CEO
Millennium bcp



Rogério Campos Henriques
CEO
FIDELIDADE



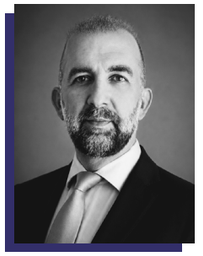
Emanuel Silva
CEO
imga



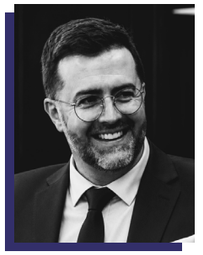
Manuel Caldeira Cabral
Former Member of the Board, ASF, and former Minister of the Economy



Inês Palma Ramalho
Partner
Pérez-Llorca



Francisco Luís Alves
Head of Legal, Insurance
FIDELIDADE



Pedro Oliveira Machado
Country Senior Director
grupo ageas portugal



André Prazeres Henriques
Chief Compliance Officer
BPI



Raquel Vila Verde
Group Head of Compliance
Caixa Geral de Depósitos



João Pedro Tavares
Programme Director
CATÓLICA LISBON
BUSINESS & ECONOMICS



Inês Caria
Chief Compliance Officer, Banco Santander Totta
Santander

Introduction

AI is transforming how financial institutions engage with their customers. Done well, it turns generic communication into something relevant, timely and genuinely useful, explaining a complex product in plain language or anticipating what someone needs. It also speeds up content production, whether that means drafting an advert for a savings account or producing personalised messages at scale.

However, the same speed and scale that make AI useful also raise the stakes for oversight. Each personalised variant is a separate piece of content the firm has to stand behind, and the landscape they travel through has fragmented across search, social, email, in-app messaging, affiliates and paid influencers. More content on more channels is produced faster than anyone can manually review it, raising the risk that something unclear or misleading reaches a customer before it is checked.

Volume is one visible problem, but not the whole of it. Whether a communication misleads depends on its context. The same claim can be fair on a product page and misleading in a social post. Firms have to assess not only what is being said, but how it is presented and where it ends up.

This changes the nature of the compliance challenge. Traditional marketing controls were designed around relatively fixed communications, like a campaign or advert reviewed before publication. AI introduces a far more dynamic model, in which content is generated, adapted and distributed continuously. That requires a different governance approach.

At the same time, AI makes a different kind of oversight possible. Instead of checking each asset once, in isolation, firms can review communications in the context of the product, audience and channel they relate to, link them to the rules they depend on, and keep checking them once they are live. They can also capture the judgement that has usually sat with a handful of reviewers, so that how the rules apply to a particular claim, product or channel stays with the firm rather than leaving with the people who worked it out. The technology does not replace that judgement, but it can make it consistent and available at the scale firms now operate.

Yet the first line of defence, the revenue-generating and customer-facing side of a firm, is adopting AI faster than the second and third lines that oversee and assure it. That was a central finding of Zango's AI governance initiative with researchers from the Oxford Martin School and the University of Glasgow, and it is a pattern we see as we deploy AI agents for risk and compliance in Portugal and beyond.

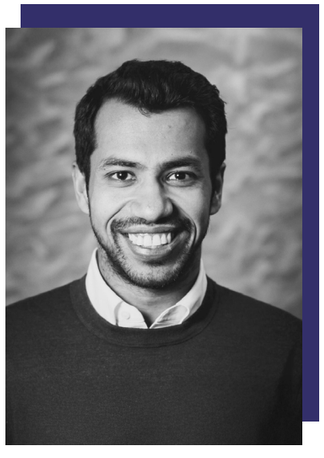
This report brings together leading voices from across Portugal's financial services sector to explore what AI means for how firms communicate with, and protect, their customers. The essays that follow ask three questions: how AI is reshaping trust across banking, asset management and insurance; how existing rules on conduct and communication apply when machines help produce content; and what governance and capabilities firms need for this new era.

Running through the essays is a shared premise. As AI takes on more of how customers are reached, priced and served, oversight has to move at the same speed, so a firm can stand behind what it produces as fast as it produces it.

In the end, this is about trust. Institutions operate on trust, and the capabilities transforming how they reach their customers have to be turned, just as deliberately, to protecting them.

Ritesh Singhania
CEO and Cofounder

zango



1
The big picture

AI, trust and the customer

ZANGO
PERSPECTIVE

The problem: keeping track of what reaches the customer

Compliance has long been built around approval at a specific point in time. Yet a firm can no longer assume that what it approved once is the accurate state of what is live. For instance, a campaign approved once may be edited by a partner months later, or a claim cleared for a product page may be rewritten and redistributed for social media.

Personalisation and the fragmentation of the media landscape are core drivers of this challenge. One campaign becomes thousands of variations, using data to make communications more targeted and often more helpful, but harder to see as a whole. Customers also increasingly form a view from content the firm did not produce, on channels it does not own, including comparison sites and finfluencers.

A further complicating issue is that whether a communication misleads depends on its context. A claim that a customer can earn up to a given rate may sound like a promise if the qualifying information is not given sufficient prominence. Effective oversight therefore requires firms to assess not only what is being said, but how it is being presented.

In the new era of marketing: compliance cannot stop at approval

Continuous assurance needs to sit alongside approval. Firms need an ongoing view of what has been published, where, and under which approval, including content distributed by partners and affiliates, backed by an audit record of who approved each asset, when, and why.

The same controls used before publication need to extend to content once it is live. If a page changes, a claim is repurposed into another channel or a relevant regulatory requirement moves, firms need to be able to identify which communications are affected rather than waiting for them to return through the approval process.

This requires a change in architecture - one with a live understanding of the regulatory context. Yet few institutions are built that way today. Firms that cannot assure continuously in this landscape risk the very trust they spent years building.

Questions to ask:

1. Could you produce, within a day, every live communication about a single product across your own, your partners' and your affiliates' channels?
2. When personalisation changes the emphasis of a message, who decides whether the balance between benefit and risk still holds?
3. Would you be able to identify which of your live assets were approved under guidance that has since changed?

Trust as a foundation of marketing in the age of AI

By Miguel Maya
CEO
Millennium bcp

Millennium bcp



For a bank, adopting AI well in marketing and customer communications is not fundamentally a technology challenge. It is a trust challenge.

AI is at its best when it helps banks understand customers better, communicate more clearly, and serve people in ways that feel more relevant, timely and human. Used well, it can transform generic communications into meaningful interactions, simplify complex financial decisions, anticipate customer needs, and make every engagement more useful and personalized.

AI raises the stakes for trust

But in banking, personalisation alone is not enough. Customers also expect transparency, fairness, privacy and human accountability. They want to know how their data is used, when AI is involved, and that important decisions remain aligned with clear ethical principles. Trust is not built because a bank uses AI. Trust is built because it uses AI responsibly.

What is at stake if it is done badly is far greater than marketing effectiveness. Poorly implemented AI can create the perception of manipulation, bias, opacity or excessive automation. A communication that feels intrusive, inaccurate or inauthentic may damage in seconds a relationship that took years to build.

As a major bank in our core markets, Millennium bcp holds a unique position in society. We manage not only money, but essentially confidence. That is why the winners in the age of AI will not necessarily be the institutions that deploy the most AI, but those that deploy it in a way that strengthens trust. I firmly believe that AI will reward those who adopt it well and unsettle those who fear it.

Trust as banking's decisive asset

Trust has always been important in banking and the cornerstone of our communications strategy, but in the age of AI in which disinformation is recurrent, it has become the decisive asset. Something that takes years to earn and only seconds to lose.

We live in a world where people have access to more information than ever before, yet often struggle to know what is true, reliable and relevant. As I have said before, disinformation

> “Whether through human interactions or digital channels, the same principles apply: accuracy, transparency, respect for privacy and a commitment to the customer’s best interests.”

is eroding trust and the quality of the decisions families and companies make. In this environment, customers increasingly value institutions that can act as trusted guides, helping them navigate complexity with confidence.

For Millennium bcp, trust is not an abstract concept. It is the foundation of every relationship. Customers entrust us with their savings, their investments, their businesses and many of their most important life decisions. Without trust, there is no meaningful relationship, no loyalty and ultimately no long-term value creation.

That reality shapes the way Millennium bcp communicates. Our objective is not simply to attract attention; it is to provide clarity. We seek to communicate in a transparent and responsible way, particularly when financial decisions are involved. Whether through human interactions or digital channels, the same principles apply: accuracy, transparency, respect for privacy and a commitment to the customer's best interests.

Innovation with responsibility

AI will transform how banks communicate, but it will not change what customers expect from them. In fact, it raises the bar. As technology becomes more powerful, trust becomes more valuable. The institutions that succeed will be those that combine innovation with responsibility, and technology with values. At Millennium bcp, we see trust not as a consequence of communication, but as its ultimate purpose.

People remain at the heart of everything we do. They are, and always will be, the reason we communicate. AI doesn't change that. What it changes is how people discover information and, increasingly, how they make decisions.

The new gateways to financial information

AI is becoming a new intermediary between brands and customers, creating both an opportunity and a responsibility. In this new paradigm, it is no longer enough to communicate with credibility through our own channels.

As AI becomes an increasingly important gateway to information, the challenge is to ensure that AI systems recognize, prioritize and surface trustworthy content, helping people make better-informed decisions. This is why Generative Engine Optimisation is becoming an important part of our marketing strategy.

It's not about optimizing for algorithms, it's about ensuring that trustworthy, accurate and authoritative financial information reaches people wherever they choose to seek answers. At Millennium bcp, we are already embedding best practices into our content strategy, alongside initiatives such as "Afinal de Contas", which reflects our long-term commitment to financial education. We believe that every opportunity to educate, clarify and simplify complex financial topics is also an opportunity to strengthen trust.

Do customers understand what they are buying?

By Rogério Campos Henriques
CEO
Fidelidade

FIDELIDADE



Insurance has a particular characteristic: we ask people to make decisions today about risks and needs that may only become apparent many years from now. When we protect a family, think about our health or start preparing for retirement, we are making decisions whose benefits may not be immediately visible and whose consequences can extend over decades.

At the same time, many insurance and long-term savings products are inherently complex. Regulation, while essential to protect consumers and ensure transparency, often brings extensive information requirements that do not necessarily make communication with customers any simpler. The challenge, therefore, is not to provide more information, but to make sure that information is understood and genuinely helps people make informed decisions.

Understanding the risks we face

Understanding what we are buying also requires us to understand the risks we face. And some of those risks are changing profoundly. Longevity is perhaps the clearest example. Living longer is excellent news, but it means longer periods in retirement and greater healthcare and assistance needs, requiring us to think earlier about how we protect ourselves and save for the future. Climate change is another example, as it alters the frequency and scale of some of the risks faced by individuals and businesses.

AI can help us address this challenge. But it can also have the opposite effect. Today, we have an unprecedented ability to understand our customers, their behaviours and, increasingly, anticipate some of their needs. Used purely from a commercial perspective, AI can make us extraordinarily effective at identifying the message, the moment or the product most likely to lead someone to buy. But selling better does not necessarily mean helping someone make a better decision.

Using AI to make the complex simpler

This is where I see one of the greatest opportunities for AI in insurance: using technology to make simple what is complex today.

> “A company's growth should never depend on a customer failing to understand what they have bought, underestimating a risk or being unaware of an important limitation of a product.”

We can explain products in clearer language, tailor information to different levels of financial knowledge, compare alternatives, simulate future scenarios and help customers understand not only what they are buying, but why they may need that protection or those savings.

Many of the decisions that will determine our quality of life 20 or 30 years from now need to be made much earlier. We cannot expect people to make good decisions if they do not understand the choices in front of them.

For an insurer, this cannot simply be a matter of regulation or compliance. Above all, it is a matter of trust. A company's growth should never depend on a customer failing to understand what they have bought, underestimating a risk or being unaware of an important limitation of a product. If a customer cannot understand what we are offering, the first question we should ask ourselves is whether we have explained it clearly enough.

Financial literacy and AI literacy must develop together

This responsibility becomes even greater with AI. As algorithms play an increasing role in the information, recommendations and propositions people receive, transparency must evolve too. People should know when they are interacting with AI and have a basic understanding of how technology can influence what is presented to them.

That is why I believe financial literacy and AI literacy will increasingly need to develop together. Responsibility is shared. Governments have a role in building financial and digital literacy from an early age. Regulators must ensure innovation takes place within a clear framework of transparency and consumer protection. And financial institutions have a particular responsibility because we design the products and increasingly develop the technology through which customers access them.

High tech must still mean high touch

But there is one dimension we cannot afford to lose. The future of insurance must be both high tech and high touch. Technology can make us faster, more efficient and better able to anticipate needs. But insurance demonstrates its true value at deeply human moments: an accident, an illness, the loss of a loved one, or simply when someone needs help navigating a difficult situation.

At those moments, technology should allow us to be closer, not create greater distance. I do not see AI as a substitute for the human dimension of insurance, but as an opportunity to strengthen it: freeing people from lower-value tasks, helping us understand customers better and allowing us to be more present when that presence really matters.

Ultimately, that may be the greatest potential of AI for our industry: to make insurance simpler and more efficient, but also more personal and, above all, more human.

AI in asset management: democratising the ability to compete

By Emanuel Silva
CEO
IMGA



AI is beginning to reshape asset management in ways that go well beyond automation. It has the potential to change how asset managers operate, develop products, process information, communicate with distributors and investors, and meet increasingly demanding regulatory requirements.

For smaller markets such as Portugal, this transformation presents a particularly important opportunity. The Portuguese asset management industry operates on a significantly smaller scale than the major European and international markets. Local asset managers inevitably have more limited resources to invest in technology, systems and specialist capabilities than the largest global institutions, which benefit from substantial economies of scale. AI has the potential to narrow that gap.

Technology as a democratising force

At IMGA, we see AI not simply as a tool to reduce costs or automate processes, but as an opportunity to strengthen our organisation and increase our ability to compete with leading international asset managers.

Access to increasingly sophisticated AI solutions is becoming progressively more democratic. Capabilities that until recently would have required substantial technology investment and large specialist teams are becoming accessible to smaller organisations.

For asset managers operating in markets such as Portugal, this matters. AI can help reduce execution times, automate lower-value activities, improve the quality and consistency of information, and allow our teams to focus on areas where experience, judgement and human expertise genuinely make a difference.

In this sense, AI can contribute not only to greater operational efficiency, but also to a democratisation of the ability to compete and, ultimately, of the investment solutions that managers are able to offer.

> “As AI increases the ability to personalise communications and investment propositions, maintaining accurate, fair and consistent representation of products across the distribution chain becomes even more important.”

From opportunity to practical application

At IMGA, this vision is already being translated into practical initiatives.

We currently have several AI-based agents and solutions in operation and are completing additional tools across different areas of the organisation. Commercial and marketing activities are naturally among them. The ability to analyse information, support content creation and improve the speed and quality of communication can significantly enhance how an asset manager interacts with its distribution partners and, through them, with investors.

But our approach extends well beyond marketing. We have developed solutions supporting onboarding and Adverse Media processes, enabling us to increase efficiency in client admission and monitoring while strengthening our ability to comply with applicable regulatory requirements.

We are also developing tools to support the preparation of fund documentation, including Key Information Documents and Prospectuses. These processes involve significant volumes of information, multiple validations and a high degree of regulatory precision. AI can substantially reduce preparation and review times while maintaining the necessary levels of human oversight and control.

These are only some examples. Our broader programme is identifying further areas where AI can improve execution, quality and organisational capability.

Innovation cannot dilute accountability

Greater use of AI also brings greater responsibility, particularly in an industry built on trust.

This is especially relevant in asset management because the relationship between the manager and the end investor is frequently intermediated by third-party distribution networks. The asset manager is responsible for the product, its investment strategy and the information produced about it, while much of the communication and client relationship may take place through distributors.

As AI increases the ability to personalise communications and investment propositions, maintaining accurate, fair and consistent representation of products across the distribution chain becomes even more important.

Technology cannot create distance between innovation and accountability. Performance and risk must continue to be communicated transparently, suitability requirements must be respected, and greater personalisation must not come at the expense of investor protection, particularly for less experienced investors.

When communication is generated or influenced by algorithms, accountability cannot simply be transferred to the algorithm. Asset managers must retain appropriate governance, validation and oversight of how their products are represented.

AI should strengthen people, not replace them

There is one principle at the centre of IMGA's approach: we do not see AI as a means of replacing people.

Our objective is to use technology to support our organisation and make our teams more effective - providing faster access to structured information, reducing repetitive tasks and allowing people to devote more time to activities where knowledge, experience and judgement are indispensable.

Asset management will remain a business built on trust: trust in the institution, in the people managing investments, in decision-making processes and in the quality of information provided to investors.

AI does not replace that trust. Properly governed, it can strengthen it.

For a Portuguese asset manager competing in an increasingly global industry, this may be one of AI's greatest opportunities: enabling smaller organisations to access sophisticated capabilities, improve their efficiency and compete on more equal terms with institutions of far greater scale.

If we combine technology with expertise, governance and accountability, AI can become more than an operational transformation. It can help democratise the ability to compete in global asset management.

AI in insurance: better risk assessment requires better supervision

By Manuel Caldeira Cabral
Former Member of the Board, ASF
Former Minister of Economy of Portugal

AI is transforming the insurance industry by reducing operational costs, allowing new forms of advertising, and by enabling a more accurate assessment of individual risk.

Traditionally, insurers relied on actuarial models based on relatively broad categories such as age, location, occupation or claims history. While these models have proved remarkably successful, they inevitably grouped together individuals with significantly different risk profiles. AI, combined with the growing availability of behavioural and personal data, has the potential to overcome these limitations by estimating risk at a much more granular level, and also by targeting advertising at the individual level.

More accurate risk assessment

Motor insurance already illustrates this transformation. Usage-based insurance products employ telematics devices or smartphone applications to monitor driving behaviour. AI models can distinguish between drivers with genuinely different levels of risk, within the same demographic group. Similar approaches are feasible in health, life, home and cyber insurance.

From an economic perspective, this development improves insurance markets. Better risk measurement reduces the information asymmetries that have long characterised insurance. Imperfect information can prevent insurance markets from reaching efficient outcomes. AI offers the possibility of narrowing these information gaps by allowing insurers to distinguish more accurately between high-risk and low-risk customers. More precise pricing may also reduce moral hazard by rewarding safer behaviour.

In principle, AI allows insurers to move closer to the ideal of actuarial fairness, where premiums closely reflect the expected losses generated by each policyholder. Yet actuarial fairness should not be confused with social fairness. From an actuarial perspective, charging different premiums to individuals with different risk profiles is both efficient and equitable.



> “The challenge is no longer limited to whether premiums are fairly calculated, but also whether consumers are being persuaded through AI-driven techniques that are transparent, proportionate and consistent with fair treatment.”

From a broader societal perspective, however, some characteristics may be considered inappropriate for pricing, because insurance also performs a social function by pooling risks. Health conditions, race, disabilities or genetic information may improve predictive accuracy, but society may legitimately decide that these factors should not determine the price individuals pay. AI cannot surpass the old policy question of asking not simply what risks can be measured, but which risks should legitimately be priced.

From risk-based pricing to personalised markets

The same technologies that improve risk assessment can also transform commercial strategies. AI models do not simply estimate the probability of a claim; they can also estimate a customer's willingness to pay, likelihood of switching insurer or sensitivity to price increases. At that point, pricing may shift from measuring risk to maximising revenue by extracting the highest possible price from each customer.

This distinction is economically important. Charging different premiums because customers present different expected risks is consistent with efficient insurance markets. Charging different prices because some customers are less informed or less likely to search for alternatives is a classic form of price discrimination. While such practices may increase profits, they do not improve economic efficiency or consumer welfare.

These concerns are reinforced by the increasing complexity of AI models. Modern machine learning systems often operate as "black boxes", generating highly accurate predictions while offering only limited explanations of how decisions are reached.

AI is also changing the way insurance products are marketed. For decades, consumer protection focused on standardised disclosures and the familiar problem of lengthy "small print". Today, digital platforms enable insurers to deliver personalised advertisements and offers based on browsing history, purchasing behaviour and social media activity.

Advertising is becoming dynamic, adaptive and socially targeted, meaning that different consumers may receive different messages about the same product. While this can improve the relevance of information, it can also exploit behavioural biases and reduce consumers' ability to make informed choices. The challenge is no longer limited to whether premiums are fairly calculated, but also whether consumers are being persuaded through AI-driven techniques that are transparent, proportionate and consistent with fair treatment.

Supervising AI with AI

These developments present major challenges for insurance supervisors. Traditional supervision relies on reviewing underwriting criteria, advertising information, actuarial assumptions and pricing methodologies. AI-based models, however, continuously learn from new data and may involve millions of parameters that cannot realistically be examined using conventional auditing techniques. Supervisors therefore face an information asymmetry of their own: insurers increasingly rely on algorithms whose complexity exceeds the analytical capacity and the legal framework of traditional supervisory methods.

The answer, however, is unlikely to be restricting AI. Instead, supervisory authorities should consider deploying AI themselves. Advanced analytical tools can help regulators audit pricing algorithms, identify the variables that most influence premiums, detect proxy discrimination, simulate pricing outcomes across different groups of consumers and distinguish between models that estimate genuine insurance risk and those designed primarily to exploit willingness to pay. The same technologies can also help monitor AI-driven advertising and distribution practices, identifying patterns of behavioural targeting.

This approach is consistent with the growing development of RegTech and SupTech - regulatory and supervisory technologies that strengthen financial oversight through artificial intelligence and advanced analytics. Rather than replacing human judgement, these tools should enhance supervisors' capacity to understand increasingly complex markets.

AI represents one of the greatest opportunities for improving insurance markets in decades. It can make pricing more accurate, reduce inefficiencies and better align premiums with individual risk. Yet these same technologies can also facilitate opaque forms of price discrimination and personalised commercial practices that move markets away from socially desirable outcomes. The challenge for policymakers is therefore not whether AI should be used in insurance, but how to ensure that it serves actuarial fairness, social fairness and consumer autonomy. Achieving this objective will require supervisory authorities to embrace AI as actively as the firms they regulate, making AI not only the object of regulation but also one of its most effective instruments.

2
Communicating compliantly

Conduct and the law

The problem: judgement lives in people, not the firm

Continuous assurance needs something to assure against, and in marketing compliance that is the firm's judgement. The underlying rules set clear requirements: information must be balanced, conditions and restrictions given appropriate prominence, risks cannot be obscured by benefits, and marketing cannot contradict required product information. Firms will often also have their own internal rules and standards on top of that. But how those principles apply to a particular claim, product or channel is often not written down.

In many institutions today, that judgement is made once and then lost. For example, a reviewer works out the right risk qualifier for a deposit campaign, which then launches, but the next reviewer facing the same claim starts again and may land somewhere different. What the firm knows sits with a handful of people, which works until they leave.

In the new era of marketing: judgement has to become institutional

AI makes it possible to apply that accumulated judgement systematically. Earlier tools largely scanned for banned phrases and missing warnings, with limited ability to account for context. A system can hold a firm's products, regulatory obligations, internal rules, approved claims and prior decisions, and assess every new asset against all of it, with a human reviewer in the loop and accountable for the final judgement.

The same capability creates an audit record of how decisions were made. Portugal's revised framework strengthened what institutions must report to their supervisor, and a firm whose decisions are recorded as they are made does not have to reconstruct them from email threads and shared drives months later.

Institutionalising that judgement is what makes it scalable. Each time a reviewer amends a claim, that amendment is captured in a record of approved claims, qualifiers and channel standards and applied to later reviews. Precedent becomes something the institution owns and keeps current, rather than something its reviewers remember.

Questions to ask:

1. When your team signs off a specific claim or disclosure, does that decision travel to the next asset, or does the next reviewer work it out again?
2. Could you show why a specific claim was approved, and on what basis, twelve months after the decision?
3. Would the same asset get the same answer from two different reviewers on your team?

How existing financial marketing rules co-exist with AI-generated content

By Inês Palma Ramalho
Partner and Head of the Financial Services Department
Pérez-Llorca

Financial institutions have always had to answer a relatively simple legal question before communicating with customers: is this message compliant? AI does not replace that question, but it certainly makes the answer considerably more complex.

A framework that already exists

Financial advertising in Portugal is already governed by a dense sectoral framework. Depending on the product and activity concerned, each of the three financial regulators (Banco de Portugal, CMVM and ASF) supervise compliance with specific rules governing communications with customers either stemming from EU legislation (namely MiFID II, the IDD or the PRIIPs Regulation) or domestic legal or regulatory regimes.

The rules differ in their detail, but they share a common principle: communications to customers must be fair, complete, current, clear and not misleading. And that principle translates into very concrete requirements.

Portuguese banking rules, for example, require an APR to be given similar prominence to a credit product's benefits highlighted in consumer-credit advertising. Under MiFID II and its Portuguese implementation legislation, references to the potential benefits of an investment must be accompanied by a fair and prominent indication of the risks; under PRIIPs international and domestic regimes, marketing communications cannot contradict (or undersell) the information contained in the product's Key Information Document (KID).

AI does not change the underlying principle

AI does not alter this common principle, its intended effect or any of its related rules. If anything, AI further enhances the importance of complying with these rules and may make a financial institution more accountable for failing to respect them.

Pérez-Llorca



> “AI may change who (or what) writes the advertisement. It does not change the fundamental proposition on which financial regulation rests: the regulated institution must remain answerable for what the customer is told.”

It is important to note that a statement does not automatically become (more) misleading because it was generated by AI rather than resulting from a marketing team's (human) brainstorming session. Nor does an institution cease to be subject to financial-marketing rules merely because AI intertwines product and customer.

When one message becomes thousands

Traditional financial advertising has generally been built around identifiable content. A television advertisement or an email campaign is produced, reviewed and eventually communicated to a specific audience. However sophisticated the process behind it may be, there is usually a single message against which compliance can be assessed.

AI, in particular generative AI, challenges this model. This technology allows institutions to move from producing one communication for a group of similar customers to producing many communications for many (different) customers simultaneously. Language, emphasis, and examples can be adapted dynamically. The communication one actually receives may therefore be materially different from that received by another.

This is an important legal shift in financial marketing rules. The question is no longer whether the campaign or the advertisement is legally compliant, but rather if the institution can ensure that the universe of communications generated by the system respects the applicable regulatory boundaries.

That distinction becomes increasingly significant when it comes to personalisation. Let's consider a scenario where an institution wants to promote the same financial product to 100,000 customers. In a traditional campaign, the starting point may be one advertisement. With AI-enabled personalisation, we could theoretically be discussing 100,000 variations, depending on all clients' preferences and past behaviors, all of which may be factually accurate. Yet each may present benefits differently, emphasize different product features or give different prominence to risk. From a regulatory perspective, however, what ultimately matters is not merely what the institution intended its campaign to say, but rather what each customer was actually told and was "pushed to believe".

This creates a challenge that is legal as much as technological. A regulated entity must be capable of standing behind the communication received by its customer, even when that communication was generated dynamically rather than drafted word by word in advance.

The AI Act applies alongside, not instead

The EU AI Act adds another layer to this framework, but it should not be mistaken for the governing regime for financial advertising. It is true that this act regulates AI practices and systems according to their characteristics and uses. Financial marketing, however, does not cease to be financial marketing (nor does it become it) because AI is involved.

The AI Act must therefore be understood alongside, rather than instead of, existing sectoral financial regulation. These regimes shall apply cumulatively with very different purposes. One governs the use of AI in the circumstances falling within its scope and the other continues to govern what regulated entities may communicate about financial products and services.

Where accountability lives

That distinction is particularly important when considering accountability. The use of AI inevitably introduces new actors: developers, tech and data providers and other intermediaries may all participate in the technological chain. But for the financial institution, the central question remains whether the communication made to the customer complies with the obligations imposed on it. The fact that particular wording originated in a model cannot, by itself, answer that question.

That is why the transition brought about by AI is more significant than simply replacing one content-creation tool with another. The traditional legal model focused principally on approving the message. The emerging one requires institutions to consider the legal consequences of governing the machine capable of producing the message.

AI may change who (or what) writes the advertisement. It does not change the fundamental proposition on which financial regulation rests: the regulated institution must remain answerable for what the customer is told or led to perceive.

Marketing compliance in the age of AI

By Raquel Vila Verde
Group Head of Compliance
Caixa Geral de Depósitos

Caixa Geral de Depositos



The advertising of financial products in Portugal has been shaped over the past decade by a simple yet demanding principle: customers must be able to understand, clearly, immediately and without ambiguity, what is being offered to them. Regulatory developments have consistently sought to reinforce this objective, keeping pace with changes in communication channels and commercial practices.

Portugal's tightening advertising regime, and its lessons

The latest revision of the applicable framework represents more than a formal update. It strengthens transparency and disclosure requirements, clarifies the distinction between different categories of advertising, and introduces more demanding standards regarding the prominence and presentation of information relevant to customers. At the same time, it reinforces institutions' reporting and information-sharing obligations vis-à-vis the supervisory authority in relation to their advertising activities.

The key lesson from this regulatory evolution is that compliance is no longer an ex post validation exercise. In an environment of increasing regulatory and supervisory scrutiny, Compliance and business teams must work together from the earliest stages of campaign development, ensuring that commercial objectives remain compatible with transparency, disclosure and customer protection requirements. Promotional claims that may appear attractive from a commercial perspective often require specific disclosures or strict conditions of use. Compliance has therefore ceased to be a final checkpoint and has become a design requirement.

The clearest opportunities for AI in marketing compliance

The use of AI represents a significant opportunity to improve the efficiency of campaign review and approval processes. A properly configured system can, for example, automatically identify potential instances of non-compliance relating to prominence requirements, the use of promotional claims or the omission of mandatory information, thereby acting as a first line of control.

> "AI can help reduce review times and improve consistency, allowing marketing and compliance teams to focus their efforts on cases that genuinely require specialist assessment."

In this context, AI can help reduce review times and improve consistency, allowing Marketing and Compliance teams to focus their efforts on cases that genuinely require specialist assessment.

However, trust in such systems requires a series of safeguards. First, sufficient explainability is needed to ensure that control functions understand the rationale underpinning the system's conclusions. Secondly, governed and auditable update mechanisms must be in place to ensure the timely incorporation of regulatory changes. Thirdly, clear human accountability remains essential. Technology may support decision-making, but responsibility for the compliance of a campaign ultimately remains with the institution and its designated accountable individuals.

Finally, systems should be subject to regular robustness and quality testing to ensure the effectiveness of mechanisms designed to detect potential instances of non-compliance. A failure to identify a breach at an early stage may result in the rapid proliferation of the same non-compliance across multiple campaigns and communication channels, significantly increasing regulatory and reputational risk exposure.

It is also important to recognise that AI does not alter the fundamental principles governing banking advertising. The obligations of transparency, accuracy and customer protection remain exactly the same. What changes is the scale, speed and complexity of the processes through which messages are created, personalised and distributed.

Personalisation: rebalancing relevance and customer protection

Personalisation is where some of the most significant challenges for the future of financial advertising emerge. Communications that are better tailored to customers' needs and characteristics can substantially improve the customer experience, reduce irrelevant information and increase the usefulness of the messages received.

However, data-driven and algorithm-based personalisation introduces additional complexity. Instead of a single campaign subject to centralised approval, institutions may produce thousands of variations of the same message, tailored to specific customer segments or, potentially, to individual customers. This reality makes oversight and control significantly more demanding.

The principal risk does not stem from personalisation itself, but from the possibility that messages become optimised solely to maximise commercial outcomes without being subject to the same compliance discipline traditionally applied to mass-market advertising. Furthermore, certain targeting practices may raise additional concerns when directed at more vulnerable customers or individuals who may be particularly susceptible to specific commercial stimuli.

The answer is neither to limit innovation nor to prevent personalisation. Rather, it is to ensure that compliance requirements are embedded within the architecture of the systems that generate and distribute communications. However, validating the system itself does not remove the need for oversight mechanisms covering the outputs it produces. As personalisation becomes more dynamic and operates at greater scale, the role of Compliance is likely to evolve from the exhaustive review of individual communications towards a model based on rules, controls, sampling methodologies and the continuous monitoring of outcomes generated by those systems.

When AI talks to AI: conduct risk and customer communications

By Francisco Luís Alves
Head of Legal, Insurance
Fidelidade Group

FIDELIDADE



The debate on AI in insurance still tends to start inside the insurer: underwriting, pricing, fraud detection, claims management and customer service. That focus is justified, but it is now incomplete. AI has also moved to the customer side of the relationship. Customers use it to interpret policies, prepare complaints and compare products. Insurers use similar tools to analyse information, draft replies and support decisions. The interaction is no longer just between a person and an organisation. Increasingly, systems sit on both sides of the exchange.

That shift changes the basic premise of consumer protection. For years, insurers had data, technology and expertise, while customers did not. That imbalance remains, but it has changed form. A complaint can now be drafted with AI, reviewed with AI and answered with AI. In some cases, the real conversation takes place through machine-generated text, with limited human judgement when it matters.

Automated decisions to automated dialogue

Financial services regulation has long focused on automated decision-making. Decisions about pricing, cover, eligibility or claims affect customers directly and must remain controlled and reviewable. AI-mediated communications create a different problem. The risk often starts before the final decision, in the framing of facts, the summary of policy terms, the explanation of exclusions and the interpretation of complaints.

Communication is not neutral. It shapes how the insurer understands the customer and how the customer experiences the insurer. A technically correct response can miss the real concern. A neat summary can hide the assumptions that drove the conclusion. Governance cannot therefore focus only on the model. It needs to look at the relationship that the model helps to create.

> “In a market where any participant can produce sophisticated language with AI, differentiation will not come from the apparent polish of the text. It will come from showing that behind the words there is responsibility, judgement and oversight.”

Responsibility does not move to the machine

Responsibility remains the central issue. AI can write the words. It cannot assume responsibility for them. In traditional customer communication, it is usually possible to identify who assessed the matter and approved the position sent to the customer. Generative AI adds an intermediate layer. The line between human judgement and machine assistance becomes harder to see. Who validated the content? What information was considered or left out? When did someone with authority review the answer?

If a chatbot explains an exclusion poorly, if an automated response misleads a customer or if an AI-generated letter omits a relevant fact, accountability still belongs to the insurer. Technology can support communication. It cannot replace accountability to customers, supervisors or courts.

The test lies in the exceptions: ambiguous wording, sensitive complaints, vulnerable customers and facts that do not fit the examples used to test the system. These are the cases where human review must be real, not decorative. The reviewer must be able to change, suspend or escalate the answer.

The illusion of quality

When AI is present on both sides of the communication, a further risk emerges: the illusion of quality. A complaint can look rigorous because it is well structured. A response can look authoritative because it uses precise language. Both can still be incomplete, misleading or wrong. The quality of language is not the same as the quality of analysis.

Generative AI reduces the effort needed to produce sophisticated communications. It does not guarantee accurate facts, sound reasoning or real understanding. A customer can submit a carefully drafted complaint with plausible references while misunderstanding the policy. An insurer can produce a polished answer that fails to address the substance of the case.

The use of AI by customers should not, by itself, be treated as suspicious. Generative AI can introduce inaccuracies, inferred facts or unsupported assumptions into a communication without any deliberate intent to mislead. Customers remain responsible for validating the information they submit, but insurers may increasingly need to distinguish intentional misrepresentation from negligent reliance on AI-generated content.

Fraud assessments should focus on the underlying facts, the customer's intent and the circumstances of the case, rather than on the mere use of AI. If AI helps a customer explain the problem, the fair treatment question is simple: has the insurer understood the issue, assessed the facts properly and answered in a way the customer can understand?

Transparency is necessary, but not enough

As systems improve, it becomes harder to distinguish text written by a person from text generated by a machine. Transparency about AI use is important, including under the EU AI Act. But disclosure is only the beginning.

Knowing that AI was used does not answer the harder questions. Who is responsible for the result? What facts were considered? What assumptions shaped the answer? What was checked? Who can correct an error? If the communication materially affects the customer, these questions matter more than the system label.

For the customer, the concern is rarely technological. It is practical. If the outcome is unfavourable, the customer wants reasons and access to someone with authority to review the case. Human oversight is therefore not a formal safeguard to be mentioned in a policy. It is a condition of trust.

A conduct risk lens for the next phase of AI

Insurers should govern AI-supported communications as part of the customer relationship, not merely as an extension of IT or model risk. That means clear ownership, escalation paths, monitoring errors and complaints, and human review where customer impact is material. It also means testing communications for clarity, fairness and suitability, not only for technical accuracy.

Insurance is built on trust. That trust requires identifiable responsibility, understandable communication and human intervention at the right moments. These are not obstacles to innovation. They are conditions for adoption.

In a market where any participant can produce sophisticated language with AI, differentiation will not come from the apparent polish of the text. It will come from showing that behind the words there is responsibility, judgement and oversight. When AI talks to AI, trust no longer depends only on the technology. It depends on the people and institutions that remain responsible for the outcome.

3
Keeping pace

Governance, capability and culture

ZANGO
PERSPECTIVE

The problem: marketing has moved faster than its oversight

Marketing is adopting AI faster than the functions that oversee it. Campaigns are increasingly drafted, varied and distributed through systems the second line had no hand in specifying and cannot readily inspect. That creates an oversight gap, which is widening.

Accountability is also harder to locate. The Three Lines of Defence model depends on clear ownership of decisions and controls. When a system contributes, it can become harder to establish who assessed the matter, what information and rules shaped the outcome, and when someone with authority reviewed it.

Few second lines are currently equipped to close either problem. Doing so calls for people who can read what a system has been given, follow how it reached a conclusion, and say whether the reasoning holds. Those are not the skills compliance has traditionally recruited for, and they sit across compliance, marketing and technology rather than in one team.

In the new era of marketing: compliance must oversee the system, not just the output

AI is transforming how the Three Lines of Defence operate in practice. The second line's job becomes setting the standards those systems run on: what they must hold, what they must never do, and when to escalate, then testing that they behave that way in practice. That changes what the function produces, from approvals to intelligence: knowing which live assets are exposed when an obligation moves, and which claims have drifted.

It is vital that the systems themselves can be challenged, not only their outputs. A reviewer who can accept or reject a recommendation but cannot question the rules that produced it is confirming decisions rather than overseeing them. That capability requires conscious investment in people, and literacy across all three lines.

Questions to ask:

1. Is your second line equipped to oversee the systems the first line is deploying, or only to review what they produce?
2. Can your reviewers question the rules a system applies, or only accept and reject its recommendations?
3. If a supervisor asked how an AI system reached its conclusion on a specific asset, who in your firm could answer?

Rethinking the Three Lines of Defence

By Pedro Oliveira Machado
Country Senior Director | Data Protection Officer
Grupo Ageas Portugal



The three lines of defence (3LoD) model has long served as a foundational architecture for risk governance across financial services and insurance. Yet that architecture is now under considerable pressure and, in important respects, being reshaped. Financial institutions are confronting an increasingly intricate constellation of risks arising from AI, including fraud, information security vulnerabilities, and the less visible consequences of algorithmic bias.

The accelerating imperative to deploy generative AI across business functions, from marketing to customer communications, further intensifies these challenges. In this environment, the once relatively clear demarcation between operational management, risk oversight, and independent assurance is becoming increasingly fluid, calling into question some of the assumptions upon which the conventional 3LoD model has been built.

Shifting to five lines

In this context, some frameworks propose a five-lines-of-defence model:

1. **First line:** responsible for managing AI-related performance and risks at an operational level, with controls embedded within regular business activities and day-to-day processes (i.e. business units).

2. **Second line:** functions responsible for specialist monitoring and for challenging decisions, providing internal oversight through frameworks and standards, policies, tools, and techniques, while ensuring an appropriate degree of objectivity and specialised competence (i.e. Risk and Compliance).

3. **Third line:** provides a high level of independent assurance regarding AI, covering the activities and controls performed by the First and Second Lines (i.e. Internal Audit).

4. **Fourth line:** responsible for managing and integrating AI-related risks into the company's strategy and overall performance. It receives relevant and reliable information from the First and Second Lines, together with assurance from the Third Line, to support appropriate management decisions (i.e. Executive Management).

5. **Fifth line:** ensures ultimate accountability and oversight of AI governance and management, acting with due diligence and appropriate delegation, while maintaining high-level oversight of material gaps and weaknesses (i.e. the Board of Directors).

> "This evolving risk landscape creates a growing need for controls that are embedded within processes and capable of operating in real time."

In addition to the classic 3LoD model, each new line brings a distinct value to the AI governance framework; the Fourth Line strengthens strategic alignment and resource allocation, while the Fifth Line reinforces ultimate accountability and external oversight.

Practical implications

For insurers and banks, this development calls for the integration of specialised AI expertise, a more robust approach to model risk management, and continued alignment with an increasingly demanding regulatory landscape. Moving from a three- to a five-line model of defence can help ensure that AI-related risks are subject to a level of scrutiny and discipline comparable to that applied to more established financial risks. At the same time, such an approach may provide a more balanced foundation for reconciling technological innovation with regulatory compliance in an increasingly complex and rapidly evolving digital environment.

Board accountability under the EU AI Act

Under the EU AI Act, the Board of Directors bears ultimate responsibility for AI governance, particularly in relation to high-risk AI systems deployed across the insurance and financial services sectors. Effective accountability at this level requires more than formal oversight; it calls for active engagement with the risks arising from AI, a sufficient degree of AI literacy, and a clear alignment between AI governance and the organisation's broader strategic objectives.

Boards should invest in meaningful AI literacy programmes, embed AI-related risk within the organisation's broader GRC framework, and draw on established governance principles and frameworks to promote responsible, consistent, and well-grounded performance.

New risks challenge the Lines of Defence (LoD)

The LoD model is increasingly being put to the test by AI-driven threats, ranging from CEO fraud enabled by voice cloning to generative content that can outpace human review in areas such as marketing and claims processing. This evolving risk landscape creates a growing need for controls that are embedded within processes and capable of operating in real time, gradually shifting the emphasis away from periodic reviews towards assurance that takes place alongside execution. In this context, the LoD model is becoming more dynamic, with AI enabling continuous monitoring and facilitating a timelier response to emerging threats, including increasingly sophisticated forms of generative fraud.

Ahead of the curve: how compliance functions get ready for AI

By André Prazeres Henriques
Chief Compliance Officer
BPI



For most of the last decade, compliance functions have become accustomed to responding to regulatory change once it is sufficiently defined. AI is different. By the time detailed supervisory expectations are fully settled, AI will already be embedded across customer journeys, operations, risk management, fraud detection and employee productivity tools.

From regulatory catch-up to regulatory readiness

The challenge is therefore not simply compliance with the AI Act; it is building the capability to supervise AI before every requirement is explicitly articulated. Recent industry discussions have consistently highlighted that good AI governance starts with clear accountability, documented use cases, explainability, auditability and effective oversight structures rather than waiting for specific supervisory findings.

The most effective compliance functions are treating AI today in the same way they once treated outsourcing, cybersecurity or model risk: as an enterprise-wide governance issue. They are creating inventories of AI use cases, assigning accountable owners, establishing risk classification frameworks and ensuring that every system can be explained, challenged and reviewed. Similar governance themes are already emerging in AI programmes across European banks, including central inventories, risk assessments, approval workflows and clear ownership structures.

A practical example is the use of AI to process customer documentation, support credit decisions or improve transaction monitoring. Even where such tools deliver better efficiency and customer outcomes, compliance should already be asking fundamental questions.

For example:

- Who owns the use case?
- What data is being used?
- How can the output be challenged?
- What controls exist if the model behaves unexpectedly?
- Can the decision be explained to a customer?

> “The lesson from finfluencers is that compliance should focus less on identifying specific future regulations and more on identifying future behaviours.”

Institutions that can answer those questions consistently will be far better prepared than those still trying to identify where AI is being used.

The finfluencer lesson: spotting risks before the rulebook arrives

The rise of financial influencers (“finfluencers”) provides a useful warning for compliance professionals. Consumer financial content exploded across social media long before most regulatory frameworks were adapted to address it. Banks, investment firms and supervisors suddenly faced new risks involving financial promotions, suitability, conflicts of interest and consumer protection. The underlying risk was not entirely new, but the delivery channel changed faster than the rules.

AI presents a similar challenge, but at a much larger scale. The lesson from finfluencers is that compliance should focus less on identifying specific future regulations and more on identifying future behaviours. Risks rarely emerge because regulation is absent; they emerge because technology changes how existing risks materialise.

For example:

- Generative AI can create marketing content at scale.
- AI agents may interact directly with customers.
- Employees can build informal AI tools without central oversight.
- Third-party AI models can influence decisions without full transparency.

These developments are creating new forms of conduct, reputational, consumer protection and operational risks long before detailed supervisory guidance is available.

The key capability is therefore horizon scanning. Compliance functions need mechanisms to identify emerging technologies, assess their potential regulatory implications and engage early with business and technology teams. This requires moving beyond traditional policy interpretation and becoming active participants in innovation discussions.

The institutions that successfully navigated emerging risks such as social media, crypto-assets and finfluencers were usually those whose compliance teams were involved early, not those who arrived after the first incident or supervisory inquiry.

What a well-prepared compliance function looks like

A mature AI-ready compliance function does not attempt to own AI. Instead, it ensures the organisation can govern AI effectively. Three characteristics stand out.

First, clear governance and accountability. Every AI use case should have a business owner, defined responsibilities and an agreed governance process. AI should not sit in a grey area between technology, risk and business teams. Effective oversight committees, approval mechanisms and reporting channels are becoming essential components of AI governance.

Second, consistent inventory and risk assessment. Many organisations still cannot confidently identify all AI tools operating across the enterprise. Shadow AI is rapidly becoming the equivalent of the “shadow IT” problem of previous years. Leading organisations maintain central inventories, classify AI use cases according to risk and apply controls proportionate to that risk.

Third, skills and challenge capability. One of the biggest gaps today is not technology but expertise. Compliance professionals do not need to become data scientists, but they do need enough understanding to challenge assumptions, assess risks and engage credibly with technical teams.

The organisations most likely to succeed are investing in AI literacy across all three lines of defence. Industry experience increasingly suggests that AI governance failures arise not because controls do not exist, but because business, technology and control functions fail to communicate effectively.

Ultimately, getting ahead of AI is not about predicting every future regulation. It is about building governance that remains effective even when technology evolves faster than the rulebook. Compliance functions that establish accountability, transparency and oversight today will be in a position to shape the conversation with supervisors tomorrow, rather than merely reacting to it.

Why AI is opening a marketing compliance skills gap in banking

By João Pedro Tavares
Programme Director
Universidade Católica Portuguesa



The financial landscape is undergoing profound change. Incumbents and new players alike are rapidly evolving their capabilities through AI, automation, and advanced technologies. Customers now rely on multiple sources of information, expect seamless digital experiences, and demand clearer understanding of financial products. At the same time, banks face talent scarcity in critical areas, rising cyber and fraud risks, hyper-personalisation challenges, trust erosion, and new supervisory requirements.

The shifting landscape of distribution and discovery

Traditional banks face increasing disintermediation of their value chain and weakening direct relationships with clients, even as they continue to push extensive product portfolios. AI is reshaping how banks advertise and sell products, and financial services are becoming embedded inside non-bank apps, reducing direct customer visibility. As a result, marketing claims must remain accurate even when distributed through partners and third parties.

Customer discovery no longer happens on bank-owned channels. Clients arrive with pre-formed opinions shaped by forums, comparison platforms, partner ecosystems, AI-generated answers, and social media. Therefore, compliance must extend to every environment where the brand appears, even those not directly controlled by the bank. Banks must monitor and validate third-party content to avoid misleading claims or regulatory breaches.

Efficiency gains leading to a new capability gap

From an industry perspective, AI is currently delivering the greatest impact in efficiency and productivity, especially in high-volume operational areas. These gains can help fund more sophisticated AI use cases in customer experience and relationship management.

The deeper transformation, however, is becoming truly AI-first and data-driven, changing not only processes but also how customers are served. Customer-facing use cases require caution, particularly around trust, privacy, fairness, explainability, and regulatory compliance. The goal is to shift from generic communication to contextual, useful engagement while ensuring accuracy, tone, and compliance.

> “Marketing knows the channels but not the regulations; compliance knows the rules but not the technology; data teams know the models but not the customer impact.”

Banks must evolve from product-oriented structures to process-driven models focused on customer service and care. Silos are no longer viable. Marketing knows the channels but not the regulations; compliance knows the rules but not the technology; data teams know the models but not the customer impact. Improvement is needed, especially as many institutions struggle with fragmented data while only a minority use their marketing data effectively.

Traditional marketing and compliance roles are insufficient in an AI-mediated environment. Banks must build teams that combine AI literacy (model behaviour, bias, explainability), regulatory fluency (GDPR, DSA, EU AI Act, financial promotion rules), digital marketing expertise (targeting, attribution, content automation), and risk & governance skills (model risk, conduct risk, disclosure risk).

Governance, ethics and building capability

New governance roles are required to supervise AI-generated marketing outputs. AI now automatically produces product descriptions, campaign profiles, segmentation models, personalised offers, and chat-based sales interactions.

This raises critical questions:

- Who validates AI-generated content?
- Who ensures fairness in targeting?
- Who monitors hallucinations or misleading claims?
- Who owns accountability when AI makes a mistake?

AI accelerates market cycles, but compliance must keep pace. Banks need a culture where compliance is embedded early, not added late.

Regulators also play a role. They are increasingly concerned about AI-generated performance claims, automated advice, biased algorithms, and certification of AI systems. Oversight must be strengthened without slowing innovation or adoption of new opportunities.

Ethics must be deeply embedded in culture. Banks must define what “responsible personalisation” means, how to avoid manipulative nudging, how to ensure transparency in automated decisions, and how to protect vulnerable customers from algorithmic pressure.

Personalisation should create clear value and use data in ways customers reasonably expect. If it goes too far, customers feel monitored or pressured. They should understand when AI is being used and be able to escalate to human support. Trust is critical in banking, requiring a balance between relevance and privacy.

Building the skills and culture banks need

Compliance must evolve from document review to digital ecosystem oversight, understanding algorithmic targeting, dynamic personalisation, real-time ad platforms, influencer marketing, and AI-generated customer interactions.

Certain human skills remain irreplaceable: integrity, curiosity, and critical judgment. People must stay open-minded, continuously learning new tools and trends while maintaining operational excellence and customer care.

Without unified data, both marketing effectiveness and compliance monitoring suffer. Compliance becomes a trust signal - accurate, consistent, understandable, and omnichannel.

Education also plays a role. Universities should embrace future-focused learning and strengthen collaboration with industry, balancing theory with practical outcomes.

Banks must invest in AI literacy for marketing teams, regulatory literacy for digital teams, cross-training between marketing and compliance, board-level education on AI governance, and scenario-based training on AI mis-selling risks. Capability building becomes a strategic differentiator.

Building a culture that adopts AI responsibly while oversight keeps pace

By Inês Caria
Chief Compliance Officer
Banco Santander Totta



People often ask whether the financial sector is ready for artificial intelligence. Yet AI is already inside our institutions. The better question is whether our capacity for oversight is evolving as quickly as the technology it must oversee.

For years, we regarded AI as a support tool for internal use. Today it interacts directly with customers, shapes commercial recommendations and is beginning to carry out actions in our systems. This is not a change of degree. It changes the nature of the risk.

Compliance will play a critical role in managing the risks that come with AI. These include data quality and integrity, privacy and data protection, bias and explainability, security, risks specific to generative AI, regulatory compliance, reputational risk, ethical and conduct risks.

Compliance will need to rely even more on a risk-based approach, proportionate controls and oversight that is independent, effective and continuous. That means questioning and assessing use cases, challenging and monitoring controls, testing compliance and helping to build robust governance.

However, there is no strong compliance culture without a robust risk culture in the first line of defence. Control functions can set policies, design controls and provide the necessary challenge. But they cannot be present in every customer conversation, every commercial decision or every new use of AI. Those who take on or generate risks in the first line are also responsible for identifying, assessing, managing and controlling it. When people see that responsibility as their own commitment rather than an obligation imposed from outside, Compliance is no longer seen as something external to the business. It becomes a partner.

That is why a culture of oversight cannot rely on controls alone. It must live in people and in the design of the systems themselves. Culture is what people do when no one is watching.

AI as weapon, target and dependency

AI is now a weapon, a target and a dependency, all at once. It is a weapon in the hands of

> “Adopting AI responsibly is not only about controlling its risks. It is about creating the conditions to innovate with confidence, without losing sight of what we must protect.”

those who attack us, through voice and video deepfakes, synthetic identities and fraud that adapts to the victim in real time. It is a target, because the models themselves can be manipulated. And it is a dependency, as shown by the concentration of suppliers and by AI use that has never been inventoried.

In banking, all three dimensions strike the same target: trust. Fraud and impersonation are not confined to a single system. Their effects spread through relationships with customers and suppliers and can reach critical services. Once trust is damaged, the problem is no longer technical. It becomes reputational, regulatory and, ultimately, systemic.

The role of compliance: from legality to legitimacy

Complying with the AI Act is only part of the answer. Regulation alone cannot settle every question that AI raises: whether a decision is fair, whether a recommendation genuinely serves the customer, or whether an interaction respects their autonomy. The question is no longer simply ‘can we do this?’. It is also ‘should we do this, and can we explain it to the customer, to internal audit or to the regulator?’.

In practice, our role is to make ethics actionable. Principles such as transparency, fairness and accountability only change behaviour when they are turned into clear boundaries. For example, AI should not be used with sensitive data without safeguards. Decisions with a material impact should not be automated without traceability, validation and human oversight.

One dimension is non-negotiable: protecting customers, particularly the most vulnerable. Models that influence decisions must be explainable, tested for bias and monitored with the same rigour we apply to financial products. We must also know where legitimate personalisation ends and where practices begin that exploit vulnerabilities or could lead to mis-selling. AI gives us an unprecedented capacity to influence. For that very reason, the ethical bar must be set higher.

Culture before control

The most important discussions about AI are not technical. They are cultural. Teams that adopted AI early quickly grasped their potential. Learning to use it responsibly takes longer. Without clear guidance, some people end up using AI without safeguards and others stop using it altogether out of fear. Neither outcome serves the institution.

Articles and studies on the gains from AI keep multiplying: productivity, speed, analytical power. Understanding what we stand to gain matters. It matters just as much to reflect on what we do not want to lose. Compliance is the function that should bring this question to the table. A responsible organisation must decide how far it is prepared to pursue efficiency when efficiency conflicts with other values. These limits are not meant to hold back innovation. They protect what is essential: the integrity of ethical decisions, the organisation’s purpose and values, sound conduct and customer protection.

Compliance also raises another discussion, perhaps the most demanding of all: people. Adopting AI responsibly depends as much on the people who use it as on technology. Training,

reskilling and upskilling are no longer merely HR initiatives. They have become matters of governance. According to the World Economic Forum's Future of Jobs Report 2025, employers expect 39 per cent of workers' core skills to change or become outdated by 2030. The task is not simply to teach people how to use new tools. It is to preserve and strengthen what no machine can replace: professional judgement, critical thinking and the ability to say no when something is not right. Leaders, too, have a decisive role to play here.

What AI can do for compliance

Compliance has a great deal to gain from AI. In preventing money laundering and fraud, it can give significant support to alert analysis, investigations of suspicious transactions, monitoring, controls and KYC processes. It can also help identify new legislation that affects business and carry out the related gap assessment. It can support the implementation of policies, the review of marketing materials and the assessment of reputational risk.

Our work rests on large volumes of information: transactions, alerts, documents and ever-changing regulations. AI enables us to handle that volume more intelligently and with greater analytical depth.

In risk management, AI also helps us move from reacting to anticipating. It separates the noise from what really matters, prioritises alerts and picks up patterns that sample-based reviews miss, such as signs of mis-selling, unequal treatment of customers or recurring complaints. We can spot risk earlier and act before it materialises. AI also brings greater consistency to the quality of decisions when cases rest on the same or similar facts.

In oversight, AI gives us wider reach and greater rigour. Rather than looking back at samples, we can monitor what happens in the first line on a continuous basis. And when systems are designed with proper logging, traceability and explainability, it becomes easier to evidence, challenge and audit what has been analysed and recommended.

Compliance by design and by default

Compliance cannot be a check at the end of the process. It must be part of the system itself. Compliance by design means that control, traceability and customer protection are built into every solution from day one.

Compliance by default starts with the fact that most people never change the settings they are given. If the default configuration is secure, verified and logged, doing the right thing becomes the path of least resistance. When Compliance is embedded in this way, it no longer holds innovation back. It is what makes innovation sustainable.

AI will continue to advance and to transform the way we work, make decisions and engage with our customers. Our challenge is to ensure that our culture of responsibility and our oversight capabilities keep pace. Adopting AI responsibly is not only about controlling its risks. It is about creating the conditions to innovate with confidence, without losing sight of what we must protect. It is that culture that will determine whether we can make the most of technology without jeopardising what allows us to go on earning our customers' trust.

Conclusion

These essays, read together, show from a number of vantage points that AI is advancing rapidly across marketing and communications, while the governance, oversight and institutional capabilities around it are still lagging behind. Whether financial services closes that gap will determine how effectively, responsibly and sustainably AI is adopted across the sector.

The first theme to emerge across the contributions is that accountability can never move to AI systems themselves. As more of what a customer sees and is offered is shaped by AI, and as the process becomes more autonomous, the question of where accountability sits becomes more complex. That has to be answered inside financial institutions, across their supply chains, and in the frameworks that hold them to account.

The second is that the conception of oversight itself has been transformed. Compliance functions were built to review a single message before it was sent; they now have to govern a system that produces thousands of personalised variations. Fragmenting media channels and the rise of finfluencers push more of what customers see beyond the firm's direct control. Firms will need to rebuild their governance around this reality, replacing point-in-time approval with continuous assurance.

Third, capability, judgement and trust must become priorities. Staff and customers alike need the skill and confidence to challenge what these systems produce. This means ensuring that, as these capabilities advance, the customers least able to navigate an AI-mediated market are protected. That is a responsibility for firms as much as governments.

AI has already changed financial marketing for the better in many ways, and it will continue to do so. Used well, it can make communications clearer and tailor them to what a customer genuinely needs, while turning everyday marketing into an opportunity to improve financial understanding. Yet the same capabilities can facilitate manipulation, exploiting behavioural biases and persuading customers through targeting they cannot see. Effective regulatory oversight has never been more necessary, and this places an obligation on supervisors to adopt these tools as readily as the firms they regulate, so that supervision itself is not left behind.

The challenge, then, is to build the institutional capacity to use AI well, where technological capability is matched with equally sophisticated governance. Firms must know where AI is being used and understand how it shapes customer outcomes, so they can intervene when those outcomes begin to diverge from what was intended. Done properly, governance should not sit in opposition to innovation. It is what allows firms to deploy these capabilities at scale without sacrificing the standards on which the industry - and public trust in it - depends.

Trust has always been the foundation of financial services. In the age of AI, it takes on an altogether different character. It becomes even more central, not just to the relationship between firm and customer, but to the functioning of the financial system itself. The capabilities used to reach and serve customers must be turned towards protecting them.

The participation of so many leaders from across Portugal's financial services industry in this report is testament to their willingness to confront that challenge. We look forward to continuing this important conversation.

Sam Green
Senior Policy & Partnerships Lead
zango



About Zango AI

Zango is the AI compliance layer for financial services. The platform uses AI agents, grounded in regulatory intelligence and human oversight, to help financial institutions monitor regulatory change, assess compliance gaps, and review marketing and product launches against regulatory rules. Founded in 2024 by Ritesh Singhania and Shashank Agarwal, Zango AI serves clients across banking, insurance, payments, and digital assets, with offices in London, Lisbon, and Bengaluru.

Regulatory Change Management:

Zango monitors regulatory sources, extracts and verifies the obligations inside each change, and routes what matters to the team that owns it. Group compliance gets one view of every open gap in a firm's policies, controls and norms.

Product Governance:

Zango scans product features, terms and disclosures against the obligations that apply in each market. When a regulation changes, affected products surface immediately, so the register stays current instead of ageing between reviews.

Marketing Compliance:

Zango reviews financial promotions against applicable rules before publication, cutting approval time while leaving audit-ready evidence for each decision. Published assets stay monitored, so a regulatory change traces straight through to the promotion it affects.

Learn more at www.zango.ai

zango

www.zango.ai

MILLENNIUM BCP	MIGUEL MAYA
FIDELIDADE	ROGÉRIO CAMPOS HENRIQUES
IMGA	EMANUEL SILVA
FORMER BOARD MEMBER, ASF	MANUEL CALDEIRA CABRAL
BANCO SANTANDER TOTTA	INÊS CARIA
BPI	ANDRÉ PRAZERES HENRIQUES
CAIXA GERAL DE DEPÓSITOS	RAQUEL VILA VERDE
CATÓLICA LISBON	JOÃO PEDRO TAVARES
FIDELIDADE	FRANCISCO LUÍS ALVES
PÉREZ-LLORCA	INÊS PALMA RAMALHO
AGEAS PORTUGAL	PEDRO OLIVEIRA MACHADO