Book Section:
Tomlinson, Joe and Thomas, Robert (2025) Judicial Review of AI in UK Government: Two Emergent Forms of Legal Risk. In: Ng, Yee-Fui and Groves, Matthew, (eds.) Automation in Governance. Hart/Bloomsbury.

Judicial Review of AI in UK Government: Two Emergent Forms of Legal Risk

Robert Thomas* and Joe Tomlinson+

January 2025

Abstract.

While the application of the general principles of administrative law is often uncertain to some extent, how these principles apply to the use of AI and automation in public administration in the UK is currently highly uncertain. In this unpredictable legal environment, legal risk is accumulating as the use of AI and automation in public administration rapidly expands. Specifically, this chapter points to how there are two distinct forms of risk developing. One legal risk is that public bodies are proliferating the use of systems that will later be found to be unlawful. This could store for the future serious and wide-scale implementation problems. A second legal risk is that systems of ambiguous legality become so embedded in bureaucratic practice that, by the time meaningful judicial review occurs, the courts will be in an invidious position. They will have to address fundamental questions of administrative law fixed with the knowledge that rulings might have wide-ranging adverse impacts across the public sector, and this could then encourage an unduly deferential jurisprudence. The chapter concludes with some observations about how—beyond the growing debate about new statutory forms of regulation—ensuring ventilation of claims, non-court redress mechanisms, and the flexibility of judicial review remedies might all have a role to play in allowing the public law system to effectively manage these emergent risks.

* Professor of Public Law, University of Manchester.
+ Professor of Public Law, University of York.

1 Introduction

AI and automation are increasingly being used by public bodies in the UK to make administrative decisions and to deliver public services.[1] While AI is becoming more widely used, we are only at the beginning of the coming wave that will transform how public administration functions.[2] General administrative law is an essential part of how the public sector's growing use of AI is regulated.[3] At present, its role is particularly important as the UK has thus far adopted a 'pro-innovation' approach to AI which, in plain language, means it is not yet creating any further legislative frameworks for AI.[4] Despite this, we have only seen two judicial review judgments from UK courts that engage meaningfully with how administrative law applies in this setting.[5] This lack of cases is likely the product of a range of factors, including procedural barriers, the opacity of the AI systems, and litigant behaviour. Whatever the cause, the outcome is that, while the application of administrative law principles in specific settings is often uncertain to some degree, there are presently strong grounds for believing there is an unusually high amount of legal uncertainty surrounding the use of AI in the public sector.

In this chapter, we set out the current situation in the UK and argue that, in this unpredictable legal environment, legal risk is accumulating as the use of these systems rapidly expands. Specifically, we observe how there are two distinct forms of risk emerging. One risk is that public bodies are proliferating the use of systems that they will later realise to be unlawful—which could create serious and wide-scale implementation problems for public bodies.[6] Another risk is that systems of ambiguous legality become so embedded in bureaucratic practice that, by the time meaningful judicial review occurs, the courts will likely be placed in an invidious position. They will be required to address basic questions of whether AI systems comply with principles of administrative law. However, the courts will be aware that their rulings will likely have wide-ranging – and potentially

[1] For evidence of current use, see the UK Government's 'Algorithmic Transparency Recording Standard Hub Repository' and the Public Law Project's 'Tracking Automated Government Register.' At present, the latter appears to be much more detailed and useful.
[2] M Suleyman, *The Coming Wave: AI, Power and our Future* (London, Vintage, 2024).
[3] We refer to public law in a broad sense here, to include statutory regimes such as the Human Rights Act 1998 as well as common law principles. For a full analysis of the applicable regimes, see: B McGurk KC and J Tomlinson, *Artificial Intelligence and Public Law* (Hart Bloomsbury, 2025).
[4] Department for Science, Innovation and Technology, *A pro-innovation approach to AI regulation* (CP 815, 2023).
[5] Our focus is on the UK in this chapter, but we understand similar issues are arising in other jurisdictions.
[6] We use the concept of legal risk in this sense to mean the chance government takes action or adopts policies/processes that will later be found to be unlawful. This is an idea less discussed in research, but one central to public law in practice, see e.g. Government Legal Department, *Guidance Note on Legal Risk* (2024).

negative and costly– impacts across the public sector. This could lead to considerable upheaval within government. Alternatively, the courts might adopt an undue level of judicial deference. After setting out this often-neglected dimension of the debate surrounding legal regulation of AI in the public sector, we conclude the chapter with some observations about how—beyond the growing debate about the possibility of new statutory forms of regulation—ensuring ventilation of claims, non-court redress mechanisms, and the flexibility of judicial review remedies might all have a role to play in allowing the public law system effectively to manage this emerging situation.

2 Judicial review cases so far

There have been only two judicial reviews of the use of AI, or more simple automated systems, to reach fully reasoned judgments. Other cases have been brought, but they have often been conceded by defendant public bodies or otherwise resolved before a full judgment is reached.[7]

The Court of Appeal's ruling *in Bridges v South Wales Police* is arguably the most extensive judgment to date.[8] The case centred around the South Wales Police's use of an Automated Facial Recognition ('AFR') system. South Wales Police had been deploying AFR technology since 2017 to identify individuals in real-time by comparing live camera feeds against pre-determined 'watchlists.' Ed Bridges, a private citizen, brought the case with the support of the campaign group Liberty, alleging that his image was captured during AFR deployments in Cardiff and Swansea without his consent. He argued that this breached his rights under Article 8 of the European Convention on Human Rights, the Data Protection Act 2018, and the Equality Act 2010. At first instance, the High Court upheld the lawfulness of the AFR system, and Mr. Bridges appealed to the Court of Appeal.[9] The Court held that the use of AFR technology interfered with Article 8 rights due to its intrusive nature.[10] Even though South Wales Police had legitimate aims, the court found the system lacked sufficient safeguards to ensure its use was proportionate: the watchlist criteria were not clearly defined or subject to rigorous scrutiny; no

[7] Arguably the most prominent settled case concerned a discriminatory visa streaming system, which was settled by the Home Office. For a full account of this case, see: J Maxwell and J Tomlinson, *Experiments in Automating Immigration Systems* (2022, Bristol University Press), ch. 4.
[8] *Bridges v South Wales Police* [2020] EWCA Civ 1058.
[9] *R (Bridges) v Chief Constable of South Wales* [2019] EWHC 2341 (Admin).
[10] *Bridges* (above n 8).

independent oversight mechanism existed to govern the police's use of AFR; and the decision-making process regarding where and when to deploy AFR was discretionary and lacked a clear legal framework. This failure to establish adequate safeguards rendered the interference unjustified under Article 8(2). The Court also found that South Wales Police had failed to meet its obligations under the Data Protection Act 2018 and the UK General Data Protection Regulation. Specifically, there was no satisfactory Data Protection Impact Assessment addressing the specific risks posed by AFR and insufficient controls were in place to minimise potential data breaches or misuse of sensitive biometric data. The Court concluded that the deployment of AFR technology did not meet the legal requirements for transparency, necessity, and proportionality under data protection law. Lastly, the Court addressed claims of discrimination under the Equality Act 2010, particularly concerns that AFR technology might disproportionately misidentify individuals based on their race or gender. While no direct evidence of bias was presented, the Court criticised South Wales Police for failing to investigate whether the system had such discriminatory impacts. This failure constituted a breach of the Public Sector Equality Duty under the Equality Act.

The second case was *R. (on the application of Johnson) v Secretary of State for Work and Pensions*.[11] *Johnson* was specifically concerned with the administration of a social security benefit –Universal Credit – and how automated payments were calculated. The claim was brought on behalf of three single working mothers by the Child Poverty Action Group, a longstanding campaigning body, and was joined with a similar case involving another single working mother. Each claimant was paid a monthly salary towards the end of the month (eg. on the 28th of each month or the 'last banking day of the month'). Each claimant also had an assessment period for their Universal Credit which started and ended close to the end of the month. However, as is the case for many employees, the claimants' actual pay dates were occasionally brought forward to take account of weekends and bank holidays. This disjuncture created a 'non-banking day salary shift' problem at the centre of this case.[12] Where the pay date was shifted earlier, claimants received double their usual income in an assessment period, leading to reduction (or complete removal) of their Universal Credit award for that month. This loss was not compensated by the system calculating their next month's benefit on the basis of no income. This created almost

[11] [2020] EWCA Civ 778. The relevant regulations were Universal Credit Regulations 2013. It is important to note the substantially similar and linked case of *R (Pantellerisco and others) v Secretary of State for Work and Pensions* [2021] EWCA Civ 1454, which we do not consider separately here.
[12] J Meers, ''Fatally upsetting the computer': universal credit, earned income, and the demands of automation' (2024) 42(4) *Journal of Social Welfare and Family Law* 520.

immediate cashflow problems and a range of other serious impacts for the claimants and many other people similarly situated. For example, the claimants fell into rent arrears, defaulted on council tax, incurred bank overdraft charges, borrowed money and became reliant on food banks to make ends meet. One of the claimants was required to decline a promotion on the basis of this issue and another gave up a job to find alternative employment with different pay dates. The claimants in *Johnson* argued that this situation was unlawful. They succeeded at first instance, on the ground that the system's operation was inconsistent with the relevant regulations.[13] This was despite arguments from the relevant government department, the Department for Work and Pensions (DWP), that the regulations ought to be interpreted in light of the 'fact that the system of universal credit was intended to be automated.'[14] To this end, the department relied on evidence from a DWP official 'indicating the importance of automation in the design of the system of universal credit and indicating that it would not be possible to make an automated change to address the issue that has arisen in this case' and thus 'any solution would have to involve a manual calculation of the amount of the award.'[15] The first instance court, the Administrative Court, did not entertain this position:

> First, this is a question of statutory interpretation. If the regulations, properly interpreted, mean that the calculation must be done in a particular way, that is what the law requires. We do not belittle the administrative inconvenience or the cost involved but the language of the regulations cannot be distorted to give effect to a design which may have proceeded on a basis which is wrong in law… Secondly, the existing regulations already contemplate manual intervention at some stages.[16]

However, the Court of Appeal overturned this construction of the regulations:

> If the Regulations were drafted with the aim of ensuring that the calculation of awards each month could be an automated process, it is legitimate to take that purpose into account when applying the canons of construction.[17]

The Court of Appeal held that the Administrative Court's approach would require the government to make an 'evaluative determination' in many different types of case about whether a person's income was earned in respect of that period.[18] It further observed that

---

[13] *R (Johnson) v Secretary of State for Work and Pensions* [2019] EWHC 23 (Admin).
[14] Ibid [59].
[15] Ibid [59].
[16] Ibid [59-60].
[17] [2020] EWCA Civ 778 [40].
[18] Ibid [39].

this ran contrary to one of the purposes of the regulations, namely: 'ensuring that the calculation of awards each month could be an automated process.'[19]

These cases demonstrate the applicability of established administrative law frameworks and principles to AI and automated systems used in the public sector. They also show that the novelty of these systems can give rise to novel legal questions about the application of administrative law—questions which require judicial resolution and to which there might not be clear answers.

These cases also show the value of reasoned judgments in clarifying the application of existing laws to AI systems thus reducing the extent of legal uncertainty for both claimants and defendant public authorities alike.[20] Arguably, *Bridges* is a paradigm example of this. It was widely accepted prior to the case that the legal position pertaining to the use of facial recognition technologies by the police was unclear. As Haddon-Cave LJ and Swift J noted in the Administrative Court judgment, the parties 'brought these proceedings before the Court in order to seek the Court's early guidance as regards the legal parameters and framework relating to AFR, whilst it is still in its trial phase, and before it is rolled-out nationally. We commend the spirit in which these proceedings have been brought.'[21] Whether one agrees with the way the Court of Appeal ultimately resolved the questions or not, the judgment represents some progress being made in clarifying what the applicable legal framework is—even for critics, such elucidation ought to be welcome.

3 Why so few judicial reviews?

Bridges and *Johnson* are important cases, but they have, to date, been the only two judicial review challenges against the use of AI and automation to have reached a final judgment by the courts. This is despite the growing use of AI and automation. The question then arises: why have there been so few judicial reviews? The apparent lack of judicial review cases reaching a full judgment, despite the proliferation of AI and automation in government, has a range of possible causes. As noted above, many emerging cases have been settled out of court—which is always a key feature of judicial review litigation.[22] For

[19] Ibid [40].

[20] This feature of judicial rulings that is widely discussed in the literature. The classic articulation is in O M Fiss, 'Against Settlement' (1984) 93 *Yale Law Journal* 1073. See also A Lahav, *In Praise of Litigation* (Oxford University Press, 2017).

[21] *Bridges* (n 8 above) [2].

[22] See generally V Bondy, L Platt, and M Sunkin, *The Value and Effects of Judicial Review: The Nature of Claims, their Outcomes and Consequences* (Public Law Project, 2015).

instance, the Home Office, which is responsible for immigration, has settled judicial reviews against the use AI in immigration decision-making.[23] Settlement can suit the interests of both claimants and government. For claimants, settlement avoids them from potentially being hit by escalating legal costs if they ultimately lose. Note that AI often affects people, such as benefit claimants who are on low incomes. Alternatively, challenges might be brought by campaigning groups, but they have limited resources and high legal fees can have strong deterrent effects. By contrast, for government, settlement enables it to protect the privacy by keeping the details of its AI systems hidden from public scrutiny. A key concern from government is that the use of releasing full details on its AI in contexts, such as immigration and benefit fraud, could enable unscrupulous people to exploit those systems. Settlement is therefore an important factor that explains the small number of legal challenges. But there are some other important factors why claims have likely failed to emerge despite AI and automation being adopted increasingly widely in the public sector.[24] We point here to four other key reasons.

Opacity. AI systems introduce unique challenges due to their inherent opacity, often resulting from the complex algorithms they use. As Jenna Burrell has noted, opacity in these systems generally falls into three categories.[25] *Intentional Opacity:* sometimes, the details of systems are withheld, e.g. to protect intellectual property or prevent manipulation. *Illiterate Opacity:* this form of opacity arises when only those with specific technical skills can interpret how these systems operate, necessitating the involvement of experts in judicial reviews. *Intrinsic Opacity:* often, systems are so complex that even their designers may struggle fully to understand them, complicating efforts to assess them for possible judicial reviews. These layers of opacity make it particularly challenging for claimants to gather sufficient evidence, as they are often unable to access or interpret the underlying algorithms. Without clear visibility into how systems reach decisions, preparing and bringing strong and well-prepared legal cases about administrative law errors becomes significantly more complex and challenging.

Challenges in fact-finding and evidence. AI systems arguably place additional burdens on claimants to gather evidence relating to a decision. Claimants have to identify

[23] See Maxwell and Tomlinson (n 7 above), ch. 4.
[24] For a full exploration of the issues considered in this section, see: J Tomlinson, K Sheridan, and A Harkens, 'Judicial Review Evidence in the Era of the Digital State' [2020] P.L. 740. See also A Sinclair, 'Automated decision-making and government opacity' [2025] P.L. 8.
[25] J Burrell, 'How the machine 'thinks': understanding opacity in machine learning algorithms' (2016) 3(1) *Big Data & Society* 1.

if the decision was in fact made by AI in the first instance. Trying to find evidence of the supposed legal error might take a claimant into highly complex systemic territory, where the issues of opacity raised above will likely occur. The operation of existing rights allowing access to official information, such as the Freedom of Information Act 2000, are also more uncertain in their operation in AI contexts.[26]

Procedural barriers in judicial review. Judicial review processes have developed over time for the purpose of scrutinising the legality of decisions made by human officials. In AI contexts, however, decision-making decisions are based (at least in part) on data processed by algorithms rather than human judgement. This means the existing procedure of judicial review might struggle to operate effectively. For instance, AI systems are often so complex that the need for expert evidence becomes central and greater candour might be required on the part of public authorities—both points made by Lord Sales in his increasingly thorough analysis of these challenges in his extrajudicial writing.[27] Another example can be seen in respect of costs. It has long been observed that the main barrier to accessing judicial review is the costs rules which, with a few exceptions, operate on a 'loser pays' basis.[28] This may disincentivise poor-quality cases from being brought. However, it also risks claimants with well-prepared and meritorious challenges can be dissuaded from pursuing them because of the risks of losing and thereby being hit with large cost orders.[29] This feature makes judicial review always risk-laden to some extent but in AI contexts, where the application of laws in more uncertain, it adds considerable uncertainty to any risk calculation in bringing a claim. Even if challenges are brought, then the risks of cost orders will inform how litigation progresses. For instance, if a public body offers to settle a judicial review challenge out of court—by for instance, withdrawing the use of AI—then procedural incentives will often all but require this to be taken up by the claimant, but it will prevent the court from ruling on the matter and thereby developing the law in that regard. As Lord Sales has noted elsewhere in this collection, 'some aspects of ordinary

---

[26] See e.g. *Public Law Project v Information Commissioner* [2024] UKUT 71 (AAC), which shows the uncertainty and complexity of the operation of these rights in even relatively simple automated systems.

[27] [This will be cross-referenced to the chapter by Lord Sales in this collection]. See also: Lord Sales, 'Algorithms, Artificial Intelligence and the Law' (2020) 25(1) *Judicial Review* 46; Lord Sales, 'Information and Automated Governance' (Institute of Directors, 24 April 2023).

[28] A detailed overview of the recent history of this issue is set out in J Tomlinson and A Pickup, 'Reforming Judicial Review Costs Rules in an Age of Austerity' in Andrew Higgins (ed.), *The Civil Procedure Rules Twenty Years On* (Oxford University Press, 2020),

[29] For a full discussion of this as part of an independent review, see: Lord Justice Jackson, *Review of Civil Litigation Costs: Final Report* (2009); Lord Justice Jackson, *Review of Civil Litigation Costs: Supplemental Report, Fixed Recoverable Costs* (2017). These concerns were further noted in the *Independent Review of Administrative Law* (CP 407, 2021), chaired by Lord Faulks.

judicial review procedure are, if not given flexible application, likely to create near-insuperable impediments in terms of securing access to justice in an ADM context.'[30]

Government's restrictive approach to transparency. Fourth, government bodies have, generally speaking, been reluctant to be transparent about their use of AI and automation. This is perhaps best demonstrated by the weak outcome of the most concerted attempt at transparency so far: the Algorithmic Transparency Recording Standard ('ATRS'). First published in November 2021, the ATRS was established by the Responsible Technology Adoption Unit of the Department for Science Innovation and Technology in conjunction with the Central Digital and Data Office, and provides a framework for accessible, open, and proactive information sharing about the use of algorithmic tools across the public sector. The standard was voluntary, meaning that a public body using AI was not legally required to report that fact under the ATRS. At the time of writing, there were only 23 completed reports recorded on the ATRS Hub notwithstanding that it is known, including as a result of responses to FOIA requests, that many more public bodies are currently using algorithms within their decision-making framework. It was said in the government's response to the consultation on the AI White Paper that a mandatory reporting requirement would be phased in, but there is currently no public sign of this despite being part of the roadmap for 2024.[31]

These factors, and potentially others, are all likely playing a significant role in why we are not seeing many judicial review judgments emerging despite wide-scale adoption of AI across the public sector. As we explore in the next part of this chapter, this lack of clear legal rulings has important implications beyond access to justice for individuals.

4 Two forms of emerging legal risk

One of the main implications of the lack of reasoned judgments about the application of administrative law to AI and automated systems is that the application of the law and legal principles to these novel systems remains highly unclear and uncertain. General administrative law principles and frameworks are stated broadly. They are applied flexibly by the courts in response to the particular contexts and issues in play. On the one hand, this is both a necessity and an advantage. On the other hand, the application of

[30] [This will be crossed referenced to the chapter by Lord Sales in this collection]
[31] Department for Science, Innovation and Technology, *A pro-innovation approach to AI regulation: government response* (CP 1019, February 2024).

administrative law principles carriesan inherent degree of uncertainty. It is important not to overstate this feature of any system of administrative law as a problem. Through both case law and the development of administrative practice and policy, certainty increases in practice over time. However, with the rapid and at-scale adoption of AI and automated in the public sector—including in a range of diverse settings for many different purposes—then new and often fundamental questions about the application of administrative law are arising.[32] The outcome is that, while the application of administrative law in specific settings is often uncertain to some degree, there are presently strong grounds for believing there is an unusually high amount of legal uncertainty surrounding the use of AI in the public sector and therefore also legal risk. This uncertainty and risk can have harmful effects for both claimants and government and in terms of the wider public interest in fair and lawful administration. Moreover, given the growing use of AI and automation by government, such harmful effects will likely increase.

A specific example demonstrates this point. Much of contemporary administrative law centres on the idea of a 'decision' being judicially reviewed, but does the output of an AI system constitute a decision for the purposes of public law? In *Pintarich v Deputy Commissioner of Taxation*,[33] a decision of the Australian Federal Court, the question was whether a letter issued by the Australian Taxation Office amounted to a 'decision,' given that it has been issued solely by a computer-generated system without human review. The majority of the Federal Court concluded that no decision had been made. A decision, applying *Semunigus v The Minister for Immigration & Multicultural Affairs*, a Federal Court decision in a non-automation context, required both 'needs to be both a mental process of reaching a conclusion and an objective manifestation of that conclusion.'[34] Kerr J dissented and emphasised that the legal concept of what constitutes a 'decision' must be flexible to accommodate the increasing use of automated decision-making systems by public authorities. He reasoned that the circumstances of the conduct must be examined in each case in order to assess whether those were within the 'normal practices of the agency' and whether the decision-making was an 'overt act would be understood by the world at large as being a decision.'

[32] AI is having similar effects on settled legal doctrines and interpretations across many areas of law, see *e.g. Thaler v Comptroller-General of Patents, Designs and Trademarks* [2023] UKSC 49.
[33] [2018] FCAFC 79.
[34] [1999] FCA 422.

Another example is the application of the duty to give reasons. There is, it has often been repeated, no general administrative law duty on public authorities to give reasons for their decisions.[35] However, the courts have, across the public sector, identified a very wide range of situations in which decision-making will be found to be unlawful in the absence of reasons, to the point where finding that a duty arises has become the norm.[36] The provision of the decision-maker's actual reasons will, of course, reveal the basis on which a decision was taken and, thus, identify whether the decision-maker committed a public law error when making the decision. As Hickinbottom LJ observed in *R (Help Refugees Ltd) v Secretary of State for the Home Department*:

> The rule of law requires effective access to justice. Therefore, generally, unless (e.g.) excluded by Parliament, there must be a proper opportunity to challenge an administrative decision in the court system. As a consequence, unless rendered impractical by operational requirements, sufficient reasons must be given for an administrative decision to allow a realistic prospect of such a challenge. Where the reasons given do not enable such a challenge, they will be legally inadequate.[37]

An AI system that failed to give reasons, or could not produce proper, adequate or intelligible reasons, would potentially frustrate the ability of an individual to identify and understand the grounds on which they might challenge a decision taken by AI. There is, at present, very little clarity on how this basic principle of administrative law applies in the context of AI but its application could have profound effects on the developing use of AI in the public sector.

These are just two examples, but they highlight the nature of the uncertainty that is reflected in many areas of public law in respect of AI systems.[38] This increased uncertainty creates two forms of legal risk for government. One risk is that public bodies are proliferating the use of systems that may well later be found to have been unlawful. Given the government's intention to use AI to drive up the efficiency and efficacy of public services and the increasingly easy availability of technology, the already wide-spread adoption of automation and AI in the public sector is only likely to continue to increase. It is possible that, in the future, a ruling could be produced from a court, on for instance, the duty to give reasons which could render not only the challenged decision unlawful but,

---

[35] See *e.g. R (JJ Management LLP) v HMRC* [2019] EWHC 2006 (Admin).
[36] *Oakley v South Cambridgeshire District Council* [2017] EWCA Civ 71.
[37] [2018] EWCA Civ 2098 [122].
[38] This is a key theme emerging from the attempt at a comprehensive analysis in McGurk KC and Tomlinson (n 3 above). See also J Cobbe, 'Administrative Law and The Machines of Government: Judicial Review of Automated Public-Sector Decision-Making' (2019) 39 *Legal Studies* 636, 647.

in principle, render unlawful many AI systems operating across thousands of public authorities. This could have dramatic costs for the government and setbacks to AI adoption strategies. It could also have significant negative impacts on, for instance, people who interact with government, such as those in receipt of social security benefits. They rely upon the benefits system operating both fairly and efficiently.

Another risk is that systems of ambiguous legality become so embedded in bureaucratic practices that, by the time meaningful judicial enforcement occurs, the courts will be in an invidious position where they have to address fundamental questions of administrative law with the knowledge that rulings might have wide-ranging impacts across the public sector. Cases like *Johnson* show that government bodies do run these arguments, stressing the possible negative impacts of legal outcomes, in context of automated systems. Further, there is a long history of cases in which government has argued that the courts should not intervene by issuing a remedy because of the negative wider consequences of doing so..[39] Equally, the courts have often been open to arguments that their intervention concerning the operation of an administrative system or scheme would, for instance, cause administrative chaos or have significant economic implications.[40] Understandably, courts can often be nervous about intervening in cases in ways that are likely to have a dramatic and potentially perverse impacts on the public sector.

We are therefore left with an inherent tension. On the one hand, judicial control of the operation of AI is necessary to ensure its lawfulness. This is required in the interests of the individuals adversely affected by AI, but it is also in the public interest and indeed of government itself – so that it has the a reasonably certain legal framework within which to design and operate fair and lawful AI systems. On the other hand, extensive judicial interventions in the operation of AI systems could themselves cause significant administrative costs and upheavals. For instance, it is inherent in the nature of large bureaucratic organisations, such as those that administer social security benefits and immigration, that they require stability in their internal operations to ensure efficiency when processing enormous caseloads. Given the cost-efficiencies promised by AI, government has invested in automation. However, judicial interventions potentially risk

[39] See e.g. *R (on the application of Elan-Cane) v Secretary of State for the Home Department* [2021] UKSC 56; R *(Arias) v Brentwood Superintendent Registrar of Marriages* [1968] 2 QB 956; *Coney v Choyce* [1975] 1 WLR 422; *R (Argyll Group plc) v Monopolies and Mergers Commission* [1986] 1 WLR 763.
[40] See e.g. *R (Hurley and Moore) v Secretary of State for Business Innovation & Skills* [2012] EWHC 201 (Admin).

upsetting this. Alternatively, the courts could retreat into non-judicial intervention thereby limiting their impact. This also seems inadequate in terms of ensuring governmental compliance with the law and legal principles. Accordingly, are there potentially better ways of managing the legal risks we have identified?

5 Managing emergent risks

There are potentially numerous ways in which our preceding observations about the heightened nature of legal risk in the public sector resulting from AI could be interpreted, especially in terms of what an appropriate response might be. The usual starting place here is to suggest new statutory regulation of AI is now required. This is an important debate which is now developing quickly in the UK.[41] However, a new set of substantive legal regulations or principles governing the use of AI in the public sector will not necessarily solve the crux of the issues considered above. Given this, we close by pointing to three ways the risks we identify here might be effectively managed.

First, government could positively facilitate the ventilation of judicial review claims pertaining to AI. While government taking action to allow for easier ventilation of grievances against AI systems might seem like it is inviting trouble for itself, the long-term systemic benefit of the ventilation of these claims seems particularly important in the development of established legal norms in this context. As we noted above, this is a long-observed virtue of litigation generally, but might be particularly important at this moment of the adoption of AI. As such, thinking about how procedural reforms or changes in judicial approach to procedure might support might be a helpful starting point. This might include promoting different approaches to the disclosure of evidence and the duty of candour in cases involving AI. It might also involve creating a protective costs regime in these cases, similar to that in operation under the Aarhus Convention for environmental judicial review claims.

[41] For instance, at the time of writing, there is a private members Bill in Parliament, introduced by Lord Clement-Jones, entitled 'Public Authority Algorithmic and Automated Decision-Making Systems Bill,' with its stated aim being to 'regulate the use of automated and algorithmic tools in decision-making processes in the public sector; to require public authorities to complete an impact assessment of automated and algorithmic decision-making systems; to ensure the adoption of transparency standards for such systems; and for connected purposes.' As a private members Bill, this Bill is unlikely to be enacted without government support, but the Bill will inform the debate concerning the regulation of AI. The government has also signalled it is more willing to consider regulation and recently became a signatory to the Council of Europe's Framework Convention on Artificial Intelligence.

The use of remedial flexibility in judicial review and particularly prospective-only quashing remedies. Judicial review remedies can generally be deployed with flexibility. This can assist a court presented with a situation where illegality has been embedded in administrative operations, and there is a concern of significant disruption resulting from a judgment. Two remedies are particularly useful in this respect. First, declarations enable the court to clarify and declare what the legal position is, with strong the expectation that the government will abide with the declaration. Declarations provide certainty to the parties as to what the court's view of the legal position is. Second, the courts can now issue retrospective quashing orders. Under the Judicial Review and Courts Act 2022, the courts have a discretionary power to suspend and to limit the retrospective effect of a quashing order. The Act provides, at s.29A(1)(a), for suspended quashing orders, whereby a quashing order will not 'take effect until a date specified in the order.' The consequence is that 'the impugned act is...upheld until the quashing takes effect'[42] and for the period it is upheld 'it is to be treated for all purposes as if its validity and force were, and always had been, unimpaired by the relevant defect.'[43] Once the quashing order takes effect, the impugned act is treated as void from the beginning.[44] At the time of writing, there has not been much case law on the use of the provision. Chamberlain J's judgement in *R (Every Child Protected Against Trafficking UK) v Kent CC and Secretary of State for the Home Department* provides early guidance.[45] He noted that:

> Until the enactment of the Judicial Review and Courts Act 2022... it was unorthodox, and arguably impermissible, for a court to grant an order permitting a public body to continue, even for a short time, to do something which it had held to be unlawful.[46]

He further observes that the new powers provide:

> [A] remedial flexibility that was previously unavailable. [It] enable[s] the court to deal with situations where one or more public authorities are engaged in conduct that is unlawful, but real harm would be caused if that conduct had to stop immediately. The court can set a date by which

[42] Senior Courts Act 1981, s.29A(3).
[43] Ibid, s.29A(5). Section 29A(8) sets out considerations the court 'must' have regard to in deciding whether to suspend a quashing order: the nature and circumstances of the relevant defect; any detriment to good administration that would result from exercising or failing to exercise the power; the interests or expectations of persons who would benefit from the quashing of the impugned act; the interests or expectations of persons who have relied on the impugned act; so far as appears to the court to be relevant, any action taken or proposed to be taken, or undertaking given, by a person with responsibility in connection with the impugned act; and any other matter that appears to the court to be relevant.
[44] Ibid, s.29A(6).
[45] [2023] EWHC 2199 (Admin).
[46] Ibid [4].

> the unlawful conduct must be brought to an end while providing, subject to conditions under s.29A(2), that the conduct may continue in the meantime.[47]

In the context of that specific case, the court concluded that 'the unlawful practice must stop. But all parties agreed that it would counterproductive if it had to stop immediately.'[48] Chamberlain J went on to explain that 'during the period while the orders were suspended there would be stringent conditions to ensure that everything was being done that could be done to minimise the extent of the continued unlawful conduct.'[49] The judge also adopted an approach which put the Court in the position of essentially monitoring the ongoing position. Though the 'normal position in judicial review is that the court determines the issues before it and then decides what relief to give on one occasion'[50] and that 'the public authority can in general be trusted to comply with the judgment,'[51] where the 'power in s.29A(1)(a) to suspend a quashing order is exercised, it may be necessary to hold a further hearing to check that the conditions for suspension have been complied with and to determine whether the suspension should be extended.'[52] These provisions might be particularly attractive to courts in the context of AI where there is concern that a ruling about the legality of a decision resulting from an AI or automated system might invalidate a wide class of decisions in a way that might cause administrative problems, or time might be required to transition to a new system or approach to decision-making.

Engaging and improving the use of non-court redress mechanisms in AI matters. We would also suggest that, given the complexities of judicial review, modes of scrutiny and redress outside of courts might have an important role to play in responding to the risks we identified. Judicial review is only one means of challenging scrutinising administration, whether it involves AI or otherwise. There is also non-judicial scrutiny by a wider range of diverse bodies, including Parliament and its select committees, public inquiries, ombudsmen, auditors, inspectorates and regulators. Indeed, the majority of the scrutiny of government is undertaken by non-judicial mechanisms. The nature of such scrutiny varies depending on the particular method involved. For instance, ombudsmen investigate complaints of maladministration whereas select committees examine the expenditure, administration, and policy of government. Non-judicial scrutiny does not

[47] Ibid [6].
[48] Ibid [7].
[49] Ibid [10].
[50] Ibid [11].
[51] Ibid [12].
[52] Ibid [14].

result in legally binding ruling, although it has advantages when compared with the courts. Non-judicial scrutiny is not limited to the relatively narrow grounds of judicial review. It encompasses concepts such as value for money, maladministration, and the effectiveness of government. Non-judicial scrutiny does not necessarily need to instigated by a claimant; it often acts on its own-initiative. There is no costs regime. Bodies such as select committees and ombudsmen are investigative and they bear the costs of scrutiny. Non-judicial scrutiny can be quicker, more responsive and more immediate than the courts. Some forms of scrutiny, such as ombudsmen, auditors and inspectorates, have direct access to government information. Select committees can directly scrutinise ministers and officials. Non-judicial scrutiny has also promoted transparency in the use of AI by revealing details about its use that would otherwise have remained hidden away. Details about the Home Office's 'streaming tool' algorithm were revealed by the immigration inspectorate and prompted a judicial review challenge, which was then settled out of court.[53] The National Audit Office has shed light on the use of machine learning by the Department for Work and Pensions (DWP).[54] In turn, the Public Accounts Committee concluded that the DWP had not done enough to understand the impact of machine learning on people and ensure public confidence that it would not result in unfair treatment.[55] Of course, the issue is not whether non-judicial scrutiny is better than that of the courts. Instead, the issue is how all types of scrutiny complement each other in the overall endeavour of making government accountable and how they will face up to the challenges presented by the use of AI by government. Given the issues emerging with the effectiveness of the judicial review procedure, focusing on how redress against AI in the public sector can be strengthened outside of the courts is imperative.

6 Conclusion

This chapter has considered judicial review challenges against the growing use of AI and automation by government in the UK. We have examined the two judicial review challenges that there reached final judgment. We have also identified the reasons why

[53] Independent Chief Inspector of Borders and Immigration, *An Inspection of Entry Clearance Processing Operations in Croydon and Istanbul* (2017); *An Inspection of Visit Visa Operations* (2023).
[54] National Audit Office, *Report on Accounts: Department for Work & Pensions* (2023).
[55] House of Commons Public Accounts Committee, *The Department for Work & Pensions Annual Report and Accounts 2022-23* (HC 290 2023-24).

there have only been two challenges - so far. Our main argument is that, in the absence of more legal challenges and the growing use of AI and automation, there is a corresponding growth of legal risk and uncertainty – for the people affected, the public as a whole and also government. The longer this situation persists, then the more legal risk accumulates and the greater the resulting uncertainty. AI and automation is now a reality. It also raises unique and unprecedented legal problems and challenges for claimants, government the courts. In this respect, we have considered ways of managing the risks involved, by way of government facilitating judicial review challenges, remedial flexibility and the use of non-judicial review redress.

Our conclusions here echo views which have emerged in the writing of Lord Sales on public law and AI. Lord Sales has suggested that non-judicial redress—in the form of an algorithm commission—might be an appropriate way to support the position of the courts.[56] He has also noted that more expansive approaches to evidence might be required by the courts when judicially reviewing AI, including taking more flexible approaches to expert evidence and the duty of candour.[57] In his contribution to this collection, he noted that '[c]ourts may have to be creative in fashioning forms of order which, while doing justice between the parties, do not have the effect that an expensive and generally effective AI system has to be scrapped and allow a public authority time to take remedial action in relation to it.'[58] In this sense, we view Lord Sales as having made a sophisticated and credible contribution to a roadmap for the pragmatic evolution of administrative law—one which ought to be taken seriously while the wider debate about the possibility of new forms of statutory regulation continues.

[56] Lord Sales, 'Algorithms, Artificial Intelligence and the Law' (2020) 25(1) *Judicial Review* 46
[57] [This will be crossed referenced to the chapter by Lord Sales in this collection]
[58] [This will be crossed referenced to the chapter by Lord Sales in this collection]