

Technology and Public Law

Popular culture abounds with narratives projecting futures where technological advances in automation and artificial intelligence (AI) have in some way outpaced humans' ability to manage them. Invariably, AI developments are designed to assist humans in some way – in *Black Mirror*, lifelike robots that simulate the personality of dead loved ones are created to assuage loneliness; in *Westworld*, theme parks with costumed human robots are created to vent human bloodthirstiness and lust, so that people are able to walk the straight and narrow in real life; in *2001: A Space Odyssey* and the *Hitchhiker's Guide to the Galaxy*, the disembodied robot Hal and the embodied but profoundly depressed robot Marvin help fly and navigate spaceships. However, the truly compelling parts of these stories are the unanticipated negative effects of these AI systems – which range from bleak to dystopian. In *Black Mirror*, there is personal disappointment and pain because a robot cannot possibly replicate a loved one in any meaningful way. In *Westworld*, the constantly brutalized robots achieve sentience and rise up against their human oppressors. In *Space Odyssey*, robot Hal chillingly decides that humans are expendable to achieve its mission and kills the astronauts on board in order to protect and continue its programmed directive.

These stories project distant and dystopian futures that may never come to be. They belie a far more mundane reality – that of gently humming computer servers under dim neon lights in drab office buildings, generating thousands of automated decisions, with civil servants casting an occasional desultory eye over the computers. Deep in the heart of the bureaucracy, machines make decisions on who should receive social security benefits, issue debt notices, and generate numerous predictions on the risk of criminal recidivism, child abuse, and welfare fraud.

These tasks were previously carried out by a cavalcade of frontline civil servants, who made individual determinations about who should get a visa or social security benefits, how much a person owes in tax, and whom to pursue for fraudulent claims. But for reasons of cost-saving, efficiency, and the desire to be at the forefront of technology, governments have been enthusiastic adopters of AI and automation, leading to the automated state that we have today.

However, where automated systems are deployed hastily and without enough safeguards, there may be significant detrimental consequences. Large-scale botched rollouts of government programs can result in injustice heaped upon hundreds of thousands of people before errors are rectified.

For example, in 2001, Republican governor Rick Snyder, a former venture capitalist who vowed to run the government like a business, decided to put into place a new automated social security system in Michigan. Unemployment agency staff were laid off to achieve cost-savings. In their place, a \$46 million automated fraud detection system was introduced: the Michigan Integrated Data Automated System (MiDAS). MiDAS was programmed to treat any data discrepancies or inconsistencies as evidence of illegal conduct and to trigger an automatic default finding against individuals. The program has been a windfall for the Michigan bottom line, collecting over \$60 million in just four years, bringing the State budget into the green.

However, it turned out that the AI system had a 93 percent error rate.¹ As a result, between 2013 and 2015, the system falsely accused more than 40,000 Michigan residents of suspected fraud. It seized their tax refunds, garnished their wages, and imposed civil penalties of 400 percent of the amount the person was owing plus interest. These people wrongfully lost access to unemployment payments and reported facing fines as high as \$100,000.

As a result of being wrongfully accused of unemployment insurance fraud, 1,100 people filed for bankruptcy, losing jobs, homes, and livelihoods. With a quadruple penalty plus interest, people went into significant debt. As their credit ratings plummeted, there was a seven-year stain on their record. Consequently, they were unable to get jobs, rent houses, or buy cars or homes. Two people even attempted suicide.

In a 2017 settlement, Michigan's unemployment agency agreed to stop using MiDAS's automated functions without human review.² The settlement also made the agency reverse and refund certain fraud determinations. In July 2022, the Michigan Supreme Court ruled in *Bauserman v Unemployment Insurance Agency* that people in Michigan could sue state agencies for monetary damages for violating their constitutional rights.³ This class action was settled for \$20 million in January 2023.⁴

¹ Michigan Office of the Auditor-General, *Michigan Integrated Data Automated System (MiDAS)*, Unemployment Insurance Agency, Department of Talent and Economic Development and Department of Technology, Management and Budget (2016), https://audgen.michigan.gov/finalpdfs/15_16/r641059315.pdf.

² *Zynda v Zimmer*, order of the United States District Court for the Eastern District of Michigan, issued 2 February 2017 (Case No. 2:15-cv-11449).

³ *Bauserman v Unemployment Insurance Agency*, 2022 WL 2965921, at 3 (Michigan 2022).

⁴ Michigan Government, *Notice of Settlement of Bauserman UIA False Fraud Class Action* (Press Release, 23 January 2023) 3, www.michigan.gov/ag/-/media/Project/Websites/AG/releases/2023/January/Notice-Settlement-Bauserman.pdf?rev=ed98484f3e4d48be8254a73c2201e611.

This is not an isolated occurrence. Across the globe in Australia, an automated system pejoratively called “Robodebt” erroneously identified overpayments and calculated debts deemed to be owed by social security beneficiaries. Errors of methodology led to incorrect or inflated debt calculations for over 600,000 individuals.⁵ Desperate people falsely accused of high debts rang the social security agency, Centrelink, only to be put on hold for eight hours.⁶ These incorrect calculations have led to grave repercussions for vulnerable low-socioeconomic debtors, including severe mental health issues and even suicide.⁷ In 2019, the Federal Court of Australia held that the basis for raising debts under the Robodebt program was irrational and therefore unlawful.⁸ Following this, the Australian government agreed to settle a class action on behalf of 600,000 persons affected by Robodebt for A\$112 million, without any admission of liability.⁹ The judge presiding over the case decried the Australian government’s handling of its online compliance systems as a “very sorry chapter in Australian public administration.”¹⁰ The subsequent Robodebt Royal Commission set up to investigate the scheme lambasted it as an “extraordinary saga” of “venality, incompetence and cowardice.”¹¹

In yet another corner of the globe, there was the Post Office Scandal in the UK. Between 2000 and 2014, 736 sub-postmasters were wrongfully prosecuted for theft, fraud, and false accounting in their branches. The Post Office Ltd, a wholly owned British government company, aggressively pursued these sub-postmasters through the courts, resulting in numerous criminal convictions, some being jailed (including a teenager and a pregnant woman), and many facing bankruptcy and financial ruin. There were at least four suicides.¹²

However, none of them had done anything wrong; the Horizon accounting software system designed by Fujitsu had produced incorrect accounting shortfalls, due to bugs, errors, and defects in the IT system that automated the post office network. Widely decried as one of the biggest miscarriages of justice in British

⁵ Senate, Community Affairs References Committee, Parliament of Australia, *Design, Scope, Cost-Benefit Analysis, Contracts Awarded and Implementation Associated with the Better Management of the Social Welfare System Initiative* (2017) 1 [1.2], www.aph.gov.au/Parliamentary_Business/Committees/Senate/Community_Affairs/SocialWelfareSystem.

⁶ *Ibid.*, ix, 33–4, 107.

⁷ *Ibid.*

⁸ Order of Davies J in *Amato v Commonwealth* (Federal Court of Australia, VID611/2019, 27 November 2019) 6 [9].

⁹ *Ibid.*

¹⁰ Luke Henriques-Gomes, “Robodebt Responsible for \$1.5bn Unlawful Debts in ‘Very Sorry Chapter’, Court Hears”, *The Guardian*, 7 May 2021, www.theguardian.com/australia-news/2021/may/07/robodebt-responsible-for-15bn-unlawful-debts-in-very-sorry-chapter-court-hears.

¹¹ Commonwealth of Australia, *Report: Royal Commission into the Robodebt Scheme* (2023) 702, <https://robodebt.royalcommission.gov.au/system/files/2023-09/rcc-accessible-full-report.PDF>.

¹² Marina Hyde, ‘Hundreds of Lives Ruined. Not a Single Person Held to Account. And Still: Silence on the Post Office Scandal’, *The Guardian*, 2 May 2023, www.theguardian.com/commentisfree/2023/may/02/post-office-horizon-scandal-inquiry?CMP=Share_AndroidApp_Other.

history, this multimillion pound IT disaster has spanned over twenty years and has generated widespread condemnation.¹³

These events happened in advanced liberal democracies, with their sophisticated bureaucracies and well-developed systems of checks and balances. How did things go so wrong? For each of these large-scale catastrophes to occur, deliberate choices were made at key decision points by elected politicians and civil servants to procure and deploy faulty technology at a large scale, and on vulnerable populations. Hundreds of thousands of people were harmed for years due to these wrongful governmental decisions. This was followed by years of challenges in courts before the governments conceded any error. Yet, in the wake of these debacles, it is hard to pinpoint responsible individuals. Such individuals hide amidst a bureaucratic behemoth, or behind ministerial orders for cost-savings through automation. Or they are one of many IT staff who procure or develop technology used by frontline civil servants, or are amongst managerial staff who authorize automated processes.

These scandals demonstrate the need to more closely scrutinize the regulation of automated decision-making across governments. We need to consider how automated systems are designed, how these technologies are rolled out, and whether appropriate auditing practices are put into place to mitigate any potential ill effects. This book thus examines the legal, political, and managerial controls that are needed to prevent large-scale disasters.

AUTOMATION AND ARTIFICIAL INTELLIGENCE

This book covers a constellation of automated technologies used by governments. They range from the more basic expert systems that merely automate decision-making based on rule-based criteria, to the more sophisticated form of machine learning systems utilizing AI, which can make predictions or decisions using machine or human-based inputs. It also covers mass surveillance technologies, such as facial recognition technology, which automatically captures facial images and matches these with government-held databases.

The global term “automated decision-making” will be used in this book to encompass three forms of technology: rule-based deterministic systems, machine learning systems, and facial recognition technology. The term “artificial intelligence” (AI) in turn is used to denote “systems that display intelligent behavior” by analyzing their environment and “taking actions – with some degree of autonomy – to achieve specific goals.”¹⁴ Examples include machine learning and facial recognition technology. AI systems can, “for a given set of human-defined objectives, make

¹³ Nick Wallis, *The Great Post Office Scandal* (Bath Publishing, 2021).

¹⁴ European Commission’s High-Level Expert Group on Artificial Intelligence, *A Definition of AI: Main Capabilities and Scientific Disciplines*, Definition Developed for the Purpose of the Deliverables of the High-Level Expert Group on AI (2018), https://ec.europa.eu/futurium/en/system/files/ged/ai_hleg_definition_of_ai_18_december_1.pdf.

predictions, recommendations or decisions influencing real or virtual environments.”¹⁵ Thus, AI is a subset of “automated decision-making” and these terms overlap.

The most basic type of technology examined in this book is rule-based deterministic systems that follow “a series of pre-programmed rules written by humans,” while predictive machine learning systems utilizing AI deploy “rules that are inferred by the system from historic data.”¹⁶ Automated systems, such as systems that match data on welfare compliance, are examples of rule-based logic systems, since they use pre-programmed rules to reach decisions about eligibility for welfare benefits. On the other hand, probabilistic predictive systems, such as machine learning systems, train models to learn from data using a variety of methods (such as neural networks and “random forests”)¹⁷ which are capable of recognizing statistical patterns in different types of data (such as numbers, text, and images).¹⁸ This “includes ‘supervised learning’, where ‘training data’ are used to develop a model with features to predict known ‘labels’ or outcomes, and ‘unsupervised learning’, where a model is trained to identify patterns in data without labels of interest.”¹⁹ An example of a predictive system would be risk assessment tools that predict child welfare cases, such as the Eckerd Rapid Safety Feedback, which is a process designed to identify child welfare cases with a high probability of serious child injury or death.²⁰ This type of decision is more problematic from a rule of law point of view, due to the opacity of such systems and the possibility of bias being baked into the system, as the training of machine learning programs has the possibility of ingraining existing biases in the AI (or even of creating new ones).²¹

GOVERNMENT AND THE TECHNOLOGY BANDWAGON

Governments in liberal democracies have jumped on to the AI bandwagon. This is in part due to the fetishization of technology, a trend that has only accelerated in the last few decades. The glittering promise of new technologies has led commentators

¹⁵ OECD AI Policy Observatory, OECD AI Principles Overview (2024), <https://oecd.ai/en/ai-principles>.

¹⁶ Australian Human Rights Commission, *Human Rights and Technology: Final Report* (2021) 37, <https://humanrights.gov.au/our-work/technology-and-human-rights/publications/final-report-human-rights-and-technology>.

¹⁷ Random forests are a ‘combination of tree predictors such that each tree depends on the values of a random vector sampled independently and with the same distribution for all trees in the forest’. Leo Breiman, ‘Random Forests’ (2001) 45 *Machine Learning* 5.

¹⁸ David Freeman Engstrom et al., *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies*, Report Submitted to the Administrative Conference of the United States (2020) 12, <https://law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf>.

¹⁹ Ibid.

²⁰ Robert Brauneis and Ellen P Goodman, ‘Algorithmic Transparency for the Smart City’ (2018) 20 *Yale Journal of Law and Technology* 103, 141–2.

²¹ This will be discussed in more depth in Chapter 5.

to breathlessly proclaim a “Fourth Industrial Revolution”²² and a “Second Machine Age,”²³ from the “computer revolution” in the 1970s to the “AI revolution” today.²⁴ Proponents of the digital revolution trumpet a new age where the acceleration of technology has fundamentally transformed human existence – in the way we live, work, and interact with others.²⁵ In this utopian vision, humans are emancipated, fully interconnected through digital networks, and able to participate fully in democratic processes.²⁶ However, the stark reality is often the opposite: the political and institutional elite have the power and resources to reap the benefits of new technologies, while those at the edges of society are left out or, worse still, further marginalized.²⁷ The use of new technologies in government has often been trialed and deployed punitively against the most vulnerable sections of society, to predict and punish criminal recidivism and unemployment fraud, or to automate social security debt notices – rather than, for example, to predict insider trading by the wealthy. Without deeper social and political change, the democratizing potential of technology is lost. Existing hierarchies in society will reassert, reinforce, and perpetuate themselves.

Of course, there are great advantages for governments in adopting new technologies. Their use will no doubt increase the speed and efficiency of decisions. Automated systems are able to process data at breathtaking volumes and velocities, far surpassing human ability. This generates cost-savings, thus reducing the expense of government programs. Automated systems also may produce decisions that are more consistent than humans, who are all too susceptible to the biological and circadian rhythms that govern all animals – bail applicants are less likely to receive bail just before lunch than any other time of the day.²⁸ Humans are limited by bounded rationality and may not be able to process all the information to incorporate it into a decision.²⁹ As a non-biological decision-making entity, AI does not suffer

²² Klaus Schwab, *The Fourth Industrial Revolution* (World Economic Forum, 2016).

²³ Erik Brynjolfsson and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (W. W. Norton & Company, 2014).

²⁴ Paul Stoneman, *Technological Diffusion and the Computer Revolution* (Cambridge University Press, 1976); Yuval Noah Harari, ‘Reboot for the AI Revolution’ (2017) 550(7676) *Nature* 324.

²⁵ JCR Licklider, ‘Computers and Government’ in Michael L Dertouzos and Joel Moses (eds.), *The Computer Age* (MIT Press, 1980) 114, 126.

²⁶ John Naisbitt, *Megatrends: Ten New Directions Transforming Our Lives* (Warner Books, 1984) 282.

²⁷ Langdon Winner, *The Whale and the Reactor: A Search for Limits in an Age of High Technology* (Chicago University Press, 1988) 107; James Danziger et al., *Computers and Politics: High Technology in American Local Governments* (Columbia University Press, 1982).

²⁸ S Danziger, J Levav, and L Avnaim-Pesso, ‘Extraneous Factors in Judicial Decisions’ (2011) 108 (17) *Proceedings of the National Academy of Sciences* 6889.

²⁹ Peter M Todd and Gerd Gigerenzer, ‘Bounding Rationality to the World’ (2003) 24 *Journal of Economic Psychology* 143, 144; Andrew D Selbst, ‘Negligence and AI’s Human Users’ (2020) 100 *Boston University Law Review* 1315, 1360.

these variations and cognitive deficiencies and so is able to produce more consistent decisions – an important feature of administrative decision-making.

Despite this, AI systems are not infallible. Machine learning systems that make probabilistic predictions based on complex algorithms can use biased inputs, leading to equally biased outcomes. For example, in certain US states, judges have adopted the controversial practice of utilizing automated decision-making tools for sentencing, drawing on historic data to infer which convicted defendants pose the highest risk of re-offending.³⁰ This practice has been criticized for perpetuating and reinforcing systemic bias. An investigation found that African Americans are more likely than Caucasians to be given a false positive score by the US sentencing tool COMPAS (Correctional Offender Management Profiling for Alternative Sanctions),³¹ although this finding was disputed by the system's creators. Facial recognition technology has been shown to misidentify Asian and African American people up to 100 times more than white men, which may mean that minority races are more likely to face false accusations of being the perpetrator of a crime.³² Accordingly, technology needs to be carefully managed, to allocate risk, apportion liability, and eliminate potential unintended and detrimental effects.

SCOPE AND CONTRIBUTION

This book will analyze the rise of automated government decision-making in the United States, United Kingdom, and Australia. It will ascertain whether the advent of automated decision-making is accommodated into the existing legal frameworks of these jurisdictions, and develop a normative technological governance framework based on legal, political, and managerial controls.

These jurisdictions were chosen as comparators because they share a liberal democratic framework and have increasingly adopted automated systems within government. But there are also significant differences that generate comparative interest. The UK operates within an unwritten constitution, with external influences incorporated into domestic law from the supranational European system, such as the European Convention on Human Rights (ECHR) and the General Data Protection Regulation (GDPR). Australia operates within a legal framework without significant constitutional or human rights protections, and so is completely reliant on domestic administrative law for the protection of individuals. The US has

³⁰ Monika Zalnieriute, Lyria Bennett Moses, and George Williams, 'The Rule of Law and Automation of Government Decision-Making' (2019) 82 *Modern Law Review* 425.

³¹ J Angwin et al., 'Machine Bias', 23 May 2016, *ProPublica*, www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

³² National Institute of Standards and Technology, 'NIST Study Evaluates Effects of Race, Age, Sex on Face Recognition Software' (19 December 2019), www.nist.gov/news-events/news/2019/12/nist-study-evaluates-effects-race-age-sex-face-recognition-software.

substantial constitutional protections that incorporate a Bill of Rights. Further, the US conception of administrative law is broader than that of the UK and Australia. In the UK and Australia, the primary focus of administrative law is adjudication, while the US has a strong focus on rulemaking in addition to adjudication.

When government decision-making is aided, supplemented, or replaced by automation, we need to consider issues of legal responsibility, the ability to launch legal challenges to defective decisions, and liability for loss. In the private sphere, there is often discussion of autonomous vehicles and autonomous weapons systems, with proposed solutions within the regulatory toolkit being the law of contract and insurance requirements to allocate risk, and the laws of tort and product liability to apportion liability.³³

This book will focus on the public sphere, which necessitates different regulatory tools in the form of constitutional, administrative, and human rights laws. Scrutiny of the operation of government is important, given the pervasiveness and powers of the State in our lives. Governments make decisions that deeply affect us in many fundamental ways, from deciding whether we are entitled to stay in the country, or if we qualify for certain benefits. Governments also have strong coercive powers to enforce their will on the populace. The normative framework that animates the recommendations throughout this book, is thus that of public law.

Certain uses of AI in government do not create angst amongst public lawyers. An example is public engagement through chatbots on websites, with pre-programmed responses that merely inform users about government policies and procedures. Likewise, public law concerns are not generated by the use of AI to facilitate administrative or procedural tasks, such as the US Postal Service using machine learning algorithms in mail-sorting equipment to “read” handwritten postal codes.³⁴ In addition, public lawyers barely raise an eyebrow where decisions do not significantly affect individual rights and freedoms, such as when government meteorologists use machine learning technologies to forecast the weather.³⁵

The crux of the public law debate about the use of AI in government thus occurs in adjudicative and enforcement settings, where individual life, liberty, and property are at stake. Accordingly, the focus of this book will be on mass adjudication of benefits by government agencies, where millions of cases are decided per year, such as in the areas of social security, tax, immigration, and veterans’ affairs. These areas of high-volume decision-making are plagued by enormous backlogs and delays. For example, in US social security cases, where the average wait time was over 605 days

³³ Simon Chesterman, *We the Robots: Regulating Artificial Intelligence and the Limits of the Law* (Cambridge University Press, 2021); Ryan Abbott, *The Reasonable Robot: Artificial Intelligence and the Law* (Cambridge University Press, 2020).

³⁴ Cary Coglianese and David Lehr, ‘Transparency and Algorithmic Governance’ (2019) 71 *Administrative Law Review* 1, 30–1.

³⁵ *Ibid.*

in 2017, 10,000 claimants died awaiting a hearing in that year alone.³⁶ In that same period, delays in immigration courts could be as long as six years.³⁷ Another issue with mass adjudication is the disparity of grant rates by adjudicators, with some decision-makers, for example, having grant rates of 3 percent, while others have grant rates of 88 percent.³⁸ Immigration researchers have dubbed this phenomenon “refugee roulette,” the suggestion being that decision-making is completely arbitrary based on the assigned adjudicator rather than the merits of the case.³⁹ This leads to significant questions about varying decisional quality and accuracy between adjudicators.⁴⁰

Due to the considerable difficulties in administering mass adjudication schemes, governments have strong incentives to streamline processes, cut costs, and ensure the consistency of decision-making. Increasingly, these goals are achieved through the process of automation. The idea is that automation of high-volume decision-making will reduce the discrepancy amongst decision-makers, by providing guidance to human decision-makers for machine-assisted decisions, or in the case of machine-made decisions, ensuring that computers with fixed rules or machine learning decide like cases in a similar fashion. Further, automation of decision-making may save costs in machine-assisted decisions by increasing the efficiency of the human decision-maker by providing guidance on whether applications should be approved or, in the case of machine-made decisions, requiring fewer human staff.

Alongside mass adjudication, this book will also consider large-scale surveillance mechanisms such as facial recognition technology, which are increasingly adopted by law enforcement authorities to identify and punish individuals for legal transgressions.⁴¹ These invasive technologies may lead to significant harm to individual autonomy through loss of control of one’s personal information and the consequent potential for manipulation and harm to human dignity, including discrimination and the loss of personal anonymity.⁴²

Although there is a burgeoning literature on AI and public law, there is surprisingly little comparative doctrinal case law analysis assessing legal grounds for challenging government automated decision-making. Most scholarship to date has been focused on a single country or region (such as the European Union), or has

³⁶ David Ames et al., ‘Due Process and Mass Adjudication: Crisis and Reform’ (2020) 72 *Stanford Law Review* 1, 16.

³⁷ *Ibid.*

³⁸ Philip G Schrag, Andrew I Schoenholtz, and Jaya Ramji-Nogales, *Refugee Roulette: Disparities in Asylum Adjudication and Proposals for Reform* (NYU Press, 2009).

³⁹ *Ibid.*

⁴⁰ Ames et al., ‘Due Process and Mass Adjudication’ (n 36) 1.

⁴¹ Legal challenges to surveillance technologies are discussed in Chapters 5 and 6.

⁴² For discussion of the harms arising from the loss of privacy, see Moira Paterson and Maeve McDonagh, ‘Data Protection in an Era of Big Data: The Challenges Posed by Big Personal Data’ (2018) 44(1) *Monash University Law Review* 1, 6–9.

speculated as to how judicial review may adapt to new technologies.⁴³ This is the first book to undertake a detailed comparative analysis of automated government decision-making across the US, UK, and Australia. It makes a contribution by conducting an in-depth analysis of the grounds of legal challenge that have been utilized across three Western liberal jurisdictions that are increasingly employing automation and AI in government. It examines legal controls across four dimensions: rationality, anti-discrimination, public sector privacy/data protection, as well as freedom of information (FOI). In doing so, the book will assess the likelihood of success of these legal avenues, and consider whether the existing laws in these jurisdictions are “fit for purpose” for challenging automated government decision-making. In doing so, the book will propose how public law will need to evolve if it is to meet the rising challenges posed by AI’s increasing application in the public sphere.

Further, the book will design a best-practice framework for technological governance that combines legal, political, and managerial controls. In doing so, the book goes beyond traditional legal mechanisms, which have been the predominant focus of legal academics, and assesses institutional arrangements of political and managerial oversight, with a view to achieving a comprehensive framework for scrutinizing automated decision-making within government. While traditional legal avenues are essential to allow for individual redress, legal accountability is an *ex post* mechanism that happens at a time distant to the initial agency decision. By the time litigation commences, many people would have already been harmed by wrongful AI decisions. Thus, *ex ante* mechanisms of managerial controls are required to effectively control agency behavior and mitigate the wide-ranging impacts of deficient AI systems. This should occur in conjunction with political oversight mechanisms. This book will thus develop a multi-faceted governance framework of legal, political, and managerial controls comprising both internal and external accountability processes.

AUTOMATION AND ADMINISTRATIVE LAW

Administrative law “defines the structural position of administrative agencies within the governmental system, specifies the decisional procedures those agencies must follow, and determines the availability and scope of review of their actions by the independent judiciary.”⁴⁴ As Stewart explains:

The traditional core of administrative law has focused on securing the rule of law and protecting liberty by ensuring that agencies follow fair and impartial decisional

⁴³ For example, Rebecca Williams, ‘Rethinking Administrative Law for Algorithmic Decision Making’ (2021) 42(2) *Oxford Journal of Legal Studies* 468; Yee-Fui Ng et al., ‘Revitalising Public Law in a Technological Era: Rights, Transparency and Administrative Justice’ (2020) 43(3) *University of New South Wales Law Journal* 1041.

⁴⁴ Richard B Stewart, ‘Administrative Law in the Twenty-First Century’ (2003) 78(2) *New York University Law Review* 437, 438.

procedures, act within the bounds of the statutory authority delegated by the legislature, and respect private rights.⁴⁵

Administrative law is thus predicated upon the control of government action, in ensuring that government acts within legal confines and is subject to the dictates of rationality, accountability, transparency, participation, and procedural fairness.⁴⁶ These safeguards protect individuals against arbitrary and unlawful government decisions. A related stream of literature on administrative justice focuses on the nature and quality of administrative decision-making by government agencies, particularly those that determine the legal entitlements of individuals, as well as the systems of redress by which people can challenge the decisions of public bodies.⁴⁷ Administrative justice aims to enable accurate, fair, and impartial administrative decision-making through internal agency procedures and redress mechanisms. This is congruent with the rule of law, which seeks to protect individuals against the use of arbitrary power.⁴⁸

Thus, as instruments of democracy, governments need to address a broad set of public law norms and values: the demand for transparency, accountability, rationality, participation, and efficiency. Automated decision-making poses challenges for each of these principles.

Transparency

Transparency in government is a democratic ideal, based on the notion that an informed citizenry is better able to participate in government, thus placing an obligation on government to provide public disclosure of information.⁴⁹ Increasing transparency in government also reduces the risk of corruption. As the saying goes: “sunlight is . . . the best of disinfectants.”⁵⁰ In terms of government decision-making, it is desirable for persons affected by a decision to know *why* it was made, including access to the reasons for decisions and underpinning principles.

However, the rise of automated decision-making in government has created significant challenges for transparency. Automated decision-making can be opaque in two ways. The first is its invisibility; people often do not realize that they are interacting with the technology, and generally know little about the programs that

⁴⁵ Ibid.

⁴⁶ Ng et al., ‘Revitalising Public Law in a Technological Era’ (n 43) 1041.

⁴⁷ Simon Halliday, ‘Administrative Justice and Street-Level Emotions: Cultures of Denial in Entitlement Decision-Making’ (2021) 4 *Public Law* 727; Marc Hertogh et al., *Oxford Handbook of Administrative Justice* (Oxford University Press, 2011) xvi.

⁴⁸ Yseult Marique, ‘The Rule of Law and Administrative Justice’ in Marc Hertogh et al. (eds.), *Oxford Handbook of Administrative Justice* (Oxford University Press, 2011) xvi.

⁴⁹ Daniel J Metcalfe, ‘The History of Government Transparency’ in Padideh Ala’i and Robert G Vaughn (eds.), *Research Handbook on Transparency* (Edward Elgar, 2014) 247, 249.

⁵⁰ Louis D Brandeis, *Other People’s Money and How the Bankers Use It* (F A Stokes, 1914) 92.

are used to make decisions about them. The second is the complexity of its functioning. This leads to what is commonly known as the “black box” problem, whereby it is possible to observe incoming data (input) and outgoing data (output) in algorithmic systems, but their internal operations are poorly understood. As highlighted by Oswald, incorporating an algorithm into decision-making “may come with the risk of creating ‘substantial’ or ‘genuine’ doubt as to why decisions were made and what conclusions were reached.”⁵¹

The demand for transparency of automated decision-making is addressed through several regulatory channels: legal controls by releasing information through FOI legislation, as well as managerial controls through the centralized disclosure of algorithms, and notification and explanation of automated decisions.⁵²

Accountability and Rationality

Apparatuses of the State have to account to the people for their actions and make reparations for any harms suffered. There are three facets of accountability that are relevant to AI decision-making. The first is the question of responsibility: Who is responsible for automated decisions? The second is the ability of individuals harmed by government actions, to obtain legal redress. The third is independent monitoring and oversight of government automated decisions.

In terms of responsibility, the underlying assumption in the administrative law context has been that an administrative decision would be made by a human or a body comprised of humans, so that in turn there would be a responsible party.⁵³ Where a responsible person can be identified, individuals harmed by a decision are able to achieve accountability by seeking legal redress for government decisions. Government decisions must be reviewable in the courts and tribunals, and reparations must be made to those aggrieved by erroneous decisions by public officials. This means that automated decisions should be subject to judicial review.

The second major requirement for executive accountability is that the use of automated systems must result in decisions that are produced rationally and in compliance with the general framework under which they are authorized. Automated decisions must be in accordance with the constitutional, statutory, and human rights framework, including through due process rights, anti-discrimination, privacy, and data protection laws.⁵⁴

The third element of accountability is independent oversight of government decisions. Like other government decisions, automated government decision-

⁵¹ Marion Oswald, ‘Algorithm-Assisted Decision-Making in the Public Sector: Framing the Issues Using Administrative Law Rules Governing Discretionary Power’ (2018) 376 *Philosophical Transactions of the Royal Society* 20170359.

⁵² This will be explored in Chapters 6 and 9.

⁵³ Oswald, ‘Algorithm-Assisted Decision-Making in the Public Sector’ (n 51) 379.

⁵⁴ These issues are further discussed in Chapters 4, 5, and 6.

making should be monitored and audited not just internally, but by external scrutiny bodies. This is aided by the proliferation of oversight bodies or officeholders to monitor the executive, such as ombudsmen, auditors, commissions, and tribunals.⁵⁵ Another source of independent scrutiny is through Congress or parliamentary committees, who provide both periodic audit-like oversight and “fire alarm” responses to political problems.⁵⁶ These scrutineers can play a significant role in investigating and exposing issues relating to automated decision-making in government agencies.⁵⁷

Participation

American administrative law is particularly focused on enhancing participation by members of the public in regulatory processes through rulemaking and adjudication.⁵⁸ Public participation in decision-making processes may promote fairness in two ways. First, the claimant would be exposed to and better understand the substantive adjudicatory norms and the decision process.⁵⁹ Participation also aims to give the claimant the opportunity to voice matters they regard as important and to rebut any adverse evidence. Participatory processes are also desirable in the design phase of new technologies to promote the dignity and empowerment of people affected by automated systems, and to allow for expert input.⁶⁰

Efficiency

Another administrative law imperative is that of efficiency of government decision-making. Achieving efficiency goes towards competency and adequacy of performance, towards attaining a desired effect.⁶¹ This in turn increases the cost-effectiveness, accuracy, and precision of decision-making. AI should be able to make speedier decisions with fewer errors than humans. Secondarily from an economic perspective, resources saved from inefficient decision-making can be applied to other worthwhile causes, reducing the burden on government coffers and ultimately the taxpayer. The fear is that too much reliance and too many tasks

⁵⁵ Christopher Hood et al., *Regulation Inside Government: Waste-Watchers, Quality Police, and Sleaze-Busters* (Oxford University Press, 1999) 11.

⁵⁶ BR Weingast, ‘Caught in the Middle: The President, Congress, and the Political-Bureaucratic System’ in JD Aberbach and MA Peterson (eds.), *The Executive Branch* (Oxford University Press, 2005) 312, 329–30.

⁵⁷ These issues will be explored in further detail in Chapter 8.

⁵⁸ Stewart, ‘Administrative Law in the Twenty-First Century’ (n 44) 441.

⁵⁹ Jerry L. Mashaw, *Bureaucratic Justice: Managing Social Security Disability Claims* (Yale University Press, 1985) 140.

⁶⁰ These issues are elaborated upon to construct a governance framework in Chapter 9.

⁶¹ Aziz Z. Huq, ‘Constitutional Rights in the Machine Learning State’ (2020) 105 *Cornell Law Review* 1875, 1915–7.

might be assigned to AI with insufficient human supervision to ensure justice is done in the administration of law – an issue identified by Chatila et al. as “the social component of the socio-technical system.”⁶² The dictates of efficiency thus have to be achieved within the constraints of administrative justice and fairness to individuals.

These public law principles and values provide a normative basis for designing a robust technological governance framework that seeks to enhance the transparency, fairness, and accountability of automated decision-making, as well as protect individual rights and freedoms.

GOVERNANCE FRAMEWORK

To achieve the administrative law norms of transparency, accountability, rationality, participation, and efficiency outlined above, this book proposes an accountability framework focusing on three main dimensions: legal, political, and managerial controls. *Legal controls* are focused on judicial review for individual redress and rationality of AI decisions,⁶³ anti-discrimination laws for challenging biased inputs to AI decisions,⁶⁴ privacy and data protection legislation for the protection of individual data,⁶⁵ and FOI legislation for transparency.⁶⁶ Second, *political controls* focus on the operation of parliamentary committees and oversight bodies such as ombudsmen, auditors-general, and commissioners.⁶⁷ These political oversight mechanisms enhance the transparency of AI decision-making, and scrutinize the rationality and efficiency of decisions. Third, *managerial controls* focus on the internal regulation within agencies through agency and central executive guidelines and internal agency measures.⁶⁸

Whilst a lawyer’s first instinct is to gravitate towards litigation, by the time automated government technologies are challenged in courts, many people have already been harmed by the wrongful use of these technologies. As such, legal controls should be the measure of last resort, to allow individuals to challenge government decision-making and seek a remedy for the harm caused.

The institutional checks and balances within a democratic system form the political controls of government decision-making. Parliamentary committees and oversight bodies are able to utilize their coercive powers to conduct investigations into specific AI scandals, or to explore policy issues and make recommendations for

⁶² R Chatila et al., “Trustworthy AI” in B Braunschweig and M Ghallab (eds.), *Reflections on Artificial Intelligence for Humanity* (Springer, 2021) 13, 15.

⁶³ This is discussed in Chapter 4.

⁶⁴ This is discussed in Chapter 5.

⁶⁵ This is discussed in Chapter 6.

⁶⁶ This is discussed in Chapter 7.

⁶⁷ This is discussed in Chapter 8.

⁶⁸ This framework will be comprehensively developed in Chapter 9.

change. In the case of large-scale controversies, stand-alone public inquiries may be established to ventilate how things have gone wrong, and consider how to prevent future recurrences of major AI disasters.

A crucial element of ensuring the proper oversight of AI in government is managerial controls, which encompass the internal regulation within agencies. These preventative measures of self-regulation within agencies form a fundamental aspect of the governance framework, as they harness internal bureaucratic discipline to self-monitor the design and deployment of their automated systems, in order to avoid harms engendered by deficient automated systems.

STRUCTURE

The first part of the book tells the story of the rise of automated decision-making in government and the effects of algorithms on vulnerable members of society. Chapter 1 introduces the book, defines key terms, and outlines the book's scope and contribution. Chapter 2 examines the historical development of the use of technology in government, from classic Weberian bureaucracies based on paper-based filing systems to the current pervasive use of automated decision-making and AI in government. It also examines the effect of new technologies on vulnerable populations, focusing on social security as a case study. It shows that the use of automation in government has often been trialed on and deployed punitively against the most vulnerable sections of society. The faulty design and implementation of these technologies have caused significant harm to these disadvantaged groups.

The second part of the book outlines the fight against the wrongful use of automated decision-making in government. It engages in a fine-grained doctrinal and comparative analysis of the mechanisms of legal accountability that apply to automated government decision-making in the United States, United Kingdom, and Australia. Chapter 3 outlines the constitutional, human rights, and administrative law frameworks in the three jurisdictions, while Chapters 4 to 7 analyze the adaptability and effectiveness of existing laws to successfully challenge automated government decision-making. Towards this end, the book scrutinizes the grounds of legal challenges to automated decision-making exhibited in the jurisdictions. There are three main forms of challenging government AI decision-making: challenges to the output, input, and use of data.

First, a person affected by AI decisions could challenge the *outputs* or results of the decision-making, if they can show that the computer produces decisions that are substantively flawed. This can be done through challenges to the rationality of the decisions in the UK and Australia, or due process review in the US, which considers the relationship between the inputs and outputs of government decision-making. This has proven to be a largely successful avenue of review, with the US adopting a more procedural approach based on due process, and the UK and Australia adopting a more substantive approach, as will be discussed in Chapter 4.

Another method is to challenge the *inputs* to the decision-making, in terms of faulty code and data. The affected person can argue based on anti-discrimination principles that the data inputted into AI systems is biased and produces discriminatory outcomes. As we will see, this has not proven to be a fruitful avenue for challenging automated government decision-making, as anti-discrimination tests based on intention, causation, and impact fail to capture the nuances of AI systems, and are difficult to prove in an AI context, as shown in Chapter 5.

Finally, it is also possible to challenge the *use of the data* in terms of data sharing and retention. A person can claim privacy protections in the government's use, dissemination, and retention of AI data. The UK is the only jurisdiction of those examined that has successfully deployed this ground of legal challenge, with specifically calibrated laws on data protection and additional protections from the European human rights framework, as will be examined in Chapter 6.

Following this, the book considers the operation of FOI laws, which allow individuals to obtain information about government automated decision-making, towards facilitating legal challenges (Chapter 7). However, FOI claims have been hindered by the trade secret exemption claimed by IT vendors. As we will see, the UK's framework is particularly effective, as it combines FOI legislation with targeted protection under the GDPR based on the right to know that a decision is automated, the requirement of disclosure of meaningful information about the logic involved, and the significance and the envisaged consequences of such processing.

The third part of the book goes beyond legal boundaries to examine other institutions and mechanisms of accountability through political and managerial controls. It will consider political accountability through the role of scrutiny bodies, such as parliamentary committees, ombudsmen, information commissioners, and auditors-general (Chapter 8). These bodies play a strong role in investigating various AI controversies, as well as issuing forward-looking policies and setting centralized standards across government.

Chapter 9 of the book develops a framework for technological governance focusing on legal, political, and managerial controls. It delves into mechanisms of managerial accountability through internal agency controls. In doing so, it provides normative recommendations for an ideal system of managerial controls comprising internal agency mechanisms, algorithmic auditing, centralized coordination, and transparency measures. The governance framework comprises a range of internal and external mechanisms of accountability, and prospective and retrospective measures of control. It is aimed at achieving the administrative law norms of transparency, accountability, rationality, participation and efficiency, as well as to ensure fairness to individuals subject to automated government decision-making.

Finally, the book concludes in Chapter 10 with a summation of the findings and recommendations, considerations for governments seeking to regulate automated decision-making in the public sector, as well as a discussion of emerging themes and issues.