THOMSON REUTERS
PRACTICAL LAW™

Grant Murtagh's insurance column: May 2024

by Grant Murtagh, Partner, Herbert Smith Freehills LLP

Status: **Law stated as at 22-May-2024** | Jurisdiction: **United Kingdom**

This document is published by Practical Law and can be found at: **uk.practicallaw.tr.com/w-043-3809**
Request a free trial and demonstration at: **uk.practicallaw.tr.com/about/freetrial**

Grant Murtagh, Partner, is a corporate insurance expert at Herbert Smith Freehills LLP. He regularly shares his thoughts on topical insurance issues with Practical Law subscribers.

Grant's column for May 2024 looks at two emerging regulatory issues. The first is an introduction to AI, and some of the legal and regulatory points that are already coming to the fore. The second is the FCA's proposal (in CP24/2) to name the subjects of ongoing FCA enforcement investigations, a proposal that has resulted in reactions from the industry, politicians and others.

AI in insurance

Beginning with a definition is boring and a cliché, and an article created by AI would almost certainly not make that kind of mistake. So, let's start with a definition.

There are almost as many definitions of AI as there are papers about it. Rather than create another, I will just borrow from the ABI's excellent guide on the topic. That guide distinguishes AI from other types of technology through two characteristics: adaptivity and autonomy. "Adaptivity" here means the system can, through training, infer new ways of getting to outcomes (and in turn produce new outcomes). "Autonomy" means the system can make decisions without ongoing human control or other interference.

IT systems that meet those criteria have been around for a while. Why then is there a relatively sudden turning of so many heads towards this part of the technology sector? The short answer is that people can now more clearly see a path towards AI fulfilling the potential that has long been talked about and that, until recently, seemed like it may never actually come to pass.

There are a number of reasons that path has only now come into view. The volume of data that is both useful and available has increased rapidly in the last few years. Computing power has increased exponentially and continues to do so.

Crucially, however, generative AI has been created and has been made accessible to all. And "accessible" here does not just mean it can be used; it also means it can be used as a base for new generative AI projects. GitHub has seen generative AI projects jump from being a fringe interest in 2017, to around 17,000 projects in 2022, to about 60,000 in 2023.

As with all new things, there is an element of hype about what is happening. While some generative AI-based systems show potential, many do not and all have varying degrees of flaws. Even when one discounts for that hype, however, it is still hard to get away from the notion that what is happening is significant. The technology can "learn" at a rate that humans cannot match, drawing otherwise invisible connections between data and producing credible (and sometimes incredible) outputs in the form of text, images, videos, code and more.

Also, the sheer range of possible uses is hard to ignore. Within insurance alone, it could have a role in asset management, product design, pricing, claims, marketing, customer care, distribution and administration, to name just a few. And that may be the most important distinction from what has gone before. Each previous piece of technology has only ever been a potential replacement for a narrow range of work types. Generative AI could have very broad application, which means huge opportunity (for example, through increased efficiency) and significant threat (such as by changing many employees' roles within a relatively short period).

Whether AI delivers on its potential will be an interesting trend to follow. Either way though, legal and regulatory teams have to react. The question for those teams is, where does one start?

As always, one will not go too far wrong by identifying, as clearly as possible, what you are looking to achieve. One goal might be to simply comply with the law. A loftier objective would be to have a more holistic programme that encompasses all of the ethical, reputational and other aspects that go with embracing AI.

Which aim one pursues will not be obvious. Merely achieving legal compliance may sound like the easier route, but that might not be the case as this can be expected to be an area of law that changes quickly (so quickly, in fact, that we may eventually need AI to keep both the design and implementation of the law up to date), and major change to a compliance programme can be quite disruptive and costly. A holistic approach will probably be more difficult and expensive to implement initially, but it would (if done well) have a better chance of being fit for purpose for longer.

The difficulty of achieving compliance over the medium term should not be underestimated. Even when one just focusses on the laws directly targeting AI (as opposed to those that are indirectly relevant, with data protection being an obvious example), it is already clear that different jurisdictions will take different approaches. China has decided to produce legislation on a technology-by-technology basis, while the EU's approach is based on the relative risk of different uses of AI.

The UK is taking a different approach again, deciding in favour of a more principles-based approach. The UK's approach (for now at least) relies on existing laws and tasks sectoral regulators with identifying gaps in the regulations for which they are responsible.

Different approaches will be hard to reconcile for companies with cross-border businesses. Even those that operate within one jurisdiction will, however, likely have to deal with conflicting laws. Much of the legislation purports to apply extra-territorially and, in the UK, the sector-by-sector approach risks different approaches being taken by different domestic regulators.

Pausing on the PRA and FCA for a moment, both responded to the government's request to identify gaps in the existing regulatory framework for which they are respectively responsible. The PRA's response and the FCA's response both suggested that we should not expect dramatic changes to address AI in the near future. This was unsurprising, given both have been looking at this area for years now (they issued, along with the Bank of England, a joint discussion paper in 2022) and neither had signalled that it thought that widespread regulatory change was needed at this point.

Regulatory change will doubtless come in time. An absence of regulatory change will not, however, immunise legal and regulatory teams from having to grapple with some interesting questions, with the following being just three examples:

- **Decision making.** One often hears of AI systems taking "decisions". The law is, however, drafted on the basis that binding decisions, such as entering contracts, are taken by legal or natural persons (section 43 of the Companies Act 2006 being one example). If an AI system (which is not a person, legal or natural) takes a decision, will it be binding?
- **Governance.** Governance systems rely on people who understand an area taking responsibility for the decisions that they make. If one introduces an AI system that can reason in a way humans cannot, who takes responsibility for the decisions that the AI system reaches? It might be that, in time, AI systems will be able to both reach decisions and explain how it reached them in terms humans can understand. Until that "explainability" is there, however, will boards and senior managers get comfortable with taking responsibility for decisions if they do not understand the inner workings of the AI system? That question will apply to companies generally, but is particularly relevant to insurers (and other financial services firms) as the senior managers and certification regime (SM&CR) attaches personal responsibility to the directors and senior managers. If those individuals cannot get comfortable with that responsibility in those circumstances, how can AI systems be adopted into businesses?
- **Intellectual property.** In very general terms, a core concept in IP law is that it protects creative and innovative inventions by the human mind. In the context of AI, that core concept can be difficult to apply. Which human mind counts? Is it the original programmer, the person who owns the data that trained the AI system, the person who gave the AI system a command or someone else? More radically, should the AI system be treated as if it were a person for these purposes, with the value of whatever the AI system creates then flowing to whoever owns that system? The courts are already facing questions like these, with the Supreme Court's decision in *Thaler v Comptroller-General of Patents, Designs and Trademarks [2023] UKSC 49* being a prominent example.

This is a fascinating area of law that can only be touched upon in a short article. It challenges legal concepts that have been around for a very long time and raises some entirely new legal issues. While it will be possible to deal with many of the issues that arise by creatively using existing laws, some judges have already queried whether that amounts to judges straying into the role that should be performed by the legislature. For now, the UK government is espousing a "wait and see" approach. How long that can last is an open question.

The FCA's "name and shame" proposal

Since my February 2024 column, the FCA has consulted in CP24/2 on changes to its enforcement approach. There is a lot to say on the consultation, but there is one area that deserves particular attention.

The FCA has proposed naming the subject of an investigation at the time that the investigation opens, subject to a new public interest test. The risks of that approach for the firms, and potentially the individuals, being investigated are significant, and would range from financial to reputational. Those risks would arise regardless of whether the subject of the investigation is innocent or guilty.

If 100% of the FCA's investigations resulted in sanction, one can see how the FCA might be able to build a justification, but that is not the case. The FCA's own numbers show that 65% of investigations close without further action. Given these are real people and businesses who would, if the FCA proposals applied, have been directly affected, it is worth pausing on that number and questioning whether the purported benefits stack up:

- The FCA suggested the approach would better protect consumers and investors. This would, however, serve to "protect" them from the 65% that the FCA decides not to take action against after investigating. How would the inevitable damage to those businesses and individuals be justifiable?
- The FCA has said that the proposal will address public concern or speculation. Doing that would inevitably mean adverse inferences would be drawn against the 65% of investigation subjects against whom no action is ultimately taken. Is that really an appropriate response to rumours?
- The FCA feels its approach would deter future rule breaches. Does publishing names at an early stage really achieve this? Details of the type of issues being investigated seems more pertinent.
- Some have argued it will encourage firms to settle sooner. That would be desirable if all of the FCA's investigations established actionable guilt, but 65% do not. Is that really likely? If names have been released, and so damage has been done, that seems like it would remove some of the incentive to settle early (and potentially increase the incentive to appeal), particularly for those who feel the FCA is pursuing a weak case.

The arguments against the proposal seem much stronger. The FCA's proposal risks creating an assumption of guilt in the mind of the public and public commentators, regardless of any caveated language the FCA announcements may include. It is questionable whether it would be seen as following due process. It seems to put too little weight on the damage the approach would have on those who are named. It does not recognise that such damage could be irreparable. It would make the UK an outlier on its approach to enforcement, and so impact the UK's competitiveness.

The reaction to the proposal has been strongly negative. This includes the House of Lords Financial Services Regulation Committee initially expressing concern in a letter to Nikhil Rathi, FCA Chief Executive, and then making clear in a subsequent letter to him that it was not satisfied with the FCA's response to its questions.

The FCA should rethink its approach. Many of the same ends could be achieved in far less controversial ways, such as the FCA using its newsletters to highlight the types of breaches it is investigating. It is hard to see how the damage the current proposal would inevitably cause merits the limited additional benefit (if any) of naming investigation subjects at the start of the investigation process.