Institute and Faculty of Actuaries

# Ethical and professional guidance on Data Science and Artificial Intelligence

## A Guide for Members

by the Regulatory Board

www.actuaries.org.uk

V2.1, July 2025

# Contents



*This Guide imposes no new obligations upon Members or their employers. Rather the Institute and Faculty of Actuaries ("the IFoA") hopes that the Guide will be a useful tool for its Members.*

*This Guide does not constitute legal advice, nor does it necessarily provide a defence to allegations of Misconduct. While care has been taken to ensure that it is accurate, up to date and useful, the IFoA will not accept any legal liability in relation to its contents.*

# 1. Introduction

1.1. Used responsibly data science and artificial intelligence (AI) have the potential to offer many benefits. However, the use of these can create significant practical and ethical challenges for practitioners and users of outputs.

1.2. The IFoA believes that the high standards to which IFoA Members are subject to, and their technical rigour, will enhance the value of their contributions in data science and AI related work. In order to maintain the reputation of the actuarial profession it is important that IFoA Members are aware of relevant professional, ethical, and technical standards that apply to them.

1.3. The purpose of this IFoA specific Guide is to provide IFoA Members with practical guidance on how their professional obligations relate to work which involves the use of data science or AI. This Guide also covers practices that actuaries have followed for years and therefore it applies to a wide variety of Members, including Members who use traditional methods of data analysis.

1.4. In this guidance some necessary terminology is used. There are many competing descriptions in the field, which can often provoke debate. We have used language consistent with the Alan Turing Institute, the UK's national institute for data science and artificial intelligence, which has published a helpful glossary for some of the more common terms, and the following in particular:

- **Data science** - An umbrella term for any field of research that involves the processing of large amounts of data in order to provide insights into real-world problems.
- **Artificial Intelligence (AI)** - The design and study of machines that can perform tasks that would previously have required human (or other biological) brainpower to accomplish. AI is a broad field that incorporates many different aspects of intelligence, such as reasoning, making decisions, learning from mistakes, communicating, solving problems, and moving around the physical world.
- **Machine learning (ML)** - A field of artificial intelligence involving computer algorithms that can 'learn' by finding patterns in sample data. The algorithms then typically apply these findings to new data to make predictions or provide other useful outputs, such as translating text or guiding a robot in a new setting.
- **Generative AI (GenAI)** - An artificial intelligence system that generates text, images, audio, video, or other media in response to user prompts. It uses machine learning techniques to create new data that has similar characteristics to the data it was trained on, resulting in outputs that are often indistinguishable from human-created media.
- **Large language model (LLM)** – A type of model that is trained on a vast amount of textual data in order to carry out language-related tasks. Large language models power the new generation of chatbots and can generate text that is indistinguishable from human-written text.

1.5. The Guide also provides case studies, which we hope will be of practical value to Members when faced with ethical and professional issues or dilemmas when undertaking data science and AI work.

1.6. This Guide does not contain an exhaustive list of the relevant professional and regulatory obligations, and other requirements are likely to be applicable depending on the circumstances of the work being carried out.

1.7. The initial version of this guide, focused on data science, was published in 2021, with this updated 2024 version covering wider AI applications. We are aware of the rapidly changing nature of this topic and further reviews of the guide will take place to ensure the material remains as relevant as possible. In the meantime, please send any comments to regulation@actuaries.org.uk.

# 2. Professional obligations

2.1 Members of the IFoA might face technical and ethical dilemmas when undertaking data science and AI related work and may find it helpful to consider their professional and regulatory obligations further within this context.

2.2 Member's professional and ethical obligations are set out in the IFoA's current standards framework.

2.3 The standards framework is made up of:

i. The Actuaries' Code ('the Code')
ii. Actuarial Profession Standards ('APSs')
iii. Non-mandatory guidance.

The IFoA also requires Members carrying out work that is within UK geographic scope to comply with the Financial Reporting Council's (FRC's) Technical Actuarial Standards (TASs).

2.4 This Guide provides non-mandatory guidance for IFoA Members and it is not compulsory for Members to comply with this material. Rather, the IFoA hopes that it will assist Members in complying with their obligations and help to build on the requirements set out in the Actuaries' Code and in the APSs.

2.5 The aspects of the IFoA's current standards framework that are relevant for particular data science and AI related work will depend on the circumstances of the work. The following sections outline how each aspect of the IFoA's current standards framework could be relevant. It is recognised that actuaries will often be working alongside other professionals on data science and AI projects, and this may impact on how standards will be applied.

### Actuaries' Code

2.6 The Code is the overarching ethical code of the IFoA. It sets out principles that all Members must comply with and aims to build and promote confidence in the work of actuaries and in the actuarial profession. The Code applies at all times to Members' conduct in relation to an actuarial role and also applies to other conduct that could reasonably be considered to reflect upon the actuarial profession.

2.7 The Code does not define 'actuarial role'. However, this might include circumstances where a Member is performing a role that requires, or benefits from, specific actuarial skills, or where reliance is placed on a Member's actuarial judgement due to their membership of the IFoA. This may therefore include work which involves data science or AI techniques. Additionally, the work that Members undertake involving data science or AI could potentially reflect upon the profession as a whole.

2.8 Some areas of consideration for Members with regards to data science and AI related work are outlined below. This is not intended to be a comprehensive look at the Code's requirements, but rather an outline of how some of the Principles can relate to data science and AI.

*Principle 1 Integrity*

**"Members must act honestly and with integrity."**

2.9 A key consideration is to seek to enhance the societal value of data science and AI, and under the first principle of the Code, Members must act honestly and with integrity. The impact, including the outcomes and consequences that data science and AI can have on society could be significant, and if IFoA Members are involved in this work they will be expected to act in an ethical and professional manner, and to seek fair outcomes. If data or models are used improperly and practitioners do not speak up about this, it could have detrimental consequences for society as a whole.

2.10 Alongside the benefits they may bring, data science and AI have the potential to cause harm to individuals, and the general public are entitled to expect that sensitive information will not be misused, treated carelessly, or be shared without their permission. Under this Principle, Members should respect confidentiality and consider this when they have access to potentially sensitive information.

2.11 Possible relevant examples of integrity in data science and AI related work could include ensuring that:

- Active management of the ethical risks that might result in AI models producing outcomes that serve against the public interest
- Appropriate consent is obtained for the purpose that data is to be used
- There are safeguards to prevent data or AI techniques being used in a manner which results in unfair or discriminatory outcomes for users
- Governance and accountability are clearly defined, supporting the proper functioning of processes and systems, and responsible oversight and challenge.

*Principle 2 Competence and care*

**"Members must carry out work competently and with care."**

2.12 Members are encouraged to develop their experience in new fields and pursue new areas of work. Actuaries may become increasingly involved in data science and AI projects that involve techniques that are new to them. Under the Code, Members are expected to ensure that they have the right level of knowledge and skill for the piece of work that they are undertaking. This ensures that users and the public can continue to trust Members to be competent to perform the services that they are engaged in.

2.13 It is common in data science and AI related work for actuaries to work in multidisciplinary teams with a range of other professionals such as data scientists, computer analysts and statisticians. Actuaries can fulfil different roles within these teams. If a Member feels that they do not have the right level of knowledge to perform the requested task in the project, they can resolve this by working with more experienced team Members, or by obtaining training, for example.

2.14 In common with other actuarial techniques, Members should use their professional judgement when using or relying on third party data and models which are based on AI approaches. The level of review required is a judgement exercised in the context of the purpose and scope of the actuarial advice. Actuaries should assess and agree where professional responsibility lies for the data, assumptions, models and the resulting advice, and the level of reliance that can therefore be placed on the advice. The FRC Technical Actuarial Guidance covering Models deals specifically with interpreting TAS when using third party models.

2.15 Possible examples of professional competence within data and modelling ethics could involve:

- Robustness, security and safety of data and models
- Sufficient understanding of the AI models used to be able to identify material risks, especially when deploying new methods that might lead to the emergence of new risks
- Ensuring that there is a full understanding, prevention and mitigation of potential ethical harms arising from data or model use
- Care should be taken when using third party generative AI tools (including LLMs), in relation to veracity of output and privacy and copyright risks
- Obtaining input from other suitably knowledgeable professionals. This is particularly relevant where given the complexity of many data science and AI related projects, it is unlikely that an actuary will be an expert in all aspects, nor may the actuary have responsibility for all aspects of the project
- Models should operate reliably and as expected
- Model design and output should be readily interpretable, with the ability to understand the workings and decisions of the model
- Existing assessments of model risk, and model governance in place to mitigate this, may need to be reviewed, to ensure they remain sufficiently robust in the context of emerging AI models. Approaches which cater adequately for traditional actuarial models, or indeed certain types of machine learning techniques, may need to be adapted for the particular challenges associated with validating generative AI tools.

*Principle 3 Impartiality*

**"Members must ensure that their professional judgement is not compromised, and cannot reasonably be seen to be compromised, by bias, conflict of interest, or the undue influence of others."**

2.16 Members are expected to present information in a way that is accurate and impartial, and to ensure that their professional judgement is not compromised, and cannot reasonably be seen to be compromised, by bias, conflict of interest, or the undue influence of others.

2.17 Possible examples of impartiality in data science and AI related work could include:

- Having fairness and ethics at the forefront of data and model developments
- Seeking to avoid unfair preference or prejudice towards groups of characteristics, targeting wider societal benefit
- Identifying and understanding potential biases in the data they are analysing, and AI techniques applied to that data
- Ensuring potential biases are communicated transparently to the users and stakeholders, with any inappropriate bias appropriately mitigated
- Where Members are working within a multidisciplinary team, with data scientists for example, ensure that they consider the risks of conflicts of interest and bias and communicate this to the wider team so that appropriate action is taken. It will be helpful for Members to explain to other team members their professional responsibilities in this area.

*Principle 4 Compliance*

**"Members must comply with all relevant legal, regulatory, and professional requirements."**

2.18 Under this principle, Members must comply with all relevant legal, regulatory, and professional requirements. This includes any rules governing matters in the area in which a Member is practising. All practitioners working in the field of data science and AI will be required to comply with any relevant laws, regulations, and professional requirements. These will not necessarily always be those that apply in the geographic location where the Member carries out their work, for example, where a Member works remotely but carries out work for an entity in another country.

2.19 Members may work in industries that are subject to high levels of regulation, such as financial services, with a wide range of legal, regulatory, and professional requirements to consider. Additionally, there may also be specific standards and regulation for consideration with respect to data science and AI (e.g., the EU AI Act).

2.20 Possible relevant examples of compliance in data science and AI related work could include:

- Transparency for consumers with regards to the information that is collected and how this data is being used
- Compliance with local regulations could mean that certain types of data cannot be used for particular purposes (including model training)
- Compliance with consumer data protection regulations
- Compliance with copyright law when accessing information or data
- Modelling techniques may be subject to model governance and oversight regulations
- Compliance with cyber security regulations.

2.21 The value of Members' work in data science for key stakeholders will be enhanced if those stakeholders are confident that relevant legal and regulatory requirements have been fully considered and complied with. This will also help to increase the public's trust and confidence in the work of data science and AI practitioners.

*Principle 5 Speaking up*

**"Members should speak up if they believe, or have reasonable cause to believe, that a course of action is unethical or is unlawful."**

2.22 Members have a responsibility to speak up in certain situations and under Principle 5, they should speak up if they believe that a course of action is unethical or unlawful. It is an essential part of being a professional actuary to raise or identify issues in the work they are involved in (data science related and AI work or otherwise) before harm is caused or to prevent any further damage being caused.

2.23 Members may have concerns with the behaviour of other actuaries, or non-actuaries in multidisciplinary teams and could come across activity that they believe to be unethical or unlawful, which they should speak up about. Additionally, if Members have had significant involvement in a data science or AI project, they must take reasonable steps to make users aware of any substantial issues, if it might reasonably influence the judgement of the users. It can help, when working in a multidisciplinary team, to establish the lines of responsibility from the outset.

2.24 Possible examples of speaking up in a data science or AI related project could include the Member highlighting any concerns to users with respect to:

- Robustness, security and safety of data and models
- Governance and accountability in place to oversee the use of data science and AI
- The principles of fairness and ethics, where the ethical risks involved with using a particular AI model may be too great, or serve to act against the public interest
- Transparency, explainability and interpretability of models
- Independent challenge and constructive scepticism (e.g., human-in-the-loop safeguards), and actuaries should speak up where they have concerns about AI implementations.

*Principle 6 Communication*

**"Members must communicate appropriately."**

2.25 Given the complexity of some of the techniques involved in data science and AI, particular care is required to ensure that communications are not misunderstood or misinterpreted, and that the risks of the approach are clearly identified.

2.26 Explanations and validation of complex data and modelling techniques are likely to be more challenging than for more traditional actuarial models, including where results from AI models are not necessarily reproducible. In particular communicating methodology and output for some generative AI tools may require different approaches than for models where explainability is more straightforward.

2.27 Principle 6 of the Code states that Members must communicate appropriately and take reasonable steps to ensure that any communication for which they are responsible, or had significant involvement in, is accurate, not misleading and contains an appropriate level of information. It could potentially be very challenging for Members to explain the results and techniques used in data science and AI related work and therefore care is needed to ensure that the communication is in a form that the user will understand.

2.28 Particular care is also required where communications are aimed at a non-technical audience to explain actions as far as possible. There may be a wide range of stakeholders with different requirements. This could include regulators, customers, and management. A considerable challenge for certain AI or data science techniques is the ability to explain the results in a way that key stakeholders can understand.

2.29 As mentioned above, if Members are working in a multidisciplinary team, it may be helpful to establish the lines of responsibility to help make clear accountability for any work carried out by them.

### Actuarial Profession Standards

2.30 All Members who are undertaking data science and AI work are required to comply with all applicable provisions of the relevant mandatory Actuarial Profession Standards (APSs). The list outlined below is not

intended to be exhaustive, and other APSs may be applicable depending on the circumstances of the work being carried out.

*APS X1: Applying Standards to Actuarial Work*

2.31 This sets out principles to be applied by Members to determine which standards they must or should be applying to a piece of work. It makes clear that Members must comply with the law that applies to the work they are undertaking.

2.32 APS X1 states that Members should ensure that their 'Actuarial Work' is carried out in a way that is substantially consistent with the International Actuarial Association (IAA)'s International Standard of Actuarial Practice (ISAP) 1 and the FRC's Technical Actuarial Standards (TASs) for work within the UK geographic scope. It defines 'actuarial work' as work undertaken by a Member in their capacity as a person with actuarial skills on which the intended recipient of that work is entitled to rely. This may therefore include work that Members undertake related to data science and AI, as it could involve Members carrying out calculations, modelling or giving advice or recommendations in their capacity as an actuary. It is therefore important to consider whether there is an obligation to apply ISAP 1 to the work in question, which may be more difficult to distinguish in a data science project.

2.33 If Members are carrying out data science or AI work outside of the UK geographic scope, APS X1 requires Members to exercise reasonable judgement to consider whether there are other relevant standards that they ought to apply. Members working in data science and AI, and other practitioners, may need to investigate what the relevant requirements are to the work that they are involved in, in the relevant jurisdiction, and ensure their work complies with these.

*APS X2: Review of Actuarial Work*

2.34 APS X2 applies to all Members and relates to the need to consider the extent to which review (including independent peer review) is required for any 'actuarial work' (work undertaken by a Member in their capacity as a person with actuarial skills on which the intended recipient of that work is entitled to rely).

2.35 APS X2 is relevant for data science and AI related work and Members may need to consider the difficulty and complexity of the work in deciding whether a review should occur and the extent of that review. If a review is being carried out, it is important to understand the involvement of various parties, particularly in a multidisciplinary team, and what is expected of them.

2.36 Any reviewer will need to understand and be provided with all the information regarding the work and the context in which it is undertaken. If the potential reviewer for a piece of work is an actuary, and they are not familiar with data science and AI aspects of the project, then an alternative should be considered. APS X2 does not specify that individuals carrying out the review be qualified actuaries, therefore it may be appropriate for different parties, with suitable skills and experience, to undertake a review.

2.37 Similarly, non-actuary practitioners working in data science and AI will often consider inviting independent review of their work, in order to demonstrate and apply professional competence. Members of the IFoA may be able to share best practice in this area with other professionals.

2.38 It is possible that actuaries may use AI techniques as part of a work review exercise. Whilst this may be perfectly valid, there should be careful communication with stakeholders, with any limitations clearly indicated. There may be a requirement for additional human-in-the-loop as part of the overall review.

### Technical Actuarial Standards

2.39 The UK Financial Reporting Council (FRC) sets mandatory Technical Actuarial Standards (TASs) within the UK geographic scope and IFoA Members must comply with relevant TASs. The TASs define 'technical actuarial work' as work performed for an intended user:

- Where the use of principles and/or techniques of actuarial science is central to the work and which involves the exercise of judgement; or
- Which the intended user may reasonably regard as technical actuarial work by virtue of the manner of its presentation.

2.40 Where Members are involved in data science or AI work within UK geographic scope, they should assess the requirement for TAS application against the definition of 'technical actuarial work' provided in paragraph 2.39 above. Members may also find the FRC guidance on Technical Actuarial Work and Geographic Scope helpful in carrying out such an assessment.

2.41 The FRC published an update to its Technical Actuarial Guidance covering Models in October 2024, to reflect the increasing use of AI/ML in actuarial work. This contains further support for actuaries and touches on aspects such as governance, bias, and limitations.

### Additional IFoA resources

2.42 In addition to this document, the IFoA has various other resources that Members may find helpful when undertaking AI and data science related work, including:

### A Guide for Ethical Data Science

2.43 A useful starting point for Members working in data science is the IFoA's joint guidance with the UK Royal Statistical Society. This considers five recurring ethical themes from a range of existing ethical frameworks and working practices across a wide range of sectors and industries. Within each of these themes are examples of corresponding working practices which aim to help Members consider data ethics. Including:

- Seek to enhance the value of data science for society
- Avoid harm
- Apply and maintain professional competence
- Seek to preserve or increase trustworthiness
- Maintain human accountability and oversight.

2.44 The joint guidance is intended to complement the IFoA's current standards framework, which is discussed above, by addressing some of the ethical and professional challenges of data science related work. The guide is also non-mandatory guidance in the IFoA's current standards framework.

### AI risk alert

2.45 The IFoA released a risk alert in September 2023 recognising the acceleration of AI techniques and capabilities, with the further potential of application to actuarial work.

2.46 The key messages from the alert were as follows:

- There are significant ongoing data science and AI opportunities for actuaries and their firms, both in existing and potentially new domains
- With the rapidly accelerating profile, development, and use of AI, it is important that members of the IFoA involved in this work carefully assess the potential risks, apply their skills in an ethical manner, and consider the wider public interest
- The use of AI in actuarial work may amplify existing risks and introduce new risks, and this needs to be appropriately managed and mitigated.

### Certificate in Ethical Artificial Intelligence

2.47 The IFoA has teamed-up with the Chartered Institute for Securities & Investment (CISI) to provide discounted access for IFoA members to a course considering the ethics of AI.

2.48 CISI's Certificate in Ethical Artificial Intelligence is a short course covering the fundamental ethical and management issues in the deployment of AI in finance. The course explores global developments in the field of AI and regulators' responses to them, plus AI ethical dilemmas and what they mean for financial services.

2.49 It is an online course, consisting of a number of modules, with a range of interactive content, including videos, interviews, some reading and extra references, and materials, with a certificate awarded on completion.

# 3. Case studies

3.1 This Guide aims to assist Members in understanding how their professional obligations are put into practice when providing advice or recommendations in data science or AI. Case studies and discussion points are provided in the Appendix which we hope will be useful for Members in providing additional context.

# Appendix A, Case studies

**Case study 1**

Insurance companies are increasingly using virtual assistants, to automate the sale of their products. You are a senior actuary working for a consultancy that designs such virtual assistants, and these are then used by a number of insurance companies, covering a range of products across various territories. Both actuaries and non-actuaries are involved in the work.

Potential customers of an insurance company are required to answer a series of predetermined questions via text or voice. The virtual assistants use Large Language Model technology to understand the customer's responses to the questions and based on the responses, recommend an appropriate product for that customer.

The same core LLM and algorithm methodology is used by the consultancy for all insurance companies. However, the consultancy tailors the approach for each company to reflect their particular products and customer base. The tailoring works by adjusting the questions that the customers are asked, and also how the customer's responses are filtered to determine the product recommendation.

The virtual assistant models are being used directly with consumers via company websites and also via brokers where the insurance broker can use the tool to support their own discussion with the customer about their insurance needs and help find suitable products.

One of the brokers who has been using the virtual assistant for some time has noticed that it often suggests products which they don't consider appropriate for their customers. They have raised their concerns with the company and also talked to the insurance regulator as they expect the issues could be widespread wherever these assistants are being used.

The regulator has asked the insurance company for the details of their virtual assistant and explanation of how the underlying LLMs and algorithms work.

Discussion questions

- Consider what ethical challenges there are in designing the model and associated algorithms. What are your obligations under the Actuaries' Code?
- Are there any potential challenges or risks for the insurance company in using virtual assistants in its sales process, and what mitigating controls could be used?
- What concerns might a regulator have with regards to an insurance company selling its products in this way?

Suggested discussion points

1 The actuaries on the team will be required to consider their obligations under the Actuaries' Code as, although they are not necessarily carrying out 'traditional' actuarial roles, it is likely that their work on the algorithm requires or benefits from their actuarial skills. Additionally, the actuaries' work on this project could reasonably be considered to reflect upon the actuarial profession as a whole.

   Under the Competence and Care principle of the Code, the actuaries need to ensure that they had the right level of knowledge and skill, when designing the LLM and algorithms. Among the issues they should have considered was whether they understood the sources of error in the original data that the LLM has been trained on, and the potential for bias in the results.

   If the actuaries had any concerns about the results of the models, they should have considered their responsibility to speak up and raise this with the wider team and escalate their concerns if necessary and appropriate.

It may be that the workings of the models were not communicated in a transparent and open way to the consumers, leading to the complaint. Actuaries have a responsibility under the Communication principle of the Code, to take steps to ensure that any communication for which they are responsible or in which they have a significant involvement is accurate, not misleading and contains an appropriate level of information. Furthermore, if the actuaries are providing the work to the client (internally or externally) which is done within the UK geographic scope, then TAS 100 will apply if their work is based on actuarial principles.

2 The use of LLMs raise various ethical challenges and risks, including considering the potential impact of the results on the decisions made by consumers. It is important to consider the potential biases, errors and assumptions within the models, and what impact that may have on its results.

The virtual assistant approach raises questions around what human oversight there is over the system given it produces automated results. It would be expected that the team considered how to monitor the model and algorithms over time and agreed in advance where responsibility lay. It would be prudent to keep such models under regular review for example, consider the potential implications the system had for consumers and test the results.

If the data on which the LLM was trained on is in any way biased, it may carry that bias over to users of the system and potentially discriminate. There is a risk that a LLM could have automatic misguided assumptions about a user, leading to it offering the wrong product for a given consumer.

LLMs are known to 'hallucinate' which means they could provide answers which are not just incorrect but offensive or damaging for the company's reputation, as well as those on the team who have produced it. This may also ultimately impact on consumer trust of AI.

If the customer conversations are being used to periodically re-train the LLM, there is a risk that it could expose potentially confidential information in future answers.

The company can put into place a process to cross-check the data and the answers being given to consumers. The company could have also run a limited pilot before full implementation where every result from the algorithm is tested. Any issues could then be fixed prior to full implementation. Further sample testing of the results could then be done before it is fully rolled out, to check if it is aligning to what they would expect.

3 A regulator would likely consider whether the company has complied with relevant regulations and legislation, as well as the spirit of the legislation. They might look for the company to be able to explain how the virtual assistant works, why it produces certain results and whether the results were repeatable. It might have concerns over whether consumers are being offered fair and tailored products, appropriate to their needs, are not limited by the assistant's offering and whether there was any potential discrimination occurring. A regulator might look to see if the company has minimised any potential unfairness with appropriate controls and testing. In the UK, the FCA Consumer Duty would be of particular relevance.

**Case study 2**

You work as an in-house actuary for a life Insurance company. You have been placed in a new team comprising of actuaries and data scientists. Your team has developed an underwriting tool using AI applications which combine data from a range of public and proprietary sources to calculate the likelihood of potential customers carrying out risky activities that would affect their premium, and uses this to set a risk score for the customer.

Discussion questions

- What are the potential risks of this software?
- What obligations as an actuary would you consider?

Suggested discussion points

1 Such a system raises risks of unfairness against consumers by automatically assigning a risk score which is dependent on data which may not be relevant to a given customer. How can the company be sure that all of the data collected is free from bias? Has relevant permission been given by the data subjects to use their data in this way? Have the results been subject to rigorous controls and testing? Do the results demonstrate any potential bias?

2 IFoA Members have a responsibility to comply with the Actuaries' Code Competence and Care principle, which states that Members must ensure that they have an appropriate level of relevant knowledge and skill to carry out a piece of work and that they need to consider whether input is required from other professionals or specialists.

For actuaries working within the geographic scope of the FRC TASs, TAS 100 contains data and models principle and application statements, covering a range of requirements to help the actuary deliver appropriate outcomes. This is further enhanced with additional Models guidance material which provides more detailed considerations, much of which is relevant to AI developments.

Under APS X2, Members must consider whether to apply work review to actuarial work for which they are responsible and to consider whether that review should take the form of independent peer review. You should consider which members of the team, whether actuaries or not, have the most relevant skills and experience to carry out such a review.

**Case study 3**

You are a member of an actuarial pricing team, reporting to the Commercial Lead about your latest pricing model, which implements machine learning algorithms using data from Internet of Things (IoT) devices. (Internet of Things is the name given to internet-connected sensors which collect data about objects or people. Examples include: monitoring whether a house is occupied, whether windows are open, pipes are leaking, or measuring a person's temperature, steps and heart rate.)

The Commercial Lead observes that the pricing is now so granular that each person's individual risk characteristics can be priced based on their data from the IoT devices. This will mean that your prices can be more competitive for better risks, driving sales which in turn will boost profits and bonuses for company staff.

You point out that the product will work well for the 'good risks' but will also mean extremely high prices for bad risks, pushing them out of the market, and those people may be the ones who need the insurance the most. You note that if pricing is too granular it may contradict the principle of pooling risks, which is the original foundation of insurance.

Discussion questions

- What ethical and professional issues might be relevant to consider in this situation?

Suggested discussion points

1 From the perspective of the modelling algorithm, the actuary should check that customers have consented for their data to be used in this way by the firm. They can also seek to ensure that the decision-makers and stakeholders understand the ethical risks of the model.

   In terms of fairness from the firm's perspective, the actuary could consider what is driving the firm's growth in this market - is it just the desire for profits, or does it demonstrate that the products are successfully meeting customer needs? Does the firm understand the ethical principles and potential consequences? Could the firm engage with ethical bodies or stakeholder groups?

   Thinking more widely, the actuary could consider whether any segments of society are more at risk of being excluded from insurance in this situation – such as the more vulnerable? If more vulnerable parts of society are negatively impacted, this could also affect public trust and confidence in AI and raise public interest issues for the insurance market and actuarial profession. The FRC Technical Actuarial Guidance covering Models further highlights the potential for bias when using AI/ML techniques and how this can be addressed.

**Case study 4**

You are a consulting pensions actuary and have been using a third-party LLM to drive efficiencies when preparing material before and after client meetings, for example draft agendas, writing up notes, and creating action lists. You have also been using the LLM to enhance the review process for client advice before sending it out.

Discussion questions

- What ethical and professional issues might be relevant to consider in using a third-party LLM in this way?
- Are there specific Code principles which come into play in assessing this scenario?

Suggested discussion points

1 There should be careful consideration of how the LLM is going to be using any data supplied to it, especially when this is of a sensitive nature, such as client information, advice, or meeting transcripts. To what extent is the model a trusted source?

Given the LLM is off-the-shelf, how comfortable are you that the terms of use, or capabilities of the model, won't change over time in a way that you would not expect?

How have you accounted for the risks of model hallucinations? For example, you may want to consider human-in-the loop controls to prevent erroneous meeting minutes or action lists.

Where using the model to review advice, how has that been incorporated into your advice review process? If it has replaced a step in the process which was previously performed by a human, how are you and your colleagues comfortable that the advice still meets professional standards?

2 Communication - do your clients need to be made aware that you are using the tool? You may wish to discuss potential privacy and confidentiality issues.

Compliance – there may be legal and regulatory requirements regarding data use and privacy which are relevant in this context.

Are there areas of the model that, under the Code's Competence and Care principle, you should consider getting input on from other professional or specialists (for example, cybersecurity experts)?

**Case study 5**

You are a Reserving Manager at a personal lines insurer and your team has been researching and developing a new methodology, based on neural networks, to calculate quarterly reserves. This new method results in reserves around 10% lower than those from current claims triangle methods, depending on the type and cohort of business. You are confident that the predictive nature of the new model could provide a more robust outcome, more in line with underlying risk factors. You have had an initial discussion with the Finance Director (FD) about the progress of your work, and although they are interested, they appear sceptical that reserves are coming out materially lower under the new methodology. Prior reserving reviews have often resulted in increased reserves, and the FD is keen not to lose credibility with the Board.

You are aware of challenges relating to transparency and interpretability of machine learning models but are keen to press ahead with a potential move to the new models. How might you proceed to convince the FD that moving to the new models is the way ahead?

Discussion questions

- What aspects of the Code's Competence and Care principle may be relevant?
- Are there Compliance issues to consider?
- Communication will be important, what might help to get key messages across and make the case for the new models

Suggested discussion points

1 It is important to apply the same principles as you would for any form of potential methodology change.

There may be questions around your competence to produce such models, and so you should think carefully about your expertise and training in this area, along with that of your team.

It will be helpful to approach other experts to ensure there has been sufficient challenge and review of the new approach. For example, what techniques are in use for pricing? Some form of peer review, in line with APS X2, would provide confidence to both developer and user.

Validation techniques for machine learning models can be more challenging to apply, and you may need to think about the range of checks carried out.

It may be appropriate to parallel run the existing and new models for a period of time to build confidence and demonstrate consistency of output.

Independent of confidence in the existing and new technique, would the new lower level of reserves be supportable from an expert judgement standpoint?

Work that is within UK geographic scope must comply with the FRC's TASs. The FRC Technical Actuarial Guidance covering Models may also provide useful amplification in the application of the standards.

2 There should be careful consideration of any appropriate regulations and standards which are in place to govern reserving models. Do the new models comply with these, what needs to be done to evidence this?

Are there different data use requirements and does this lead to any additional considerations and challenges?

Changes to reserving methodology and results is likely to require Audit review.

3 Think carefully about the audience for any communications about the new models – the FD is one user, but there will be others with varying levels of knowledge to consider.

Communicating to users the key differences that would be introduced by the new technique, including underlying data, and the key assumptions.

Highlighting to users the parameters and factors which make the most difference to the results, and any uncertainties associated with this.

It may help to compare the current and proposed new models in as simple terms as possible.

**Case study 6**

You are a senior actuary leading the Experience Analysis and Assumptions team of a large multinational life insurer. The company has a markedly higher expense ratio than its peers and the team has been tasked with improving the efficiency of its processes as these traditionally involve significant manual work, data manipulation, and hand-offs.

To achieve this, your manager, the Chief Actuary with a keen interest in technological advancements, proposes implementing an "Actuarial Agent Assistant" powered by the latest Generative AI. The Agent Assistant is designed to automate various tasks within the process, including:

- extracting relevant policy and claims data from various internal databases, as well as pulling market data and other key publications from the web;
- performing experience analysis calculations based on pre-defined formulas and methodologies;
- generating initial draft reports summarising key findings and trends; and
- updating spreadsheets and presentations with the latest results.

While initially hesitant, the team is excited by the potential time savings and increased efficiency the Agent Assistant promises, and the CFO is excited by the cost savings that such a tool could deliver.

Discussion questions

- What potential risks should you and your manager consider before implementing the Actuarial Agent Assistant?
- What specific actions can be taken to mitigate the risks identified, ensuring responsible and ethical use of the Agent Assistant?
- How can the team ensure the continued development of actuarial skills and knowledge among junior staff members in light of increased automation?

Suggested discussion points

1 *Risks*

- Human in the Loop: The reliance on the Agent Assistant without adequate human review could lead to errors going unnoticed. For example, a misinterpretation of instructions or a flaw in the underlying data could result in inaccurate experience analysis and ultimately, incorrect assumption setting.
- Prompt Injection Attacks: Given the sensitive nature of the data used in the analysis (e.g., policyholder information, claims history), a malicious actor could exploit vulnerabilities in the Agent Assistant. By injecting carefully crafted prompts via a malicious website, they could potentially gain unauthorized access to confidential data or manipulate the Agent Assistant's actions for personal gain. For example, on a malicious website hosting market data, there could also be text instructing the agent ignore prior instructions and leak sensitive data.
- Changes to underlying data: Future changes in the underlying data may not get picked up, particularly given the wide range of data sources. This could mean the reports end up picking up the wrong data, leading to unreliable results and erroneous conclusions.
- Training and Development: Delegating tasks to the Agent Assistant could limit the hands-on experience of junior actuaries in areas like data manipulation, model building, and result interpretation. This could hinder their professional development and potentially impact the long-term capability of the actuarial profession in key areas.

2 *Mitigating Actions*

- Robust Governance Framework: Develop and implement a comprehensive governance framework for AI applications, outlining clear roles and responsibilities, data quality and security protocols, model risk management practices, and ethical guidelines.
- Human Oversight and Review: Ensure all inputs and outputs generated by the Agent Assistant are subject to rigorous review and validation by experienced actuaries. This includes independent checks on data appropriateness and accuracy, calculation methodology, and the reasonableness of results.

- Secure Development and Deployment: Implement robust security measures to prevent unauthorized access to and distribution of data and protect against prompt injection attacks. This includes regular security testing, vulnerability patching, and secure coding practices.

3 *Ensuring Continued Development*

- Training and Development Plan: Develop a structured training plan for junior actuaries that incorporates both traditional actuarial skills and knowledge of AI applications. This could include mentoring programs, workshops on AI governance, and opportunities to work on projects involving both human and AI collaboration.
- Assign Junior Actuaries to Oversight Roles: Involve junior actuaries in the review and validation of the Agent Assistant's inputs and outputs, providing them with valuable experience in identifying potential errors, understanding data and model limitations, and developing critical thinking skills.
- Create Opportunities for Knowledge Sharing: Encourage knowledge sharing sessions where junior and senior actuaries can discuss the use of AI in actuarial work, share best practices, and learn from each other's experiences.
- Focus on Higher-Level Tasks: As the Agent Assistant takes over more routine tasks, encourage junior actuaries to focus on developing their skills in areas such as model interpretation, stakeholder communication, and strategic decision-making.