FOCUS

AI in Retail Banking: The Ultimate Challenge.

Christophe Baniol
Nicolas Miart

Circle strategy

Circle strategy

AI in Retail Banking: The Ultimate Challenge.

Authors: Christophe Baniol and Nicolas Miart

September 2024

ACRONYMS

AI	Artificial intelligence
LLM	Large language model
GPU	Graphics processing unit
GDPR	General Data Protection Regulation
AI Act	Artificial Intelligence Act
LPU	Language Processing Unit
TPU	Tensor Processing Unit
SLM	Small language model
AML	Anti-money laundering
NBO	Next Best Offer
CRC	Customer relations centre
KYC	Know your customer
BO	Back office
Q&A	Questions and answers
WMA	Wealth management advisor
EPD	Energy performance diagnosis
ADR	Automatic document recognition
BU	Business unit
KPI	Key performance indicator
POC	Proof of concept
NBI	Net banking income
FTE	Full-time equivalent
ROI	Return on investment
SaaS	Software as a service

CONTENTS

AI: A MAJOR DIFFERENTIATING FACTOR FOR RETAIL BANKING

BANKING MEETS ALL THE CRITERIA FOR MAXIMUM 'AI-SATION'.

After fifty years in the wilderness during the 20th century following the Turing test and some interesting but targeted breakthroughs in the early 21st century, AI entered a new era at the end of 2022 with the public launch of ChatGPT, promising exceptional advances across an immense field of application. The algorithms are mature and increasingly accessible to everyone (see the proliferation of open-source libraries, including on LLMs), and the necessary computing power (GPU) is available and constantly increasing (Moore's law continues to hold). Its cost is becoming more affordable by the day, particularly with a move to the cloud (pay-per-use billing). The regulatory framework is clear or in the process of becoming so (GDPR and AI Act). Finally, acceptance by society is reaching new levels, liberated by the ChatGPT revolution. The only downside is the energy consumption of these models, which are inherently very energy intensive. The industry is working on this, whether through the development of specialised chips (LPUs and TPUs) that consume less energy than current GPUs, or through the use of SLMs for case of targeted use.

AI is therefore set to have an impact on every sector. The extent to which it does so will, however, vary.
For a number of reasons, those sectors **particularly reliant** on **data, humans, digital technology** and that are intrinsically **complex** will be the most affected. **This is precisely the case in the banking industry**. As far as data is concerned, no other sector can boast such a large, varied, and relevant mass of data. What is more, because of the intangible nature of its products, banking is an industry that can be fully digitalised and is already digitalised to a large extent, which will make it easier for AI to be embedded in processes. From a human point of view, this sector involves millions of customers and tens of thousands of employees, all potentially 'augmented' by AI in the future. Finally, the complexity factor of the banking industry is high. Whether in terms of actual complexity (see the regulations) or complexity as perceived by customers (undesirable goods and an uninviting subject matter), there is much to be simplified. This is a task that AI, and in particular generative AI – a veritable engine for digesting complexity –, performs better than anyone or anything else. Scoring top in all the indicators in our 'AI-sation' index, retail banking is therefore set to experience the AI revolution in full.

In short, **AI will be THE next transformative technology for the financial sector, just as digital technology has been since the beginning of the 21st century**. There is one difference, however: the digital revolution has taken twenty years. The AI revolution will be much faster.

AI WILL BECOME THE KEY COMPETITIVE FACTOR FOR THE BANKING SECTOR BY THE END OF THE DECADE.

The AI shock will affect an industry that is, despite its civilised appearances, the scene of a merciless war of competition. While relatively slow, often resembling a tactical war of position, the manoeuvres have actually been rather violent. Over the years, they have claimed a number of victims: regional banks that no longer had the necessary critical mass (Crédit du Nord and, before it, the CIC Group), foreign banks unable to exploit the cross-border synergies they had hoped for (Barclays and HSBC), online banks that struggled to win over a quality customer base (Orange Bank, ING, and Ma French Bank), not to mention players that failed to survive a number of systemic crises (LCL and Dexia, in their day). The consolidators themselves have paid a high price for these clashes, which have weighed heavily on their revenues. Despite the consolidation that has taken place in recent years, competition is set to continue unabated in the years ahead. AI will open up a new front in this merciless battle, one that will take the form of **a war of movement** based on the exponential pace of technological progress. And this will be all-out war. It will affect all products in the range, all professions (front, back, risk, marketing, etc.), all markets (individuals, professionals, and businesses), and all performance drivers (productivity, revenue, risk, customer and employee experience). In short, retail banking in 2030 will either be augmented by AI … or it won't.

Without claiming to be exhaustive, we present below what we consider will be the major impacts on the **banking profession** (without dwelling on those linked to functional support, such as the human resources, finance and information systems departments, which are not specific to the banking sector).

Figure 1. **AI impact matrix**

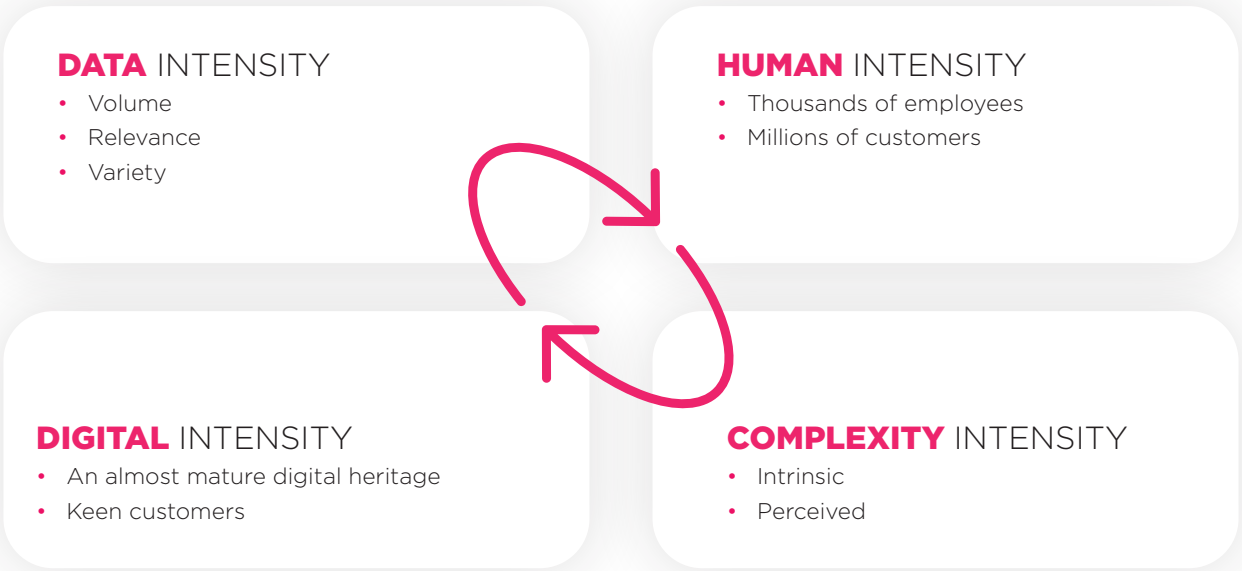

THE USE CASES WILL APPLY TO ALL PROFESSIONS AND MODELS

THE USE CASES WILL ADDRESS EVERY LINK IN THE BANKING VALUE CHAIN.

After exhaustive mapping, it will be clear that all aspects of performance will be affected. Predictive AI has already addressed many productivity and revenue issues. Generative AI will also address customer and employee experience issues which will, in turn, have an impact on productivity. Figure 2 below illustrates the interweaving of AI generations and their impacts, classified by type of issue.

Notwithstanding the strategic and technological contexts specific to each institution, we believe that the priorities in terms of AI initiatives should be based on a dozen or so areas of use and according to a Maslow pyramid of AI use cases, as represented in Figure 3 (see next page).

Figure 2. **AI, technologies, and banking issues**

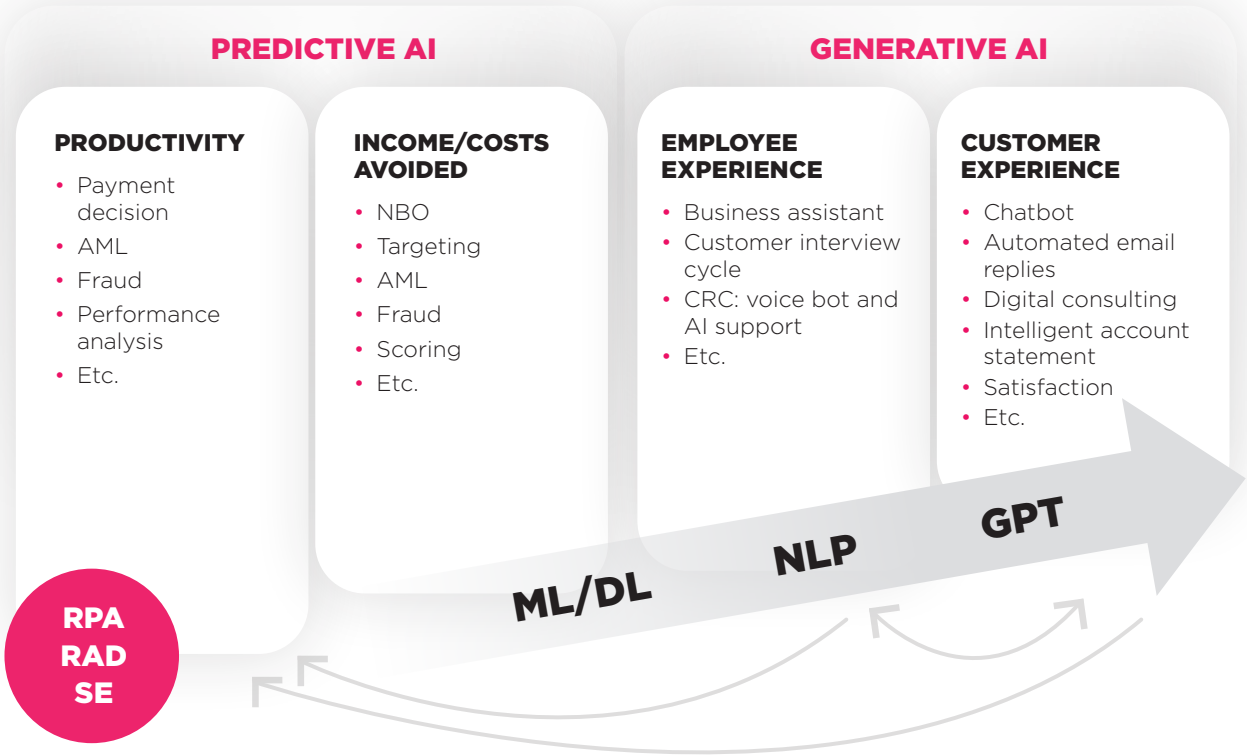

Figure 3. **Maslow pyramid of AI use cases**

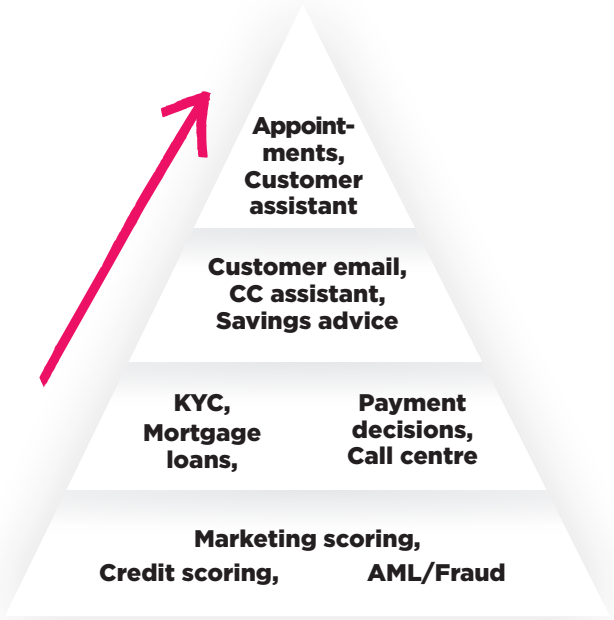

Many of these macro-cases of use will give rise to the development of in-house solutions. That said, specialist off-the-shelf solutions are likely to emerge rapidly. The question of *make or buy* will become increasingly acute in the years to come, as the market offering becomes richer. The answer will not be a global LLM that covers everything, but rather the development of solutions dedicated to a specific need in a sector that will gradually be addressed by a number of start-ups.

THE DISTRIBUTION MODEL WILL BE IMPACTED BY THE EMERGENCE OF CONVERSATIONAL BANKING.

Today, omnichannel banking is limited by the expertise available across the various channels. While all products can be sold in branches, not all can be sold via digital channels or call centres. This is essentially due to their complexity. The direct channels of today's omnichannel bank are, for example, incapable of providing the level of advice needed to sell any financial savings product that is even slightly sophisticated.

With the exception of the agency channel, channels are specialised for precisely this reason. Some simple products, labelled 'omnichannel', such as consumer credit and car insurance, are struggling to go fully digital by more than 10 to 15%. This is due to the lack of qualitative dialogue possible in a digital funnel. For example, a customer seeking reassurance on the question of possible early repayment penalties when taking out a consumer loan will not find the answer in their transformation funnel. **Today's digital channel is dry, devoid of that grey area of understanding and conviction that is so crucial to the act of selling**. With a genuine conversational agent powered by generative AI and integrated into the process, this hurdle will be overcome. This is where the digital channel will come into its own.

Generative AI should therefore enable us to move to a 3.0 distribution model: a conversational bank to take the place of the omnichannel bank built up over the past twenty-five years thanks to the advent of digital technology. Within this new bank, we will be able to engage in high-quality dialogue via **equally omniscient channels**. In effect, each channel (advisor, digital, and call centre) will be augmented by a powerful conversational agent that will give it a much wider reach than it has today, if not total.

We can imagine this evolution culminating in **an advisor who is not simply augmented but accompanied by a personalised 'virtual teammate'**, known and recognised by the customer. This avatar will respond directly to email requests. Initially, it will process the simplest requests with the advisor's approval. It will then specialise in more complex applications, without the need for human intervention. It will be ever-present on all digital sales and after-sales screens, ready to respond to the customer at all times. Powered by an ultra-high-performance natural language interface, it will also be able to respond directly to certain requests over the phone, or at least categorise them before directing the customer to the most qualified human (in the branch or in the call centre), whom it will then help with the customer interaction.

Figure 4. **Evolution of distribution models**

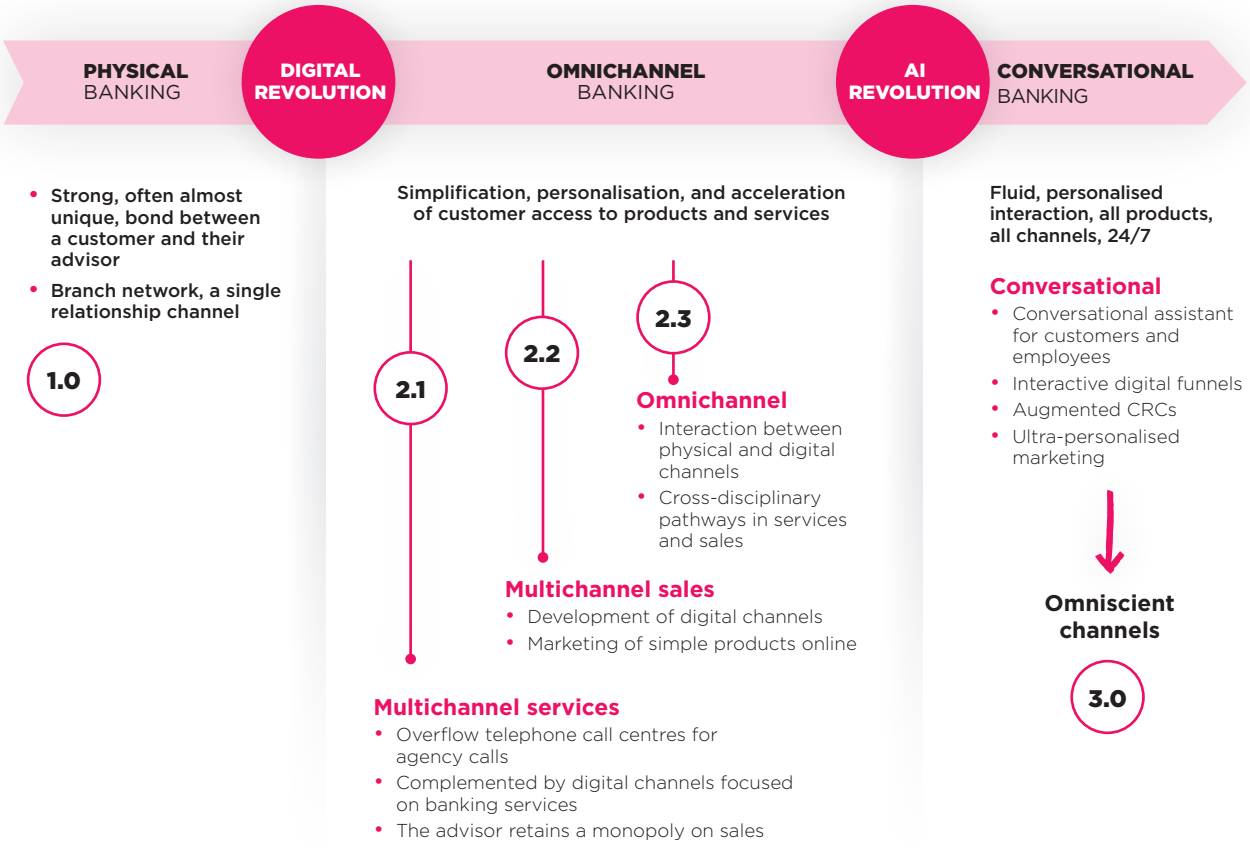

Also designed to support the advisor, this avatar will answer even the most technical questions, finding an efficient way through the ever-increasing documentation. It will also offer **interview preparation sheets that really tell the story** of past interactions between the customer and the bank and the opportunities they present today, without being limited to a 'flat' vision of the customer's identity and products. Finally, it will record the interviews and provide an automatic summary report, saving the advisor a considerable amount of time and enriching their knowledge base for the next sales round.

It is important to note that this AI-powered virtual advisor will sooner or later lead to a **rethink of the digital pathways** themselves. The sequence of screens we see today, the completion of which is the responsibility of the customer, will gradually be replaced by interfaces that perfectly simulate a dialogue with an advisor. When customers enter a digital pathway, they will be met by a virtual advisor instead of the current data input interfaces.

The customer will explain their needs orally, respond to the avatar's requests for clarification and the avatar will 'enter' the answers into the system. The conversational agent will, of course, respond to any questions and objections before presenting the customer with a summary of their request and offering an electronic signature to complete the transaction. Equipped with a

sentiment analysis module that will enable it to tailor its message to meet customer expectations as closely as possible and trained in the best commercial practices desired by the company, its impact will be immense. This is because contrary to what we have come to believe through our extensive use of office automation, **human beings are not designed to wield a mouse**, click on a drop-down menu, or fill in a form. They are born to speak, see, and hear. **In the future, virtual advisors will enable us to respond once again to these fundamental human needs. Virtual though they may be, they will offer a very human response**.

A final element to be integrated into conversational banking is **the management of customer complaints and feedback**. This mass of textual data (verbatim feedback from on-the-spot 'appointment surveys', Google reviews, social network posts, and formal complaints received through various channels) is currently poorly exploited in a number of ways. On the one hand, insufficient responsiveness disappoints the customer. On the other hand, the absence of statistical analysis within a reasonable timeframe prevents the implementation of a continuous improvement cycle. Generative AI is changing the game in this area, with its ability to understand the subject in detail, identify feelings, and classify in a relevant way. It can, for example, be used to spot a faulty product process (linked to an IT anomaly, a BO backlog, etc.) at an early stage, or to identify an agency or group that is collecting very negative or positive feedback.

For all these reasons, AI is ushering in a new era in customer relations.

CONVERSATIONAL BANKING WILL ALSO BRING ABOUT A REVOLUTION IN CUSTOMER INFORMATION.

Advances in AI offer a fantastic opportunity to make information sent to customers more meaningful.

This primarily concerns day-to-day banking. Banking transactions have always been punctuated by periodic account statements. Thanks to digital technology, these now flow continuously into applications and web spaces, often in real time. In both cases, the customer comes face-to-face simply with a list of operations, without any perspective. Most banks have tried to make these statements more meaningful with budget management functions that categorise customer spending over flexible time horizons in a semi-automatic way. However sophisticated these functions may be, they are perceived as complex and are rarely used. Generative AI will revitalise the account statement and 'current accounts' sections of applications in the form of an 'intelligent account statement'. Each month or quarter, a transaction analysis report will be sent to the customer. This report will be highly personalised, highlighting underlying trends and significant transactions, pointing out differences in behaviour with peers (in the professional customer sector, for example), and warning of potential risks (future cash flow problems, for example). Lastly, this report will offer suggestions for optimisation and personalised advice (not forgetting, with regard to the digital component, calls to action along the path to purchase and Q&A).

The same progress will also take place in the area of financial savings. After the initial investment proposal, the customer is usually left to fend for themselves (except in the case of discretionary management contracts). Whether the financial environment changes or whether their returns rise or fall, the bank takes no action, at least not until the next date in the sales calendar.

As a result, customers are neither alerted to opportunities nor warned if their positions are exposed. They simply benefit from reporting enabling them to analyse the performance of their investments and draw their own conclusions. **In the future, generative AI will be able to transform these financial savings reports** by enriching them with an analysis section followed by duly justified allocation proposals. These analyses

will be carried out at regular intervals (quarterly, for example) or triggered by exceptional market events. This reporting will boost portfolio turnover by improving performance and, in so doing, increase the bank's revenues.

Finally, in wealth and private banking, AI could considerably **speed up the production of personalised wealth studies and summaries** delivered to customers. These are based on a fairly stable fiscal and regulatory framework and a standard structure (a summary of the customer's personal situation, assets, projects and investment options). There is therefore scope for automating the production of these summaries (under the ultimate control of the private banker) so that the wealth management advisor or private banker can concentrate solely on tailored levers.

CUSTOMER-FACING PROFESSIONS WILL EVOLVE UNDER THE PRESSURE OF NEW EXPECTATIONS.

In recent years, **banks have delegated many of the tasks traditionally carried out by advisors to tools**. They have done this for both regulatory and efficiency reasons. Whereas in the past advisors were autonomous in their sales process and even in their risk analysis, they are now almost systematically supported or supervised by the tools they use.

Few sales transactions escape an ultra-guided process involving a succession of screens ensuring that the information required is entered in an orderly fashion, that in-house commercial practices are properly monitored, and that compliance and risk requirements are met. What was for a long time the prerogative of mortgage underwriting – given the complexity of the product – has extended into savings (driven by the requirements of the duty to advise), protection (driven by the expansion of the range and its complexity), and management of the process of entering into a relationship with customers (under the impact of increasingly demanding KYC requirements). Against this backdrop, the advisor's know-how and technical and commercial 'tricks' have been undermined. The most spectacular case is undoubtedly that of wealth management advisors (WMAs), who in the past designed investment strategies for their clients but today simply 'pass on' a stereotyped client discovery and asset allocation proposals that are strictly controlled by the machine and can only be subject to marginal modifications.

Conversely, a new area of expertise will need to be developed: the ability to interact with the AI tools made available to the advisor. As Ethan Mollick rightly explains in his reference work *Co-Intelligence*, these knowledge workers will have to use AI in 'centaur' or 'cyborg' mode. In the first case, the AI does part of the work, while the professional takes care of the rest (for example, my avatar suggests an email response to send to a customer, which I personalise before sending). In cyborg mode, the professional interacts back and forth with AI before arriving at a satisfactory solution (for example, simulations of investment strategies).

As well as know-how, the arrival of generative AI is also calling into question **the skills associated with knowledge as such**. Product knowledge, for example, is set to become part of our general knowledge, as access to an exhaustive description of product capabilities will soon be at the fingertips of the advisor or customer alike, offering an unrivalled level of completeness and relevance.

Having an opinion on the future performance of the financial markets as part of a savings strategy will be just as incongruous, as access to the bank's best analyses, also contextualised, will be immediate and less risky in regulatory terms. These developments will even affect private banking, where a generative AI-powered robot will very quickly be able to propose a wealth strategy.

Valued less and less for their knowledge and expertise, what will be left for the advisor? **More than ever**, everything will be about people skills! For example, they may have to rephrase an explanation given by a new-generation conversational agent that a customer has not understood. Hence, the advisor will personalise the reassurance that many customers will continue to value. They will also have to intervene at certain 'key moments', where the added value of human interaction will be decisive. Some will be rewarding such as discovering the customer and their needs before the sale, while some will be more difficult such as in the after-sales phase (in the case of inheritance, for example). The advisor will then have to embody the brand through their attitude and customer skills. Obviously, they will also have to be very responsive when dealing with malfunctions, when they occur. In an 'AI-ed' bank, the advisor's knowledge and know-how are augmented and their interpersonal skills are given a higher profile. More human than ever, **and a repository of the brand's 'emotional added value'**, responsive when necessary, they will play a key role in brand loyalty by remaining empathetic during the sales process and reassuring in after-sales. They will interact less frequently with the customer, but these interactions should have a strong emotional impact. These developments need to be **taken into account when recruiting**. From now on, recruitment needs to be directed towards **profiles with strong interpersonal skills, to the detriment of technicians**.

SALES MANAGEMENT WILL HAVE TO TAKE THE NEW PARADIGM IN WHICH ADVISORS WILL BE OPERATING INTO ACCOUNT.

Many sales managers currently derive their managerial legitimacy from their expertise. In this sense, they are mentors for younger staff and resources for the more experienced and in more complex cases. **If AI delivers on its promises, employees' reference point for expertise will become their virtual teammate rather than their manager. This will mean refocusing on the fundamentals of sales management:** setting the course, collective and individual sales promotion, fair objectification and evaluation, active listening to individual situations, and oversight.

The job of sales manager will be impacted by AI in no small way. Sales managers today spend a considerable amount of time analysing the performance of their business unit and their staff. Working with individual customers is a good example of this complexity, as it involves setting targets for business, quality, compliance, risk and financial results for a range of 30 product lines to be marketed. In this area, **generative AI could offer off-the-shelf 360-degree analysis of the performance of a sales** unit (branch, group, or region) or an employee, in natural language. This summary would cover all performance levers (commercial, financial, risk, compliance, from a multichannel perspective) and would include suggestions for improvement. These standardised documents would form the basis for managerial discussion at each level of the structure as part of the various performance reviews. More prosaically, it would provide each sales manager with a clear, well-argued, and structured view of the strengths and weaknesses of their sales outlet. The expected gains take several forms:

- Time-saving in performance analysis by managers;
- Time-saving in preparing performance reviews;
- Standardising the performance culture within a network.

BANKING MARKETING WILL MOVE TOWARDS ULTRA-PERSONALISATION.

Given the sheer volume of data at its disposal, banking is a highly advanced industry when it comes to segmentation. In this area too, AI is set to change the game. Traditional value segmentation (customers segmented according to what they bring in) will remain

useful for distributive segmentation, allocating customers in the portfolio to advisors with expertise commensurate with the customer's value. This segmentation must, however, be based on stable criteria, as a network cannot be continually resized, nor a customer change advisor too frequently. It is also limited in terms of the number of segments, as it is not feasible to multiply the number of professions in agencies or call centres.

Conversely, **AI will revolutionise relationship marketing segmentation**, which is used to organise requests pushed to the customer, whatever the channel used. We're talking here about targeting, which determines the push of a particular offer to a particular customer. Today, these target types are designed by marketers on the basis of common sense criteria (non-holding, borrowing or saving capacity, risk profile, customer value, socioprofessional category, etc.). These models are stable over time and have the advantage of being easy to explain. As a result, they do not perform very well (campaign success rates are usually well below 1%). Tomorrow's models will escape the marketers: they will be delegated to AI. Based on machine learning, they will combine criteria far too numerous to be imagined and manipulated by humans. They will be fine-tuned by the machine almost in real time and will be adjusted for each campaign.

They could be enhanced by the use of data external to the bank such as the EPD rating according to address, when targeting for an energy retrofitting loan. **In this sense, for banks, AI will mark the end of marketing segmentation as it currently stands and will enable the advent of the long-promised one-to-one marketing**.

In addition, while predictive AI will enable us to target better and better, generative AI will enable us to **personalise the message to an extreme extent**. The push marketing customers will receive tomorrow will not only be more accurate in nature and timing, it will also be more convincing in form. Generative AI will be able to write an ultra-personalised message explaining to

Figure 5. **The impact of AI on segmentation**

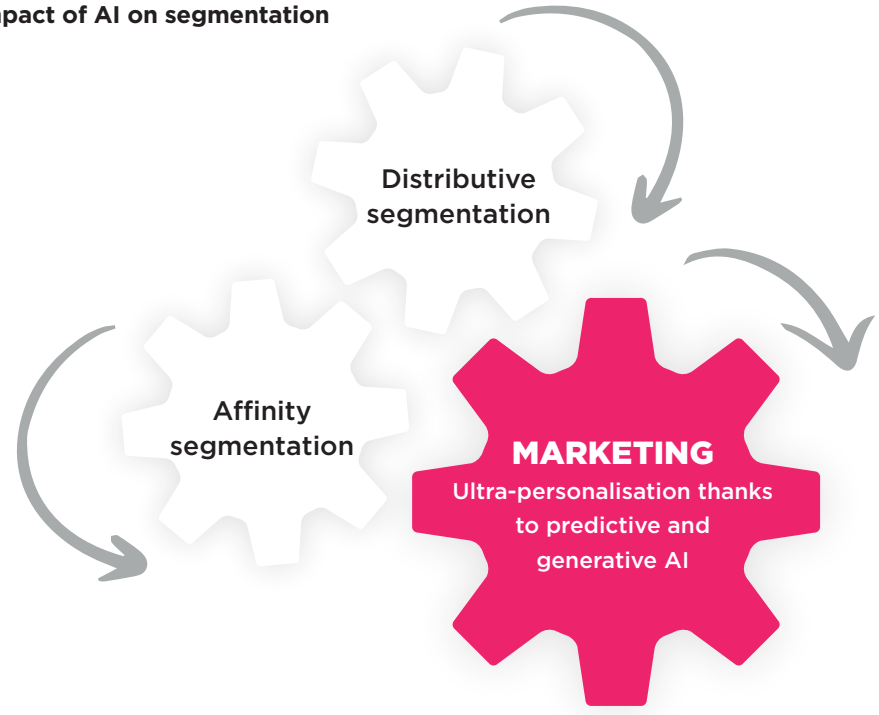

the customer why they are being targeted. Instead of 'The holidays are fast approaching, sign up for an international option card', the customer (or the advisor on behalf of the customer) will receive a message such as 'You have spent a lot of money outside the euro zone in recent weeks, incurring significant costs. By taking out our international option, you will be exempt from this type of charge in future'. **The combination of predictive and generative AI** will also make it possible to personalise financial savings advice, in particular by making it consistent over time, beyond the initial investment. Personalised investment proposals can be put to the customer on a regular basis.

THE OPERATING MODEL WILL ACCELERATE THE MOVE TOWARDS THE FRONTALISATION OF CREDIT.

Whether in the retail or corporate markets, the role of back offices in the credit field is essentially one of document verification, with a dual front-end control approach. In this respect, the case of mortgages is spectacular. The mission of back offices is to secure the issue of an offer or a disbursement in two ways:

1. Check that all the documents required for the dossier to be compliant are present (the list can be long);

2. Ensure that the information entered into the bank's tools by the advisor (or sometimes directly by the customer) is consistent with the information contained in the documents on file (e.g. address of the property, surface area, EPD or, quite simply, the spelling of the sellers' and buyers' names given in the offer to sell compared with those entered in the bank's instruction tool). The maturity of RAD tools coupled with generative AI means that these tasks can be fully automated. At the end of the assessment process, the advisor would then receive an endorsement to issue the loan (or even disburse the funds once the deed of sale has been received) or a non-completion alert that they would have to resolve before being able to issue the loan (or disburse the funds).

Such a development means **that back-office intervention would be simply eliminated and that the credit offer would in fact be made from the front**. In the end, customers will appreciate the significant reduction in processing time achieved by the much-awaited empowerment of advisors.

In addition to granting or disbursing credit, AI will also have an impact on the life of the loan. To this end, it will once again rely on its document management capabilities: automating the payment of invoices in the context of an off-plan sale, automating the processing of a request for partial or total early repayment (simply by reading the customer's request on the bank's secure messaging system, for example), or automatically proposing the best collection strategy.

FINALLY, AI WILL FURTHER IMPROVE, SECURE, AND AUTOMATE BANKING RISK MANAGEMENT.

Risk management lies at the heart of a banking franchise's performance, both in terms of credit and operational risks. In this area, AI is already a great help. Future advances in models and computing power, on the one hand, and breakthroughs in generative AI, on the other, will make it possible to continue improving existing uses, but also to tackle other areas of application.

In terms of **fraud management**, AI is currently doing a great deal to detect 'remarkable transactions' as part of the **fight against money laundering. Reducing the number of false positives** will be the next challenge, as this is a costly phenomenon in terms of staff time and customer satisfaction. Yet fraud prevention also means identifying false documents, whether at the point of contact (foreign identity documents, for example) or

when granting credit to individuals (forged pay slips or bank statements, for example). **Cheque fraud is another example**. In this case, AI can considerably help by suspending payment of a cheque deemed suspicious until it is cleared, based on models of past scams.

In terms of risk in the strict sense of the term, AI will have a considerable impact in two areas. Upstream, **risk scorings** are set to be further refined through much more advanced use of machine learning. Financial analysis, which is a prerequisite for corporate credit commitments, will also be considerably automated. During the life of the loan, much more sophisticated monitoring of customer risk profiles can also be envisaged, through the emergence of advanced indicators based on account operating data as well as on data that is becoming increasingly available in open source, particularly for businesses.

Finally, the age-old practice of making daily payment decisions at the customer advisor's discretion is likely to be overturned. There is no doubt that an AI is better able to make, or propose, a better decision in this area than a human (access to much more data and absence of affect).

Finally, **debt collection practices** are set to change considerably. In terms of prevention, AI can carry out real-time behavioural analyses (spending habits and changes in financial flows) to predict payment defaults and propose restructuring plans at an early stage. Downstream, by understanding the unique circumstances of each borrower, AI can propose tailored repayment plans to increase the likelihood of recovery while improving productivity. As in marketing, it will also be able to suggest the most appropriate contact channels for a given situation, and the right form to adopt.

HOW DO YOU STRUCTURE AND MANAGE AN AI PROGRAMME?

With AI, particularly generative AI, banks are caught between two contradictory pressures: the fear of being right too soon, on the one hand, and the fear of missing a transformative wave on the other. The first pressure could lead them to embark on adventurous projects involving technologies that are highly promising but not yet mature and still risky, particularly in a field of activity as exposed as their own. The second risks them falling considerably behind and never being able to catch up on THE transformative technology of the decade.

In these uncertain conditions, while the option of doing nothing is to be ruled out, AI must nevertheless be approached methodically. The idea is to move forward as quickly as possible, while controlling the risks and supporting the strategic project as effectively as possible. This means addressing the issue through **a corporate programme steered by the General Management, but … for which empowerment at the lowest level of the pyramid** is just as key.

THE GOVERNANCE OF SUCH A PROGRAMME MUST BE BASED ON TECHNOLOGICAL CENTRALISATION AND DECENTRALISATION OF USE.

Steering the AI programme at the highest level by the General Management will secure the following:

- The major strategic goals of the programme (overall or by BU) and its KPIs: optimisation of the target operating/distribution model, top line, customer and employee experience, etc.;
- The volume and nature of technological investments (cloud vs on-premises, proprietary or open LLM, budget envelope);
- The organisation of technological capabilities and the extent to which they are shared;
- The management ethical and cybersecurity considerations;
- The methodology for prioritising and selecting use cases in a transformative logic.

That said, in the context of a technology that will affect all businesses at every level, **this top-down approach must be closely paired with a bottom-up approach too**. As this issue is set to affect everyone's working practices – managers, support functions, network, back office –, a hyper-collaborative approach must be adopted at every level of the company. First of all, employees on the ground must be listened to and heard, and their pain points and those of their customers addressed. The effective widespread use of AI will therefore require a form of AI democracy. Each department must organise a structured ideation, carry out a prioritisation exercise and then manage a backlog. This should be regularly challenged to incorporate the latest developments in a rapidly advancing technology.

Piling up business cases without a transformative vision is, however, to be avoided. Once an initial inventory of use cases has been made and the first POCs launched, steps must be taken to move beyond a simple inventory of significant advances and look at the overall organisational impact. The integration of AI will not just mean faster and more efficient completion of tasks.

It must also enable them to be carried out differently, sometimes even accomplishing others, changing their missions, their distribution or production model, and sometimes even their very *raison d'être*. **Act local then think global** is a key motto for getting the best out of an AI programme.

PRIORITISING USE CASES WILL THEN REQUIRE AN ANALYSIS OF THE SPECIFIC VALUE.

In addition to a broad, precise and qualified mapping of business opportunities, there must be prioritisation. An analysis of the value in relation to the costs and risks will be necessary. On a subject like AI, however, this approach will require certain adjustments.

In terms of value, some projects will offer easily quantifiable gains in terms of avoided costs (AI for risk management), productivity (AI for coding), or even NBI (ultra-personalisation of digital marketing offers). For others, particularly those in the field of generative AI, this will be more complex.

This is because their added value will lie in the area of customer or employee experience, where the benefits are often difficult to measure and cannot always be harnessed. For example, an effective virtual assistant for an advisor could save 10 to 15 minutes of productivity per employee per day (with two to three complex requests per day). In a network of ten thousand advisors, the theoretical FTE gains would be enormous. However, their dispersal would make them impossible to harness. This would not necessarily justify abandoning such a project. In fact, effectively dealing with these two or three painful work situations may be enough to turn a frustrating day into a satisfying one for the advisor. At a time when the commitment of customer advisors is declining everywhere, this is certainly something to consider. The ROI in this case could, however, be approached, for example, through the estimated impact on the rate of empty chairs in a network caused by resignations, the cost of recruiting and onboarding a new associate, or, more simply, the number of additional sales appointments made possible by this time saved, coupled with the average commercial effectiveness of such appointments.

Figure 6. **Risks specific to AI**

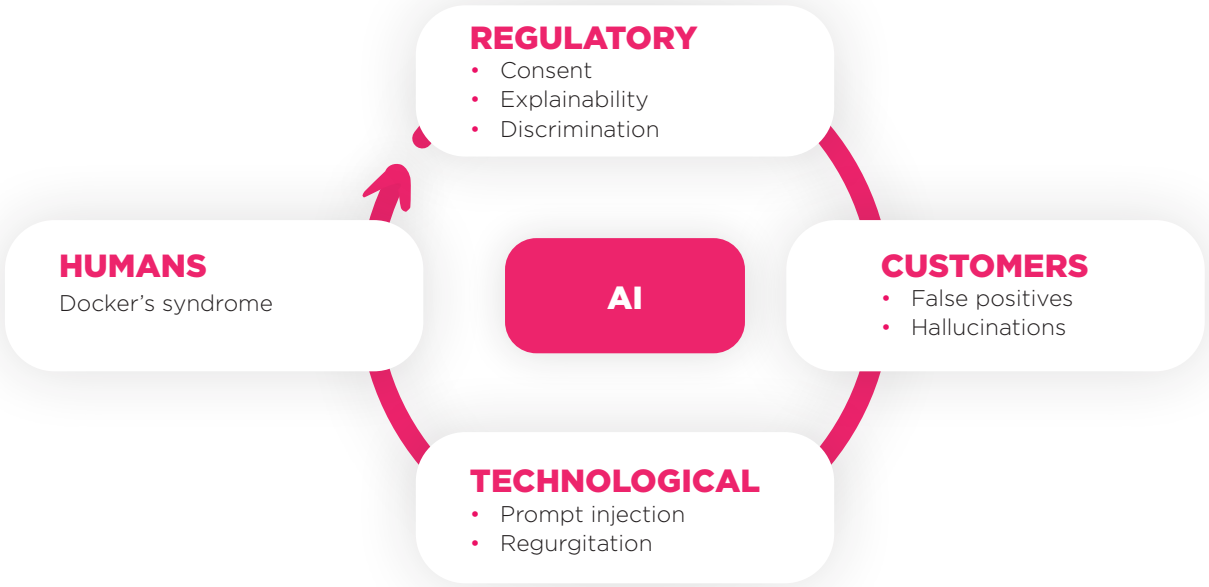

As we can see, many use cases will require a relatively complex value analysis.

From a cost point of view, the difficulty lies less in the intrinsic cost of the project than in **the significant infrastructure costs** required by AI and the data it embeds. Acquiring GPUs to process generative AI use cases is therefore particularly costly and not easy to scale (investment scalability). Use of the cloud and hyperscale models (or emerging sovereign solutions) will make it easier to calculate the cost of using the solution, which will be invoiced on a fee-for-service basis in SaaS mode. It is important to note, however, that the tenths of a euro cent charged per request will have little impact on a model that works for employees, but that **extending it to millions of retail customers could cause the bill to increase dramatically. Some projects will, moreover, be used for more than one use case**. The best example is applied generative AI. A project of this type can be used for an employee assistant, a customer chatbot, and to automate call handling in a customer relations centre. ROI must therefore be calculated with caution.

AS WELL AS THE VALUE OF THE PROJECTS, AN AI RISK ANALYSIS MUST ALSO BE CONDUCTED.

Once the value has been established, it will be necessary to assess the project based on another criterion and one that is far from trivial when it comes to AI: risk. **Four types of risk need to be examined: regulatory risk, customer risk, technological risk, and human risk**.

From a regulatory point of view, two aspects specific to AI projects must be examined: compliance with the GDPR and compliance with the AI Act. As far as the GDPR is concerned, the issue of customer consent will be the most sensitive to resolve understood here in the broadest sense, i.e. from consent upstream to the ability to change choices downstream. For the AI Act, it is undoubtedly the explainability of the models that will pose the greatest problem, particularly in view of the increasing complexity of ever-earlier risk prediction models. Customer scoring models, although not explicitly classified as high risk by the AI Act, must take care to avoid any discrimination and comply with the obligations of transparency and non-discrimination, especially when they have a significant impact on the rights of individuals (AI Act, Appendix 3, Section 5).

As far as customer risk is concerned, two areas need to be examined, starting with the matter of false positives. What percentage of customers would agree to being wrongly penalised by a model which, for example, suspended the payment of a cheque for several days until it was cleared on suspicion of fraud? Then, as far as generative AI is concerned, it is hallucinations that must be tracked. A risk/benefit assessment will also be necessary in this area. What is an acceptable level of error and instability, for example, for an automatic response system to certain customer emails?

As far as technological risk is concerned, IT departments will, of course, be keeping a close eye on all applications that are open to the outside world. By this we mean models requiring the pairing of external data with internal data in order to be fully measurable (the use of public data to approximate the EPD of a property, for example). This also applies to the generative AI solutions that will be made available to customers (chatbots that will finally be fluent) in order to become an omnichannel conversational bank. The risks of inappropriate content, prompt injection and regurgitation of internal data are taken very seriously by cybersecurity teams. For these reasons, the projects selected will initially focus on internal customers.

Finally, **there is a human risk**. We need to keep a close eye on 'docker's syndrome', which can affect operational staff in branches or the back office. When the first forklift trucks arrived in the 1920s, some dockers were worried that their jobs would disappear. Others were delighted at the prospect of having less backache and riding these new, fun machines. Employees will there-

fore need to be supported to make the most of the benefits that AI can bring to their jobs. Automating the processing of customer emails, in particular, will come under this heading. This task is tedious for customer advisors, but its disappearance could destabilise a number of players who also see it as a reassuring link with their customers. At managerial level, AI will scare the experts, those who derive their legitimacy from their knowledge of procedures, tools, and products, as well as their ability to make decisions (credit applications and tariff exemptions). These individuals will have to shift the focus of their skills towards leadership and pure management.

In short, the notion of **responsible AI** will become increasingly important. It is driven by the AI Act, but will also be increasingly driven by public opinion. In this context, it would appear essential to establish close governance of AI risks within each establishment.

FROM A TECHNOLOGICAL POINT OF VIEW, IT WOULD APPEAR ESSENTIAL TO POOL HUMAN RESOURCES WITHIN AN AI FACTORY.

From a technological point of view, there are two particular constraints relating to AI that need to be taken into account:

- Skilled resources are scarce and will remain so for a long time: coders, data scientists, but also prompters;
- Progress is incredibly swift, often promising and sometimes deceptive.

These arguments tend to favour the creation of an AI Factory: a centralised technological skills centre where the company's best experts will be pooled. Its main tasks will be to:

- Build and operate the technological infrastructure: computing power, base models (LLM, machine

Figure 7. **Knowledge management: the foundation of generative AI**

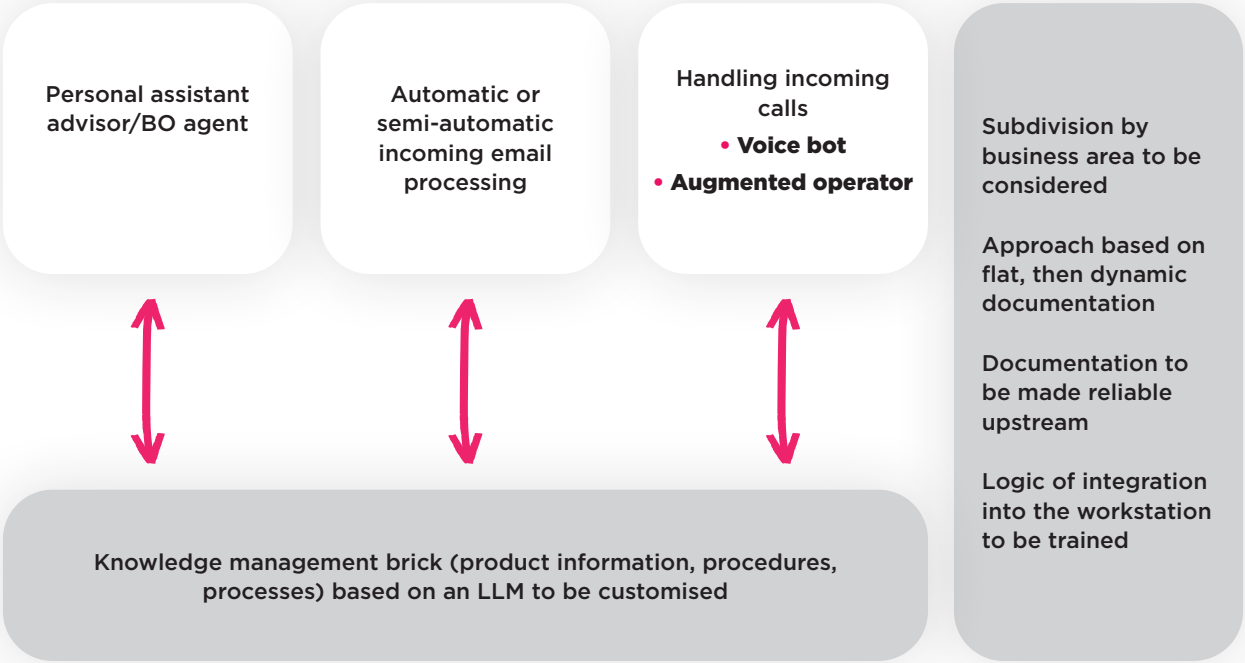

learning), development and test environment (particularly in terms of data);
- Develop in-house methodologies;
- Assist business units in drawing up their AI roadmap (agile mode);
- Develop, implement, and maintain the projects selected;
- Assist the business lines in their change management approach prior to implementation;
- Conduct a technology watch.

TACKLING KNOWLEDGE MANAGEMENT IS THE NUMBER ONE PRIORITY.

Notwithstanding the strategic priorities of one group or another, the field of knowledge management stands out as one in which it is in the banking industry's interest to invest. While we have known for several years about the importance of the quality of structured digital data thanks to predictive AI, **we are discovering with generative AI the key importance of a new unstructured data environment, that of documentation**. The quality of documentation is crucial to many key AI uses, including almost all those linked to conversational banking. In order to interact with a customer in a fluent and relevant manner, the generative AI algorithm needs to use reliable unstructured data: product characteristics (the 'what') and all the bank's processes (the 'how'), in an omnichannel environment. Investing in knowledge management means consolidating a base layer that will be essential to the quality of a chatbot or a tool for handling emails, customer complaints, or customer calls.

This data is, however, outside the scope of the IT or data departments and belongs to users. Its governance is generally unstructured as its importance has, up until now, been secondary. More than simply a new kind of data cleansing operation, this is a new way of organising knowledge management.

Figure 8. **AI for the better,** *People @ Work* magazine *#8*

FINALLY, THE CHANGE MANAGEMENT ASPECT WILL BE COMPLEX AND MUST BE ADDRESSED WITH CARE.

AI is perceived by employees as a subject that is both anxiety-provoking and complex. Things must therefore be properly planned if they are to go well, based on a fairly simple aim: overcome fears to encourage adoption. **There is a great risk of developing expensive tools that are not used by employees or customers**. The semi-automatic tools for processing customer emails deployed in certain networks in recent years have, for example, proved disappointing. At customer level, the investments made in budget management solutions on websites or apps have not yielded any ROI either.

Considering change in terms of AI means first and foremost **educating** employees on a massive scale about this technology, its nature, its fields of application, and its limits, right down to the lowest level of the hierarchy.

General concerns must then swiftly give way to **tailored thinking for each profession**. This is the purpose of calls for projects organised by certain chains, for example. Their aim is to get employees to think about the potential benefits of AI for their jobs and their day-to-day work.

Once the employee has been given the opportunity to learn, to find out, and then to plan, the r**ationale for the project must then be explained and** linked to the bank's corporate purpose. At Crédit Mutuel Alliance Fédérale, the statutory promise to use technology solely 'for the benefit of our employees and our customers' makes it easier to express a convincing case for AI and its applications.

Finally, change management must take into account the expectations of customers and employees. Some are obvious, such as the need for reassurance. Others will need to be investigated to better define their specific characteristics.

Figure 9. **The keys to embracing AI**

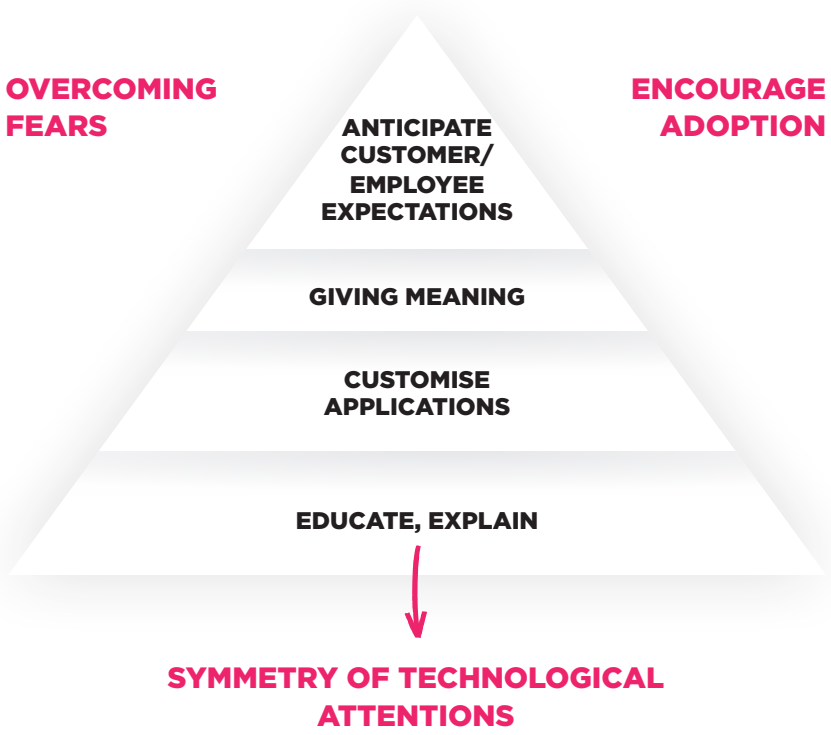

CONCLUSION
TOWARDS A CHANGE IN CULTURE AND THE EMERGENCE OF A FORM OF URGENCY

AI is therefore set to profoundly distort the banking industry's distribution, production, and risk models. The challenge will be as much technological (data, infrastructure, and languages) as human and organisational. The breadth of the scope and scale of both the potential gains and the risks mean that this issue needs to be turned into a real corporate project driven by the General Management, while being genuinely taken in hand by the business departments.

Over the years, corporate culture itself is likely to be affected by these profound changes. The enduring administrative culture, particularly in retail networks, could be diluted by the promises that AI will offer in terms of operational efficiency and customer relations. AI will then have transformed the banking industry once and for all.

In the meantime, European banks, and French banks in particular, seem to be lagging behind. According to a November 2023 analysis by US research firm Evident Insights, the top five North American banks now account for more than 67% of the industry's AI research publications worldwide and have filed 94% of the patents. No French bank appears in the top 10 of Evident's Bank AI Index, and only one European bank (ING) features on it. While, as is often the case, the winners in this emerging revolution will be those who move fastest (the famous 'first-mover advantage'), European banks should be feeling a certain sense of urgency. They need to accelerate and put **AI at the heart of their strategic plan**.

Circle Strategy is a strategy consulting firm rooted in entrepreneurial culture and unique in the dual profile of its consultants, who have a combination of consulting and field experience.

Circle Strategy assists senior management with growth strategy issues, responsible business models, and digital and operational transformation.

Circle Strategy is characterised by an assertive blend of pragmatism and strategy, illustrated by its motto ‘Strat for impact’ and its corporate purpose: ‘We shape strategies for bold and sustainable futures.’

Circle Strategy is part of Square Management and operates worldwide from its offices in France, Spain, Belgium, and Luxembourg.

www.circle-strategy.com

CONTACT

CHRISTOPHE BANIOL

Partner

+33 6 50 47 12 85

christophe.baniol@circle-strategy.com

NICOLAS MIART

Partner

+33 6 34 58 35 19

nicolas.miart@circle-strategy.com

Cover image: Freepik


Circle
strategy