AI in regulation

This note considers the use of AI in regulation – not the regulation of AI, but rather how AI might be used by regulators (particularly competition regulators) in their work. Our focus is on use-cases specific to regulation, rather than on the general AI productivity improvements common to many organisations.

AI is a diverse field, but we will primarily consider generative AI (GenAI): models that generate coherent output (such as text, images or code) based on learned patters from training data. ChatGPT, Claude and Gemini are examples.

Current constraints of AI

GenAI is a powerful technology, and we discuss below some of the substantial impacts it may have on regulation. However, it also has some material constraints, some general and some particular to the regulatory context.

As a general matter, current models can be given to hallucinations – simply making up facts. For instance, the Trump administration's recent 'Make America Healthy Again' report appears to have been at least in part AI-generated, and includes numerous citations of studies that do not in fact exist.[1]

Benedict Evans has described AI (at least today) as giving you 'infinite interns'.[2] As he notes, one of the problems with interns is that you have to check their work. For some types of problem, checking the work takes as long as doing it, and AI may be less suited to such tasks (again, at least for the time being).

For regulation in particular, AI's limits include:

- AI cannot be held responsible for its decisions. You might reweight a model after bad outputs, but you cannot in any meaningful way penalise it for failure
- In a regulatory context, decisions must be explicable and justifiable, both for legal defensibility and for public trust. AI findings can be 'black boxes', and fail this test (though more recent models are getting better at showing their reasoning)
- AI may be a poor tool for complex, occasional decisions, where training data is lacking, and where decisions need to be 'right', not just statistically better. (This contrasts to the use of AI to interpret medical imaging, say, where accuracy that is 10% better that that of human radiographers is immensely valuable)
- While AI may support the analytical side of regulatory decisions, it is less well placed to handle the political judgement and wider contextual trade-offs that are often an important part of such decisions

Thus AI is unlikely to be allowed to take (material) regulatory decisions any time soon. But it can nonetheless have enormous value.

Contribution to the regulatory decision-making process

While AI may not make decisions, it can support the process that leads to them. It can:

- Find useful precedents. Specialist regulators know their own industry well, but may be less familiar with other markets that have equivalent economic and competitive structures. AI can easily identify precedents from these
- Provide an outline or rough first draft of papers, providing a useful strawman to prompt early thinking
- Summarise responses to consultations. The Scottish government has trialled this in a consultation on non-surgical cosmetic procedures, and found that the AI

[1] Washington Post, *White House MAHA Report may have garbled science by using AI, experts say*, 30 May 2025
[2] Benedict Evans, *Generative AI and intellectual property*, 27 August 2023

delivered very similar results to human experts, at far less cost.[3] Ofcom too has been investigating this[4]

- Anticipate challenges (to draft findings, for instance). An AI might be asked "Imagine you are the incumbent - what would you object to in this document and how would you argue against it?"
- Counter group-think by acting as a devil's advocate ("How would you critique this proposal?")
- Provide multiple perspectives ("Imagine in turn you are a consumer of these ten types. How do you view this proposal?")
- Undertake quick post implementation reviews. While major decisions will benefit from a traditional PIR, AI can enable 'quick and dirty' reviews, to allow assessment of a broader range of decisions, and to identify those that may merit deeper investigation. AIs have the advantage that they have no conflict of interest in doing a PIR

One could imagine a standard set of prompts, to be used at different stages of a regulator's decision-making process.

Contribution to work planning

Quick and dirty PIRs may point to regulation that needs revisiting. AI can also be used more generally to test whether regulation remains current: "Review this determination from 2020; identify the key assumptions that underpinned the decision; determine whether these assumptions are still valid; and if not, assess whether updated assumptions would be likely to change the decision".

Periodic queries of this type may help identify areas for future work.

Contribution to monitoring and enforcement

AI has the potential to pick up monitoring tasks that would otherwise be too labour-intensive or impractical.

One early example is the CMA trial of AI to detect bid rigging in public procurement.[5] Another is Ofcom's development of Machine Learning techniques to monitor digital platforms' compliance with online safety requirements.[6]

Looking ahead, many other applications are possible.

Media regulators typically do not monitor the output of broadcasters to check for compliance with content codes, but rather rely on complaints to prompt investigations. However, in future it may be possible to have 'watching bots' that flag up potential breaches. This might also allow for better detection of patterns of behaviour than ad-hoc complaints do.

AI might monitor consumer offers, to flag those that may appear to be unfair in some way. For example, automated review of small print might identify burdensome clauses.

AI could also be used to review notifications of smaller mergers (that might otherwise be ignored on the basis of scale) to identify those that are potentially problematic and merit further review.

Approaches like these might lead to changes to regulator data requests. For instance, AI review of Ts & Cs would likely be coupled with a requirement on companies to notify the regulator of changes to those terms, if these were not already in the public domain.

[3] DSIT, *Government-built "Humphrey" AI tool reviews responses to consultation for first time, in bid to save millions*, 14 May 2025
[4] Ofcom, *Ofcom's strategic approach to AI, 2025/26*, 6 June 2025
[5] Procurement Magazine, *CMA: Tackling 'Significant Risk of Bid-Rigging' with AI*, 22 January 2025
[6] Ofcom, *Online safety: our research agenda*, 15 April 2024

Contribution to regulatory burden

While information requests to support regulatory AI analysis might increase regulatory burden, there are other areas where AI may enable a reduction.

For example, when considering a planning application, a consumer can get informal guidance from planning officers before making a formal submission. If the project is flawed, this guidance helps avoid spend on plans and applications that are likely to fail. Could an AI – trained on a regulator's internal data – provide indicative guidance on whether a proposed merger was likely to be approved, or whether a licence application would be successful? It might also identify areas of concern.

Clearly this could not be definitive guidance. It might be couched as a steer that the application was 'more likely' or 'less likely' to be successful. But it might – like the planning guidance – offer substantial savings in wasted effort avoided.

Changed expectations of regulatees

Impacts on regulation will not simply be from AI use within regulators. AI will also change regulators' legitimate expectations of regulatees, and their analysis of the costs and benefits of regulation.

For example, AI is increasingly employed in fraud prevention techniques (as it is in fraud itself). Norwegian banks saw a 90% reduction in phishing losses after implementation of AI-based detection tools.[7] Other companies are using AI 'victims'- such as O2's 'dAIsy' - to take up fraudsters' time and learn their techniques.[8]

The availability of such tools potentially changes the calculus for regulators considering incentives for financial institutions and other organisations to reduce fraud rates on their platforms. Preventable fraud may be an externality for many platforms, but if more fraud is preventable, then there is greater logic for structures that internalise the fraud cost to the platform, creating appropriate incentives.

There are many other potential examples. Financial firms have a duty of care to ensure consumers are sold suitable products. If AI allows tailored advice at scale, then the bar for 'suitable' might legitimately be set higher.

For customer service, if AI agents can effectively handle a large number of queries, then call waiting times should fall – and regulators may be dissatisfied if they do not.

These examples are within the framework of existing regulation. But it also possible to imagine AI leading to fundamentally different regulatory or policy approaches.

For example, countries with minority languages often support them with minority language broadcasters (MLBs), such as S4C broadcasting in Welsh in the UK, and TG4 in Irish in Ireland.

However, AI tools such as TrueSync now allow comparatively low cost 'visual dubbing' (where both the audio and the lip movements of speakers are converted to another language).[9] This may create the potential for an entirely different policy approach in the future – a mandate on some or all national broadcasters to provide a minority language feed of their programming. This could provide far more minority language programming, at potentially much lower cost than a standalone MLB. (That said, such an approach would not address the distinct cultural needs of minority language audiences).

Changed expectations of consumers

Much regulatory intervention is designed to redress imbalances of power and information between companies and consumers. For

[7] Visa, *Visa Boosts AI Capabilities to Further Reduce Fraud*, 3 April 2025
[8] O2, *Meet dAIsy, the scam-fighting AI bot*, 14 April 2025
[9] Flawless, *TrueSync* [accessed 30 May 2025]

instance, Ofcom has made numerous interventions to prompt consumers to consider switching suppliers, and to provide them the best information when they do so.

However, AI has the potential to (effectively) make consumers considerably more sophisticated. For example, Trim is an AI-powered financial assistant that scans your bank account to identify areas for savings, which can also renegotiate telecoms bills on your behalf.

Not all consumers will use AI purchase agents, but as more do, the cost/benefit trade-off of interventions to support consumers in their purchase decisions will change.

More generally, information asymmetry and the advantages it may give providers over consumers are an important area of regulation. If AI support for consumers helps reduce this asymmetry, it may in time create the potential for reduce regulation in a number of areas.

Risks

Use cases like those above may recalibrate regulation, and make it more efficient. However, there are also risks to consider.

To pursue some of these opportunities will require specialist IT skills, and these are in high demand. Regulators will need to consider which skills they need in-house.

As with other organisations, regulators may need an active push to shift non-specialist staff to using AI (where appropriate), since it requires new skills and modes of working. To encourage AI usage, companies are providing training, prompt libraries, employee innovation incentives, time off for innovation, celebration of AI methods adopted, and so on.

As we have noted, regulators also have a particular need for transparency and replicability in some aspects of their work. Generative AI platforms are largely a black box for users, and a dynamic and customised one at that. Thus even if a regulator publishes the prompts it used in a given context, a different body on a different day using the exact same prompts night get a different result.

Confidentiality concerns will also be important for regulators, particularly if models are be to trained with internal data. This may argue for the use of local rather than hosted models. Care will also be needed to avoid inadvertent exposure of training data to any external users.

There is also a fundamental issue of the biases that may be inherent in AI because of the set of training data. For instance, regulators typically publish extensively when they intervene in a market. By contrast a 'do nothing' decision may not be publicly documented at all. Thus training data may include much justification for regulation, but less on the benefits of forbearing.

This is another reason to be cautious of using AI for regulatory decisions, rather than regulatory analysis. It also highlights the importance of 'devil's advocate' prompts that explicitly ask the AI to consider the case against an intervention.

Further, by its nature, training data is historic, and thus AI is less likely to be 'cutting edge' in its thinking. For instance, there is a long history of regulators being wary of '4 to 3' mergers in mobile telecoms. Recently, regulators have been less hostile, but an AI trained on the corpus of regulatory decisions is perhaps likely to hew to the historic line.

Conclusions

Regulators – like most organisations – are still feeling their way in the deployment of AI. However, it has the potential not just to reduce costs, but also to enhance decision making and to reshape engagement between regulators and their stakeholders.

Robert Kenny
rob@commcham.com

June 2025