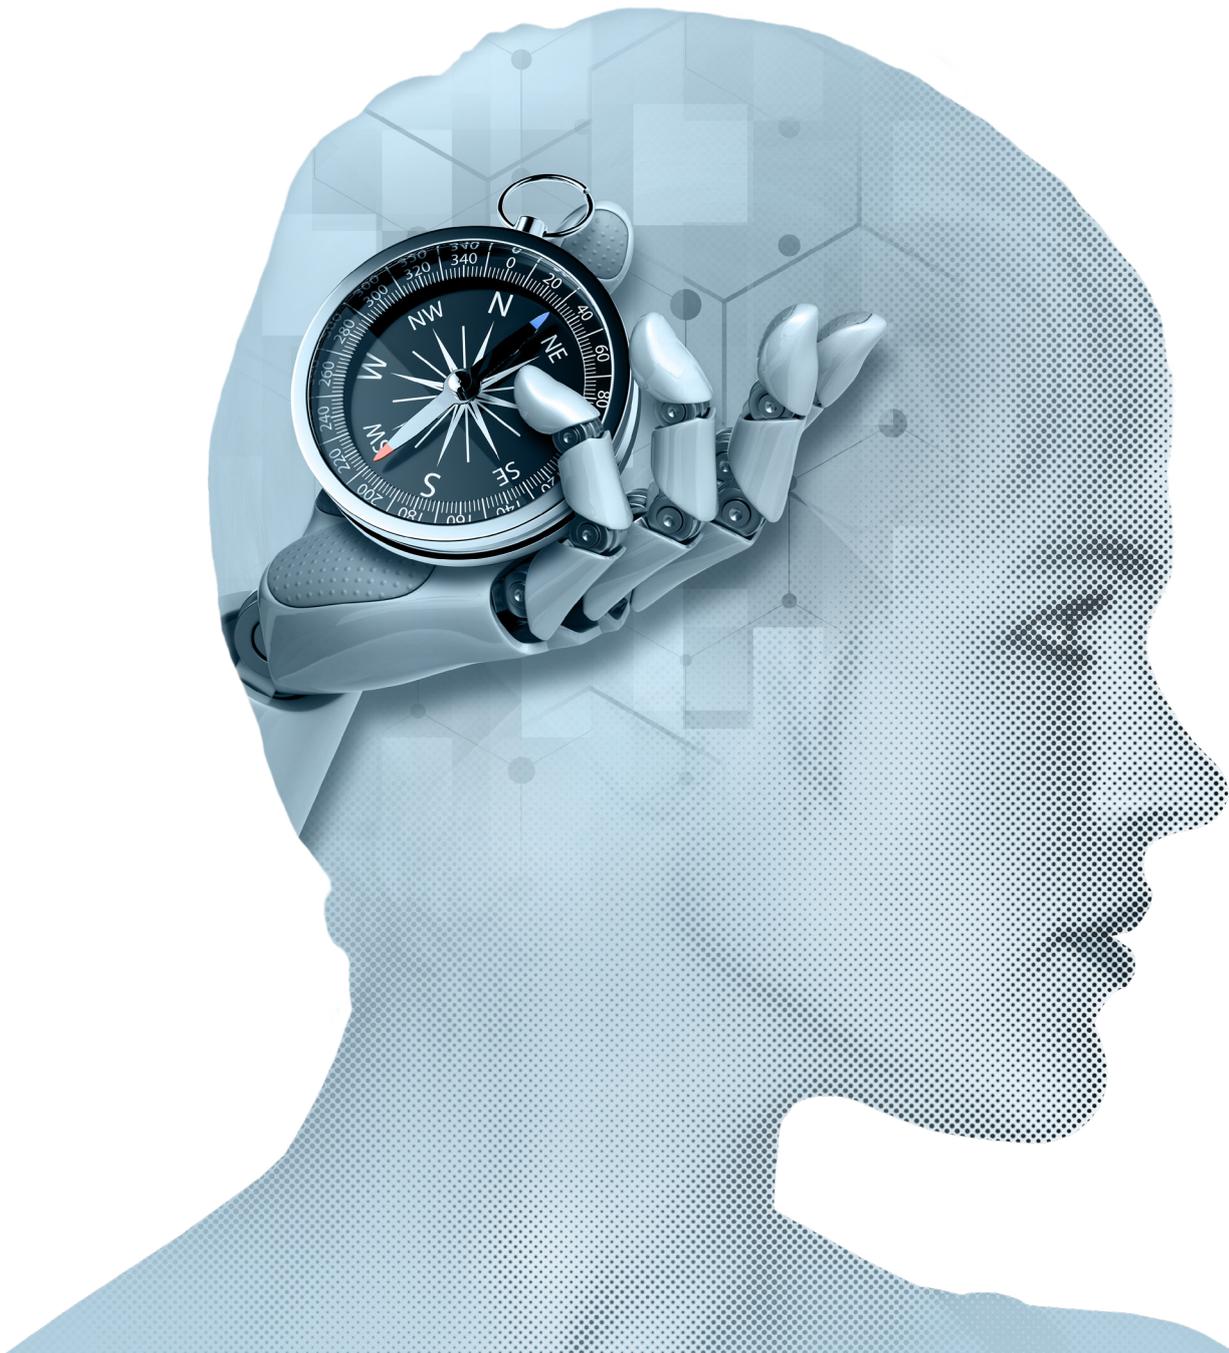


# *New technologies, ethics and accountability*



# Contents

EXECUTIVE SUMMARY .....	1
INTRODUCTION .....	2
BUILDING CONFIDENCE IN TECHNOLOGY .....	3
Understanding the problem .....	3
The role of ethics .....	4
The role of accountability .....	5
Regulatory support and intervention .....	6
ETHICS FRAMEWORKS .....	7
Identifying ethics principles .....	7
Practical meaning of key principles .....	8
Ethical dilemmas, tensions and trade-offs .....	10
The risks of bias and unfair outcomes .....	12
ACCOUNTABILITY FRAMEWORKS .....	14
Maintaining human accountability .....	14
Explaining and justifying decisions .....	15
Organisational accountability, governance and oversight .....	16
Personal responsibility and culture .....	18
PRACTICAL APPROACHES .....	19
Businesses: identifying good practices .....	19
Businesses: implementation challenges .....	21
Accountants: updating professional ethics .....	22
Accountants: broader contribution of the profession .....	23
Policymakers: establishing incentives and regulation .....	23
Policymakers: building social norms and consensus .....	25
NEXT STEPS .....	27
Practical guidance and knowledge sharing .....	27
Regulatory and professional thinking .....	27
Public engagement and debate .....	28
Engage with ICAEW .....	28
APPENDIX - QUESTIONS FOR BOARDS ABOUT NEW TECHNOLOGIES, ETHICS AND ACCOUNTABILITY .....	29
FURTHER READING .....	30
END NOTES .....	31

## *Executive summary*

### **Ethics and accountability underpin confidence in business use of new technologies**

Technology can transform people's lives for the better. But it can also be used in ways that harm people, or have unexpected and undesirable consequences. Concern about these harms is reducing confidence in business use of technology, and raising risks that the full benefits will not be realised.

Taking an ethical approach can help businesses increase confidence that they are 'doing the right things' with technology. Ethics translates moral philosophy into practical principles so it can help businesses to focus on the good they can do with technology and avoid potential harms.

Strong frameworks for accountability provide vital support for ethics. Accountability means that someone needs to explain what they did, and it increases transparency, encourages good behaviour and provides someone to blame when things go wrong. Therefore, it helps to translate good intentions into meaningful action.

### **Technology creates new tensions and pressures for ethics and accountability**

There is growing consensus around the key ethics principles for technology. However, translating these principles into practice raises difficult questions about the specific meaning of concepts like fairness and justice, for example, or how to identify harms caused by technology. Technology also heightens or creates new tensions between ethics principles, between the interests of different groups, including shareholders, customers and employees, and between diverse cultural values.

Furthermore, technology puts pressure on traditional accountability structures. Maintaining human control and ensuring decisions can be explained is getting more complex. Accountability needs to be shared effectively between technology specialists, business users and senior management, avoiding gaps and ensuring clear hand-overs. To support the right behaviours, leaders also need to encourage an ethical culture around technology, as well as a diverse workforce.

### **Practical action by businesses is in the early stages and can be supported by regulators**

Implementing ethics and accountability frameworks in practice is proving to be a complex challenge for businesses. There are some emerging areas of good practice that can help, such as Ethics Advisory Boards and Algorithm Impact Assessments. The context-specific nature of ethics, though, makes it difficult to develop and adopt generic tools.

Regulation can sometimes be needed to support ethics and accountability in businesses, providing red lines, creating certainty and driving behavioural change. But good regulation is difficult where we do not have clear settled social norms or where innovation is happening quickly. Furthermore, good business behaviour is generally grounded in an ethical culture rather than being driven by compliance activities.

### **New technologies create ethics challenges and opportunities for the accountancy profession**

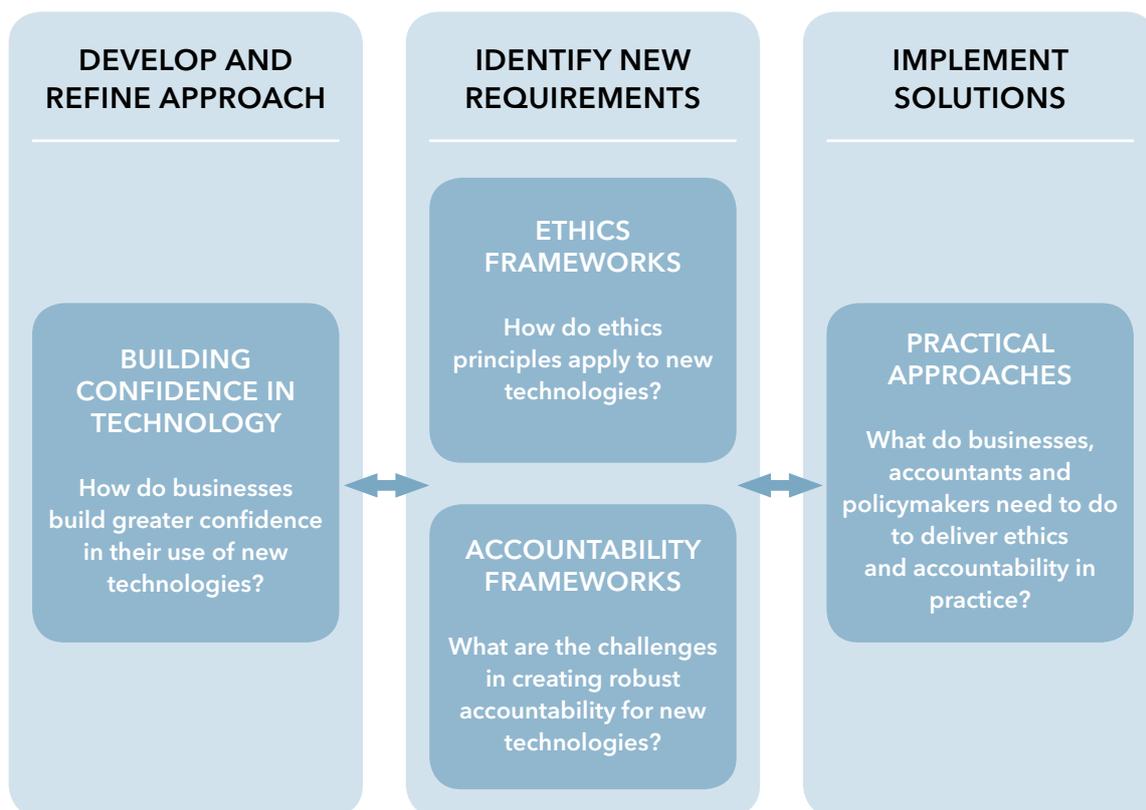
It is critical for the accountancy profession to maintain confidence in its ethical use of technology. This presents many new questions for the profession around the practical application of its ethics code, how threats to ethical behaviour are changing and the nature of professional judgement. The profession should also consider how its long history of ethics and accountability around financial data can help businesses improve practices in other areas of data and technology use.

## *Introduction*

Technology can help businesses be more efficient, deliver better customer services and generate greater shareholder and social value. It can enable new products and services, and can ultimately transform people's lives for the better.

But technology can also be used in ways that harm people, or have unexpected and undesirable consequences, and concern about these risks is reducing confidence in business use of technology. If businesses are to make the most of these new opportunities, they need to show people that they are 'doing the right things' with technology.

This report aims to support practical thinking about how to address these growing concerns and is based on two key concepts - ethics and accountability. Our approach builds on the accountancy profession's grounding in ethics and accountability around financial data, as well as ICAEW's obligation to work in the public interest. It also underpins a wide range of work across ICAEW to develop practical support for businesses and the profession in specific areas, and is shown below.



## *Building confidence in technology*

Businesses need to build more confidence in how they are using new technologies. While the notion of confidence incorporates many elements, such as functionality, convenience, safety and reliability, 'doing the right things' with technology is an increasingly important part. This section outlines some key building blocks to increasing confidence in this area.

### UNDERSTANDING THE PROBLEM

Technology is raising new challenges and ethics dilemmas that touch on some fundamental questions about society and human nature. Among these are:

- **Privacy and autonomy:** extensive gathering of personal data and growing surveillance is reducing our ability to maintain private thoughts and actions. Furthermore, high levels of personal targeting enables online manipulation and may limit our choices or access to services.
- **Bias, exclusion and social mobility:** increasing reliance on algorithms and data in many areas of decision-making risks mirroring and entrenching bias in societies, increasing exclusion from services, delivering unfair outcomes, and reducing social mobility in the process.
- **Growing up and online harms:** the challenges of growing up are being amplified by trends in social media, with reports of online bullying, mental health problems, radicalisation, grooming and teenage suicides often linked to harmful images or online interactions.
- **Undermining democracy:** fake news, harassment and abuse over social media, micro-targeting of political messaging and interference by foreign governments all contribute to concerns about the impact of technology on political debate and democracy.
- **Jobs and inequality:** trends in automation, robotics and artificial intelligence (AI) are fuelling worries that machines will take over large amounts of jobs that will not easily be replaced, leading to high levels of unemployment or greater inequality in societies.
- **Automated warfare:** the threat of autonomous drones being used in warfare raises fears over machines making life and death decisions, and humans ultimately losing control.

Concerns about things such as unequal societies, negative influences on teenagers and innovative forms of warfare are not new. So why are concerns growing today, and what new problems are technologies actually creating?

In many cases, it is the sheer power, scale, and reach of new technologies that creates or amplifies potential harm. Social media is a good example of this transformation. Privacy breaches or harassment are not new problems but where people are constantly online, it becomes impossible to get away from bullying or abusive behaviour. As a result, the harm caused may be qualitatively transformed by these new technologies, and established coping mechanisms may no longer work.

Data-centred technologies – including big data and AI – also raise specific issues. They can enable more efficient decisions grounded in hard evidence or accurate predictions, rather than assumptions or guesswork. However, this is not always a comfortable process and it can involve difficult (and new) choices about how to treat people.

Algorithmic decision-making, especially where it is automated, has many strengths in terms of speed, efficiency, accuracy, consistency and objectivity, but it typically reduces the extent to which we can show discretion and recognise individual circumstances. AI techniques, such as machine learning, put less emphasis on programming explicit rules into computers – computers learn from patterns in data to produce outputs, rather than following pre-defined rules.<sup>1</sup> Such processes place greater reliance on technical specialists, making it harder for non-experts to understand how decisions are being made, and ultimately take decision-making further away from explicit human control.

Emphasising the specific impact of technology helps to contain discussions and focus responses on new problems to be solved. In many cases, technology is heightening long-standing tensions or stressing traditional ways of resolving issues, rather than creating brand new problems. It is important to recognise this broader context and link to existing thinking about problems from other business areas where needed.

### THE ROLE OF ETHICS

Taking an ethical approach can help businesses increase confidence that they are 'doing the right things' with technology. Ethics translates moral philosophy into practical principles so it can help businesses to focus on the good they can do with technology, as well as avoid potential harms or unintended consequences. This can operate at a number of levels, including:

- business ethics, which influences the decisions and actions of people working in businesses;
- professional ethics, which focuses on the ethical requirements for professionals such as accountants; and
- personal ethics, which reflects our own moral beliefs and values.

The best known example of ethics principles is in medicine, with the Hippocratic Oath taken by doctors grounded in ancient Greek philosophy. Medical ethics puts four core principles at the heart of any decision or action taken by a doctor:

- beneficence - promoting the wellbeing of others;
- non-maleficence - doing no harm to others;
- respecting human autonomy - respecting the right of individuals to make their own choices; and
- respecting justice - distributing scarce (medical) resources in the best or fairest way.

These principles capture core moral values and reflect the highly personal decisions that doctors have to make, including life and death decisions. They are also common to many areas of ethics.

Professions such as accountancy and law have strong ethical codes. While they are grounded in moral values, these ethical codes often operate in practice as regulatory instruments that are enforced through practicing licences, exclusion from membership of professional bodies and fines. They reflect the reliance placed on professionals by non-professionals, and the knowledge and skills asymmetries between them.

There are five core principles that underpin the role of a professionally-qualified accountant:<sup>2</sup>

- integrity - being straightforward and honest in all professional and business relationships;
- objectivity - not allowing bias, conflict of interest or undue influence of others to override professional or business judgements;
- professional competence and due care - duty to maintain professional knowledge and skill at the level required;
- confidentiality - respect the confidentiality of information acquired as a result of professional and business relationships; and
- professional behaviour - comply with the relevant laws and regulations and avoid any action that discredits the profession.

Therefore, professionally-qualified accountants – whether giving tax advice, signing audit opinions, providing financial or management reports or inputting into strategic business decisions – must ensure that their decisions and actions adhere to the fundamental ethics principles. Alongside these fundamental principles, the ethics code also details threats to ethical decision-making or conduct, such as self-interest, self-review, advocacy, familiarity and intimidation.

While these principles are grounded in the professional responsibilities of accountants, they also reflect the nature of accountancy as an information-based profession. For example, handling confidential information appropriately and resisting pressures to manipulate data and report overly favourable results are common ethical challenges for accountants to deal with.

### **Example: improving accounting ethics through technology**

Most discussions frame ethics primarily as a way of stopping or limiting harm from new technologies. However, technology can play a role in improving ethics in practice and reducing opportunities for ethical breaches. By having better models and more accurate forecasts, accountants may be better placed to resist business pressures to massage results. Greater use of technology in audit could improve independence and reduce risks of familiarity or intimidation by clients.

Regulators can also use technology to improve their monitoring of businesses. Tax authorities, for example, are often well advanced in their use of data and analytics to identify potential tax evasion or aggressive avoidance by businesses.

Technology could lead to an ‘arms race’ in this context. Regulators can get better at spotting tax irregularities but businesses or advisers may be able to devise more sophisticated tax planning or avoidance schemes through technology too. By using sophisticated text analytics, capital markets regulators may be able to analyse narrative statements better to identify risks of misstatement. But businesses could use these technologies to craft statements more carefully in the first place. As a result, we need to recognise both the opportunities and risks in this area.

## **THE ROLE OF ACCOUNTABILITY**

Strong frameworks for accountability provide vital support for ethics. Accountability means that someone has to give an account of what has happened and can help to translate good intentions into meaningful action. It plays a number of roles:

- by making people explain their actions and choices to others, it is closely linked to transparency and openness;
- the need to explain actions encourages people to behave well to avoid later shame or embarrassment – no-one wants to have to explain why they did something bad; and
- it ultimately provides someone to blame when things go wrong, and can be linked to sanctions where needed.

Businesses have formal structures of accountability, which all ultimately lead to the board. When defining accountability, a business needs to consider:

- who is accountable for particular decisions, actions or outcomes – to be most effective, this ultimately needs to sit with specific individuals;

- who they are accountable to - this could be layers of management, or it could be external parties, including shareholders, customers, regulators or others;
- how they will be held accountable - the process by which explanations are demanded and given; and
- any consequences or sanctions - where appropriate, sanctions can be attached for actions or decisions that are unacceptable.

Accountability therefore sits alongside the practical application of ethics, and together they can provide a robust approach to building greater confidence. This is distinct from focusing on compliance with legal obligations. In many cases they lead to the same outcome - good regulation is likely to be built on ethics and accepted social norms. But occasionally, they can diverge, for example, where laws require the disclosure of confidential client information, forcing a breach of professional ethics.

### REGULATORY SUPPORT AND INTERVENTION

While ethics and accountability form the bedrock of greater confidence in businesses, they sometimes need to be supported by regulation. Regulation can provide red lines and ban or limit activities where significant harm is being caused. It can provide greater certainty about what actions are acceptable or not. It can drive behavioural change more quickly and deeply.

Stronger intervention may also be needed because the economic incentives do not drive the desired business decisions. For example, common business models around technology and data tend to exacerbate ethical tensions. Most internet businesses are heavily driven by advertising and the incentives drive companies to gather more data so as to maximise the efficiency of advertising. Business models often benefit from network effects, resulting in first mover advantage. This emphasises speed and a 'winner takes all' approach, and underpins the Silicon Valley approach of trying things, breaking things and fixing things quickly.

Regulators do not have all the answers, though. Regulation is always going to lag behind innovation and people in businesses will need to make their own decisions about the right thing to do when they are at the leading edge. Regulation focuses attention on compliance activities, potentially at the expense of broader accountabilities to customers, users, employees and others. This can lead to a culture that places undue reliance on processes, checklists and rules rather than underlying values.

Furthermore, there may not be clear settled norms or social expectations about different ethical choices, making good regulation very difficult. Indeed, this puts the onus on regulators to define these norms and choices, which may not always be appropriate. As a result, regulation supports rather than replaces effective ethics and accountability in businesses. However, it may sometimes be essential to driving behavioural change and setting clear boundaries.

## *Ethics frameworks*

Ethics principles reflect basic moral judgements about what is right and wrong. They guide our decisions and actions so that our behaviour accords with our moral values. This section highlights emerging ethics principles around technology, and outlines some of the specific challenges to be considered when applying them in practice.

### IDENTIFYING ETHICS PRINCIPLES

Discussions about ethics in fields such as robotics have a long history. However, recent developments in AI in particular have propelled ethics high up the agenda. This is reflected in a proliferation of new ethics frameworks specific to technology, and some examples are shown below.<sup>3</sup>

#### Example ethics frameworks for AI

Source	Principles
AI4People, a multi-stakeholder initiative from the EU which outlines principles that underpin a 'Good AI Society'	<ul style="list-style-type: none"> <li>• Beneficence</li> <li>• Non-maleficence</li> <li>• Autonomy</li> <li>• Justice</li> <li>• Explicability</li> </ul>
UK House of Lords' proposed principles for AI, outlined in their inquiry into AI	<ul style="list-style-type: none"> <li>• Should be a force for common good and the benefit of humanity</li> <li>• Should operate on the principles of fairness and intelligibility</li> <li>• Should respect data protection and not diminish privacy rights</li> <li>• People should have the right to be educated and to flourish alongside AI</li> <li>• Should never be given the power to hurt or destroy human beings</li> </ul>
Google's AI development principles	<ul style="list-style-type: none"> <li>• Socially beneficial</li> <li>• Avoid creating or reinforcing bias</li> <li>• Built and tested for safety</li> <li>• Accountable to people</li> <li>• Follow privacy design principles</li> </ul>
Microsoft's AI development principles	<ul style="list-style-type: none"> <li>• Fairness</li> <li>• Reliability and safety</li> <li>• Privacy and security</li> <li>• Inclusiveness</li> <li>• Transparency</li> <li>• Accountability</li> </ul>

Principles like these all aim to guide business thinking about how to design and use technology. There is an alternative approach to ethics that focuses on personal responsibility. This approach aims to professionalise the field of programming, or data science, and expand or develop individual codes of ethics to be similar to those in professions such as accountancy. The Association for Computer Machinery (ACM) code of ethics, for example, lays out detailed principles for its members, who are computer programmers, around general ethics (eg, do no harm), professional responsibilities (eg, competence), leadership and compliance.

While there are differences between approaches, there is growing consensus around the key ethics principles, which are consistent with the core principles highlighted earlier - doing good, avoiding harm, human autonomy and justice. Some approaches add principles that relate to the design and development of technology. Some are grounded in universal human rights, whereas others are more driven by the specific nature of the technology, such as AI.

A variety of detailed approaches is to be expected, given the range of interest in the subject, the variety of technology uses and the scale of the potential issues. However, it does ultimately risk confusion, duplication of effort and a lack of focus, and there is a growing need for some level of consistency or mapping of core principles and terminology.

Another challenge stems from the global nature of technology. While ethics principles based around individual human rights may sit easily in Western countries, Asian and African countries often have cultures with stronger roots in group-centred philosophies. Many commentators see three evolving but distinct approaches to ethics and regulation in technology.

- A model that prioritises innovation and market-based approaches to shape behaviour, as well as rights to free speech. This is typical of Silicon Valley and has led to a small number of lightly regulated technology and social media companies.
- A model that sees a greater role for regulation in building trust. The EU has already taken a global lead in defining rules around privacy through the Global Data Protection Regulation (GDPR). In contrast to Silicon Valley, New York has also taken a very strong regulatory position around technology in financial services.
- A model that gives greater government control over technology and use of data. This is the approach in China, and given the Chinese government's goal to become the global leader in AI in particular, China is becoming increasingly influential in technology development and standards.

Consequently, different cultures and approaches present significant challenges to defining a universal set of ethics principles. This is not unique to the ethics around technology - managing across cultures is a long-standing business challenge. But the global nature of technology often raises tensions and cultural clashes. Furthermore, technology may create different ethical questions in different cultures, so businesses need to be increasingly sensitive to these issues.

### **PRACTICAL MEANING OF KEY PRINCIPLES**

Identifying high-level principles is an important foundation. However, translating these principles into practice often raises difficult questions about their specific meaning.

All of the approaches highlighted start with the principle of beneficence and emphasise the need to use technologies such as AI to benefit people or society more broadly. But who is defining what that means and how are they determining what is beneficial?

Defining what is good or right is a fundamental concern of moral philosophy that has been debated for thousands of years. There are many different approaches to defining what constitutes good or right actions. Some approaches draw on religious concepts or universal values. Some focus on reasoning and logic.

This results in a wide range of theories that can be used to justify specific decisions. For example, consequentialist approaches, such as utilitarianism, focus on outcomes - they would compare the benefits that the use of technology could bring with the harms it could cause, and make a judgement about the balance. Deontological approaches focus on the duties we have as human beings, and underpin human rights. They might look at how technologies threaten human rights, or how they help us fulfil particular moral obligations.

Defining 'good' can also be approached from a political or economic standpoint. New technologies are usually framed in the context of encouraging economic growth and innovation. But is the ultimate objective to maximise wealth or maximise the distribution of wealth? And how does this fit with concerns about sustainability and the impact of some types of economic growth on the environment?

The principles of non-maleficence, or causing no harm, can also be more complex than it first seems. It is easy to recognise direct financial harms caused to individuals, such as being denied access to credit, or physical harm from rogue autonomous vehicles. However, technology is often involved in less tangible harms, which are harder to universally identify.

For example, breaches of privacy are viewed differently by different people or in different cultures. Some harms just impact specific groups. Representational harms, for instance, reinforce stereotypes of particular groups (eg, the search results for 'engineer' or 'CEOs' predominantly reflecting white males). Then there are more general societal harms, such as damaging democracy through the proliferation of fake news. Many of these harms are hard to detect and measure, may be unexpected and take a long time to manifest themselves.

### **Example: identifying the harm caused by reduced autonomy**

Many argue that autonomy, or our ability to exercise free will, is being diminished by increasing use of big data and algorithms. While personalisation of advertising may be more effective, it means that our choices are increasingly being limited to what companies think we are interested in. Similarly, social media feeds reflect what we have clicked on previously and whatever else is fed into algorithms. This can lead to a narrowing of experience, knowledge and interaction, as well as the loss of serendipity.

We are also at risk of greater manipulation. Very high levels of message tailoring can exploit emotions or biases and work to unconsciously 'nudge' our behaviour. Although advertising has always worked on this basis, the amount of detailed personal data available takes it to a new level and can be seen to reduce our ability to make truly free choices.

These kinds of developments may cause significant long-term harm to individuals and societies. It can exclude people from services, lead us to make sub-optimal choices and reinforce our own beliefs and prejudices. However, as we cannot see what others are being offered or what opportunities we are missing, the harm can be difficult to spot.

Justice, another core ethics principle, is about the application of law. Justice generally goes beyond law, though, and incorporates a more abstract notion of fairness. This makes it a particularly contentious ethical concept. It is also a very important principle for technology, given how a combination of data, algorithms and AI are increasingly driving decisions about the allocation of resources. While these decisions may reflect economic efficiency and mathematical optimisation, do they accord with ideas of fairness and justice?

This is a complex issue for a number of reasons. Fairness often gets used interchangeably with equality, but they are not the same. Equality has a sense of equivalence between people and can ultimately be seen as a mathematical concept that algorithms can cope with. Fairness is a moral principle that is harder to reflect in algorithms. Furthermore, notions of fairness or equality can focus on different things:

- opportunity - this focuses on everyone having the same opportunities to do things and this tends to emphasise treating people the same;
- outcome - this focuses on what actually happens in practice, and would potentially support different treatment of people in order to achieve greater equality or fairness in outcomes; and
- process - this focuses on the procedure followed - the means is more important than the ends.

### Example: fairness and bias in recruitment algorithms

The importance of these debates is highlighted by the growing use of data, algorithms and AI in recruitment practices. Increasingly, the scanning of CVs and initial interviews are conducted through technology, and algorithms sift out the strongest candidates. Using technology in this way should reduce human bias in decision-making and select candidates in a more objective way.

This has not necessarily been the case in practice though. Frequently, systems have reflected an implicit assumption that current employees are the best qualified and so aim to replicate the existing staff base. However, workforces reflect all kinds of human biases (conscious and unconscious), which systems simply repeat and amplify. For example, the **system used by Amazon** had learned to penalise the word 'women's' (as in playing for a women's sports team), and downgraded graduates of some all-female colleges. This reflected the dominance of men in the company, especially in technical roles.<sup>4</sup>

To some extent, these kinds of problems are the result of poorly designed and executed systems, not ethical challenges. Indeed, clear discrimination against particular groups is likely to be illegal in many cases. There may be underlying questions about the ultimate goals of such systems, though, which businesses need to make explicitly. Should organisations be aiming to achieve particular targets to different groups in their workforce? Should their focus be on an open and transparent process so that individuals can understand how decisions have been made?

As a result, high-level principles need to be translated into more detailed guidance about their practical meaning if they are to be useful and this can require difficult and contentious choices. Nevertheless, bringing these issues and choices to light is essential for building a good ethics framework.

## ETHICAL DILEMMAS, TENSIONS AND TRADE-OFFS

Real life decisions often reflect conflicts or tensions between different principles or interests. An ethical dilemma happens when two different outcomes satisfy different principles and the decision-maker needs to choose one option. Trade-offs occur when more of one thing means less of another.

Ethical dilemmas are common in medicine, such as conflicts between the principles of autonomy (consent), beneficence (helping people) and non-maleficence (doing no harm). Should doctors do things to help people when they can't or won't consent? Or do things that harm the patient, even though the patient has chosen it? Euthanasia is a classic illustration of this kind of ethical dilemma - how does a doctor weigh the autonomy of the patient to choose to die against the principle of doing no harm?

Technology often heightens or creates new tensions between ethics principles, between the interests of different groups, including shareholders, customers and employees, and between diverse cultural values.

For example, there are often strong tensions between autonomy and privacy, two common ethics principles in technology. Although consent isn't everything, privacy rules generally put a strong emphasis on it – as long as we choose for organisations to have our personal data, we accept the implications this has for our privacy. However, as the potential harms from privacy breaches increase, this balance might change. Furthermore, the notion of consent is becoming less meaningful as we have less understanding of how our data might be used and we frequently need to consent to data policies in order to access essential services. As a result, the idea of autonomy may become less central and other considerations may gain further weight.

Indeed, the concept of privacy has inherent tensions with many other principles and rights. Privacy prioritises the right of individuals to keep information about themselves private. But there can be many benefits from such data being shared or used by others, such as in medical research. Security is another example. Protecting citizens would typically be viewed as an important ethical goal too – should governments prioritise using personal data to improve security ahead of privacy?

These principles are not necessarily in conflict – well-designed security should also protect privacy; data can be accessed for medical research in ways that don't identify individuals and breach privacy. Nevertheless, these debates and tensions are likely to intensify with technologies such as facial recognition, emotion recognition and biometrics. As these become ever more accurate and widely used, businesses and governments will need to make more difficult choices about how they balance and prioritise these different benefits. This will need structured thinking to enable consistent and justifiable decisions.

### **Example: tensions between rights, harms and shareholder interests**

Social media highlights some of the clear tensions between different interests. Fake news, harmful comments about teenagers or content that glamourises violence, among others, have the potential to lead to harm for individuals or societies. As a result, under a principle of doing no harm, such content should be removed. However, taking down such content conflicts with freedom of speech, especially where the content itself is not illegal – people should be able to share whatever content they like, even if it offends others. Furthermore, anything that generates more interest and clicks serves the ultimate customers (advertisers) better, and creates more benefits for shareholders. How should businesses balance these different interests?

These tensions are also seen as technology increasingly impacts on the quality of people's jobs and lives. For example, shift patterns can be managed far more efficiently from a business point of view, but this can have a detrimental impact on employees' lives, with less regular working or unsociable shifts. Micro-managing productivity with stricter surveillance of employee activities may be more efficient for a business but can add enormous stress for employees, with few breaks, unrelenting targets, and privacy breaches. And replacing jobs with automated systems, robots and AI will create further tensions, as businesses balance the potential benefits to shareholders of reduced costs with the harmful impact on staff.

Of course, balancing business models with ethical considerations is not a new challenge. Indeed, long-standing debates about the primacy of shareholders versus the interests of other stakeholders reflect these tensions. Technology is simply bringing them into sharper relief.

### THE RISKS OF BIAS AND UNFAIR OUTCOMES

As highlighted earlier, principles of justice or fairness are critical in the distribution of resources. Algorithms and data have the potential to allocate resources more objectively and in an optimal manner. In practice, though, this is proving to be a very challenging issue, with many instances of bias and discrimination in models. These include:

- criminal justice, where **probation models** incorrectly predicted that African American offenders were far more likely to reoffend than white offenders;<sup>5</sup> and
- recruitment, where tools, such as that used by Amazon, discriminated against women and favoured male applicants.

Some of these cases simply reflect badly designed or executed systems, or poor quality of data, which can be corrected with good practice, effective data governance and appropriate testing. In other cases, issues relate to data selection and reflect statistical understandings of bias. Where data is not representative of entire communities, for example, the results will reflect that bias. Therefore making data more representative of entire communities is another part of the solution.

Developers can ensure that they remove specific sensitive features from the dataset, such as gender and ethnicity, to avoid those being used in decision-making. However, in practice, other features can be closely correlated with sensitive features and act as proxies - postcode is often a proxy for ethnicity for example. Common proxies can also be taken out of the dataset, but clearly this starts to get more complex to understand and manage. Furthermore, sensitive features can sometimes be important and need to be included, such as in medical diagnosis systems.

Then, there are deeper social elements. Classifications of data reflect social values and definitions that are very hard to change, or even identify and agree on. Indeed, bias exists in models because bias is deeply embedded in societies today. This raises particular issues for machine learning techniques in AI, as they are based on learning from the real world. While developers can do a lot to correct for biases, technology cannot solve these deeper problems. Indeed, we need to build better understanding of bias in society today and how it is being reflected, changed or entrenched by new technologies.

There can also be hard, ethical choices about how we want resources to be allocated. All decisions about resource allocation involve some element of differentiation. New sources of data allow businesses to be more precise in how they do this, based on accurate predictions or hard evidence around past behaviour, for example.

However, this can result in outcomes that may be seen as unfair. Models may penalise or reward people on the basis of things that are out of their control. There may be concerns about disproportionality - small mistakes in one context causing significant disadvantage elsewhere. Models are also making predictions based on the behaviour of a closely correlated group of people, rather than on the basis of that specific individual. All of these can result in seemingly unfair outcomes, even if they are economically rational.

### Example: risks of data excluding more people from insurance

Insurance is a concept that relies on shared risk – we don't know who will be unlucky and need to make a claim and therefore we pool our funds to protect ourselves. As insurance companies get more and more data, they increasingly can predict individual risk with far greater accuracy – from the likelihood of getting cancer to that of crashing your car. Biometric or genetic data, as well as data from fitness trackers or other lifestyle indicators, can be powerful health predictors. Telemetric data, which shows how individuals actually drive their cars, can also provide detailed evidence about the risk level of individuals. This moves insurance to a model that is more economically rational and optimised – those with higher risks pay more, and those with lower risks pay less. But is this fair?

We have no control over our genetics, so if we are unlucky to have poor genes and high risks of particular diseases, is it fair to pay more, possibly to the extent of being uninsurable? Driving late at night can be a high-risk indicator, but is that fair on someone coming home from unsociable shift working? As we get more and more data, these kinds of tensions – between what is economically efficient and what might seem 'fair' – are going to grow.<sup>6</sup>

## *Accountability frameworks*

Holding someone to account for their decisions and actions underpins business and human relationships. It means that someone needs to explain why they did something, and it increases transparency, encourages good behaviour, and provides someone to blame when things go wrong. However, technology puts pressure on traditional accountability structures and this section outlines some of these challenges.

### **MAINTAINING HUMAN ACCOUNTABILITY**

Computers, at least today and in the foreseeable future, do not 'think' in the way that human beings think. They may produce outputs based on logic or pattern recognition, like humans, but they have no sense of moral values. Therefore to think of machines or algorithms making ethical decisions, like humans do, is not appropriate. The human must remain the moral agent and accountable for decisions made by computers. But, how do you ensure that a computer's decisions follow our ethics principles?

There are different schools of thought here. Some propose essentially a rules-based approach, deciding what the right answer is in any given circumstance and ensuring the computer follows that. So, for instance, some researchers are looking at the 'trolley car' dilemma for autonomous vehicles, as outlined in the box below, and trying to define rules for the algorithms to follow in the event of a crash. Others want to teach computers ethical behaviours from the 'bottom up' by learning through data and previous examples, just as we would teach a child to be ethical.

However many of these approaches are still in early stages and, therefore, today we still need to rely on humans maintaining meaningful control over algorithmic decision-making. There are many different models for keeping a 'human-in-the-loop', involving elements such as designing and testing the system, clear and robust controls to identify and correct errors, and providing explanations for why decisions have been made in the way they have. However, as more and more decisions become automated by algorithms, maintaining meaningful control will become harder.

#### **Example: maintaining human control over autonomous vehicles and warfare**

Decisions over life and death must be the most profound ones that human beings make. The development of autonomous vehicles is raising questions about how we can ensure that they make the right moral choices. It seems that autonomous vehicles would be safer than human-operated ones, given all the human errors that happen on roads every year. But still, accidents will happen, and there is substantial research around how an autonomous machine would resolve difficult moral choices.

The trolley-car dilemma is a thought experiment that considers what action should be taken in the event of an accident where someone will be harmed - who should the driver save and who should they hit? While this is an old question, handing control of cars to computers is putting greater focus on it. After all, if we want computers to make the 'right' choices in these cases, we need to work out what the right choices are. However working out what the right answer might be in different circumstances or different cultures is very difficult. And is it even possible to devise a set of rules for every possible situation?

Autonomous weapons also have the potential to make life and death judgements. Using facial recognition to match people to 'targets' may make them more accurate decision makers. However, the idea that machines take their own life and death decisions is deeply uncomfortable for most people, and fully autonomous weapons is not a route that any country has yet gone down. In all cases, there should be a human operator who makes the final decision and there are growing pressures for international rules to ensure that this remains the case.

This raises interesting questions about the interplay between human judgement and algorithms. Algorithmic decision-making is entirely consistent in its application of the rules. But there are times when we want to exercise discretion and bend the rules to accommodate particular circumstances. When should human discretion overrule an algorithm?

Models generally make predictions - what is the process for translating that prediction into a business decision? What are the other inputs, and how much do we rely on the algorithm?

Furthermore, being able to sense when something looks wrong based on experience and knowledge is a core part of professional judgements in areas such as accounting. So what happens when an algorithm gives an unexpected output? Does the decision maker just ignore it because it doesn't feel right? Or follow the computer because it said so? Consequently, working out when we rely on computers, when we rely on human judgement and how the two work together is a **key area to consider**.<sup>7</sup>

### EXPLAINING AND JUSTIFYING DECISIONS

Accountability is based on the need to explain and justify decisions. But most business decisions makers, such as accountants, are not computer scientists. So, where they are relying on AI or other models to inform or indeed make their judgements, how much do they need to understand the technology behind it?

Of course, business decision makers and accountants have always used tools to inform judgements - the fact that they're not programmers hasn't detracted from their accountability. Nonetheless the growing complexity of tools and reliance on the output raises the risks that people won't be able to explain how they've come to make a judgement.

Challenging decisions or management assumptions is another core part of professional judgement. If these decisions are increasingly made by machines, how does an accountant meaningfully challenge management and exercise scepticism? Will there be a need to question the programmer to understand what the algorithm is really doing?

AI raises even more difficult questions. We would expect the programmer to understand what the system is doing, even if the user doesn't have the expertise to do so. However, in areas of AI such as deep learning, where deep neural network algorithms are used, it may not be possible even for the programmer to explain why the computer has come up with a particular answer. These algorithms are effectively 'black boxes' - they are given an input and calculate an output, but there is no way to understand how the algorithm is weighing up different factors in the process. This problem is well recognised in the technology community and a lot of work is underway to improve the understandability of models, often referred to as Explainable AI (XAI).

It is also frequently the case that the more accurate the predictive model is, the less explainable it is. This creates an ethical dilemma in itself - do businesses use the algorithm that provides the most accurate or optimised output, or the algorithm that they can explain?

In financial services, there are examples of companies erring on the side of explainability. BlackRock decided to shelve **AI-based liquidity models** for forecasting market volume and redemption risk, even though they appeared to be significantly outperforming other approaches, because the developers could not explain how they worked to senior managers.<sup>8</sup> Healthcare is likely to pose similar dilemmas - we may not want doctors to rely on algorithms they don't understand, but what if the algorithm result is highly accurate and could help to save lives?

### Practical approaches: GDPR requirements

The General Data Protection Regulation is the main piece of privacy regulation across the European Union (EU) and requires businesses to protect the personal data of EU citizens and personal data processed in the EU. One of the new provisions it brought in was the right for citizens to demand an explanation for how decisions about them are made – what personal data was used and the criteria for decision-making. This, therefore, provides a clear challenge for businesses that are using algorithms to explain their decisions. It remains to be seen how it will be applied in practice and the extent to which it will limit the use of unexplainable algorithms. However, it is being watched closely.

There is a risk that we are setting too high a bar for computers. After all, we don't fully understand how the human brain works, or how we make decisions. Should we worry so much about not being able to understand how computers produce their outputs? Addressing these concerns, defining when unexplainable algorithms can be used and identifying appropriate governance mechanisms will be essential to maximising the benefits of these technologies.

### ORGANISATIONAL ACCOUNTABILITY, GOVERNANCE AND OVERSIGHT

Accountability needs to be allocated to the right person and the number of people involved in technology and systems can make this a complex challenge. For example, the developer should be accountable for delivering a product that meets the quality and ethical standards set and delivers what it has been asked. There are a raft of technical controls, peer review and tests that can be put in place to ensure that work is delivered correctly, whether done by a third-party supplier or an internal technology team.

However, the business user will ultimately be held responsible for how they use the system. For example, if they are relying on system outputs to make decisions, and the outputs are biased, the business user cannot simply blame the software and take no responsibility.

This raises questions of how accountability is handed over between technology specialists and the users, and how much each needs to understand about the other. There are risks of accountability falling through the cracks between different people or driving dysfunctional behaviour to pass the buck and avoid personal blame. Identifying exactly where problems lie, between data, programming and users, can be very difficult.

Boards, of course, have ultimate accountability for the things that happen in their organisations. They have to be able to explain the actions of the business not only to shareholders, but increasingly to regulators and politicians, the media and the public more generally.

Boards have both collective and individual accountabilities so an important question to consider is who leads on the issues around technology-related ethics at a board level? Would it be the CIO, Chief Data Officer or someone similar who comes from a technology background? Or the CEO, the CFO or whoever is responsible for risk and other ethical issues? Maybe every board member has some responsibility for the ethical use of technology and data in their functional area?'

### Practical approaches: Ethics Advisory Boards

There has been a growing interest in Ethics Advisory Boards, as a way of focusing on ethics at a high level. These have historically existed in science, research and healthcare sectors but have become more relevant to technology companies in recent years.

Such boards may consider whether new products or uses of technology fit with agreed ethics principles, discuss any specific dilemmas or tensions, ask difficult questions or challenge actions they think contravene ethics principles. They may also provide an avenue for employees to discuss any concerns, enabling whistleblowing. However these groups are typically only advisory in nature, and therefore have limited power in practice to hold boards to account.

Accountability also needs to have some consequences - good or bad - to drive the right behaviour. In this sense, businesses need to move beyond good intentions, and aspirational high-level principles, to build processes that have teeth. This could include measures in performance management or rewards systems, or disciplinary processes where needed.

The need to move beyond good intentions has been highlighted by direct employee action in companies such as Google and Microsoft, protesting against what they perceive to be unethical actions by the company - [working in China through censoring search results](#), and [working on autonomous weapons respectively](#).<sup>10</sup> While these companies have published ethics principles for developing AI, many employees clearly felt that they were not abiding with them, or had different views on what they meant in practice.

Finally, there is a need to consider the limits of organisational accountability and where other forms of accountability can play a role. Businesses are not accountable for everything that occurs as a result of using their product or service. But equally, they cannot walk away in all cases. Legal frameworks tend to rely on concepts such as 'reasonable foreseeability' when determining liability for harm. Ethics frameworks may lead to higher expectations of what businesses should do to prevent harm. But where is the line?

### Example: business accountability for teenage suicide

Social media businesses have been very clear from the beginning that they do not see themselves as publishing businesses - users are responsible for the content that they publish and the business is simply providing a platform. While this remains the predominant legal position, it is under intensifying pressure as many people push for social media companies to monitor the content of their platforms more closely, take more action and be more responsible for the harm that results.

The growth in cases of teenage suicides highlights the debate here. Campaigners argue that being able to view distressing images or instructions on how to commit suicide have been instrumental in encouraging some teenagers to take their own lives. They want to see social media companies have greater accountability for these consequences, which would force the companies to be more aggressive in what content they take down. Many of these companies employ sophisticated tools, moderators and reporting systems to take down offending material, and recognise some responsibility in this context. However, for many people, this doesn't go far enough and countries such as the UK are looking at [stronger regulation](#).<sup>11</sup>

Putting more responsibility onto businesses raises fresh questions about transparency and power. To what extent should private companies be the main arbiters of what people see on the internet or social media applications? On what basis are they making those decisions, and how can people challenge them? Another example is payment companies, and suggestions that they should not process transactions related to 'unethical' businesses, such as [essay-writing services](#).<sup>12</sup> Are payment companies the right organisations to make judgements about what is an unethical (albeit legal) business?

### PERSONAL RESPONSIBILITY AND CULTURE

While processes and incentives are important, they need to align with our personal values to be most effective. Otherwise they can become tick-box exercises, done for the sake of it rather than because people believe in what they're doing. Furthermore, people can game targets, producing dysfunctional and ultimately poor ethical behaviour.

There are a number of elements to this idea of alignment with personal values. One is thinking in terms of organisational culture and doing the right things because that is simply how people in the organisation behave - it becomes part of the organisational DNA, rather than driven through rewards or sanctions. Creating an ethical culture within businesses is an essential part of achieving real behavioural change, encouraging employees to challenge unethical actions and making ethics a core part of any business activity.

Another important cultural issue is the diversity of staff. It is well known that the technology sector is heavily dominated by men. As a result, the decisions and values captured in systems reflect the experience of men, and often fail to recognise the perspective of women or minority groups. This is recognised as a key contributing factor in the issues around bias as many of the risks simply didn't occur to the predominantly male developer community. They didn't have the personal experience to recognise where bias might be occurring. In order to build solutions that reflect the needs, perspectives and values of society as a whole, and therefore deliver on ethical commitments, the technology community needs to reflect that diversity in its makeup.

Professional codes of ethics for technology specialists also have an important role to play in this regard. Not only do such codes provide principles that underpin use of technology, they provide a framework for individual accountability. However, this needs to sit within a wider professional structure of membership bodies, qualifications and disciplinary processes. While there have been a variety of attempts over the years to professionalise the IT sector, it is still in early stages, with many different bodies and specialist areas.

## *Practical approaches*

Building confidence around technology, based on ethics and accountability, requires a lot of practical action. This section considers some of the key questions and areas for action for three specific groups – businesses themselves, the accountancy profession and policymakers.

### **BUSINESSES: IDENTIFYING GOOD PRACTICES**

The starting point for any business is building a structured approach to consider the ethical elements around technology. This involves:

- defining high-level ethics principles to support decision-making across the organisation;
- developing more detailed guidance around the interpretation and application of principles in specific cases; and
- establishing mechanisms for identifying, debating and resolving conflicts or new ethical questions.

There are a number of sets of ethics principles that have been developed around technology, as highlighted earlier in this report. These provide a range of suggestions for businesses, especially in terms of high-level principles.

Once this high-level framework is agreed, a range of good practices can help businesses to implement it in practice. These could include the following.

- High-level governance processes, for example:
  - establishing written core ethics principles against which the business, its processes and its employees can be judged;
  - setting up an Ethics Advisory Board, which reports into the main board or risk committee;
  - clarifying accountability for technology-related ethics at board level; and
  - establishing whistleblowing processes for individuals who have concerns about ethics.
- Increasing transparency and engagement, for example:
  - explaining to stakeholders the business's core ethics principles and what it is doing to protect them; and
  - testing the ethics principles with its stakeholders and the wider public, and modifying where needed to meet societal demands.
- Detailed processes around accountability, for example:
  - clarifying end-to-end accountability for the use of data, algorithms and AI;
  - establishing appropriate levels of 'human in the loop' for automated decisions; and
  - setting up processes to enable the challenge or correction of decisions, actions or data where needed.
- Controls, reviews and assessments, for example:
  - mapping how algorithms, AI and data are being used in decision-making processes;
  - undertaking Algorithm or AI Impact Assessments where new uses are proposed, and rejecting those systems that do not meet the business's core ethics principles;
  - conducting assurance activities around algorithms, data and relevant governance processes to provide comfort that they are operating as expected and risks are being managed; and
  - implementing appropriate controls regarding the purchase of technology from third parties to ensure that it meets ethics requirements.

- People-related processes, for example:
  - fostering a corporate culture that supports ethics;
  - building ethical considerations into recruitment, performance management and reward structures; and
  - conducting training and raising awareness for staff around relevant ethical issues.
- Technology-specific processes or tools, for example:
  - tools that can test for, and correct, biases; and
  - tools to help with privacy protection, where appropriate.

Good practices around privacy have evolved in recent years, as regulation has strengthened, which provide good starting points for many of these activities. Businesses may also have many existing processes that can be extended to look at some of the specific ethical concerns around technology.

### Practical approaches: Algorithm Impact Assessments and algorithm assurance

An Algorithm Impact Assessment (AIA) is similar in concept to a Privacy Impact Assessment and provides a structured and consistent way of understanding the implications of specific algorithmic decision-making systems. This includes increasing the capacity of employees to understand the issues, and provides greater transparency to affected communities. The assessments can be done for new or existing systems and have particular resonance in the public sector, given the impact of systems on often vulnerable citizens and the potential harm that can be caused.<sup>13</sup>

Algorithm assurance (or audit) is another important component of good business practices, and has been discussed most frequently as a way of managing the risks of bias. There are many elements around algorithms that can be independently checked, including the data inputs, the algorithm itself, its outputs and the governance and controls around the system, as well as the appropriateness of the use and reliance on the algorithm. Such audits can help to increase transparency around the use of algorithms as well as improve accountability. However, the field is still in the early stages, and there is a need to develop a holistic and multi-disciplinary approach that can incorporate these different elements.

One common demand is for checklists of things to do or questions to ask at different points during the development of technology. The UK Digital Catapult, which supports innovation, has an [AI Ethics Framework](#), for example, which provides a detailed set of questions for innovators to ask as they develop new technology-based products and services.<sup>14</sup> [The EU Ethics Guidelines for Trustworthy AI](#)<sup>15</sup> has a comprehensive assessment checklist which businesses can use, alongside high-level guidelines.

Although useful, such lists tend to focus on the technology sector and specialists who are designing and building systems. There is limited resource for company boards, for example, to outline the questions they should be starting to consider in terms of their risk and governance. The Appendix of this report suggests some questions for boards to consider in this context.

### BUSINESSES: IMPLEMENTATION CHALLENGES

Implementing ethics and accountability frameworks in practice is proving to be a complex challenge for businesses, though. Ethics is highly context specific and while general principles can have broad relevance, their detailed application is very dependent on the particular circumstances. There can also be many trade-offs and choices – some minor, some more significant – that are specific to the situation.

Approaches will therefore vary according to the business size, maturity and sector. For start-up businesses, for example, the attitude of the founders and investors is likely to be a critical point – how do they view these ethical questions, and how do they fit with the business model of the organisation? They may then look for simple tools or checklists to prompt discussion about key questions, and demonstrate good practices to investors or customers.

More mature businesses may already have established processes around ethics or data governance that can be expanded to incorporate some of these new issues, especially in sectors like financial services or healthcare. The challenges here are more likely to focus on scaling ethics frameworks, ensuring consistency across complex businesses and embedding ethical considerations into governance and oversight processes.

Taken as a whole, this can represent a significant amount of business change that will require time and resources. Consequently, building a strong case for the benefits of ethical behaviour is crucial. Without commitment from the top of the organisation, and a belief that an ethical approach is the right one, it is unlikely to get traction across the business. This is particularly important as it is often necessary to overcome the perception that ethics hinders innovation.

Then there is the practical challenge of understanding the views of customers or users about what behaviour is acceptable and unacceptable. Where innovations are in early stages, the reactions of customers may be unpredictable or inconsistent, and businesses will need to look to their basic ethical code for assistance. Even where the product or service is well-developed, changing societal views will require a flexible response that recognises concerns and works with stakeholders to mitigate any harms. Successful businesses in the long term are likely to be outward-looking and engaging with peers, regulators, customers, academics or others to gain different perspectives and develop effective feedback loops on their products and services.

Furthermore, the ethical debates around technology have to date been driven by technology specialists. This is not surprising given that they are closest to the issues. However, most of the issues highlighted in this report are not technology issues – they involve business (and societal) choices that need broad debate. Looking at these issues purely through a technology lens risks missing existing good practices that can be extended to cover these issues, or links to other debates about good governance in businesses.

#### Talking points for businesses

- **Strategy:** Why is an ethical approach important? How committed is the board to ethics around technology? What are the risks of failing to consider ethics? How does the business model balance ethical considerations with other interests?
- **Specific technology uses:** How do you incorporate ethical considerations into decisions about technology? How are tensions between different interests resolved? To what extent do you check that systems are doing what is expected and not producing biased outputs?
- **Accountability:** How does the board approach ethics issues and who is responsible for leading in this area? How is accountability spread between users and technology specialists? How are employees empowered to stand up to what they see as unethical practices? How do you communicate your ethics to stakeholders and wider society and gain feedback on it?

### ACCOUNTANTS: UPDATING PROFESSIONAL ETHICS

The accountancy profession already has an established approach to ethics and is therefore building on an existing base of thinking and practice. As accountants start to make more use of technologies such as AI, these new issues provide opportunities to strengthen and reinvigorate the profession's thinking.

More emphasis on the ethical goals of new technologies could help the profession articulate its vision for the future, and how technologies such as AI can help it to deliver greater social value. Accountants may need to be more aware of the ethical dimensions of new uses of data or consider the risks of bias in their use of algorithms, prompting fresh discussion or engagement on ethics.

A key question is whether the profession's current ethics principles are fit for purpose in this context. Are any new principles needed to respond to concerns about technology? Does their detailed application need to be rethought?

**Initial discussions** have concluded that the principles are still appropriate, and no specific need for new principles has been identified at this point.<sup>16</sup> However, detailed guidance about the practical application of the principles is likely to evolve. For example, questions have been raised about the following.

- **Competence** - this could expand to include greater knowledge of data, statistics and algorithms. These may be considered useful business and professional skills at present, but as accountants rely more and more on technology, they could increasingly be seen as essential for an accountant to do their job and therefore part of their ethical responsibilities.
- **Objectivity** - given the particular risks of bias in data and models, greater emphasis may be needed in this area, as well as more discussion of the notion of fairness.
- **Confidentiality** - as accountancy firms gather more and more data from clients, which they may be able to use for other purposes, or which could be shared for wider benefit to society, what 'confidentiality' really means is being increasingly questioned.

There may be challenges to the current approach of threats and safeguards, which reflects the actions of humans, not computers. For instance, familiarity is a threat to independence - if an auditor has worked with a client for too long, they may lose the independent viewpoint, creating the need to rotate audit partners. But it makes no sense for a computer to be 'too familiar'. The meaning of 'supervision' may change where work is increasingly automated and done by computers, not junior accountants. Does supervision in this context mean anything different, or can computers be treated the same as a junior accountant?

Accountability for professional judgements is at the heart of the accountancy profession and growing reliance on machines presents profound questions in this context. AI systems cannot be held accountable for their decisions, even where they are explainable. Nor is there any meaningful sanction that can be made if decisions are wrong - a computer cannot be fined. Moral agency and therefore accountability, at least at present, remains with the people using the system. As a result, accountability in the context of growing reliance on complex and sometimes non-explainable models needs careful consideration.

Furthermore, while the core principles may meet current needs, this could change as accountants refine and evolve their roles. Greater involvement with non-financial data, especially personal data, or taking on a broader role around data governance, could raise new questions. Involvement in strategic decision-making will increasingly include issues related to technology. As a result, reviewing the ethics principles must be grounded in the real-life roles of accountants and reflect any new ethical challenges that arise.

### ACCOUNTANTS: BROADER CONTRIBUTION OF THE PROFESSION

Professions such as accountancy can provide useful lessons to help the technology community embed ethics in professional conduct. Such codes of conduct sit within much deeper structures of ongoing professional development and disciplinary actions. Ethics is also deeply embedded in qualifications and a core part of professional education. Without this broader institutional framework, codes of conduct are unlikely to get the traction that they need. Indeed, research shows that where programmers simply read codes of ethics, it makes no difference to their actions.<sup>17</sup>

Accountants can bring their skills and experience to help organisations build broader ethics and accountability frameworks around technology. Accountancy, after all, is an information-based profession, and accounting ethics principles such as objectivity, integrity and confidentiality are relevant in these wider debates. Accountants' professional scepticism can encourage businesses to ask difficult questions, especially at board level, around the ethics of new digital goods or services, or particular uses of data.

Audit or assurance activities over algorithms are likely to be an important part of practical solutions and there is an opportunity for the profession to work with data scientists and other technology specialists to develop holistic and robust approaches in this area, which may ultimately provide comfort to a range of stakeholders. There could be new certifications or standards in this area which would be more valuable if properly assured.

Most businesses would also benefit from stronger data governance, and accountants can be well placed to take on further responsibilities in this area. The disciplined and structured approach of the profession to financial data could help other business areas to improve data quality and standardisation, with ethics an integral component.

#### Talking points for accountants

- **Accounting ethics:** Are the current ethics principles fit for purpose? To what extent does detailed guidance need to be updated? Does the threats and safeguards approach reflect changing ways of working?
- **Judgement and accountability:** To what extent should accountants make use of 'unexplainable' models? How do accountants challenge models and exercise scepticism over their outputs? To what extent should they overrule machines when it goes against their instincts?
- **Broader contribution:** How can accountants best contribute to wider business discussions on ethics and technology? What key lessons can they share about the operation of ethics codes of conduct?

### POLICYMAKERS: ESTABLISHING INCENTIVES AND REGULATION

While good intentions from businesses are encouraging, action may also be needed from policy makers and regulators. Existing approaches rely heavily on self-regulation with regard to ethics principles and there are growing pressures for regulators to take a more forceful approach. Furthermore, some decisions cannot simply be left to businesses - they need governments to set the right frameworks and red lines.

Softer measures often aim to help investors and customers push businesses to do the right thing through their market power, with, for example:

- requiring companies to make disclosures to shareholders and markets about their ethical approach to technology; and
- developing kitemarks or certifications which signify good practices in this area.

But it is generally hard for customers or shareholder to judge whether a business is ethical or not.

Customer behaviour can be inconsistent in many cases, and customers have different values and priorities. Indeed, some may behave in ways that others would view as unethical. Therefore, relying on customers and markets to hold businesses to account is often unrealistic without further government intervention.

There is broad agreement that regulating a specific technology such as AI or algorithms is problematic. The real world impact on people will vary substantially depending on the sector and how the technology is being used. As a result, it makes most sense for sector regulators each to consider how technology is impacting on their area and take appropriate action through practical guidance or new regulatory principles, for example. This could be within a broader framework set by governments.<sup>18</sup>

Good regulation is difficult though. In many cases, regulators lack the skills and expertise to understand the technology and how it may impact their specific sector or sectors outside their own expertise – unintended consequences become particularly hard to establish where society is highly interconnected. Furthermore, the pace of change makes it hard to keep up to date and understand what is happening in practice, raising risks of regulation lagging behind practice.

There are deeper questions around the relationship between regulation and innovation. On the one hand, regulation may be necessary to build trust and ensure that risks to consumers or others are managed. On the other hand, regulating too soon can stifle innovation by putting a high burden on small companies, and limiting the opportunity to understand the potential benefits and risks of new services. Getting the balance right requires a flexible and open approach from regulators.

### Practical approaches: regulation and innovation

The financial services sector has been at the forefront of thinking about regulation and innovation in fintech. The regulatory sandbox provides a way to try out new innovations, understand how they fit with existing regulation and identify how they might need to be regulated in the real world.

One example of financial services innovation is robo-advising, which has strong links into the responsibilities of financial advisers. Robo-advising is where algorithms recommend financial investments for customers, based on a risk-reward profile. This can be a much cheaper model and enable a wider range of people to get access to financial advice. How do we ensure that robo-advisers are looking after the interests of the clients, though, and not acting purely in the interests of the company? The sandbox enables new products in this space to operate, but regulators can then compare them against the recommendations that a human adviser would have made.

The UK regulator, the FCA, has released some common themes on [the experience of robo-advising](#) which highlight the continuing responsibility of the service provider – it is not possible to blame a third-party technology provider – and a strong focus on outcomes, whether they be generated by human, computer or a combination of the two.<sup>19</sup>

This kind of careful and structured consideration will be increasingly important for all regulators in order to avoid knee-jerk reactions and bad regulation.

Then, there are questions about who regulators want to hold accountable, and who should be making difficult judgement calls – companies, regulators themselves or courts. Consumers may also have responsibilities to ask questions and understand products or services, and regulators may intervene to support consumers in making better decisions. Therefore, regulators need to look at the wider eco-system of stakeholders and decide on the best way to meet their objectives.<sup>20</sup>

### **POLICYMAKERS: BUILDING SOCIAL NORMS AND CONSENSUS**

Businesses need to build their activities on broad social expectations and ethical choices. In many cases of technology innovation, these expectations and choices are not necessarily clear or consistent as technology enables new activities. But developing and understanding these changing expectations can be very difficult. People are busy, not generally engaged in many of the issues outlined and the decisions involved are often complex. What people think in theory does not necessarily align with how they act in practice. There are also likely to be diverse views, leading to a need to reconcile views, build consensus or find ways to recognise opposing opinions.

#### **Practical approaches: citizens' juries**

Citizens' juries are one emerging way of getting direct input about what people think. The MIT Moral Machine looks at autonomous vehicles and invites users to make moral choices about different scenarios through an online platform. Users are presented with specific examples about who an autonomous car should hit or avoid (the 'trolley car' problem highlighted earlier), and asked to make a choice. Research that summarised the responses received to date was published in October 2018,<sup>21</sup> and it particularly highlighted the variances between different countries, suggesting a single set of global principles and applied rules will be difficult to achieve.

These challenges underline the need for robust evidence and research on a multi-disciplinary basis, including social sciences, computer sciences and maths, economics, psychology and others, to support policy makers and businesses. Research that could be useful includes:

- Case studies – case studies about particular dilemmas are an important part of translating ethics into practice. Case studies can help people learn and think through the decisions they would make in specific circumstances, based on a structured and coherent framework.
- Empirical research – there is a need for more empirical research on the impact and harms associated with new technologies. This would help policy makers to develop robust evidence bases and support decision-making around regulatory options. Empirical research can also help to identify what business practices work or don't work, for example the best mechanisms for pushing ethical values through a company, or when written ethics codes are most effective.

### Talking points for policy makers

- **Impact on existing regulation:** Where is the use of technology creating new harms or risks in the sector? Are current approaches to regulation still relevant? Are new business models or services by-passing existing regulatory regimes?
- **Setting incentives:** To what extent does the market drive the right behaviour? Are customers, investors and other stakeholders empowered to challenge businesses? Is stronger regulation needed?
- **Social norms:** How can people be more engaged in these debates? How can different views and priorities be reconciled or recognised? How can academic research support good policy-making?

## *Next steps*

This report maps out a wide range of issues relating to the ethics around new technologies, especially social media, data, algorithms and AI. It also highlights some of the practical implications of these issues and argues that ethics principles need to be embedded in meaningful accountability frameworks across businesses if they are to have any impact.

However, much work is needed to progress current debates and deliver more ethical decision-making and behaviour around technology in practice. ICAEW will lead and support a range of activities, based on its role as a professional body acting in the public interest and the breadth of its member experience across all sectors of business.

### **PRACTICAL GUIDANCE AND KNOWLEDGE SHARING**

There is a clear need for a wide range of practical guidance around ethics, good practices and effective governance in this space. There is also the opportunity to learn about what good companies in this space are doing, and to share key lessons more broadly. Examples could include practical experience with Ethics Advisory Boards, understanding of key ethics principles being applied or lessons in cultural change to embed ethical thinking about technology across businesses.

ICAEW is developing practical thinking for accountants and businesses on these issues across a range of different areas, including the following.

- Accountancy profession: working with other professional bodies and the international standard setters to identify any issues with the current code of ethics and keep members up-to-date.
- Financial services: identifying specific challenges related to financial services, and drawing on the experience of members working in the field to develop new principles and guidance around the ethical use of big data in particular.
- Audit and assurance: developing fresh thinking about assurance over new technologies, including algorithms and AI, to enable assurance to become an important tool in raising confidence in the appropriate use of technologies.
- Corporate finance: sharing knowledge and experience around the use of AI technologies in the corporate finance and advisory sector, and identifying any specific ethical questions that this raises.

We will share this work through publications, online resources and events. We also welcome opportunities to collaborate with others in order to learn, build common understanding of the issues and develop practical approaches.

### **REGULATORY AND PROFESSIONAL THINKING**

While good business practice is important, there is a clear role for regulators, professional bodies and others in setting the right incentives and shaping good behaviour. This needs careful consideration, robust evidence and clarity of purpose.

ICAEW is keen to work with all stakeholders to contribute ICAEW's and our members' experience to relevant discussions. In particular, debates should be grounded in the broad business context, and avoid getting stuck in a 'bubble' of technical specialists. Many of the challenges around technology are not specific to technology - businesses always have to balance how they make money with a wide range of considerations. They operate in complex international environments with different cultures and change is often difficult to control. ICAEW can support debates with input from a broad range of perspectives, from boards to investors to SMEs.

Professions can also lead thinking in some of these areas. Ethics is core to the notion of a profession and therefore considering the ethical implications of adopting new technologies is essential. ICAEW is keen to work with other professional bodies to share experience of how professional ethics is evolving to reflect new technologies. This includes issues such as:

- How notions of professional judgement may need to evolve as more decisions are made by technology;
- How to manage concerns about the explainability of models and resulting decisions;
- Implications for training models and curricula from new technologies; and
- The need for new types of regulation around software, where professionals are being replaced by new business models that are not subject to the same ethical standards.

The AuditFutures programme collaborates with many of the largest accountancy practices and leading universities on the future of the audit profession. This examines what it will mean to be a 'professional' in the future, with ethics as a central feature of this debate. Its work with the University of Edinburgh, for example, is considering how AI systems can be designed and adopted in ways that are aligned with the profession's ethical values.

ICAEW's Ethics Standards Committee is also running a series of roundtables on both specific ethics issues (such as financial inducements) and more general issues (such as the future of professional ethics, the use of AI in accountancy and ethics for directors) in order to develop ICAEW's policy thinking in these areas.

### **PUBLIC ENGAGEMENT AND DEBATE**

We need to engage widely across society on some of the ethical choices that need to be made and support informed debate. ICAEW will do this with its own members, as well as working with other bodies, to raise awareness of the issues, provoke thinking and encourage debate. This includes events, articles and online resources, as well as roundtable discussions and debates.

Furthermore, we are keen to work with academics, think tanks and others to support engagement activities, especially across businesses, and promote debate and good practice.

### **ENGAGE WITH ICAEW**

For more information on ICAEW's activities or to collaborate with us on any of these activities, please visit [www.icaew.com/ethicsandtech](http://www.icaew.com/ethicsandtech) or contact [techfac@icaew.com](mailto:techfac@icaew.com)

## ***Appendix – questions for boards about new technologies, ethics and accountability***

This appendix suggests some questions that boards of all types of business could consider. It is not comprehensive but aims to help boards get started in thinking about the ethical risks of technology and how they should approach this area.

### **1. What are the core ethical values of the business?**

- a) Are these appropriate given new technologies or are any new principles needed?
- b) How do these ethics principles fit with the business model and strategy?
- c) How important are these values to the board?

### **2. Does the business have an ethical culture around technology?**

- a) Is the board conveying the right tone from the top about the importance of ethics?
- b) How are new staff taught about the ethical values of the business?
- c) Is the organisation suitably diverse and able to recognise different perspectives when developing technology?

### **3. How is ethics around technology included in board governance?**

- a) Who is accountable at board level for these issues?
- b) How often is ethics and technology discussed by the board?
- c) Does the board have the right skills and knowledge to consider the risks and issues?

### **4. Are you comfortable that any ethical risks around technology are being managed?**

- a) Do you understand how data, algorithms and other technologies are being used in the business, especially to make key decisions?
- b) What checks or assurance is done to make sure that any key models are working as expected and not producing biased results, for example?
- c) Is ethics considered when reviewing or signing off technology based projects or risks?

### **5. Are appropriate accountabilities in place?**

- a) How does accountability between business and technology specialists fit together?
- b) Where there is automated decision-making, to what extent have controls been reviewed to ensure that there is sufficient 'human in the loop'?
- c) To what extent are 'unexplainable' models relied on in decision-making?

### **6. Is the business open and engaged with key stakeholders around ethics and technology?**

- a) Are there mechanisms for employees to raise concerns about ethical questions, such as whistleblowing processes?
- b) To what extent has the business published its ethical approach and engaged customers and others in discussions and feedback?
- c) Where the business is regulated, to what extent is it engaged in discussions with regulators about any changing requirements here?

## *Further reading*

AI4People (2018), *Ethical Framework for a Good AI Society*

AI Now Institute (2018), *AI Now report 2018*

AI Now Institute (2017), *AI Now report 2017*

AI Now Institute (2016), *AI Now report 2016*

Bostrom, N. and Yudkowsky, E. (2014), 'The Ethics of Artificial Intelligence' in Frankish, K. and Ramsey, W. (eds), *Cambridge Handbook of Artificial Intelligence*, New York: Cambridge University Press

Cath, C. (2018), 'Governing artificial intelligence: ethical, legal and technical opportunities and challenges', *Phil. Trans. R. Soc.* 376 (2133) November 2018 (and rest of the special edition)

Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M. and Floridi, F. (2017), 'Artificial Intelligence and the 'Good Society': the US, EU, and UK approach', *Science and Engineering Ethics*, 24(7625) March 2017

European Commission (2019), *Artificial Intelligence: a European Perspective*

House of Commons Science and Technology Committee (2018), 'Algorithms in decisionmaking', *Fourth Report of Session 2017-19*

House of Lords Select Committee on Artificial Intelligence (2017), 'AI in the UK: ready, willing and able?', *Report of Session 2017-19*, HL Paper 100

'ICAEW Code of Ethics', available at [icaew.com](http://icaew.com)

ICAEW (2019) *AI in Corporate Advisory - investment, M&A and transaction services*

ICAEW (2019), *Ethical use of big data in financial services*

ICAEW (2017), *Artificial Intelligence and the Future of Accountancy*

Institute of Business Ethics (2018), *Business Ethics and Artificial Intelligence*

Mittelstadt, B., Allo, P., Taddeo, M., Wachter, S. and Floridi, L. (2016), 'The ethics of algorithms: Mapping the debate', *Big Data and Society*, 3:2, December 1, 2016

Nuffield Foundation and Leverhulme Centre for the Future of Intelligence (2019), *Ethical and societal implications of algorithms, data, and artificial intelligence: a roadmap for research*

RSA (2018), *Artificial Intelligence: Real Public Engagement*

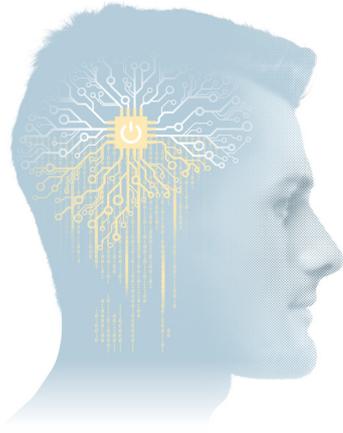
Russell, S., Dewey, D. and Tegmark, M. (2015), 'Research Priorities for Robust and Beneficial Artificial Intelligence', *Association for the Advancement of Artificial Intelligence*, Winter 2015, pp.105-114

## *End notes*

- <sup>1</sup> For more discussion of the strengths and limits of machine learning, see ICAEW (2017), *Artificial intelligence and the future of accountancy*.
- <sup>2</sup> ICAEW Code of Ethics, which is based on the Code of Ethics for Professional Accountants of the International Ethics Standards Board for Accountants (IESBA).
- <sup>3</sup> For a comprehensive list of existing approaches, see Nuffield Foundation and Leverhulme Centre for the Future of Intelligence (2019), *Ethical and societal implications of algorithms, data, and artificial intelligence: a roadmap for research*.
- <sup>4</sup> See 'Amazon scraps secret AI recruiting tool that showed bias against women' at [reuters.com](https://www.reuters.com)
- <sup>5</sup> See the results of the Prorepublica investigation of the Compas software in the US, published as 'Machine Bias' at [propublica.org](https://www.propublica.org)
- <sup>6</sup> ICAEW (2016), *Audit insights: insurance*.
- <sup>7</sup> For example, see a recent military-based study on when people trust AI outputs or their own judgements: 'Most people overlook artificial intelligence despite flawless advice' at [phys.org](https://www.phys.org)
- <sup>8</sup> See 'BlackRock shelves unexplainable AI liquidity models' at [risk.net](https://www.risk.net)
- <sup>9</sup> As an example, ICAEW's *Principles for ethical use of big data in financial services* recommends that the Chief Data Officer (or similar) becomes a regulated role under the Senior Managers and Certification Regime and takes the lead in this area.
- <sup>10</sup> See for example, Adi Robertson, 'Google employees push to cancel Chinese search engine in new letter', at [theverge.com](https://www.theverge.com) and Julia Carrie Wong, 'We won't be war profiteers': Microsoft workers protest \$480m army contract' at [theguardian.com](https://www.theguardian.com)
- <sup>11</sup> See the 'Online harms white paper' at [gov.uk](https://www.gov.uk)
- <sup>12</sup> See 'PayPal urged to block essay firm cheats' at [bbc.co.uk](https://www.bbc.co.uk)
- <sup>13</sup> AINow Institute (2018), *Algorithmic impact assessments: a practical framework for public agency accountability*.
- <sup>14</sup> See [migarage.ai/ethics-framework](https://migarage.ai/ethics-framework)
- <sup>15</sup> See 'Ethics guidelines for trustworthy AI' at [ec.europa.eu](https://ec.europa.eu)
- <sup>16</sup> See, for example, ICAEW's 2018 roundtable on some of these questions, *New technologies and their implications on the code of ethics*, available at [icaew.com](https://www.icaew.com)
- <sup>17</sup> Justin Smith and Andrew McNamara, *Code of ethics doesn't influence decisions of software developers*, presented 7 November 2018 at the ACM Symposium on the Foundations of Software Engineering.
- <sup>18</sup> For example, a key role of the UK Centre for Data Ethics and Innovation is to act as an expert body to support regulators. While not impinging on the specific remit of individual regulators, it aims to help regulators understand the issues and develop consistent approaches where that is appropriate.
- <sup>19</sup> See [fca.org.uk/news/speeches/robo-advice-fca-perspective](https://www.fca.org.uk/news/speeches/robo-advice-fca-perspective)
- <sup>20</sup> National Audit Office (2019), 'Regulating to protect consumers: Utilities, communications and financial services markets'.
- <sup>21</sup> E. Awad, S. Dsouza, R. Kim, J. Schulz, J. Henrich, A. Shariff, J.F. Bonnefon, I. Rahwan (2018), 'The Moral Machine experiment' *Nature* 563, pp 59-64.

## STAY AHEAD WITH ICAEW THOUGHT LEADERSHIP

ICAEW is required by its Royal Charter to advance the theory and practice of accountancy in all its aspects. One way we do this is through papers published in our ICAEW Thought Leadership series. Some papers are more practical in nature and aimed at helping members do their work, others consider how law and standards might develop for the future professional.



### *Artificial intelligence and the future of accountancy*

#### ASKING INTELLIGENT QUESTIONS

AI brings many opportunities for accountants to improve their efficiency, provide additional insight and deliver more value to businesses. But, what is the long-term vision for AI and accountancy and how can artificial and human intelligence work together?

[icaew.com/Alreport](https://www.icaew.com/Alreport)



### *Blockchain and the future of accountancy*

#### EXAMINING IMPACT AND POTENTIAL

Blockchain has the potential to increase the efficiency of the process of accounting for transactions and assets, operating as a system of universal entry bookkeeping. The paper explains how the technology differs from the familiar, and how these features drive the potential applications of blockchain.

[icaew.com/blockchainreport](https://www.icaew.com/blockchainreport)

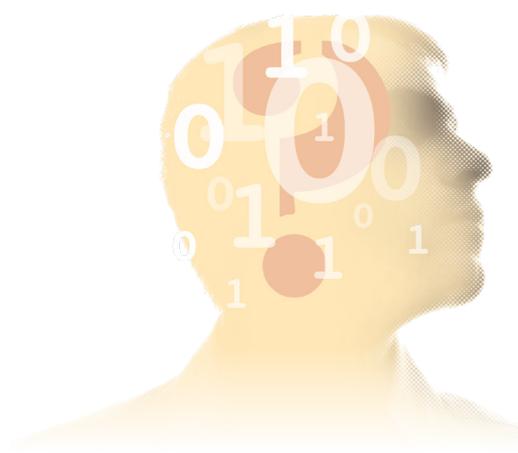


### *Audit insights: Cyber security*

#### COPING WITH INCREASING COMPLEXITY

Many businesses have onerous and complex legacy IT environments with fragmented and non-standardised systems. These present challenges and make good cyber security much harder. What can businesses do to cope with this complexity and be safe?

[icaew.com/Alcyber](https://www.icaew.com/Alcyber)



### *Big Data and analytics – what's new?*

#### INFORMING DECISION MAKERS

There has been a lot of talk about big data and analytics in recent years, but what is really new about it? What's creating big data? What are the opportunities and risks, and how do we exploit it?

[icaew.com/bigdatareport](https://www.icaew.com/bigdatareport)

The ICAEW Tech Faculty represents chartered accountants' tech-related interests and expertise, contributes to tech-related public affairs and helps those in business to keep up to date with tech issues and developments. The faculty also works to further the study of the application of tech to business and accountancy, including the development of thought leadership and research. As an independent body, the Tech Faculty is able to take a truly objective view and get past the hype surrounding tech, leading and shaping debate, challenging common assumptions and clarifying arguments.

For more information on the Tech Faculty and how to get involved, visit [icaew.com/techfac](https://icaew.com/techfac), mail [techfac@icaew.com](mailto:techfac@icaew.com), or call +44(0)20 7920 8635.

---

There are over 1.8m chartered accountants and students around the world – talented, ethical and committed professionals who use their expertise to ensure we have a successful and sustainable future.

Over 180,000 of these are ICAEW Chartered Accountants and students. We train, develop and support each one of them so that they have the knowledge and values to help build local and global economies that are sustainable, accountable and fair.

We've been at the heart of the accountancy profession since we were founded in 1880 to ensure trust in business. We share our knowledge and insight with governments, regulators and business leaders worldwide as we believe accountancy is a force for positive economic change across the world.

[www.charteredaccountantsworldwide.com](https://www.charteredaccountantsworldwide.com)  
[www.globalaccountingalliance.com](https://www.globalaccountingalliance.com)

## ICAEW

Chartered Accountants' Hall  
Moorgate Place  
London  
EC2R 6EA  
UK

T +44 (0)20 7920 8100  
E [generalenquiries@icaew.com](mailto:generalenquiries@icaew.com)  
[icaew.com](https://icaew.com)

