



Automated Public Decision Making and the Need for Regulation

ANDREW MURRAY 

RESEARCH



ABSTRACT

Our relationship with the state is unique. Only the state can compel us to use its services and be bound by its decisions. Today our relationship with the state is being complexified by the adoption, by public decision-makers, of automated decision-making technologies which can be subject to bias and error, and which are often intransparent. We need to re-orientate our relationship with the state and public decision-makers to ensure such tools meet standards of fairness, justice and transparency. This paper suggests we need a policy on the use of automated decision making in public services. This policy must be robust and must provide sufficient regulatory oversight to ensure that sufficient safeguards are adopted to ensure that the individual is respected and protected in their dealings with public decision-makers who deploy automated decision-making tools.

CORRESPONDING AUTHOR:

Andrew Murray

Professor of Law, LSE Law School, London School of Economics, UK

a.murray@lse.ac.uk

KEYWORDS:

Automated decision making; Artificial Intelligence; Governance; Regulation; Public Sector; Public Services; Fraud Detection Algorithm; Sham Marriage Triage Tool

TO CITE THIS ARTICLE:

Murray A. Automated Public Decision Making and the Need for Regulation. *LSE Public Policy Review*. 2024; 3(3): 3, pp. 1–10. DOI: <https://doi.org/10.31389/lsepr.110>

PUBLIC DECISION MAKING, THE PUBLIC, AND THE RISK OF HARM

The state occupies a unique position in our lives. We cannot choose not to engage with it. The largest of corporations may have greater day-to-day impact on our lives but we can, in theory, choose not to have Apple products or use Google's services, not to buy from Amazon or Tesco, or not to travel with British Airways or LNER. We cannot, though, choose not to pay our taxes, to drive or travel overseas without official documentation, we cannot settle in another state without visas or the right to work, and we cannot have our marriage or civil partnership recorded without the authorisation of the state.

Where we appear to have a choice not to engage with the state, that is normally a privilege restricted to a fortunate few. We might say we have the choice to have private healthcare, but emergency care in the UK is always provided by NHS hospitals; we might prefer private landlords over local authority housing, but in times of homelessness it is local authorities who have responsibility for emergency and temporary housing; we might benefit from driving our private cars but we do so on roads built and maintained by the state (either National Highways or local authority).

The unique position public authorities play in our lives has long been recognised in law. State and public authorities (including executive agencies) face specific legal checks and balances. The most obvious is the requirement that public authorities submit to judicial review. The power of the High Court to review the decisions of public bodies is 'inherent', meaning that it is based on the common law and is not awarded to the courts by Parliament (1). While recent focus on judicial review has been on the constitutional implications of judges ruling on the relationship between the executive and the legislature in cases such as *R (on the application of Miller and another) v Secretary of State for Exiting the European Union* [2017] UKSC 5 and *R (on the application of AAA (Syria) and others) v Secretary of State for the Home Department* [2023] UKSC 42, the quotidian work of the courts is in reviewing decisions of local authorities, government departments and other public officers.

The procedural rules for judicial review show that the heart of the judicial review process is not just the checks and balances of the separation of powers and the rule of law (often cited as the jurisprudential foundations of judicial review and central to the constitutional settlement between the courts and Parliament) (2), but also a reflection of the role of common law in protecting the weak and the vulnerable from injustice (3). By part 54 of the Civil Procedure Rules 'a "claim for judicial review" means a claim to review the lawfulness of – (i) an enactment; or (ii) a decision, action or failure to act in relation to the exercise of a public function' (CPR 54.1(2)(a)). It is only the first of these, a claim to review the lawfulness of an enactment, which gets to the heart of the relationship between Parliament and the courts and thereby to the separation of powers and rule of law. The second, a claim to review the lawfulness of a decision, action or failure to act in relation to the exercise of a public function, is designed to protect the individual from the outcomes of bias, error or failure. It is there to protect us from capriciousness so that we do not end up in a similar situation to Josef K in *The Trial*. It is this second strand, the protection of individuals in relation to the exercise of a public function, that is the basis of most judicial review cases. Data from the Office for National Statistics showed that in 2023 most applications for judicial review were made against the Home Office (many of which were asylum and immigration claims), with local authorities being the source of the second-highest number of applications (4). These claims would overwhelmingly be made under CPR 54.1(2)(a)(ii). Reviewing decisions of public bodies and officers is the bread and butter of judicial review.

PUBLIC AUTOMATED DECISION-MAKING

The exercise of the public function is costly and often inefficient. The Office for Budget Responsibility predict in 2023–24 that the overall cost of public services, state pensions and debt interest is £1,189 billion, the equivalent to around £42,000 per household or 46.2% of national income. Much of that reflects the costs of the welfare state and other public services such as the NHS and defence, which we might call fixed or sunk, but we can see why with figures such as this successive governments have argued for 'austerity' or 'fiscal restraint'. To parse out how much of this spend represents the administration of the public function is complex, but one

statistic offered by the National Audit Office tells us that total civil service salary costs in 2022 were £16.6bn. Of course, staff costs do not equate to inefficient spending, but these figures show the scale of just one element of the delivery of public services. A reasonable suggestion of the cost of inefficiency in the delivery of the public function may be the announcement in March 2022 by then Chancellor Rishi Sunak of a ‘new crackdown on cross-Whitehall waste to drive efficiency, effectiveness, and economy across government’. This efficiency drive, it was hoped, would save ‘£5.5 billion worth of waste’. These data only reflect the likely inefficiency of central government. Further inefficiencies may be found in local government (5) and in executive agencies (6).

One way to reduce cost and inefficiency is to replace human employees, ordinarily the most expensive recurring cost in any organisation, with an algorithmic system that can replicate the tasks of these employees such as decision-making in public services. The introduction of an algorithm to replace a human worker offers multiple benefits. Firstly, ongoing costs are reduced. While there may be ongoing licensing, servicing, or service provision costs, these will be much lower than the recurring salary and related costs of employees replaced (7). Secondly, automated systems can work continually without the need for breaks or limits on working hours, and performance of automated systems does not drop off during the day. Algorithms do not get tired or distracted. Thirdly, automated decision-making does not make subjective errors or mistakes. In theory, at least, every decision made by an automated system will follow the same logic pattern and lead to the same outcome. It will not be swayed by emotion, attachment, or personal experience. While it is true that automated systems may be imbued with systemic bias (8), in theory at least, such systemic bias can be identified and dealt with more easily than a multitude of individual, subjective biases or errors of individual employees (9). Fourthly, automated decision-making tools do not require expensive estate costs, they do not need city centre office space, parking spaces or related overheads. In 2022, the Cabinet Office reported that in 12 years, the size of the central general-purpose estate had been reduced by 30%, which had ‘cut annual running costs by £1.6 billion’. This suggests the overall cost of the Central Government estate in 2022 remained around £3.7 billion (10). Increased automation could help reduce that cost by further reducing the size of the physical estate, though such estate savings must be offset against the costs of operating the digital infrastructure, such as paying for cloud provision. Finally, automated systems are less likely to give rise to actionable claims, and when they do, these may be insurable or may be underwritten by the developer. Nor will they make claims for personal injury, unfair dismissal or other employment claims, and although there may be contractual exit or termination fees, they do not require redundancy payments.

With these possible efficiency savings in mind, it is not a surprise that the UK government has sought to deploy algorithmic decision-making systems in public decision-making. In doing so they made some bold pronouncements. For example, the Department of Work and Pensions suggested that the use of four algorithmic models to assess and flag suspected fraudulent claims around Universal Credit could ‘generate savings of around £1.6bn by 2030–31’,¹ while in a speech given in February 2024, then Deputy Prime Minister Oliver Dowden suggested that AI algorithms could help to move asylum claimants out of hotels more efficiently, saving millions of pounds, and could be used to flag fraud and error in pharmacies, which costs the taxpayers £1 billion annually. He reported that figures from the Institute for Public Policy Research suggested AI and other algorithmic tools had ‘the potential to save £24 billion each year’.²

To capture these savings, the UK public sector has rolled out several algorithmic decision-making systems. In addition to the Universal Credit fraud detection algorithm mentioned, the Public Law Project on 30 August 2024 listed 55 applications in its Tracking Automated Government (TAG) Register (11). One challenge highlighted by the Public Law Project is the so-called ‘black box problem’, the problem that while observers can witness the inputs and outputs of complex, non-linear processes, they cannot observe or study the inner workings (12). They note that of the 55 applications registered, 38 were classified as being of low

1 <https://www.publictechnology.net/2023/07/12/society-and-welfare/dwp-commits-70m-to-algorithms-and-analytics-to-tackle-benefit-fraud/>.

2 <https://www.gov.uk/government/speeches/deputy-prime-minister-speech-on-ai-for-public-good>.

transparency, 16 as being of medium transparency, and only 1 as being highly transparent.³ Further, the Public Law Project declared that 83.6% of these tools were only uncovered or fully understood through the submission of Freedom of Information requests, while only 49.1% of the tools have publicly available government assessments on their impact on the protected characteristics of individuals, normally a legal requirement under s.149 of the Equality Act 2010. The data gathered by the Public Law Project shows heavy reliance on algorithmic tools by just two government departments – the Home Office and the Department for Work and Pensions. Tools listed include the sham marriage triage tool (which we shall return to below) used to decide whether a couple who has given notice to be married should be investigated to identify or rule out sham marriage activity; the border risk and targeting capability tool, used to detect and combat fraud, crime and illegal migration at UK borders; the visa streaming algorithm, designed to determine which grade of decision-maker should handle a given visa application, and the level of scrutiny to be applied; and the self-employment fraud and error detection tool used to identify potentially fraudulent or erroneous benefits claims involving self-employment for review.

While the government is seeking to capture these efficiency savings by speeding up the roll-out of algorithmic and automated decision-making (13), mistakes, errors, biases, and examples of injustice emerge. In January 2024, Neil Couling, a Director General within the Department for Work and Pensions and responsible for Universal Credit, revealed to the Commons Work and Pensions Committee that there were problems with the Universal Credit algorithmic models deployed by the Department. While being questioned by Peter Grant MP, he disclosed that the Department had stopped routinely suspending benefit claims flagged by its models. When asked by Mr. Grant what happened when the algorithm flagged a claim, he revealed that ‘we actually changed our approach in the light of feedback from claimants and elected representatives. We used to suspend all the cases, and now we don’t suspend. We go in and do the checking as quickly as we can’. He was then pressed by Mr. Grant on bias and errors ‘If the machine or algorithm were to have any unintended inclination towards bias—for example, gender or ethnic origin—that would put you in breach of the law, wouldn’t it?’ In reply, Mr. Couling revealed perhaps more than he intended, and certainly more than the media reported at the time. He first admitted that bias exists in the system, noting that ‘the systems do have biases in; the issue is whether they are biases that are not allowed in the law’. He then went on to both clarify that some biases are intended ‘because you have to bias to catch fraudsters’ before confirming that the Department had systems to check for unintended bias (14). This is extremely interesting as it is an admission of deliberate bias in the algorithm. Why might the Department feel such bias is necessary? It is not clear, and the Department is unlikely to say as revealing exactly what the bias is would assist fraudulent claimants; but, at least one organisation, campaign group Foxglove, is extremely sceptical, believing this bias represents likely discrimination against disabled people (15).

Another tool heavily critiqued by human rights groups and news outlets is the Home Office’s sham marriage triage tool. It launched in 2015 as part of the new investigation process for sham marriages introduced by Part 4 of the Immigration Act 2014. Originally called ‘the dial’, it was a Heath Robinson system⁴ operated on Excel spreadsheets by Marriage Referral Assessment Unit staff. It used data points such as records of immigration offences, absconding or individuals otherwise in breach of leave to remain; illegal entrants; and individuals with a criminal conviction or evidence of links to criminality to profile whether there was evidence of potential harm posed to the UK by categories of individuals. It was recorded by David Bolt, Independent Chief Inspector of Borders and Immigration, in his 2016 report that ‘the overall process involved multiple handovers between different IT systems, with each human intervention providing an opportunity for inputting error’ (16). As a result, the Home Office sought a more robust system, all the while ensuring that nationality (a protected characteristic) was not a relevant factor in their calculations. By 2019, that system was in place. An Equality Impact Assessment

³ If the UK government were bound by the new EU Artificial Intelligence Act, then as public sector algorithms are classified as ‘high risk systems’ under Art.6(2) and Annex III of the Act, the government would be required under Art.13 of the Act to ensure that the system ‘be designed and developed in such a way as to ensure that their operation is sufficiently transparent’. The Act does not apply in the UK post Brexit.

⁴ William Heath Robinson was an English cartoonist, illustrator and artist, best known for his drawings of elaborate machines to achieve simple objectives. A ‘Heath Robinson system’ may be seen to be an unnecessarily complex system to achieve simple aims.

(EIA) released under a Freedom of Information Act request revealed that an automated triage tool came into operation in April 2019. As the EIA reported, ‘this automated triage process, developed in accordance with Analytical Quality Assessment (AQA), [utilises] best practice to leverage historic outcomes and associated data on sham marriages in order to identify referrals of couples where there is an indication of potential sham activity’ (17). The EIA noted that the tool was introduced, in part, as an efficiency saving resource, noting that the number of Immigration Intelligence staff required for consideration of referrals was reduced from 25 to 4 and that it had reduced the time of an initial assessment from 14 days to 3 or 4 days. In compliance with the Public Sector Equality Duty under s.149 of the Equality Act 2010, eight protected characteristics were examined (although, unhelpfully, the document is redacted and does not give the criteria used in triage). We do know that marriage notifications, which are used to populate data, include information on three protected characteristics; nationality, age, and sex, but not on the other five; maternity status, gender reassignment status, disability, sexual orientation, or religion/belief. The EIA was signed off as being in compliance with s.149 in November 2020, but doubts remain, and, as we shall see, the legality of the sham marriage triage tool has been subject to a legal challenge by the Public Law Project.

Both the DWP’s fraud algorithm and the Home Office sham marriage triage tool are recorded on the Public Law Project’s TAG Register as being of low transparency. In relation to the fraud algorithm, they record that, ‘The DWP has conducted a “fairness analysis” but noted some bias regarding age and noted the difference between the size of the training dataset and the population of claimants it applies to’. They further note that, ‘the National Audit Office noted an “inherent risk of bias” in the use of such machine learning tools. The DWP stated that its capacity to test for bias is limited due to claimants not providing optional demographic information and DWP segregating personal data for security reasons’ (11). In relation to the sham marriage triage tool, they recorded that the algorithm currently uses six ‘risk factors’ - there were eight in 2020. In 2020, we were aware of two risk factors: shared travel events and age difference between the couple. It is not known whether these still apply - the Home Office has refused to disclose the criteria, and that it ‘flags Albanian, Bulgarian, Romanian, Greek couples at a higher rate. Under the Equality Act 2010, race is a protected characteristic and refers to a group of people defined by their race, colour, and nationality (including citizenship) ethnic or national origins’ (11).

The risk of bias, error or simple administrative mistake in these systems was clear, and recent activity by civil society and media organisations has uncovered what looks to be further evidence of errors and bias. You will remember that in January 2024, a Director General within the Department for Work and Pensions revealed to the Commons Work and Pensions Committee that the department had changed its approach towards cases flagged by the algorithm. Whereas originally they had suspended all flagged claims before investigating them, they now investigated flagged cases without suspension of claims. This was due to pressure from, among others, the Public Law Project, who made a number of Freedom of Information Act requests; The Chartered Institute for IT, which published an article in October 2022 arguing that there was a high risk of discriminatory decisions and that incorrect decisions would challenge public trust in AI;⁵ and adverse warnings in summer 2023 in mainstream media including the BBC⁶ and the Guardian.⁷ To date, there is no indication that the DWP plans to abandon the algorithm altogether; all they have publicly stated is the change in policy announced quietly in committee. One concern is that this mirrors the Post Office Horizon scandal, where initial reports in technical journals and blogs and pressure from civil liberty groups eventually lead to mainstream media coverage before the full scale was revealed. In committee Conservative MP Sir Desmond Swayne asked Peter Schofield, Permanent Secretary of the DWP whether there were ‘shades of Horizon?’ In reply Mr. Schofield replied, perhaps with less conviction than we might like, ‘I really hope not’ (13).

Similarly, the sham marriage triage tool has come under further scrutiny. In 2023, the Guardian reported the ‘algorithm used by the Home Office [] has been disproportionately selecting people

5 <https://www.bcs.org/articles-opinion-and-research/dwp-machine-algorithm-and-universal-credit/>.

6 <https://www.bbc.co.uk/news/uk-politics-66133665>.

7 <https://www.theguardian.com/technology/2023/oct/23/uk-risks-scandal-over-bias-in-ai-tools-in-use-across-public-sector>.

of certain nationalities⁸ while in 2021 the Public Law Project reported that ‘the Home Office’s documents show that some nationalities have their marriages rated “Red” at a much higher rate than others’ and that ‘the Home Office has – so far – refused to disclose all of the “risk factors” used by the algorithm to rate a case’.⁹ Following a decision of the First Tier Tribunal (Information Rights) that there was the potential for bias and an apparent discriminatory effect in the algorithm (18), the Public Law Project launched a judicial review in February 2023. Unfortunately, that challenge seems not to have progressed.

THE NEED FOR A POLICY ON THE USE OF AUTOMATED DECISION MAKING IN PUBLIC SERVICES

Thus far, four things have been established. First, our relationship with the state is different from our relationship with private actors. We are required, sometimes by law, to engage with the state in a way we cannot be compelled to do with private actors. Second, the unique power of the state has been recognised in law through the principle of judicial review, which is designed to protect the individual from the overwhelming power of the state. Third, public decision-makers are looking to algorithms and other forms of automated decision-making to improve efficiency and to reduce costs. Fourth, these systems are used by public decision-makers in ways which are untransparent, subject to bias and error, controversial, and possibly illegal.

To this we might add that our near neighbours in Europe have recognised the particular risks in the deployment of automated decision-making in public services and have taken steps to regulate them. The recently passed EU Artificial Intelligence Act (AI Act) makes specific provisions for the use of AI in the public sector. The Act classifies AI systems into one of four classifications: unacceptable risk (prohibited), high-risk, limited risk and minimal risk. AI systems intended to be used by public authorities or on behalf of public authorities to evaluate the eligibility of natural persons for essential public assistance benefits and services, including healthcare services, as well as to grant, reduce, revoke, or reclaim such benefits and services; as well as AI systems intended to be used by or on behalf of competent public authorities or by Union institutions, bodies, offices or agencies to assist competent public authorities for the examination of applications for asylum, visa or residence permits and for associated complaints with regard to the eligibility of the natural persons applying for a status, including related assessments of the reliability of evidence are listed in Annex III of the Act as examples of high-risk systems. This means that, in terms of the AI Act, both the Universal Credit fraud detection algorithm and the Home Office sham marriage triage tool are high-risk AI systems. As a result, under the AI Act they would carry a significant legal-regulatory burden, including from Article 13 a transparency requirement and from Article 15 a requirement that the systems be designed and developed in such a way that they achieve an appropriate level of accuracy, robustness and cybersecurity.

Perhaps most tellingly, if the AI Act applied to these tools they would be required to have a risk management system, which under Article 9(2)(a) of the Act can identify and analyse the known and reasonably foreseeable risks that the high-risk AI system can pose to health, safety or fundamental rights when the high-risk AI system is used in accordance with its intended purpose. According to Article 9(2)(b), this risk management system would be one that can also estimate and evaluate the concerns that may emerge when the high-risk AI system is used in accordance with its intended purpose, and under conditions of reasonably foreseeable misuse. In short, any developer or supplier of such an AI system used by the UK government would be legally required under the Act to ensure that any automated decision-making system deployed by the public sector had adequate risk management protocols in place to protect against risks to fundamental rights (rights recognised under the Human Rights Act 1998). Of course, none of this applies post Brexit.

The UK looks immediately out of step, therefore, when compared to its neighbours. But even if one were to ignore the evidence of positive action elsewhere there are other compelling

8 <https://www.theguardian.com/technology/2023/oct/23/uk-officials-use-ai-to-decide-on-issues-from-benefits-to-marriage-licences>.

9 <https://publiclawproject.org.uk/latest/sham-marriages-and-algorithmic-decision-making-in-the-home-office/>.

reasons for the UK government to develop a robust policy on the use of automated decision making in UK public services.

Following years of political action, the state has become diminished. Small-state conservatism has flourished in the United Kingdom. Thatcherism has driven this diminution of the state. It has seen the privatisation of public utilities and introduced fiscal deregulation and the increase of private-public partnerships and outsourcing of public services to private providers. The state is arguably underdelivering (unless one takes a libertarian stance) while the populous pay high levels of taxation and are arguably suffering historically high limitations of individual liberty. It is an oxymoron that we both need to rely on the state, and we cannot rely on the state. The need to rely on the state is being driven by an aging and growing population; we have high levels of relative and absolute poverty and rising levels of food insecurity. More people report a disability or chronic illness and perhaps as a result we have experienced growing waiting lists for NHS care. People in Britain report feeling poorer, and statistics show that reports of stress and anxiety are at an almost all-time high with 1 in 14 UK adults (7%) reporting they feel stressed every single day.¹⁰

Against the backdrop of this increasing demand on state services, the increased deployment of algorithms and automated decision-making by the state is driving user frustration. In 2022, Tanja Sophie Gesk and Michael Leyer examined the response of the public to the deployment of AI in public services (19). Although they found a generally high level of acceptance of AI by the public, their data showed some interesting results. For example, they found that with specific services (such as decision-making), there was a significant difference between public acceptance of AI vs human services with acceptance for human decision making reported at 94% as against only 72.1% for AI tools. They further found that ‘risk and trust are important aspects that influence acceptance of AI systems’ and that, for AI systems, part of that loss of trust may be explained by the finding that ‘citizens feel that they do not have much of a choice regarding the introduction of AI [and that] respondents do not see AI as being able to handle exceptional concerns’. This led them to observe that ‘citizens see specific public services as more personally relevant, and they put more value on individual requests and requirements’.

It is to be hoped that with a recent change of government, the new administration recognises this and other evidence that the UK public sector is failing in its responsibilities to its citizens due to the intransparent, perhaps unjust nature of algorithmic decision-making tools (20). The public lacks trust for these tools and values individualised (subjective) decision-making over detached objective algorithmic decision-making. To regain public trust, it is time for a robust and transparent policy on the adoption and deployment of automated decision making tools.

DESIGNING A POLICY ON THE USE OF AUTOMATED DECISION MAKING IN PUBLIC SERVICES

To begin, the government must undertake an urgent a review of the development and deployment of algorithmic and other automated decision-making tools (including AI) by authorities and agencies exercising all forms of public decision-making authority (including executive agencies); they must produce a policy for their use; and they must implement systems of review and appeal against actions and decisions made by or through such tools.

Such a review should be focussed on five key areas of risk of harm:

1. Intransparency and lack of democratic accountability
2. Bias, error and discrimination
3. Procedural unfairness and injustice
4. Algorithmic confirmation bias (computer says no)
5. Automation

Let me briefly frame these five risks and why we need to produce effective governance strategies for them.

10 <https://championhealth.co.uk/insights/stress-statistics/>.

Intransparency and lack of democratic accountability are hopefully self-explanatory. Algorithms make decisions that are often opaque and unclear. We need to ensure some basic principles are met when public decision-makers deploy automated decision-making tools. A simple solution would be to adopt a transparency impact assessment (TIA) to be completed before an automated decision-making tool is deployed. This would require the organisation deploying the tool to answer questions about the explainability and accountability of an algorithmic tool as against a human decision-maker. To ensure they remain relevant and up-to-date, TIA's would require regular reviews and updates. Secondly, it would be simple to require the designation of a decisional transparency officer (DTO) in relevant organisations. This would be a senior member of staff responsible for ensuring compliance and responding to requests from the public. Thirdly, a system of reporting and auditing should be developed. Public sector decision-makers should be required to publish annual transparency reports, and these should be made subject to external audit. Finally, there must be an accountability and appeals process (see below).

To protect against bias, error and discrimination there must be a further decisional impact assessment (DIA) that examines and identifies risks of bias or error. Public sector decision-makers must allow researchers access to training data and any decisional logic (this may be done confidentially where sensitive information is processed). There should be specific protection for whistle-blowers and a specific timetable for responses to any claims or challenges for bias, error or discrimination. Decision-makers should be subject to an external regulator who has power to inspect and investigate any evidence or claims of bias, error or discrimination.

To meet procedural unfairness and injustice we need a new process and rights for citizens. As with judicial review, we disempowered individuals need a shield against the application of public automated decision-making; this shield, however, needs to be more responsive and targeted than judicial review. I will say more on this below.

The challenge of algorithmic confirmation bias is that even when a human is in the loop, they will often defer to algorithmic recommender systems on the basis that the computer cannot be wrong (see the Post Office Horizon Inquiry). We need to ensure that the human remains empowered. This is likely to start with algorithmic literacy campaigns for public decision-makers, but in terms of governance means that the law ensures human decision-makers are encouraged to challenge algorithmic recommendations, for example, by requiring them to give reasons for their decisions, protecting whistle-blowers (see above) and periodically auditing or reviewing the actions of decision-makers.

Finally, by automation I mean the automation of enforcement, or the automatic action of decisions taken by algorithms. These might be to automatically stop payment of benefits, to remove people from a housing register, to withdraw a visa or similar. We need governance rules to ensure that automation is only permissible where certain levels of accuracy are established by external audit, where the automated decision can be appealed and reviewed by a human in a defined (short) period and that there is in general a presumption against automation.

In addition to these specific safeguards, we need a new governance framework that reflects principles of natural justice, fairness and due process. Unfortunately, we cannot rely on the common law and the current judicial review framework to protect citizens. The courts are already overwhelmed, with around 250 new judicial review claims being raised each month (around 8 per day), with cases often taking more than a year to be resolved (4). To add thousands of potential claims driven by automated decision-making would simply add to the workload of the court and cause inevitable delays. This is why we need a new dedicated governance and regulatory framework for the use of automated decision making in public services.

The new governance framework, like judicial review, should apply to anyone fulfilling a public decision-making function, whether they are employed by a central or local government agency or authority or if they are providing a public function under contract or service agreement. Equally, rights granted under the framework must be enforceable by all who are impacted by the decision – whether they are UK citizens or not, whether they be adults or minors (who may be represented by a parent or guardian), whether they are claiming because of their membership of a particular class, group or characteristic. There should be a regulator or authority to whom people can complain, and that regulator should have powers of investigation, search and entry. They should authorise the forms of TIAs and DIAs and should keep a register of DTOs. They


should have the power to require public decision-makers to immediately suspend the use of any algorithmic decision-making tool they have concerns about and the power to require any decision made by an algorithmic decision-making tool to be retaken *de novo* by a human decision-maker. Finally, to support them in their work there should be a dedicated Automated Decisions Tribunal (as part of the First-Tier tribunals system) which allows a route into the courts.

With the current pressure on public finances, it is no surprise that some groups and organisations, such as the Tony Blair Institute (TBI), are pressing the government to move quickly to adopt AI tools, claiming that ‘adopting AI in the public sector is a question of strategic prioritisation that supersedes everything else’ and offering savings of ‘up to £40 billion a year’, something the TBI suggests is ‘an opportunity the [] government cannot afford to let slip’ (21). It is vital, though, that we do not rush to adopt without sufficient safeguards in place. The risk is that the harm caused by the use of AI in public decision making, although potentially extensive, is likely to be invisible. No government would allow self-driving cars on our roads without sufficient safeguards, as the risk of people being physically harmed or even killed are clear and obvious. Similarly, we would not allow AI to treat cancer patients with radiotherapy or chemotherapy or to develop new pharmaceuticals without appropriate safeguards. Indeed, this is the root of the safety model of the EU AI Act. We will, though, it appears, allow AI and other algorithms to make decisions on access to benefits, visa status or marriage provision with little independent oversight. It may be thought that, unlike self-driving cars or medical AI, these systems cause little harm, and unlike self-driving cars or medical interventions any actual harm can be easily reversed on appeal or review. Anyone holding this ‘lack of harm’ view may like to speak to Chermaine Leysner, a Dutch citizen who received an erroneous €100,000 tax bill in 2012 due to errors in a self-learning algorithm used by Dutch tax authorities to create risk profiles to spot childcare benefits fraud. In the nine years that it took Leysner to clear her name, the stress caused by the tax bill drove her into depression and burnout, causing her to separate from her children’s father.¹¹ Her case, and the case of tens of thousands of victims of what the Dutch have dubbed the *toeslagenaffaire*, or the child care benefits scandal, led directly to the inclusion of public sector tools as being of ‘high risk’ in the European AI Act and caused European Commission Executive Vice President Margrethe Vestager to observe that this ‘is exactly what every government should be scared of’. We should, as the TBI suggests, harness the benefits of AI in public services, but **only** when we have sufficient safeguards in place.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR AFFILIATIONS

Andrew Murray  orcid.org/0000-0002-5550-7250
Professor of Law, LSE Law School, London School of Economics, UK

REFERENCES

1. **Le Sueur A, Sunkin M, Murkens J.** Public law: text, cases, and materials. 5th ed. Oxford: OUP; 2023. DOI: <https://doi.org/10.1093/he/9780192870612.001.0001>
2. **Elliott M.** The constitutional foundations of judicial review. Oxford: Hart Publishing; 2001.
3. **Allan TRS.** Law, liberty, and justice: the legal foundations of British constitutionalism. Oxford: OUP; 1994. DOI: <https://doi.org/10.1093/acprof:oso/9780198259916.001.0001>
4. **Office for National Statistics.** Civil justice statistics quarterly: January to March 2024 [Internet]. 2024 [cited 2024 September 2]. Available from: <https://www.gov.uk/government/statistics/civil-justice-statistics-quarterly-july-to-september-2023/civil-justice-statistics-quarterly-july-to-september-2023>.
5. **Chester Buxton R, Radnor Z.** How do they do it? Understanding back office efficiency savings made by English councils. International Journal of Public Sector Management. 2012; 25:118–132. DOI: <https://doi.org/10.1108/09513551211223776>

11 <https://www.politico.eu/article/dutch-scandal-serves-as-a-warning-for-europe-over-risks-of-using-algorithms/>.

6. **James O.** Evaluating executive agencies in UK government. *Public Policy and Administration*. 2001; 16:24–52. DOI: <https://doi.org/10.1177/095207670101600303>
7. **Acemoglu D, Restrepo P.** Robots and jobs: evidence from US labor markets. NBER Working Paper No. 23285, Cambridge, MA; 2018. DOI: <https://doi.org/10.3386/w23285>
8. **Silberg J, Manyika J.** Notes from the AI frontier: Tackling bias in AI (and in humans). McKinsey Global Institute. 2019; 1:1–31.
9. **Schwartz R, Vassilev A, Greene K, Perine L, Burt A, Hall P.** Towards a standard for identifying and managing bias in artificial intelligence. NIST Special Publication 1270, US Department of Commerce, National Institute of Standards and Technology; 2022. DOI: <https://doi.org/10.6028/NIST.SP.1270>
10. **House of Commons Debates.** 8 December 2022 vol 724, col 499.
11. **The Public Law Project.** Tracking Automated Government (TAG) Register. [cited 2024 September 2]. Available from: <https://trackautomatedgovernment.shinyapps.io/register/>.
12. **von Eschenbach WJ.** Transparency and the black box problem: why we do not trust AI. *Philosophy & Technology*. 2021; 34:1607–1622. DOI: <https://doi.org/10.1007/s13347-021-00477-0>
13. **Infrastructure and Projects Authority.** Data analytics and AI in government project delivery [Internet]. 2024 [cited 2024 September 2]. Available from: <https://www.gov.uk/government/publications/data-analytics-and-ai-in-government-project-delivery>.
14. **Work and Pensions Committee.** Oral evidence: DWP's annual report and accounts 2022–23 HC 417 [Internet]. 2024 [cited 2024 September 2]. Available from: <https://committees.parliament.uk/oralevidence/14052/html/>.
15. **Foxglove.** DWP boss admits there is bias in algorithm targeting disabled people for benefit fraud investigations [Internet]. 2024 [cited 2024 September 2]. Available from: <https://www.foxglove.org.uk/2024/02/12/dwp-boss-bias-algorithm-disabled-people/>.
16. **Bolt D.** The implementation of the 2014 'hostile environment' provisions for tackling sham marriage [Internet]. Independent Chief Inspector of Borders and Immigration. 2016 [cited 2024 September 2]. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/577880/Sham_Marriage_report.pdf.
17. **What do they know?** [Internet]. Jack Maxwell FOI Request. 2020 [cited 2024 September 2]. Available from: https://www.whatdotheyknow.com/request/sham_marriages_8/response/1693677/attach/3/61422%20Maxwell%20Annex%20C%20Redacted.pdf.
18. Public Law Project v The Information Commissioner [2023] UKFTT 102 (GRC).
19. **Gesk TS, Leyer M.** Artificial intelligence in public services: When and why citizens accept its usage. *Government Information Quarterly Article*. 2022; 39:101704. DOI: <https://doi.org/10.1016/j.giq.2022.101704>
20. **Office for National Statistics.** Public awareness, opinions and expectations about artificial intelligence: July to October 2023 [Internet]. 2023 [cited 2024 September 2]. Available from: <https://www.ons.gov.uk/businessindustryandtrade/itandinternetindustry/articles/publicawarenessopinionsandexpectationsaboutartificialintelligence/julytooctober2023#positive-and-negative-impacts-of-ai>.
21. **Tony Blair Institute for Global Change.** Governing in the age of AI: a new model to transform the state [Internet]. 2024 [cited 2024 September 2]. Available from: <https://institute.global/insights/politics-and-governance/governing-in-the-age-of-ai-a-new-model-to-transform-the-state>.

TO CITE THIS ARTICLE:

Murray A. Automated Public Decision Making and the Need for Regulation. *LSE Public Policy Review*. 2024; 3(3): 3, pp. 1–10. DOI: <https://doi.org/10.31389/lseppr.110>

Submitted: 04 June 2024

Accepted: 23 September 2024

Published: 04 November 2024

COPYRIGHT:

© 2024 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

LSE Public Policy Review is a peer-reviewed open access journal published by LSE Press.