


AI in tax administration: Harder, Better, Faster, Stronger

General Features

Management of taxes



23 June 2025

We consider the growing role of artificial intelligence in tax administration and its use in automated decision-making by tax authorities.

Artificial intelligence (AI) has been the recent buzz. Whether it's a self-released deepfake video of Emmanuel Macron, market chaos because of DeepSeek or the new acquisition of a data centre, AI has become mainstream news now. Private enterprises and public bodies alike are adopting AI on a larger scale. The advantages of adopting AI are clear: enhanced speed and efficiency, cost savings, the ability to find correlations that may have been previously undetectable to the human mind, systematic and consistent decision-making.

People define AI differently and this tends to cause some confusion. In this article, AI refers to software that simulates elements of human behaviour such as learning, reasoning and classification; in particular AI in this article refers to machine learning. AI, in this context, does not include software developed through traditional decision-

trees (in other words traditional algorithmic tools).

In September 2025, the Institute of Fiscal Studies' Tax Law Review Committee (TLRC) published a paper entitled 'Artificial intelligence in automated decision-making in tax administration: the case for legal, justiciable and enforceable safeguards', written by the author of this article. The paper sets out the author's view in support of the deployment of AI for automated decision-making in tax administration, noting that this deployment is inevitable, and given the benefits of AI could bring significant efficiencies.

Benefits and questions

From a tax administration perspective, there are also additional potential benefits of using AI, such as detecting undetectable or hidden correlations, suspicious activity, trends and indicators of tax loss, etc. (sometimes in real time) and facilitating the use of pre-emptive or defensive measures. All of these ultimately could have the effect of bringing down the tax gap (which has been a key focus of recent governments).

However, as noted in the paper, the current legal tax framework does not properly facilitate the use of AI to make discretionary decisions in tax administration, such as determining the amount of certain penalties to be imposed. There are questions around the legality of the use of AI in tax administration and the availability of adequate safeguards for taxpayers. These questions derive from the fact that the existing tax administrative framework was set up in a world where decisions were primarily made by human HMRC officers.

In a world where AI is deployed to make discretionary decisions, the primary decision-maker would be the AI technology (and not the human). The 'black box' nature of machine learning means that the technology does not provide an explanation for why a decision was arrived at, and programmers are unable to explain with certainty why a decision was arrived at. This means that any human subsequently attempting to explain the rationale for a decision made by AI would in effect be reverse engineering such a decision. Any explanation offered by the human would be of why they **think** a decision was arrived at, rather than why a decision was **actually** arrived at. This inevitably involves a degree of speculation.

Since publication of this paper, the TLRC has received further comments and has held discussions on the topic, including with the HMRC's Professional Standards Committee. The advisory Professional Standards Committee provides oversight on how HMRC administers the tax system and offers critical challenge to how HMRC exercises its powers, supporting good practice in the use of its powers and safeguards.

A summary of the meeting at which the TLRC paper was presented is published at: tinyurl.com/28ndax7k. The rest of this article summarises some of the key points that have emerged from the paper and the discussions around it.

Make it

There is no publicly available list that sets out the processes in which AI is currently deployed by HMRC. From discussions with HMRC, it is understood that the key areas in which AI is currently being used by HMRC are compliance risking (for example, risking taxpayers at the time they file tax returns or make applications for repayment) and detecting VAT fraud through VAT risking.

HMRC is also in the process of developing an internal large language model (LLM) which would be made available to HMRC officers to help them answer questions asked by taxpayers in real time based on HMRC's guidance. Canada has a similar chatbot that it developed (niftily) named Charlie the Chatbot, although this chatbot is made available directly to taxpayers.

It seems that HMRC has no plans to make its LLM available to taxpayers at this stage. HMRC's view is that the development and use of AI tools for automated decision-making is still a few years away; however, given that there has been a general trend towards optimising public sector efficiency, both within the UK and outside, these developments may come sooner than anticipated.

AI has already been deployed by public bodies around the world with varying uses, including the US, Australia, the Netherlands, France, Slovakia, India and others. Even within the UK, departments such as the Department for Work and Pensions have been reported to be using AI to assist in decision-making. In fact, since December 2024, there have been a large number of entries published on the Algorithmic Transparency Standards Hub (ATS) (see tinyurl.com/6au3ynjs), which discloses different government bodies developing technology involving AI (and machine

learning) tools to varying degrees to assist with tasks.

While these tasks are in many cases administrative and non-discretionary, the disclosures published show an increased use of AI in UK public administration. It is unclear whether the ATS represents all the uses of AI (and machine learning) tools in public administration as of yet, but some clarity on the point would be welcome.

The TLRC's view is that it is not practical to assume that AI will not be more widely deployed by HMRC in the relatively near future for automated decision-making. There has been an overall push to bring down the tax gap and improve efficiencies in HMRC (especially given pressures on public finances), and AI can serve as a helpful tool to do so. HMRC already has an internal data sciences team that works on AI solutions. Once an internal or governmental decision is made to develop an AI tool, the pace of development could therefore be quick (as it is understood that HMRC generally would look to develop these tools internally).

It would, however, take some time to socialise, legislate for and develop a new HMRC AI Charter (along with corresponding changes to existing legislation). Therefore, mitigating the risks posed to taxpayers by the use of AI in automated decision-making, and ensuring that taxpayers have justiciable rights, requires pre-emptive action rather than retrospective action, as explained later in this article.

As the reported minutes of the Professional Standards Committee state:

'The members acknowledged that although HMRC is a long way from using AI in the manner hypothesised in the paper, it is important to begin considering the risks and the appropriate protections in any proposed future use. Members were reassured that many of the suggestions in the paper such as best practice model input processes, technical standards and extensive model testing pre-, during and post-deployment, were already in place.'

This is encouraging and it would be good to have more transparency around this on the HMRC website.

Work it

It is important that taxpayer safeguards are brought into effect pre-emptively, before AI is deployed by HMRC in automated decision-making.

As discussed in the report, there are questions around the legitimacy of using AI in tax administration without further legislation facilitating this. This view is yielded further weight by the judgment in the Court of Appeal case *Peter Marano v HMRC* [2024] EWCA Civ 876, where Asplin LJ notes with respect to Finance Act 2020 s 103 that: 'Section 103 is not intended to authorise the use of artificial intelligence.'

Further, the objective of using AI in automated decision-making in an institution is to replace human decision-making with AI decision-making. Given the scale of deployment of AI, any biases or errors in the technology can have significant effects on a large number of taxpayers before detection, even where there is a human element or overview. In view of the black box nature of machine learning, detection (at least based on current levels of technology) can be difficult, which enhances the risk posed to taxpayers.

For example, in the Dutch *toeslagenaffaire* 11,000 parents were subjected to audits on a discriminatory basis due to their non-Dutch nationality. In the Australian *Robodebt* matter (which technically did not involve AI but similarly involved the use of algorithmic tools that replaced human decision-making), AUS \$746 million was erroneously recovered from 381,000 people (with AUS \$1.751 billion of debt having to be written off).

Supplementary tools that are intended to provide an explanation of the reasons for decision made by the AI (also known as 'explainable AI' or 'XAI') are not developed enough to provide certainty and present elements of unreliability.

Some argue that humans can also have biases and make errors; and therefore there is no need for any additional safeguards where AI is being deployed. This argument does not appreciate the scale of automated decision-making as demonstrated above. Given the systematic and institutional deployment of AI (and speed of decision-making), the number of decisions made by AI in a few years (or less) would often outweigh the number of decisions made by a human HMRC officer or group of HMRC officers in a career.

Further, to train an automated decision-making system using machine learning, the system would (under current methods of development) need to be trained on large volumes of historic data. However, human approaches to decision-making and

biases evolve (and have evolved) over years. Therefore, safeguards are needed to ensure that a system is not being trained using historic data that imports outmoded biases into decision-making.

The *COMPAS* risk management system in the US is used to predict the likelihood of offenders committing future crimes (i.e. recidivism). It was found to discriminate against Black American offenders more than White American offenders as the system had a higher false positive rate for Black American offenders than White American offenders (i.e. it incorrectly predicted that Black American offenders would reoffend more than for White American offenders). Some research suggested that the reason for this was because the system used number of arrests as one of the factors to predict recidivism. Given the historic difference in law enforcement approaches between the two races (e.g. Black Americans historically were arrested for marijuana offences more than White Americans, even though both use marijuana at approximately equal rates), this resulted in higher false positives for Black Americans (thereby importing bias into the system). The literature on AI deployment in public administration generally coalesces around having robust training data in place to ensure that the AI developed is not biased.

As seen from its published minutes, members of the Professional Standards Committee raised an important point about building a culture within tax authorities such that where AI is deployed (or is in the process of being developed): 'HMRC ensure that any proposed systems were supported by a culture that included healthy consultation, high levels of scrutiny and to be accepting of the need to have defined routes for errors to be escalated, acknowledged, and corrected.'

The TLRC takes this to mean that if HMRC officers (or other teams at HMRC) spot errors in decision-making, logic or the process of development, the internal culture should foster the ability of the individual to report such issues, knowing that will be without any backlash; in other words, 'whistleblowers' should be adequately protected. The Professional Standards Committee's view was that building this culture will help to mitigate some of the risks posed by AI by enabling robust development, deployment and vetting systems for AI. The TLRC agrees with this, and the importance of culture cannot be understated.

However, culture in and of itself is unchallengeable. HMRC already has certain internal checklists and processes in place that it follows when developing AI. Although this is a step in the right direction, none of the internal checklists and

processes have been disclosed to the public. Further, internal checklists, processes and culture are not challengeable and are unjusticiable. If for any reason individuals at HMRC didn't adequately follow these checklists or processes or ignored them, taxpayers would be without proper redress.

Do it

The discussion paper recommended that an HMRC AI Charter should be legislated for and that this HMRC AI Charter should set out some of the key parameters necessary for HMRC when deploying AI technology in the context of automated decision-making. The implementing legislation should be drafted such that the HMRC AI Charter constitutes affirmative obligations of HMRC rather than mere intentions, which are unchallengeable.

The paper's view is that legislation enabling the use of AI in tax administration with robust remedies and protections for taxpayers, coupled with an internal culture protecting the safe development of AI, would provide HMRC with the legitimacy to deploy AI in tax administration, not least because this would follow due democratic process.

The Australian National Audit Office (ANAO) published a report analysing the governance of AI at the Australian Tax Office (ATO) and found similar shortcomings as expressed in the TLRC paper, making similar recommendations (see [tinyurl.com/3m52p9ej](https://www.tinyurl.com/3m52p9ej)). The ATO agreed to the recommendations by the ANAO and have confirmed that they are working towards implementing them. HMRC should not lag behind.

Summary

In summary (and hopefully, 90s electronic music fans have caught on already), I need not say more than Daft Punk already has:

Work it, make it

Do it, makes us

Harder, better

Faster, stronger

Like Daft Punk did with electronic music in the 90s, HMRC has the opportunity to take the lead in fundamentally changing the face of tax administration and making it fit for the AI era.

© Getty images/iStockphoto