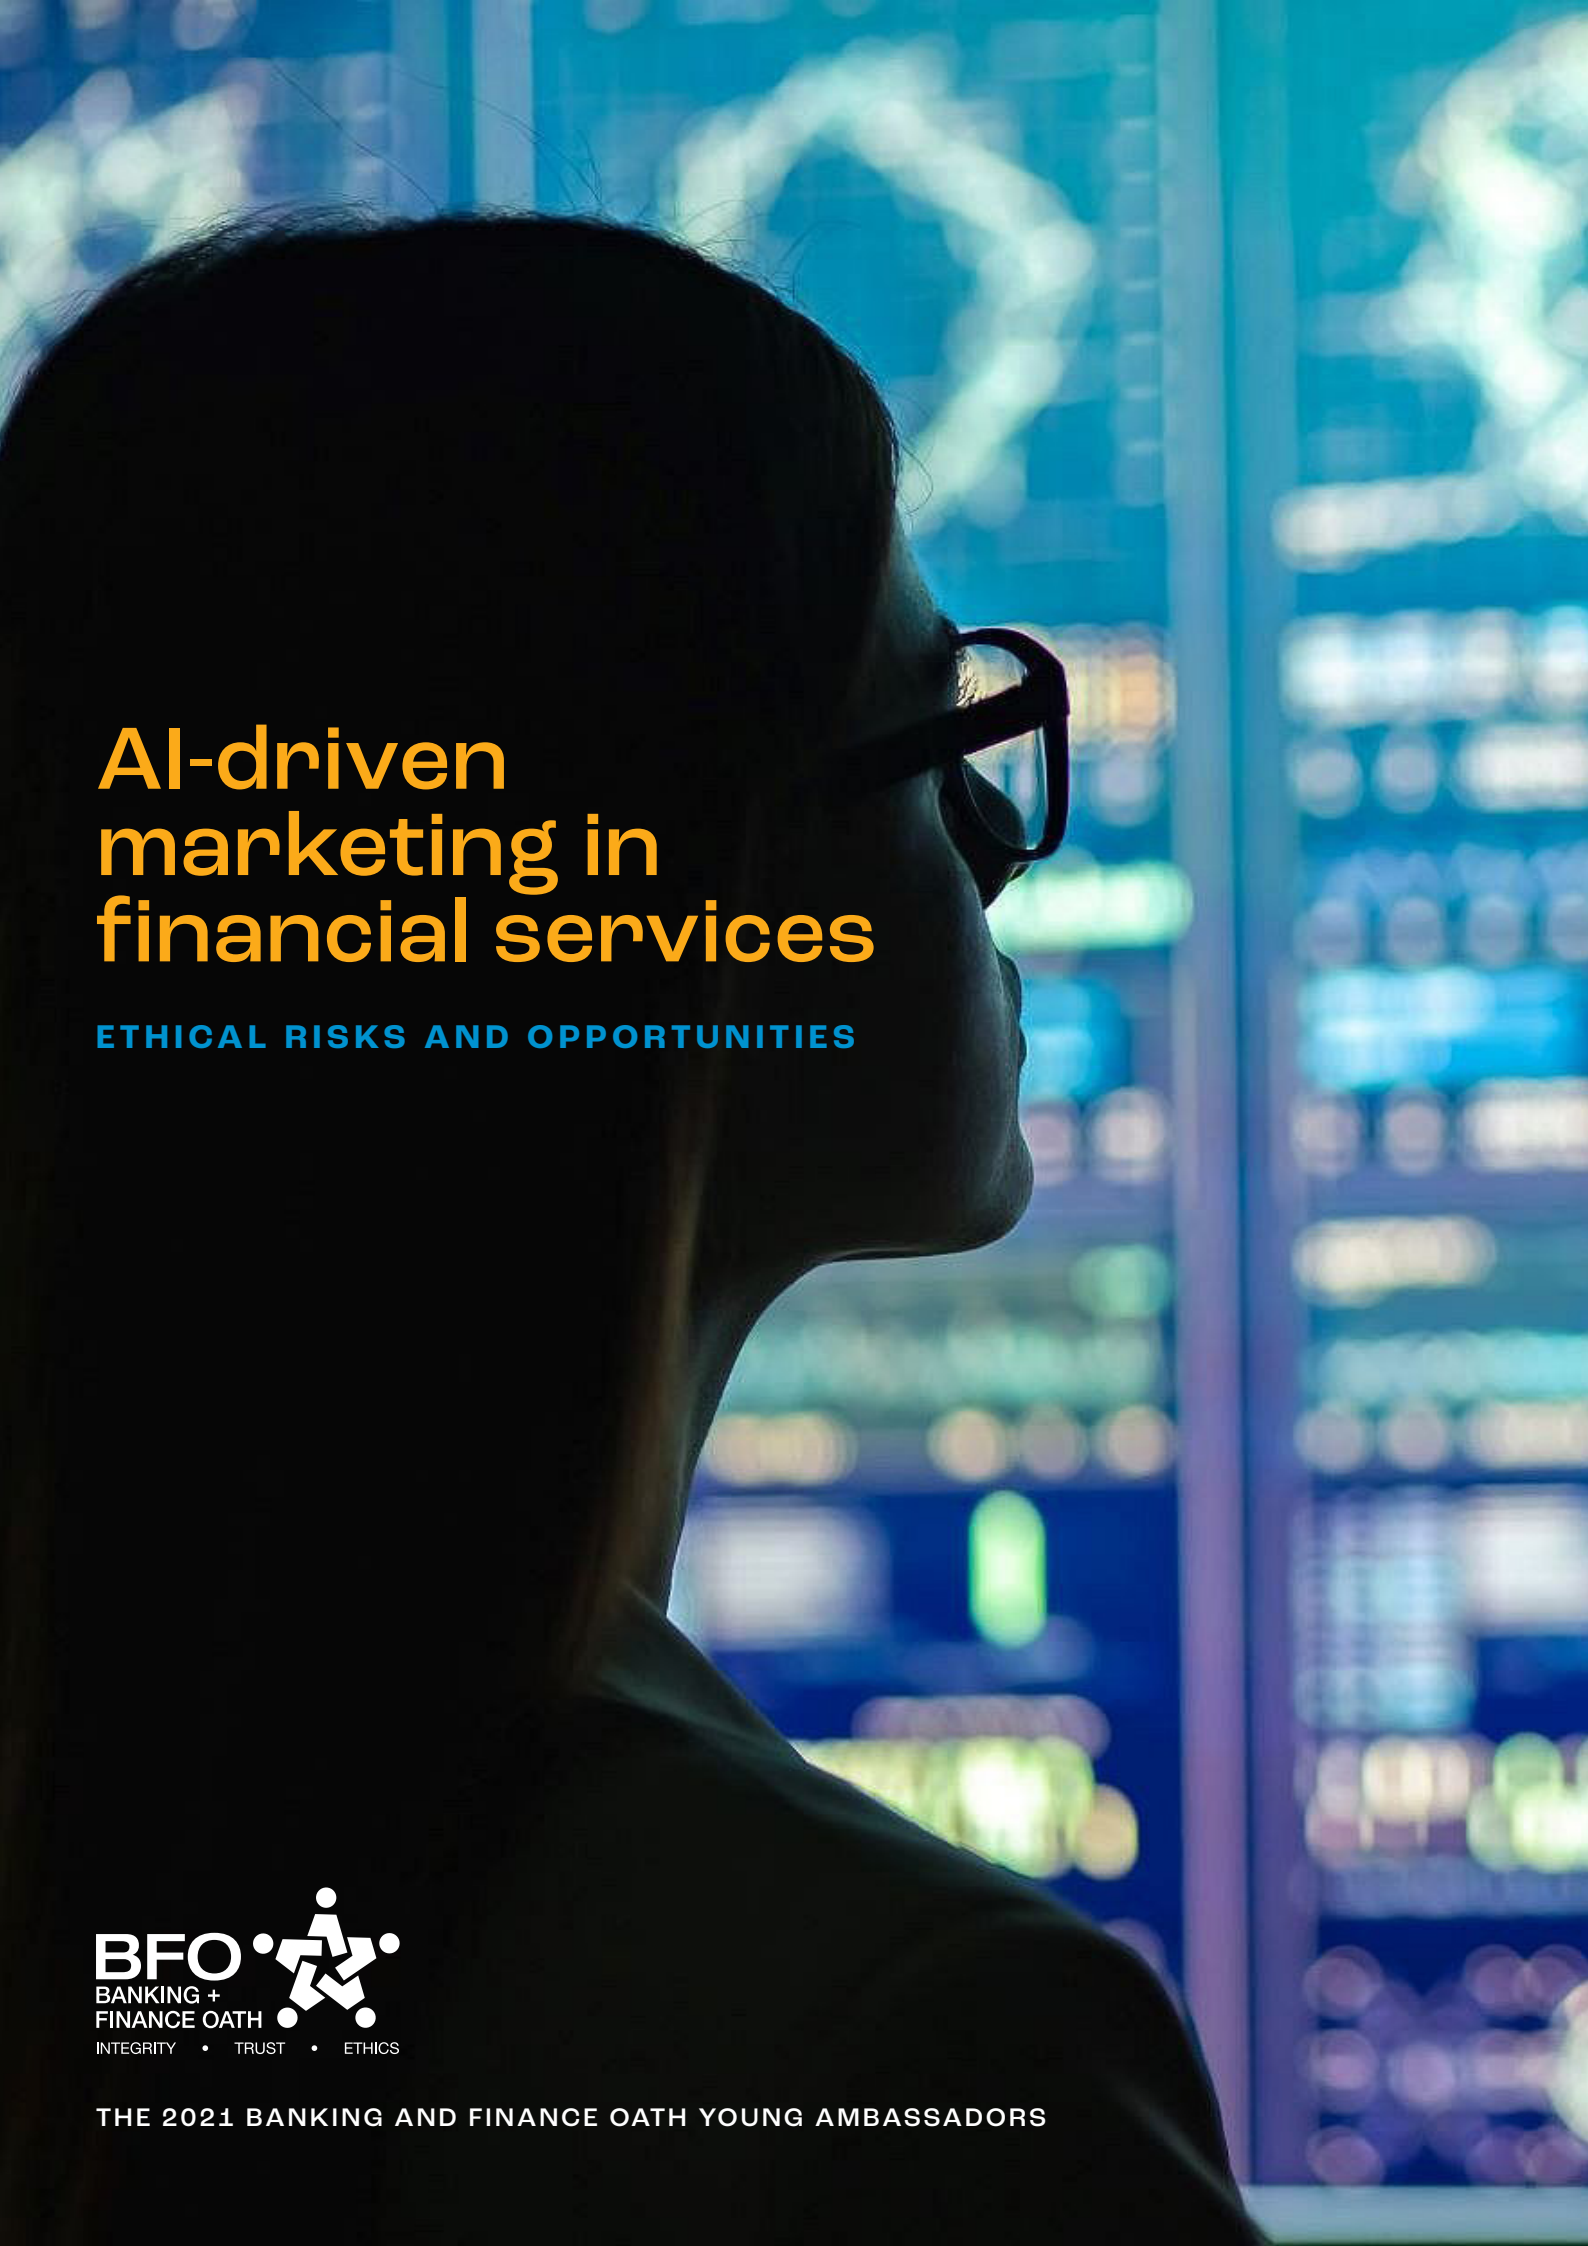

AI-driven marketing in financial services

ETHICAL RISKS AND OPPORTUNITIES

BFO
BANKING + FINANCE OATH
INTEGRITY • TRUST • ETHICS

THE 2021 BANKING AND FINANCE OATH YOUNG AMBASSADORS

A report by
The Banking and Finance Oath
Young Ambassadors

This report was written by The Banking and Finance Oath (BFO) 2021 Young Ambassadors. Created in 2013, The BFO is a voluntary oath that financial services professionals can take to demonstrate their commitment to ethical practices and to drive positive change for the industry and society more broadly.

Each year, The BFO chooses a group of young industry professionals to promote the oath, and supports them in strengthening ethical practices.

This year, the Young Ambassadors chose to explore the growing issue of the ethics of using AI in financial services. After speaking to people within their own organisations and across the financial services industry and government, they worked with Gradient Institute – an independent non-profit working to develop ethics, accountability and transparency in AI – to refine their research question and identify an area within financial services in need of a focus on AI ethics.

They found that while there is a growing awareness of AI ethics in the broader industry, and in some specific areas like credit scoring, little attention has been given to issues around AI-driven marketing, and the related harms and benefits.

To fill this gap, they researched and wrote this report.

Contents

Executive summary

Artificial intelligence (AI) is reshaping the way Australian financial services organisations do business and market their products, but not without real social, economic and ethical implications.

Before AI, humans and human teams decided what to market, who to market to and how to achieve their business objectives. But now AI, big data and machine learning are increasingly integral to how banks, financial planners, insurance companies and superannuation funds make decisions that profoundly affect their customers' lives.

While these technologies are streamlining front-line processes, making it easier for financial services organisations to predict credit risk, and enabling them to better tailor products to customers' needs, they are also raising questions. For instance, should AI-based marketing be used to target home loan advertising at people undergoing a divorce? Do human prejudices stemming from racial and socio-economic profiling influence AI's algorithms when it comes to who is targeted for specific products? How does the use of data in AI-based marketing affect customers' privacy? And, perhaps most important of all, are there positive ways we can harness the potential of this new technology?

In the financial services industry, where the practical impacts of our products on customer's lives are over-sized, the potential impacts of AI – both positive and negative – are similarly large.

But there are ways to mitigate the risks. So far more than 60 countries – including Australia – have adopted an AI policy or instrument, or have developed ethics frameworks for the responsible design and deployment of AI.[01] Organisations can also head off some of the risks in using an AI-driven marketing approach by:

- **maintaining good governance systems** that avoid data siloes and the use of inaccurate data.
- **investing in training in responsible AI** for AI developers, system owners, system integrators, business leads and boards, to ensure they have the expertise and awareness to effectively govern AI systems.
- **properly demarcating responsibilities for using data and AI-driven marketin**g, so that people and not AI ultimately retain accountability.

By adopting ethical frameworks such as that described in this report, organisations can also build ethics into the design of AI systems from the start.

Developing comprehensive understanding of the ethics involved in business decisions around AI will assist financial services leaders to avoid real harm to customers and society, maintain trust and legitimacy, and prove AI can be a force for good.

01 OECD.AI Policy Observatory (2021), 'National AI policies & strategies'. Retrieved from https://oecd.ai/en/dashboards.

Introduction

AI is already reshaping the Australian financial services landscape.

Australia may be a laggard when it comes to artificial intelligence (AI)[02], especially responsible AI[03], but from making front-line processes more efficient to better predicting credit risk and guiding personalised product recommendations, AI has the potential to unlock a vast amount of latent value.

This use of AI and related data carries well-documented risks, including in relation to transparency, privacy, bias, discrimination and reinforcing entrenched inequalities. Business leaders need to develop a comprehensive understanding of these issues so they can avoid real harm to customers and society, maintain trust and legitimacy, and on balance prove AI to be a force for good.

To avoid undoing years of recent progress following the Hayne Royal Commission, the financial services sector must focus on ensuring AI technologies are designed and deployed in ways that maintain core corporate and democratic values.

Scope and audience

This report sheds light on the growing ethical risks and opportunities of AI-driven marketing in financial services. We believe this focus is an effective way to unravel the complexities of AI's impact in one aspect of the industry, establishing a framework that can be applied to other contexts within financial services, then across other industries.

We all have a responsibility to better understand the technologies that shape our lives and choices, and those of the people around us. However, this report is specifically aimed at:

- **marketers** working in financial services, and AI developers designing marketing products.
- **professionals and leaders in financial services** and other sectors that are already using AI or may use it in the future.
- **general readers who want to know how AI works**, how it is applied in a specific context and some of the ethical risks to consider.

02 Cilento, M., 'Australia needs a more co-ordinated approach to digital technology', CEDA website, 26 March 2022. Retrieved from https://www.ceda.com.au/NewsAndResources/Opinion/Technology-Innovation/Australia-needs-a-more-co-ordinated-approach-to-di; Thomas, W., 'Why Is Australia an AI Laggard?, CDO Trends, 23 August 2021. Retrieved from https://www.cdotrends.com/story/15815/why-australia-ai-laggard.

03 Nuttall, S. 'The 2021 Responsible AI Index', Fifth Quadrant, 2021. Retrieved from https://www.fifthquadrant.com.au/2021-responsible-ai-index

PART 01

What is AI-driven marketing?

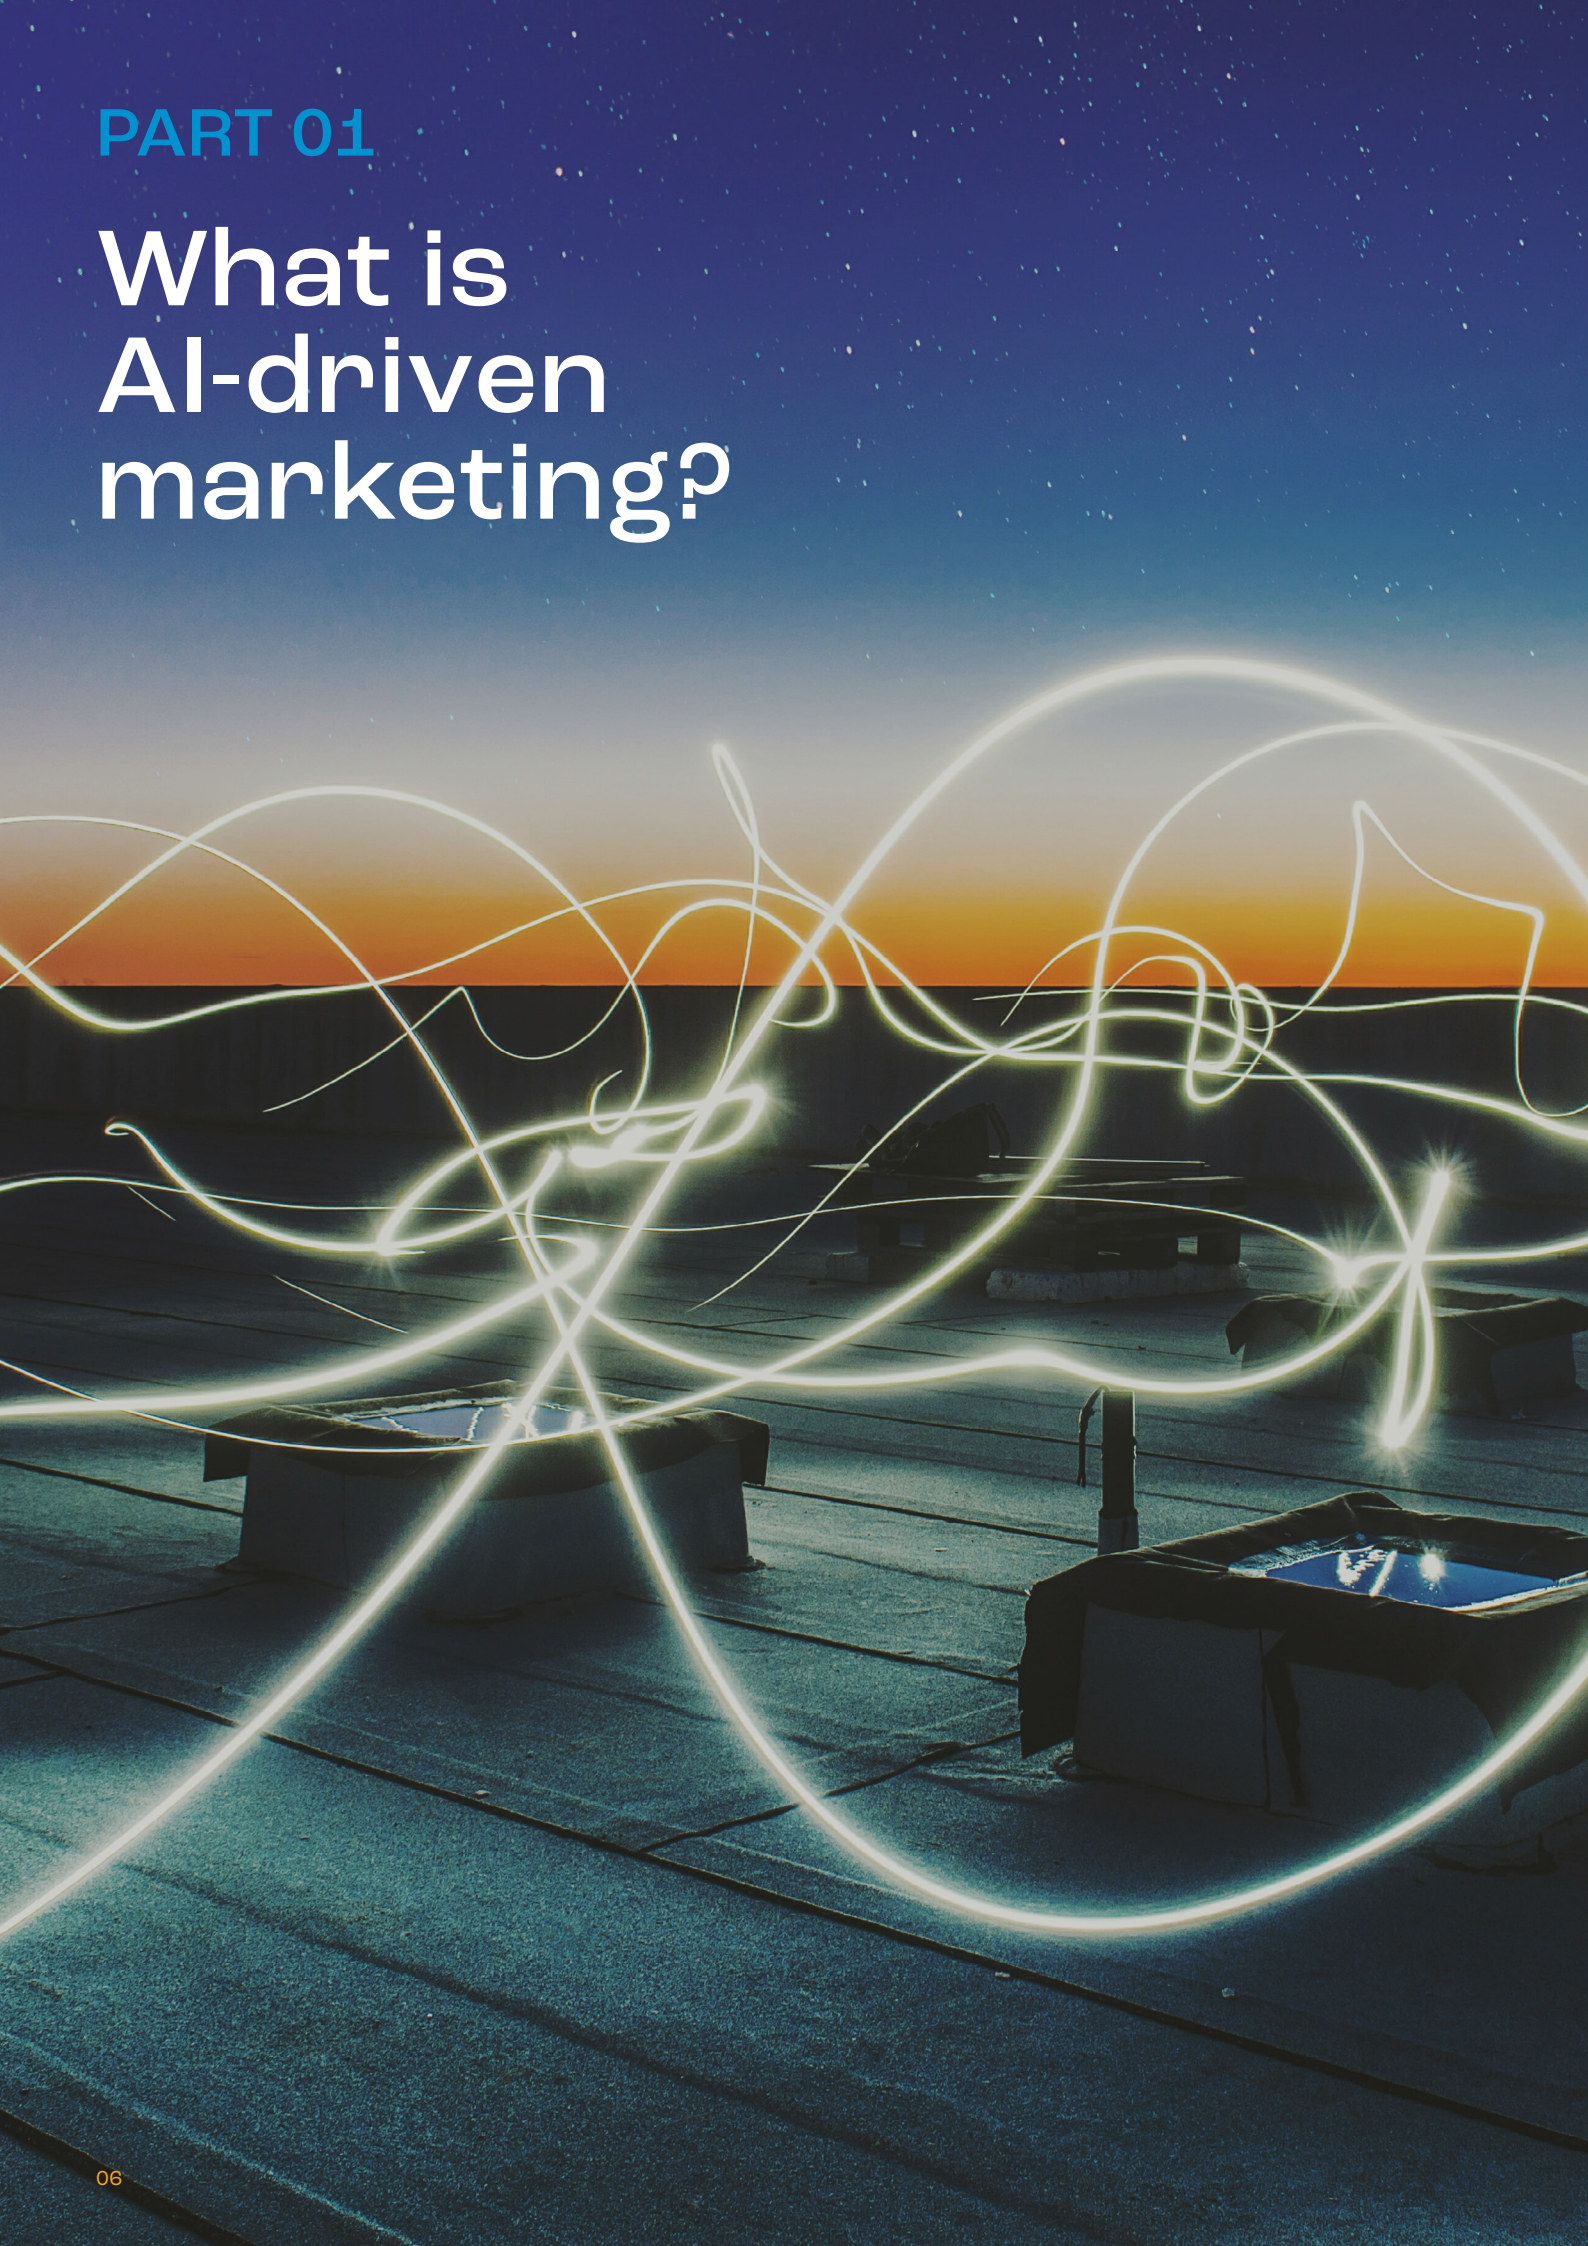

"Given the amount of consumer data available – and the number of channels, segments and competitors – marketers have to use AI and algorithms if they are to operate successfully in an online environment."

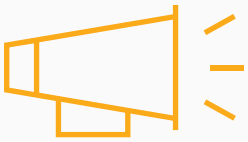

Marketing before AI

Marketing is about designing, promoting and selling goods and services.[04] In a market economy, it facilitates the interaction between consumers and producers: consumers receive information that leads them to buy things that meet their demands, and producers can increase profits by suitably designing and promoting their products to meet those demands.

Marketers have always tried innovative ways to reach customers. Newspapers, radio and television gave marketers powerful tools for targeting certain customer groups. For example, businesses have for decades paid a premium to advertise goods and services on television during evening 'prime time' viewing hours, when most people are watching. With the advent of big data, AI and online marketing, marketers can now target specific demographics in ways that were previously unimaginable. Indeed, given the amount of consumer data available – and the number of channels, segments and competitors – marketers have to use AI and algorithms if they are to operate successfully in an online environment.

04 Sharp, B. (2013), *Marketing: theory, evidence, practice*, Oxford University Press, South Melbourne.

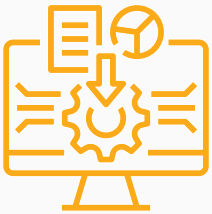

Marketing after AI

Before AI, humans and human teams decided what to market, who to market to and how to achieve these business objectives. **Now, AI, big data and machine learning can facilitate the connection between consumers and producers**. Although humans are still involved in the process, data inputs and decision-making processes are fundamentally changed.

Examples of AI-driven marketing include:[05]

- **content generation:**
  generating personalised content for customers and showing personalised offers that are relevant to the customer.
- **web and app personalisation:**
  leveraging historical and real-time data to customise the web and app pages customers see at different purchase stages, based on their interests.
- **targeting offers:**
  determining the likelihood of a customer purchasing a product if they receive a conditional marketing offer such as a discount.
- **omnichannel assistance:**
  supporting a seamless purchasing journey across multiple channels, such as from mobile phone to personal computer to an in-person meeting in a bricks-and-mortar business outlet.

How does this work?

Customer data + Business objectives → Machine learning algorithm → AI system applied to new cases

05 Veritas Document 2 (2020), FEAT Fairness Principles Assessment Case Studies, p. 8. Retrieved from https://www.mas.gov.sg/-/media/MAS-Media-Library/news/media-releases/2022/Veritas-Document-4---FEAT-Principles-Assessment-Case-Studies.pdf

Model building

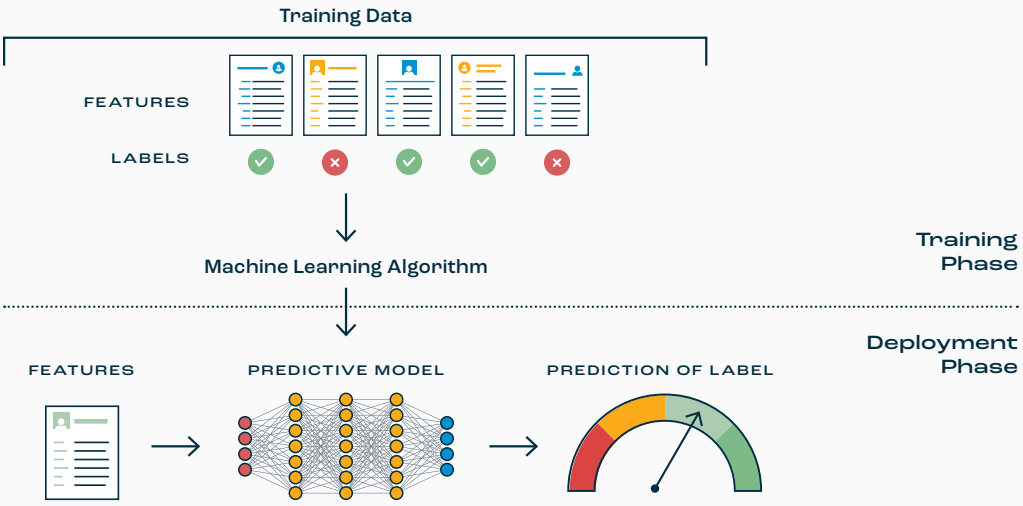

Financial services organisations, like organisations in many other industries, tend to follow a standard process when using AI and customer data.

First, they draw on a **pool of data**, such as customers' credit card transaction histories or social media click-throughs. They then match this data with a **business objective**, for example increasing the rate of uptake of a new product by showing advertisements to interested customers on social media.

Data scientists then use a **machine learning algorithm** to analyse trends in the data. For example, an algorithm could use historical customer data and transaction patterns to determine the likelihood of customers being interested in certain products. Once the algorithm has learned these patterns, it can be fed current customer data, enabling it to **predict the likelihood** of a customer being interested in a product – and how likely they are to click on a social media advertisement for it.

An algorithm like this would form the core of an AI-driven marketing system, along with decision rules setting the threshold of likelihood that must be reached before a customer sees the advertisement.

Many organisations already use this process, powered by large amounts of data plus the computing capability and growing sophistication of machine learning models. However, there are ethical risks to navigate at every step – from the quality of the data used to train the model, to how well AI developers and business managers understand the business objectives. It is also vital to consider how the model is audited and checked for unintended consequences. The following section delves into the marketing-specific risks of this process.

CASE STUDY

AI-driven marketing in online banking

RappiCard began offering a credit card to customers in Latin America in 2021. The market was flooded, so it was hard to engage consumers. Consumers were unresponsive to messages and online push notifications.

To reach customers, RappiCard used AI-generated marketing content.

To do this, it partnered with Persado, an AI content generation and decision-making platform. The platform is powered by a large language knowledge base. It combines natural language generation, machine learning and experimental design to deliver customisable advertisements to consumers.

RappiCard provided control data to Persado to use in its machine learning engine. Persado categorised RappiCard's messaging into components to identify the specific words, phrases and emotions that would entice RappiCard customers. It then created multiple messages to deliver to consumers. This process was iterated over and over again, narrowing down a customised message that would most effectively secure sales. Throughout the campaign, RappiCard discovered the different emotions and headlines that resonated best with certain customer segments.

RappiCard generated a 179 per cent uplift in conversion rates by using this AI to communicate with consumers more effectively and in a more customised way, based on their perceived emotional response to various messages.

As more content is pushed through the platform, the more data the AI processes and the more powerful and accurate RappiCard's marketing campaigns will be.

RappiCard is still working to extract value from its huge store of first-party data, even with dedicated data scientists on the job. The company wants to improve its operating margins, which it could do by using the data to better manage risk, increase conversion and reduce acquisition costs.

Marketing has traditionally played on consumers' emotions to guide their purchasing decisions. Compared to traditional marketing methods, AI-driven marketing is an extremely effective and efficient way for companies to do this. These types of marketing tools will become ever more powerful as companies learn how to use consumer data in their pursuit of better margins.

Marketers might ask themselves, does this example have ramifications for consumer autonomy? How do you feel about companies using powerful AI to predict your emotional response and entice you to purchase a product? Do you think this approach is ethical?

HYPOTHETICAL USE CASE

Data governance issues in home loan advertising

"Incomplete data can result in ineligible or vulnerable customers being targeted by advertisers for products that they would otherwise be ineligible for."

Home loan advertising is not a perfect science. Not every customer reached will be a right match for the product, and not all people who meet the loan criteria will want to take out a loan at that time.

In a crowded market with dozens of options, marketing can become quite complicated. Whether marketing directly to existing customers or prospecting for new customers, data gaps that can mean some people receive advertising that isn't appropriate for them.

Any AI solution is only as good as the data it has been fed, so AI alone cannot fix this problem.

An example might be a home loan refinance campaign targeted at current customers who, based on the available data, 'look like' they have a home loan at a different financial institution. The offer piques a customer's interest and they begin the application process. During the application process it becomes apparent that the customer holds more risk in their financial portfolio than the data showed, since it didn't include the risk held with other financial institutions. As a result, the customer's application is declined.

At the same time, a broader advertising campaign using a similar data set is promoting a new home loan rate. Based on the available data, the declined customer also fits into the target group for this campaign. The customer receives advertising for this other home loan product, even though they have just been declined for the previous loan.

Incomplete data can result in ineligible or vulnerable customers being targeted by advertisers for products that they would otherwise be ineligible for. This example shows how internal processes, data silos and organisational structures around AI and data systems can negatively impact the customer experience when not managed correctly. And if that customer's bid for a new home loan fails, they are within their rights to submit a formal complaint to the business.

Effective and responsible AI relies on having good quality data. As many banks and organisations are beginning to use AI, there are challenges to overcome surrounding legacy systems, data silos and organisational processes.

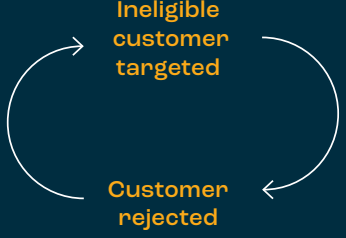

PART 02

The ethical risks and opportunities of AI-driven marketing

Ethics is the discipline of thinking about how we act in the world. This includes the duties we have to others, the impact we have on others and what type of person we want to be. Ethical considerations are just as relevant for individuals as they are for organisations – ethics deals with fundamental questions of how we ought to act in any given situation, especially when there aren't clear laws or guidance to follow.

Risks of AI-driven marketing

01 Organisations preying on people in vulnerable circumstances

A key tenet of marketing ethics is that it is morally wrong to market products to people in ways that take advantage of their vulnerability[06] – for instance, offering risky loans or mortgages to people who don't have the cognitive capacity to make a fully informed decision. Although the Design and Distribution Obligations (DDO – see page 17) aim to prevent explicit discrimination, there is still a risk of AI unintentionally drawing correlations between vulnerable demographics and risky financial products.

Everyone is vulnerable to some extent, and people sometimes buy products against their better judgement, but there are clear ways that AI could be used to intentionally or unintentionally prey on vulnerabilities in ways that most people would find objectionable.

AI has no moral oversight or awareness of its actions; it simply makes connections based on patterns in the data. So while most humans would reject the idea of targeting customers with particularly risky financial products when they are at their most vulnerable – for example, when they are going through a divorce, or even at sensitive times of the day or week – an AI-driven marketing system with no controls in place would see these as events with a high likelihood of sales conversion.

The ubiquity of internet banking, smartphone finance apps and omnichannel assistance means banks have access to a massive amount of financial and personal data. As such, banks and other financial institutions have a responsibility to not encourage people to buy things they don't need or that could harm them.

02 Data misuse and threats to privacy

Consumers have a right to know what data is being used to target them with product advertising. If the data is their own, they should have access to it and knowledge about how it is being used. If the algorithm has learned a user's general demographic features and is extrapolating from this, the user has a right to know what type of demographic category they've been placed in.

Biased data sets used to train AI models and problematic inferences between demographic categories can result in products being targeted in inappropriate ways. For example, if the data you're using to train the AI model is out of date, inaccurate or missing key information, it may incorrectly target the wrong people.

Marketing companies must follow common data ethics principles about privacy and data protection and transparency to ensure they don't misuse personal data and that they do respect individuals' privacy.

06 Brenkert, R. (1998), 'Marketing and the vulnerable', *Business Ethics Quarterly*. Retrieved from https://philosophia.uncg.edu/media/phi361-metivier/readings/Brenkert-Marketing%20and%20the%20Vulnerable.pdf.

An AI tool could deliver personalised information about a suite of financial services, including its purpose, uses and benefits.

Opportunities of AI-driven marketing

There are opportunities for organisations wanting to use AI to advance their marketing goals, alongside the ethical risks. It is important to highlight some of the benefits and not focus only on the downsides.

01 **Improved financial advice and product recommendations**
Implemented in the right way, AI may be able to connect consumers with the financial products and advice most suited to their needs. This may also benefit people who cannot afford personalised financial advice.

02 **Cost reductions**
AI can reduce costs across the marketing life cycle of research, strategy and implementation. For example, the use of AI in marketing campaigns can deliver enhanced analytics, insights and predictions. The resulting increased return on investment and more cost-effective marketing can in turn reduce the cost of goods and services for the consumer.

03 **Enhanced customer service**
Customers are increasingly interacting with customer service AI chatbots for digital assistance. Chatbots are available 24/7 and cater to customers who enjoy the self-service aspect of interacting with AI. Meanwhile, this frees up employees to focus their efforts on more challenging customer concerns or issues.

04 **Improved financial literacy**
Australians have some of the highest levels of per capita wealth in the world but relatively low levels of financial literacy. An AI tool could deliver personalised information about a suite of financial services, including its purpose, uses and benefits. This could significantly improve general consumers' understanding of the services on offer. Greater choice and understanding can lead to better purchasing decisions.

05

Data privacy vs AI reliability – a point of tension for AI-driven marketing

The ethical issues outlined previously point to a major concern for AI-driven marketing: algorithms could draw objectionable correlations between data inputs that no human decision maker could or would draw.

The more data an AI system has, the more accurate its predictions and targeting. Likewise, the more data the system's developers have, the more informed the AI system will be, reducing the risk of sending inappropriate or discriminatory marketing. However, this raises the issue of data privacy, and how much personal data private institutions, like banks, should be allowed to gather and use. Is there a limit to the volume and type of data that can be used in AI systems, even if it is to avoid negative outcomes?

Returning to the home loan example, an advertising campaign might begin with analysing psychographic data for the target audience, then creating personalised content based on this. A person who loves dogs can be assumed to react positively to an advertisement that shows people with a dog enjoying their new home. A person who loves big parties and socialising would receive advertising that depicts people enjoying social gatherings. These examples are largely uncontroversial, but what if AI drew a correlation between somebody battling a medical condition, unemployment or addiction and their increased likelihood of needing a mortgage?[07]

AI systems lack the basic moral restraint of human decision makers. This makes it all the more important to start with accurate data on people's circumstances, so that developers can build technical parameters into AI systems.

07 CHOICE staff, Mortgage brokers rewarded for preying on the vulnerable: consumer groups, 17 July 2017. Retrieved from https://www.choice.com.au/money/property/buying/articles/brokers-targeting-vulnerable-170717.

AI regulation

Ethical issues emerge when individual companies choose to assess and monitor their own AI systems. Even if an organisation has the expertise required to do so, there is an inevitable conflict of interest and lack of independent perspective.

One way to mitigate the novel risks of AI-driven marketing is through specific regulation, which is currently in its infancy.

Current AI regulation

Different countries are taking different approaches to regulating AI-driven marketing; so far more than 60 countries worldwide have adopted an AI policy or instrument, or have developed ethics frameworks for the responsible design and deployment of AI.[08] Currently, the most formalised piece of AI regulation is the European Union's proposed Artificial Intelligence Act, which classifies different AI systems into one of four categories:

01 unacceptable risk
02 high risk
03 limited risk
04 minimal to no risk

Australia has begun designing its own frameworks and building capabilities to support the development and adaptation of responsible AI. For instance, the Australian Human Rights Commission's *Human Rights and Technology Final Report* recommends checks and processes that will ensure technologies are developed and used in ways that are inclusive and accountable, and with robust human rights safeguards.[09] The Australian Government's Department of Industry, Science, Energy and Resources launched a voluntary AI Ethics Framework comprising eight principles to help organisations and governments build AI systems based on good governance and ethics.[10]

08 OECD.AI (2021), 'National AI policies & strategies'. Retrieved from: https://oecd.ai/en/dashboards.

09 Australian Human Rights Commission (2021), *Human Rights and Technology Final Report*. Retrieved from: https://humanrights.gov.au/our-work/rights-and-freedoms/publications/human-rights-and-technology-final-report-2021.

10 Department of Industry, Science, Energy and Resources, Artificial intelligence, 31 March 2022. Retrieved from: https://www.industry.gov.au/policies-and-initiatives/helping-industry-and-businesses-

AI easily reaches beyond national and even regional borders, so evolving regulations must incorporate a global perspective. The US-EU Trade and Technology Council discussed AI policy and issues at its inaugural meeting in September 2021, reaching a consensus that AI regulation should follow a risk-based approach.[11] The Global Partnership on Artificial intelligence brings together 25 nations, including Australia, aiming "to bridge the gap between theory and practice on AI by supporting cutting-edge research and applied activities on AI-related priorities".[12]

The safe development and integration of AI into our society is a clear priority globally. But although laws and regulations such as these are essential for AI governance, it is also important to consider the ethical issues of using AI and the continuing role of human judgement.

Design and Distribution Obligations

The recently updated Design and Distribution Obligations (DDO), found in part 7.8A of the *Corporations Act 2001* (Cth), require issuers to "identify in advance the consumers for whom their products are appropriate and direct distribution to that target market."[13]

In effect, this means issuers and distributors of financial products must develop a Target Market Determination (TMD) for each product they make publicly available. The TMD must specify the target market for the product, any restrictions on distributing the product, and other features including product review periods and complaints reporting procedures.

Companies can use AI to efficiently identify which customers fall into the TMD target market for a specific product, but using AI for this purpose involves some ethical risks. If the AI system is not coded in a way that aligns with the sentiment or needs of the targeted client specified in the TMD, the company may inadvertently target unsuitable – and potentially vulnerable – individuals. Using biased data to identify these target groups carries the same risk. There are many factors to consider when determining who is a target client: would an AI make the same decision as a human, and if so how often? It is important to consider these issues and realise that AI, even when operating in accordance with the law, can cause harm if risks aren't considered and addressed.

Although introducing the DDO is a step in the right direction, AI developers and business managers still need to apply judgement in applying the technology. For this legislation to be effective, it is essential for these developers and managers to have a robust understanding of the ethical risks of AI-driven marketing, and the steps they can take to mitigate those risks.

---

harness-technology/artificial-intelligence.

11 European Commission (2021), EU-US Trade and Technology Council Inaugural Joint Statement. Retrieved from: https://ec.europa.eu/commission/presscorner/detail/en/STATEMENT_21_4951.

12 The Global Partnership on Artificial Intelligence (2022). Retrieved from https://gpai.ai.

13 ASIC's approach to enforcement after the Royal Commission | ASIC - Australian Securities and Investments Commission

PART 03

Sources of AI-related ethical risks

A few ‘bad apples’ like rogue programmers or data managers, cannot be held responsible for an AI-driven marketing campaign’s ethical failings. Instead, problems can generally be traced back to issues with governance. Major sources of risk include issues with the management of data and poor communication between AI developers and business managers.

Data governance issues

AI technology is still nascent. As organisations begin to implement AI-driven marketing solutions, they may encounter significant data governance issues. An AI solution is only ever as good as the data it has been fed, so marketers must address the quality of data. At the same time, they must focus on how that data is managed to optimise the success of the system and reduce potential risks – all the while adhering to data privacy laws.

Good data governance identifies who is responsible for data, AI algorithms and reporting, and ensures all data is appropriate, accurate, timely and truthful. Following are some common data governance issues organisations should consider and address.

- **Siloed databases:**
  These are databases that do not 'talk to each other' or share definitions. They can exist when different legacy systems are required to work together. While siloing can be a problem for all financial institutions to some extent, it is often a larger problem for smaller organisations that have fewer resources available to 'de-silo' or better integrate their data. Common solutions include establishing agreed definitions; centralising and cleaning up databases to create one primary data store; and determining which teams should have access to a database.

- **Data accuracy:**
  When working with multiple databases, application programming interfaces and data sources, a clear picture can quickly become clouded. Something as simple as a dataset not refreshing quickly enough can cause an AI algorithm to incorrectly assume a person is or is not in the target audience for a product. It is important for data to be accurate and up to date at all times to avoid these sorts of problems.

- **Incorrect or missing customer flags:**
  When flags are wrong or missing, customers may be targeted for products or services that are inconsistent with their needs, or may miss out on offers that are appropriate. Automation raises questions about how to exclude customers from certain product marketing and for how long, including whether these flags should be permanently hardcoded or only applied for a defined time period (such as during periods of financial hardship, or while individuals are under a certain age).

- **Deferral of responsibility:**
  When too many people are designated as data owners or managers and responsibility is split, there is a risk of one person (or multiple people) assuming that someone else is in charge of making decisions. The same issue exists when decision making is split across AI systems, or between people and AI. Establishing and communicating clear lines of responsibility – especially among business managers and AI developers – ensures that data management duties are clear, and that people not AI ultimately remain accountable.

Control and monitoring gaps[14]

The data governance challenges on the previous page highlight the types of gaps in control and monitoring that exist in organisations starting to use AI. One of the fundamental issues with AI in any context is the 'explicability' of its decisions: it is often difficult or impossible to know exactly how the machine learning algorithm has made its correlations. This lack of explicability represents a gap in control and therefore accountability. The issue is particularly acute in situations where the reasoning behind a decision has traditionally been important, such as in professions like medicine.

One of the fundamental issues with AI in any context is the 'explicability' of its decisions: it is often difficult or impossible to know exactly how the machine learning algorithm has made its correlations.

In decisions made without AI, there is a traditional flow of authority: all decisions follow a human-to-human chain of responsibility. A board of directors may set the strategy and high-level policies for an organisation, delegating key functions to management. Management then establishes business objectives based on the strategy, setting detailed rules, policies and procedures, and delegating execution to the front-line team, along with the authority to exercise discretion when the rules are silent or unclear. The front-line team interacts with customers by following those established rules, policies and procedures, and using discretion when required.

In AI decision making, the front-line team is replaced by two 'teams': a team of AI developers tasked with encoding machines with available data and objectives articulated by management, and a 'team of machines' (the front-line automated systems) that learn precise rules from the data, then make decisions aimed at most effectively realising the set objectives. In this sense, front-line automated systems are empowered to both generate and execute rules.

14 The analysis in this section has been provided by Gradient Institute (2022), referencing *De-Risking Automated Decisions: Practical Guidance for AI Governance*. Retrieved from: https://gradientinstitute.org/wp-content/uploads/2022/03/gradient_minderoo_report.pdf.

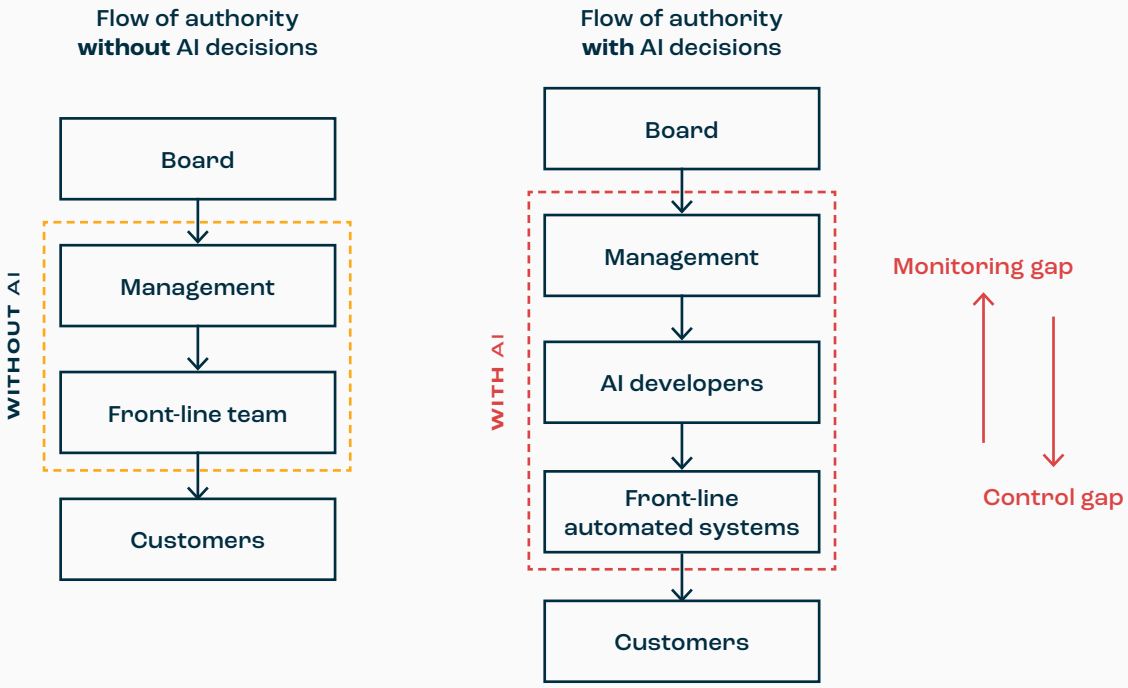

This new decision-making structure changes and challenges the traditional governance approach, introducing two main gaps:

- **Control gaps**
  - Management may be ill suited to delegating to AI developers, especially if they have insufficient technical understanding of AI to effectively instruct them.
  - AI developers may be unfamiliar with the business domain, so they may misunderstand instructions from management, including industry-specific jargon.
- **Monitoring gaps**
  - The highly technical nature of AI development work may lead to a communication gap between developers and management.
  - AI systems see a data point, not a person. The data representing a customer is a very crude approximation of a real person, which AI systems are not equipped to understand.

Although this is not an exhaustive list of the problems inherent in AI-driven marketing, it can help pinpoint the source of some key issues. To better understand their risk exposure, organisations should engage in deeper, context-specific analysis.

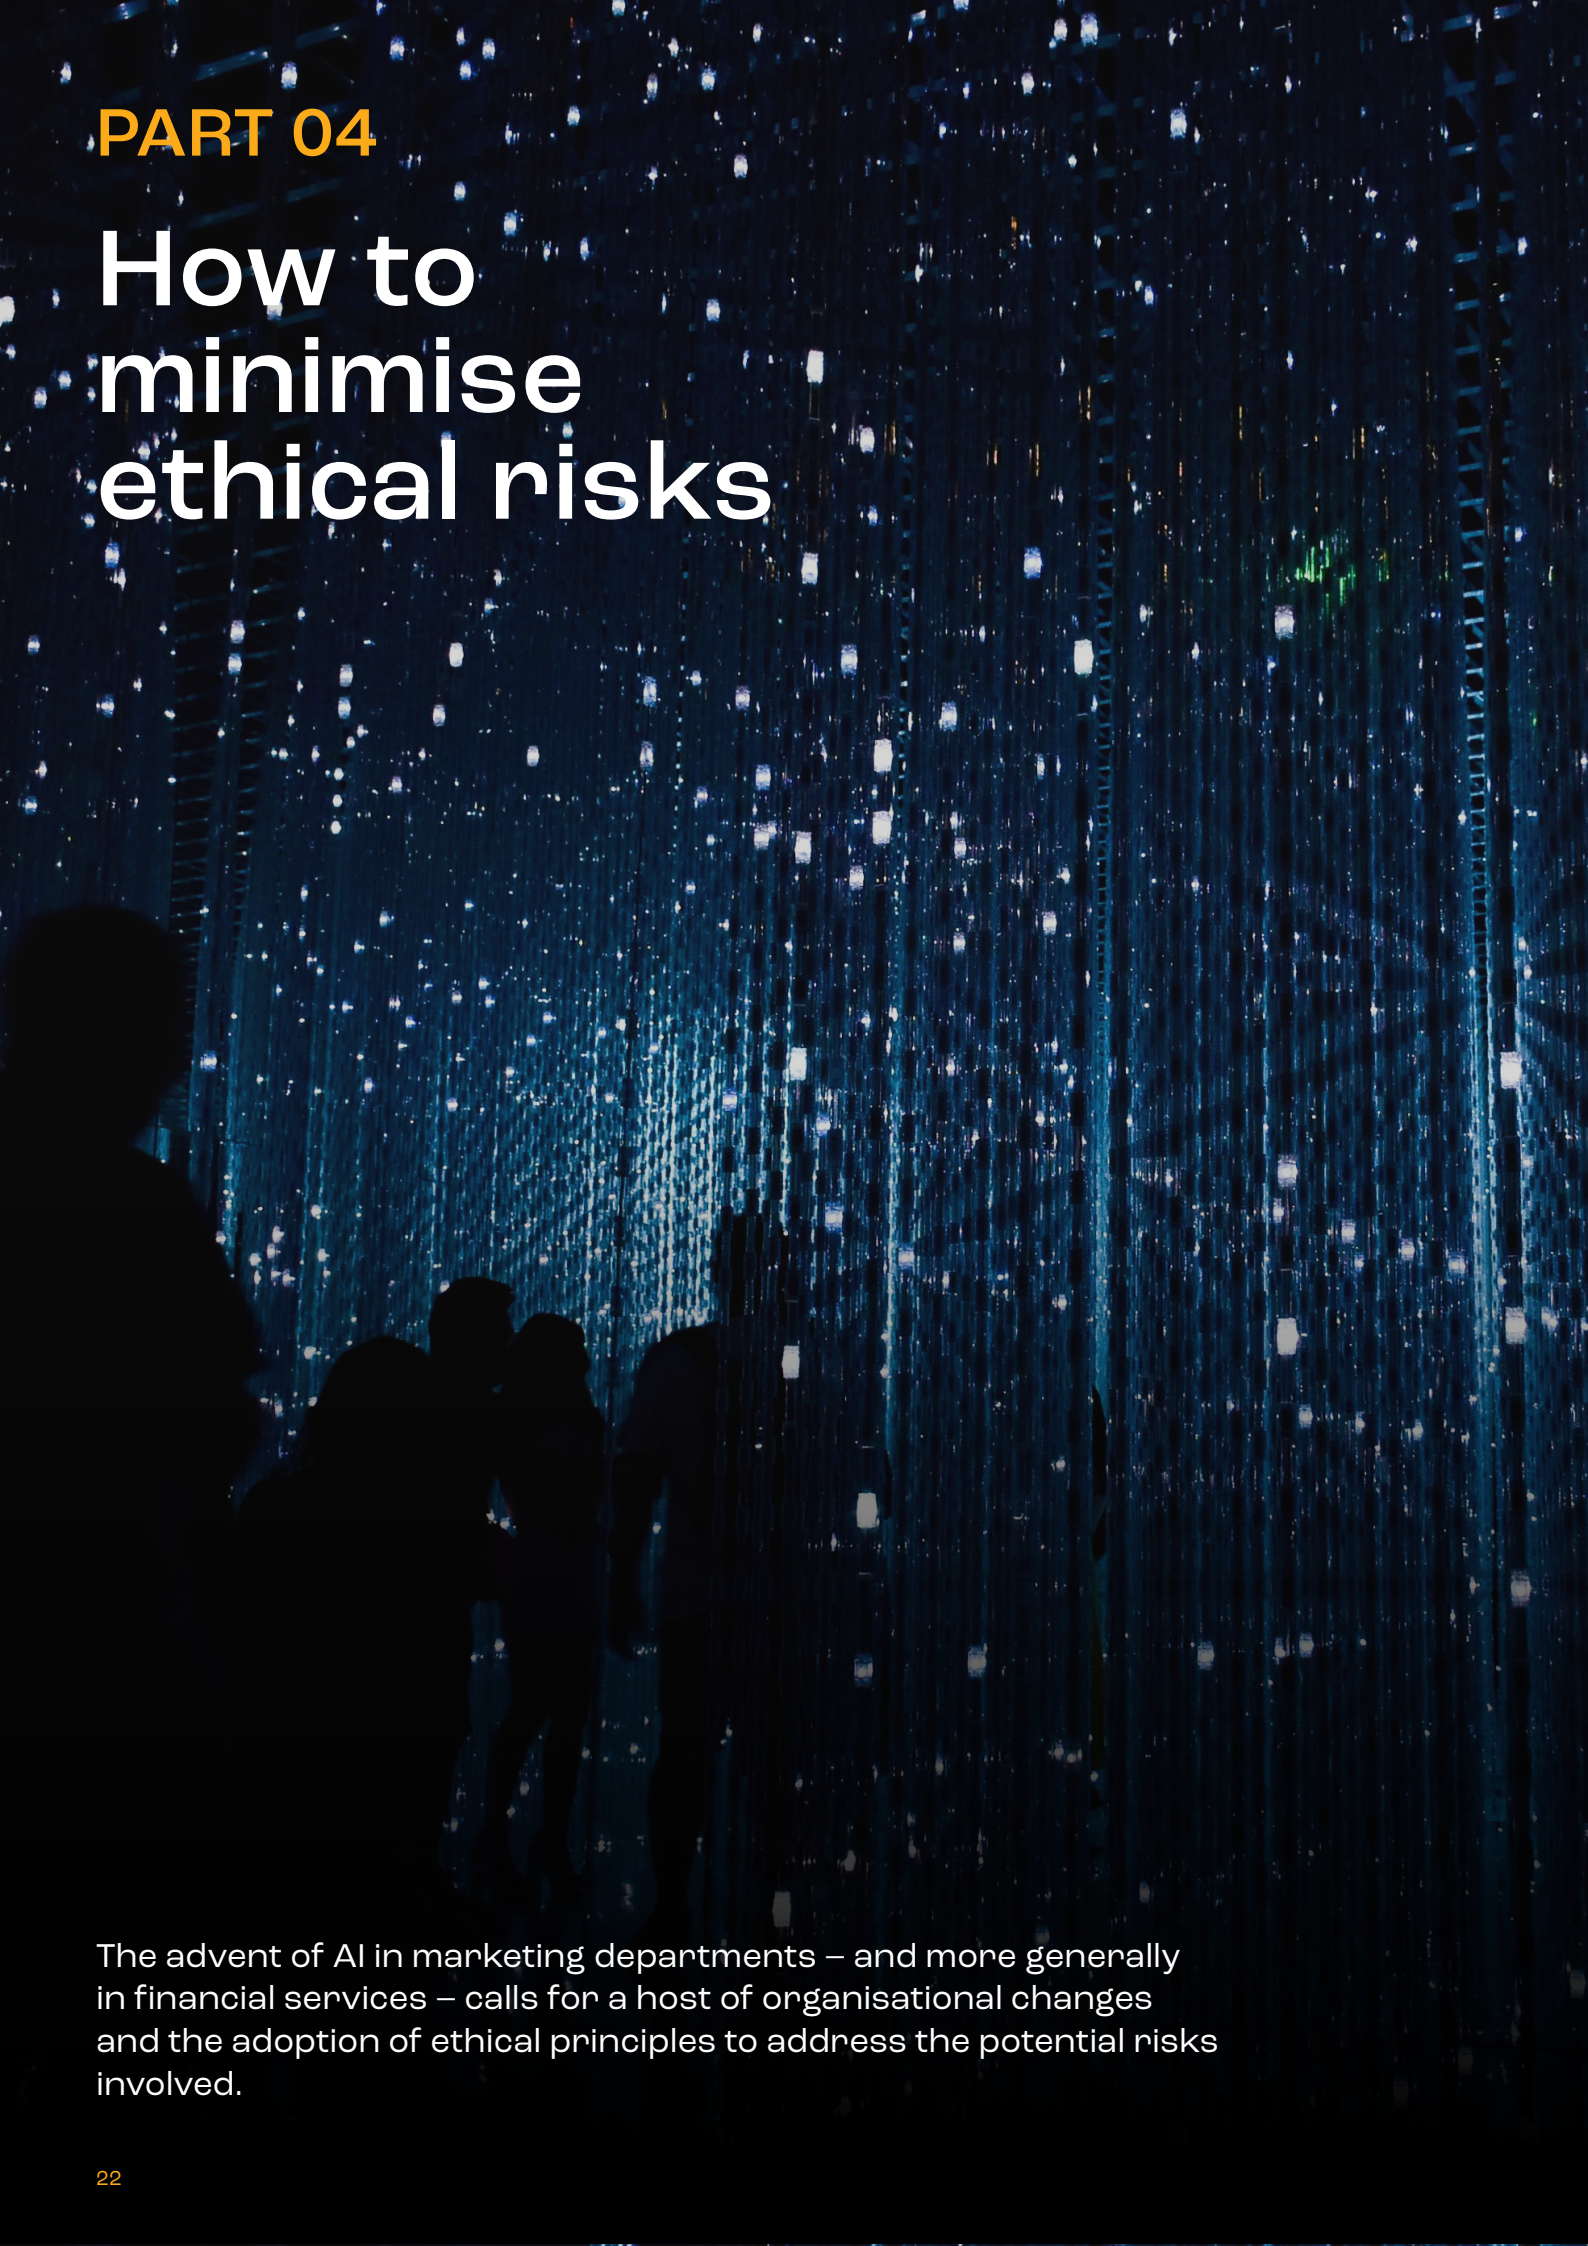

PART 04

How to minimise ethical risks

The advent of AI in marketing departments – and more generally in financial services – calls for a host of organisational changes and the adoption of ethical principles to address the potential risks involved.

Organisational changes

The following recommendations have been selected from Gradient Institute's *De-Risking Automated Decisions: Practical Guidance for AI Governance* report.[15] They are relevant but not specific to the marketing industry.

People and culture

- **Incentives:** Organisations should consider ways to incentivise development of and adherence to responsible AI practices among product owners, senior management and AI developers.
- **Training:** Organisations should provide training in responsible AI to AI developers, system owners, system integrators, business leads and boards, ensuring they have the expertise and awareness required to effectively govern AI systems. See Appendix 1 for a list of the skills needed in each role.
- **Stakeholder engagement and co-design:** To foster greater legitimacy, organisations should give the diverse stakeholders who will be potentially impacted by an AI system the opportunity to contribute to its design and operation.

Routines and processes

- **Audits:** Independent technical audits can unearth technical flaws in the design, construction or operation of an AI system; give confidence to non-technical senior management or board members that systems have an appropriate level of technical oversight; and provide a mechanism for organisations to receive advice or guidance on better practices.
- **Questionnaires and checklists:** A large element of risk management involves identifying potential problems that could occur. Questionnaires and checklists provide a standardised starting point for pinpointing and reflecting on these potential issues. When used well, they can also encourage imagination and creativity that might identify blind spots within the development team, and cultivate a more active and engaged attitude towards the risk management process. Some examples include the Open Ethics Canvas[16], MIT's AI Blindspot tool[17] and the Ethical OS Toolkit.[18]

Technical practices and tools

- **Dashboards and control panels:** The people responsible for a system must understand its impacts if they are to correctly monitor and control it. As these impacts can change over time, managers may benefit from automated visualisation of real-world impacts via a software dashboard that includes the system's impacts and ways to control them.

---

15 Gradient Institute (2022), *De-Risking Automated Decisions: Practical Guidance for AI Governance*. Retrieved from https://gradientinstitute.org/wp-content/uploads/2022/03/gradient_minderoo_report.pdf.

16 A tool for developers and system owners to initiate conversations, aimed at revealing blind spots in the development process: https://openethics.ai/canvas/.

17 A resource designed to spot unconscious biases and structural inequalities in AI systems: https://aiblindspot.media.mit.edu.

18 An aid for organising workshops aimed at identifying risks in AI systems: https://ethicalos.org/.

Adoption of ethical principles

Many AI-related ethical principles exist in industry and government contexts, often focused on specific sectors like financial services and healthcare, and even unique use cases like self-driving cars. The framework presented here is based on a meta-analysis of 84 of the most well-known AI ethics frameworks worldwide, adapted to the field of marketing.[19] It takes a multi-stakeholder perspective, aiming to frame each principle in terms of the company, customer and societal perspective.

Multi-stakeholder model of AI ethics in marketing

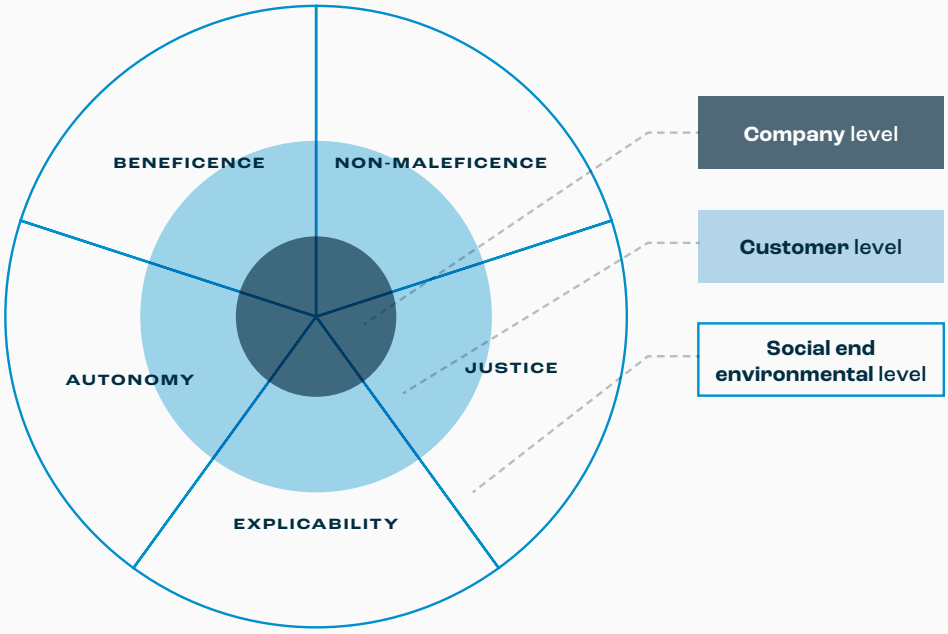

Source: Hermann (2021)

This framework is designed to remind developers, business managers and leaders of the ethical significance of the products they are creating. To avoid the risk of 'AI ethics washing' – where companies espouse high-level principles without much practical implementation – these principles need to be read in the context of the organisational structures and technical requirements that they are applied to in practice.

19 Hermann, E. (2021), 'Leveraging Artificial Intelligence in Marketing for Social Good—An Ethical Perspective', *Journal of Business Ethics*. Retrieved from https://doi.org/10.1007/s10551-021-04843-y.

- **Beneficence (do good)**
  AI has the potential to personalise marketing in a way previously unimaginable. This personalisation can be used to improve the financial outcomes of customers by directing them towards products and services suited to their needs, and improving their financial literacy by delivering customised advice and suggestions.
- **Non-maleficence (avoid harm)**
  AI-driven marketing has the potential to target customers in inappropriate or harmful ways. This can arise due to poor data governance, or when the underlying data used in algorithms that contain biases or inaccuracies.
- **Justice**
  AI-powered personalisation strategies in marketing can discriminate based on demographics – including racial background, socio-economic status and individual psychological factors. AI systems are particularly at risk of reinforcing gender, age and racial disparities, prejudices and stereotypes.
- **Explicability**
  It should be possible to interpret and explain how an AI system works. With knowledge of the system's intentions, data inputs and sources, and the relation between inputs and outputs, human experts should be able to understand its results, predictions, classifications and recommendations.
- **Autonomy**
  At a company level, governance mechanisms should ensure humans are kept 'in the loop' in appropriate ways. At a consumer level, AI applications should be designed so that they support responsible decision-making processes through personalisation and recommendation tools.

This framework can help organisations build ethics into the AI system design process. With greater communication between AI developers and business managers, developers might be able to provide evidence of how they have embedded certain organisational or AI-specific values into the system. These could then be tested and confirmed periodically by forms of auditing and assurance mentioned in the section above.

With greater communication between AI developers and business managers, developers might be able to provide evidence of how they have embedded certain organisational or AI-specific values into the system.

Questions to ask when deploying AI

Although the following questions are not specific to AI-driven marketing in financial services, they are a useful starting point for addressing the ethical issues that arise in this field and in AI generally.

- Have you **consulted widely** with stakeholders to identify the potential harms the system may cause?
- Have you **determined the ethical objectives** you need to pursue to control these harms, and how to ensure the AI system achieves them?
- Have you identified the **people at risk** of being systematically disadvantaged by the AI system and ensured that actions are taken to protect them?
- Are there processes in place to **document the AI systems** that affect people's lives, including their purpose, risks, key design decisions and justifications, performance, and oversight responsibilities?
- Are there processes for **continually monitoring** AI systems; measuring them against business and ethical objectives; searching for unintended harms; and building in mechanisms for risk review, resolution and mitigation?
- Do staff have adequate **training** in the novel risks of AI systems, and do they understand their role in controlling those risks?
- Have existing **risk management frameworks** been extended to incorporate risks introduced or potentially amplified by the use of AI systems?

Conclusion

Marketing is not ethically neutral, nor is it intrinsically bad. The same can be said for AI. Ethical risks have always existed in marketing, including in products that harm people (such as tobacco), products that may lead to future harm (such as loans that may lead to default), and misleading or manipulative advertising. But there is a significant difference between human teams and AI systems when it comes to marketing. Humans, in general, have access to common sense and an intrinsic understanding of right and wrong. AI is based entirely on the rules it is programmed to follow; it lacks any moral constraint, unless that restraint is explicitly programmed.

The risks of AI-driven marketing in financial services are very real – as are the opportunities. And while this report is directed specifically at the use of AI in financial services marketing, it can serve as a template for other functions within financial services organisations that are using AI to improve products, processes and customer experiences.

By acknowledging the ethical issues across the financial services spectrum and taking steps to understand and address them, organisations can seize the opportunities while mitigating the risks, benefitting their customers as well as their own commercial objectives.

APPENDIX

AI skills needed in different roles

The below table, adapted from Gradient Institute's De-Risking Automated Decisions report[20], summarises the types of skills needed in different roles within organisations.

Role	Skills needed
Developer Data scientists and data engineers developing, deploying and validating AI systems	– Understand the link between design decisions and the impacts of AI systems. – Communicate the implications of design decisions to non-technical audiences. – Quantify non-commercial ethical objectives into mathematical measures, understand their limitations and integrate them into the design of AI systems. – Identify unintended behaviour in AI systems, investigate causes and apply mitigation strategies. – Understand the ways in which AI systems can ingrain existing societal biases. – Understand the business domain in which they are operating.
Business or system owner and integrator Business leaders responsible for the operation of AI systems (whether developed in house or procured externally)	– Understand the novel risks introduced when AI systems replace manual systems. – Define business and ethical objectives for AI systems. – Decide how to balance those objectives when they compete. – Understand and communicate the implications of the system's impact on technical design decisions.
System review committee A cross-disciplinary committee that ensures system owners follow responsible AI processes in designing, developing and deploying AI systems	– Assess the degree to which a system's design and operation aligns with its stated objectives, and the values and priorities of the organisation. – Understand the limitations of the system as implemented and its potential unintended impacts. – Ensure the organisation maintains an effective process for responsible AI governance.
Board of directors The governing body that sets AI system management and monitoring policies	– Understand the risks introduced when AI systems replace manual systems. – Set the strategic direction on AI use (including when not to use AI). – Ensure the organisation maintains an effective process for responsible AI governance.

20 Gradient Institute (2022), *De-Risking Automated Decisions: Practical Guidance for AI Governance*. Retrieved from https://gradientinstitute.org/wp-content/uploads/2022/03/gradient_minderoo_report.pdf.

A report by
The Banking and Finance Oath
Young Ambassadors

This report was written by The Banking and Finance Oath (BFO) 2021 Young Ambassadors. Created in 2013, The BFO is a voluntary oath that financial services professionals can take to demonstrate their commitment to ethical practices and to drive positive change for the industry and society more broadly.

Each year, The BFO chooses a group of young industry professionals to promote the oath, and supports them in strengthening ethical practices.

To find out more about The BFO and how you can take the oath, visit thebfo.org.au.

Report authors

Aoife Donnelly
MetLife Australia

Ben Robinson
Deloitte

Emma Roberts
Australian Prudential Regulation Authority

Jack Thompson
Bank Australia

Nancy Loomba
National Australia Bank

Acknowledgements

Thanks to Gradient Institute for guidance, discussion and review through the project.

GRADIENT INSTITUTE

For more information about Gradient Institute and its ongoing work to ensure responsible use of AI in the finance industry, see here.

AN INITIATIVE OF

thebfo.org / ethics.org.au