

Authorisation and accountability of automated government decisions under Australian administrative law

Samuel White*

The delegation of government decision-making can be a ‘practical administrative necessity’¹ for effective governance, due to the time and resource constraints that public agencies face on a day-to-day basis. Due to these very same resource constraints, there is significant interest in deploying automated tools to assist or take over government decision-making processes. The promise of automation, in theory, includes increased efficiencies and cost savings for government, as well as more prompt service for the public.²

Automated decision-making processes have been noted to have caught administrative lawyers off-guard,³ and significant public attention has been placed on these tools in Australia due to questions about the fairness and legality of automated processes — concerns that are central to the concept of administrative justice.⁴ This public concern was perhaps most obvious in response to Centrelink’s ‘Robodebt’ online compliance initiative, in which automated debt collection notices were sent to social security payment recipients who purportedly under-declared their income. The calculation of these debts was found to be deficient in a test case that found a Robodebt decision to be unlawful on grounds of irrationality.⁵ The Australian Government, subsequently, has made the decision to refund all debts levied under the scheme at a cost of \$720 million.⁶

There has been much academic work on the automation of administrative decisions in Australia. The highly pertinent work of Ng and O’Sullivan⁷ considered the Federal Court’s decision in *Pintarich v Deputy Commissioner of Taxation*⁸ (*Pintarich*) and argued that there must be a ‘modern interpretation of existing administrative law principles’ that does not exclude automated decisions from the scope of review under the *Administrative*

* Samuel White is a Captain in the Australian Army Legal Corps. He is currently posted to the Directorate of Operations and International Law. This posting sees the author advise on delegated decision-making and machine learning with respect to military decision-making processes both with kinetic and non-kinetic operations. The author also served as Associate to Justice John Logan RFD of the Federal Court of Australia and is grateful for the exposure to alternative perspectives of administrative law. Finally, many thanks to Mr Thomas Mackie of the Attorney-General’s Department for significant input on this article. All opinions and errors remain the author’s alone and do not reflect those of the Australian Defence Force or Department of Defence.

1 *Re Reference; ex parte D-G Social Services* (1979) 2 ALD 86.

2 Yee-Fui Ng and Maria O’Sullivan, ‘Deliberation and Automation — When is a Decision a “Decision”?’ (2019) 26 *Australian Journal of Administrative Law* 21, 21.

3 Justice Melissa Perry, ‘iDecide: Administrative Decision-Making in the Digital World’ (2017) 91 *Australian Law Journal* 29–34, which developed on the earlier work of Justice Melissa Perry, ‘iDecide’ (Speech, Cambridge University, Cambridge Centre for Public Law Conference, 2014); see further considering the specific issue of the delegation of administrative decisions to an automated decision-maker and whether current statutory authorisations to employ automated systems stand up to existing Australian administrative law principles.

4 Robin Creyke, ‘Administrative Justice — Towards Integrity in Government’ (2007) 31 *Melbourne University Law Review* 705.

5 Orders of 27 November 2019 by Davies J in *Amato v Commonwealth* [2019] FCA.

6 Luke Henriques-Gomes, ‘Robodebt: Government to Refund 470,000 Unlawful Centrelink Debts Worth \$721m’, *The Guardian* (online, 29 May 2020) <<https://www.theguardian.com/australia-news/2020/may/29/robodebt-government-to-repay-470000-unlawful-centrelink-debts-worth-721m>>.

7 Ng and O’Sullivan, above n 2.

8 *Pintarich v Deputy Commissioner of Taxation* [2018] FCAFC 79.

Decisions (Judicial Review) Act 1977 (Cth). Dominique Hogan-Doran SC⁹ and Justice Melissa Perry¹⁰ have also written on the key principles of administrative law that will apply to automated decisions and the potential grounds of review that could be made against an automated decision.

This article considers the specific issue of whether current statutory authorisations for the use of automated systems stand up to existing Australian administrative law principles, as well as potential risks posed by the current approach. It is necessary to expand upon how automated decisions are made before considering the legal principles that govern delegations and authorisations. The article examines existing attempts to deal with the use of automated systems, which deem an automated system to be one made by a senior government official; and considers how these provisions will be construed and whether any risks exist. Finally, the article presents observations and conclusions.

Automated decision systems in government

Automated decision-making systems are becoming more prevalent in government processes around the world, in areas as diverse as the administration of social security, taxation, criminal sentencing and migration.¹¹ These systems are most likely to be deployed in branches of government that must cope with a high caseload volume, as well as repetitive assessments against prescriptive criteria.

However, as will be shown below, automated systems can vary in nature, which is likely to have implications for the manner in which they are authorised or delegated, as well as the risks that might be posed by indiscriminate use of those systems.

How are automated decision systems used in government?

Automated systems can be designed in different forms and may employ decision-making processes that rely on explicitly coded logic or logic that is developed through machine learning techniques.

Expert systems have been used for a number of decades¹² and typically consist of pre-programmed rules¹³ that recommend a particular outcome to users of the system when a set of facts are input. For example, medical expert systems may allow a physician to search or input certain observations of a patient into the system, and the system can provide a provisional diagnosis based on those input facts. A legal expert system may similarly provide a legal 'diagnosis', based on input facts, by identifying whether any legal issues arise or whether the elements of a cause of action are likely to be met.

9 Dominique Hogan-Doran SC, 'Computer Says "No": Automation, Algorithms and Artificial Intelligence in Government Decision-Making' (2017) 13 *The Judicial Review* 1, 9.

10 Perry, above n 3, 29, 30.

11 Monika Zalnierute, Lyria Bennett Moses and George Williams, 'The Rule of Law and Automation of Government Decision-Making' (2019) 82(3) *Modern Law Review* 425.

12 See Kevin Ashley, 'Case-Based Reasoning and its Implications for Legal Expert Systems' (1992) 1 *Artificial Intelligence and Law* 113; Paul Hynes, 'Doctors, Devices and Defects: Products Liability for Defective Medical Expert Systems in Australia' (2004) 15 *Journal of Law, Information and Science* 7.

13 Zalnierute, Moses and Williams, above n 11, 433.

Typically, this generation of expert system relied on pre-programmed rules or symbolic logic¹⁴ that operated in a deterministic manner. Critically, the logic which these expert systems followed can usually be traced back to the explicit coding choices of the software developer.¹⁵

By contrast, software that employs machine learning techniques is able to make a prediction or decision about something in the world without the need for human intervention.¹⁶ Using this programming technique, software engineers enable the algorithm to learn from historical data, such as past cases, fed to it by example and through its own trial and error experience. On the basis of this learned experience, the algorithm programs its own internal decision logic as to what the optimal way to complete a task is.

Machine learning is very good for very specific tasks. Machine learning algorithms use probabilistic reasoning and can be implemented so that they recommend or choose a particular course of action based on a particular degree of confidence.

These are algorithms that have an adaptive quality,¹⁷ as they can develop their recommendations, and improve their outputs, over time as they are exposed to more training data. Through the development of the algorithm, a large amount of data is used to teach the algorithm how to come to a decision while reducing the likelihood of predicting or recommending a false positive or false negative.

Machine learning software is particularly powerful at increasing the efficiency of decision-making processes and has been implemented by governments outside of Australia. One particularly well-known example is COMPAS, which is a machine learning system that conducts a recidivism risk assessment for the purposes of criminal sentencing by predicting an individual's likelihood of reoffending.¹⁸ The factors which lead to a prediction that an individual has a high likelihood of reoffending are notoriously opaque and restricted by intellectual property rights;¹⁹ however, the system has received widespread criticism for correlating factors such as race or postcode with a risk of recidivism, without taking into account the causal factors and systemic inequalities that may lead to recidivism rates within a minority community.²⁰

14 Hynes, above n 12.

15 David Vladeck, 'Machines Without Principals: Liability Rules and Artificial Intelligence' (2014) 89(1) *Washington Law Review* 117.

16 Andreas Matthias, 'The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata' (2004) 6 *Ethics and Information Technology* 175.

17 Emad Dahiyat, 'Intelligent Agents and Liability: Is It a Doctrinal Problem or Merely a Problem of Explanation?' (2010) 18(1) *Artificial Intelligence & Law* 103, 106.

18 Cynthia Rudin, Caroline Wang and Beau Coker, 'The Age of Secrecy and Unfairness in Recidivism Prediction' (2020) 2(1) *Harvard Data Science Review* 1.

19 Zalnieriute, Moses and Williams, above n 11, 441.

20 Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, 'Machine Bias', *ProPublica* (online, 23 May 2016) <<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>>.

However, another critical factor that may determine the risk profile of using automated systems for government decisions is the *manner* in which the software is deployed and the extent to which a human has input into the decision-making process.²¹

Automated systems may assist government officials to make administrative decisions to a variety of different degrees. For example, their level of input may include making the decision; recommending a decision to the decision-maker; guiding a user through relevant facts, legislation and policy; and providing useful commentary as a decision support system.²²

A general distinction may be drawn between systems that are deployed to provide mere decision support to a human decision-maker and systems that are deployed so that there is no human input into the decision at all. Of course, systems that arrive at an administrative decision or levy a penalty without human input have been in existence for some time. For example, speed monitoring cameras have been used on our roads for many years. They automatically detect potential speeding offences, and fines may be subsequently issued to the address of a vehicle's registered owner. This use of an automated system is largely uncontroversial and, while there have been occasional malfunctions, this does not cause a significant risk to the rights and interests of individuals in the general public.

Perhaps of more concern is where automated systems are implemented without human oversight in realms of significant consequence to the individual concerned. For instance, in 2018 the Parliamentary Joint Committee on Human Rights raised concerns about the use of automation to make decisions about who is a 'non-citizen' and application of the 'public interest test' under the *Migration Act 1956* (Cth).²³ In the Parliamentary Joint Committee on Human Rights Report 7 of 2018, the committee noted:

it appears that under the 2018 instrument some matters which could be subject to decision by computer program may involve complex or discretionary considerations. Specifically, for the minister to determine whether a person is an 'eligible non-citizen' involves a decision as to whether the minister thinks such a determination would be in the 'public interest'. By contrast, it is noted that, in relation to other provisions of the Migration Act that involve consideration of the 'public interest', the Migration Act has exempted such determinations from being 'designated migration law' (that is, the decision cannot be made by computer). It is unclear why subsection 72(2)(e) of the Migration Act is not similarly exempted from the 'designated migration law' or excluded from the 2018 instrument.²⁴

The Minister for Immigration responded by clarifying that the Minister's personal decision-making powers are not automated through departmental computer programs and that the computer program could only *grant* the relevant visa — it could not make a decision to *refuse*. The Minister clarified that:

21 Anna Huggins, 'We Need Human Oversight of Machine Decisions to Stop Robo-debt Drama', *The Conversation* (online, 2 July 2019) <<https://theconversation.com/we-need-human-oversight-of-machine-decisions-to-stop-robo-debt-drama-118691>>.

22 Australian Government, *Automated Assistance in Administrative Decision-Making: Better Practice Guide* (2007) 4.

23 Parliamentary Joint Committee on Human Rights, *Human Rights Scrutiny Report* (Report 7 of 2018); Parliamentary Joint Committee on Human Rights, *Human Rights Scrutiny Report* (Report 11 of 2018).

24 Parliamentary Joint Committee on Human Rights, *Human Rights Scrutiny Report* (Report 7 of 2018) 7–8, [1.40]–[1.41].

In instances where the online application 'hits' against risk systems, or where binary responses provided by an applicant do not support an immediate auto-grant decision, the computer program will refer the BV [bridging visa] application to a departmental decision maker to manually decide upon the application. The computer program is designed to grant BVs in association with substantive applications in the majority of straightforward cases. Instances in which the BV application cannot be immediately granted by the computer program, including where there are public interest considerations, are always considered by a delegate or the Minister.²⁵

This position is still maintained by the Department of Home Affairs.²⁶ It does not seem far-fetched, however, that automated decision-making can and will extend to deny rights, such as an automatic visa cancellation for not being of good character.²⁷ In such high-stakes cases, there should arguably be a human that is charged with overseeing an automated system recommendation and, following careful consideration, decides to affirm or reject the recommendation. This human intervention is critical to minimise risk and to provide a degree of accountability for automated decisions so that affected individuals may have a right to challenge the decision or, at a minimum, receive an explanation as to the reasoning behind an administrative decision. This contention is discussed further below.

However, there are definitive challenges with respect to machine learning. One such challenge is the nature in which machine learning is to be trained. The human mind is a complex and barely understood machine. Being able to crack how the human mind learns and being able to apply that process to algorithms is an ongoing challenge. A person is uniquely designed to be able to take a small piece of data and be able to extrapolate that data to identify similarities. If a child was shown a picture of a kangaroo (even a non-realistic one), they can either immediately or quickly learn to identify another example of a kangaroo. Current machine learning algorithms require large amounts of verified data to be able to do the same thing to the degree of a child. This creates a problem with designing algorithms — there is a need to have large datasets, and a human to verify that dataset, to properly train and verify the output of the algorithm.

The previously mentioned supervised learning is an intensive method of teaching an algorithm as there are four classifications²⁸ of the algorithm identifying inputs:

- **True positive** — correct identification of a correct input
- **True negative** — correct identification of an incorrect input

25 Parliamentary Joint Committee on Human Rights, *Human Rights Scrutiny Report* (Report 11 of 2018) — responses from legislation proponents, [1.43].

26 Australian Government, Department of Home Affairs, 'The Administration of the Immigration and Citizenship Program' (Background Paper, 4th ed, February 2020) 28 [173] states: 'Importantly, no adverse visa decision is ever made by a machine. ... The officer might be prompted and assisted by the latest technology and automated analytical tools, but it is a person who will be the decision-maker'. This so-called golden rule is discussed in Jake Goldenfein, 'Algorithmic Transparency and Decision-Making Accountability: Thoughts for Buying Machine Learning Algorithms' in Cliff Bertram, Asher Gibson and Adriana Nugent (eds), *Closer to the Machine: Technical, Social, and Legal Aspects of AI* (Office of the Victorian Information Commissioner, August 2019) 41, 48.

27 Such as those made under s 501 of the *Migration Act 1958* (Cth); see Samuel White, 'Godlike Powers: Unfettered Ministerial Discretion' (2020) 41(1) *Adelaide Law Review* 1–38.

28 Machine Learning Crash Course, 'Classification: True vs False and Positive vs Negative' (Web Page) <<https://developers.google.com/machine-learning/crash-course/classification/true-false-positive-negative>>.

-
- **False positive** — incorrect identification a correct input
 - **False negative** — incorrect identification of an incorrect input.

The false negative and positive are the key areas that must be verified by the supervisor. This is usually done through the use of validated datasets and identifying how often the algorithm produces a false response. Depending on the classification method used (such as decision tree or Bayesian²⁹) different levels of input will be required. As an example, a decision tree is simple and fast and supports incremental learning, but it requires very accurate data and a long training time to get an effective output. A machine learning technique that is trained on particular labelled datasets or data domain may not be suitable for another dataset or data domain given that the classification may not be robust over different datasets or data domains.³⁰ This means that there would need to be an algorithm developed for each area, especially if there is a narrow output required.

Kline and Kahneman created a theory on the validation of the environment when it came to being able to intuitively predict an outcome in an environment based on the regularity of variables.³¹ This idea breaks down intuition and how it can and cannot be applied to different environments. Two extremes of this scale would be firefighting and the share market. A firefighter with many years of experience can use their intuition to determine whether it is safe to enter a building or even when to stop fighting a fire. This can be based on the number of variables that determine how a fire acts — this is easier to validate and a person who experiences a large number of fires can learn to see what variables must exist to determine how it will act. This is where fire modelling is used to determine flashpoints and how a fire will act³² and can determine the action taken. This can be classified as a high validation environment. The stock market, on the other hand, would be considered low validation, as there are so many variables from the economic to human behaviour it is currently impossible to develop intuition about the market — guessing is just as accurate as experience. Applying this idea to machine learning brings up the question: can we ensure that all the variables that a human would consider can be plugged into an algorithm to give us the best decision?

Equally, bias in machine learning is a significant issue that can have long-term effects on the organisation.³³ An algorithm is developed with a particular outcome in mind, but the bias of those who develop the algorithm through the design process and how it is trained can affect

29 H Bhavsar and A Ganatra, 'A Comparative Study of Training Algorithms for Supervised Machine Learning' (2012) 2(4) *International Journal of Soft Computing and Engineering* 2231–2307 <https://www.researchgate.net/profile/Amit_Ganatra2/publication/265068741_A_Comparative_Study_of_Training_Algorithms_for_Supervised_Machine_Learning/links/5780b65f08ae9485a43ba431.pdf>.

30 S Suthaharan, 'Big Data Classification: Problems and Challenges in Network Intrusion Prediction with Machine Learning' (2014) 41(4) *ACM SIGMETRICS Performance Evaluation Review* 70–73 <http://delivery.acm.org.wwwproxy1.library.unsw.edu.au/10.1145/2630000/2627557/p70-suthaharan.pdf?ip=149.171.67.148&id=2627557&acc=ACTIVE%20SERVICE&key=65D80644F295BC0D%2EB-811333C2AA88C82%2E4D4702B0C3E38B35%2E4D4702B0C3E38B35&__acm__=1575083898-ec66f9123f4f58ee8fa5f9d7406aa94c>.

31 'Kahneman and Klein On Expertise', *Judgment and Decision Making* (Blog, 28 July 2013) <<https://j-dm.org/archives/793>>.

32 E Ronchi and D Nilsson, 'Fire Evacuation in High-rise Buildings: A Review of Human Behaviour and Modelling Research' (2013) 2(1) *Fire Science Reviews* 7.

33 Suraj Acharya, 'Tackling Bias in Machine Learning', *Insight Data Science* (online), 19 March 2019 <<https://blog.insightdatascience.com/tackling-discrimination-in-machine-learning-5c95fde95e95>>.

how it produces an output. If it is trained on only data that soldiers who commit assault are not promoted and have been in for three years then it will have a bias against soldiers who meet that criteria regardless of other variables.

There is opportunity in bias if it is deliberately introduced in a controlled way. Inductive bias can be used to help develop an algorithm that can deal with new situations. This is a human trait — we can come to a conclusion without knowing all the information about a situation. For a machine to do this, an inductive leap would need to be possible³⁴ whereby it can deliberately invoke biases for choosing one generalisation of the situation over another.

A final issue with the utilisation of machine learning is the transparency and consistency of the data that is being used. This can make it hard to challenge, especially for individuals who lack technical or hardware knowledge.³⁵ Unless you are a software engineer, how are you going to understand how the algorithm got the decision it did? The current system of human decision-making can often be perplexing, but you are able to find the individual and unpick their logic. It is much harder to unpick the logic of hundreds of lines of computer code. This creates a fear of the unknown — after all, who do you hold to account for a decision that was created by an algorithm? These are issues that deserve thorough attention, although they are outside the scope of this article.

Authorised automated decision-making under Australian administrative law

It is critical to ensure that automated tools which assist the government decision-making process are designed in a way which will ensure proper outcomes under the principles of administrative law. There are requirements under Australian administrative law that govern who may lawfully make an administrative decision. But before looking at who is permitted to make an administrative decision, it is necessary to look at just exactly what a decision is.

Pintarich and the reviewability of automated decisions

Decisions 'of an administrative character' are sometimes difficult to distinguish from legislative decision. As the Full Court of the Federal Court noted in *Federal Airports Corporation v Aerolineas Argentinas*,³⁶ general tests for characterisation of acts — as either administrative, or legislative — are unfortunately of limited utility and viability. Administrative and legislative acts can be difficult to differentiate. Accordingly, they must be characterised on *context and subject matter*. The relevance of this is that, where decisions are made by artificial intelligence (AI) and increasingly are made by AI, it might be that there is initial difficulty in the justiciability of such decisions.

34 TM Mitchell, 'The Need for Biases in Learning Generalizations' (Department of Computer Science, Laboratory for Computer Science Research, Rutgers University, 1980) 184–191 <http://dml.cs.byu.edu/~cgc/docs/mldm_tools/Reading/Need%20for%20Bias.pdf>.

35 See Australian Administrative Review Council, *Automated Assistance in Administrative Decision-making* (2004). See further M Perry and A Smith, 'iDecide: The Legal Implications of Automated Decision-making' [2014] *Federal Judicial Scholarship* 17 <www.austlii.edu.au/au/journals/FedJSchol/2014/17.html>.

36 (1997) 76 FCR 582.

It would be remiss to not reflect on the topic of government automated decisions without discussing the recent *Pintarich*³⁷ decision by a Full Court of the Federal Court of Australia. The majority found that a decision had not been made for the purposes of the *Administrative Decisions (Judicial Review) Act 1977* (Cth) because, even though an automated tool has issued an outcome letter to a taxpayer, the Deputy Commissioner had not undertaken a process of deliberation, assessment or analysis.³⁸ However, we expect that this case is likely to be distinguished or departed from over time, as government administration will increasingly rely on automated systems and decisions.

A detailed discussion of the *Pintarich* decision is outside of the scope of this article. However, Ng and O'Sullivan³⁹ have considered the *Pintarich* decision in depth. They argue that there must be a 'modern interpretation of existing administrative law principles' that does not exclude automated decisions from the scope of review under the *Administrative Decisions (Judicial Review) Act 1977* (Cth). The author agrees with the arguments of Ng and O'Sullivan — administrative decisions should not be immune from review solely because they have been delivered through a computerised medium.

Current Australian principles governing authorised decision-makers

There are clear requirements in Australian administrative law that govern who is permitted to exercise administrative power. These requirements stem from the principle of legality — the principle that government agents require a positive justification, or legal authority, for any action undertaken, particularly where such action has a detrimental effect on the rights or legal interests of an individual.⁴⁰ This principle has a number of functions, including the fostering of accountability and transparency of government action; and upholding the rule of law.

The validity of an administrative decision, in part, depends on whether it has been made by a particular person that is authorised in statute.⁴¹ The historical principle is encapsulated in the Latin maxim, *delegatus non potest delegare* — one who is vested with a statutory power must exercise it personally rather than delegate it.⁴² The underlying rationale for this principle is to ensure that only individuals whom Parliament has empowered with decision-making authority exercise the relevant public power.

The implications for failing to understand the self-imposed legislative limitations could result in an *ultra vires* decision and the decision being challenged in judicial review proceedings. Specifically, under s 5(1)(c) of the *Administrative Decisions (Judicial Review) Act*,⁴³ a decision may be challenged on the ground that 'the person who purported to make the decision did

37 *Pintarich v Deputy Commissioner of Taxation* [2018] FCAFC 79 (Moshinsky and Derrington JJ; Kerr J dissenting).

38 *Ibid* [56].

39 Ng and O'Sullivan, above n 2.

40 Perry, above n 3, 31. See also Robin Creyke and John McMillan, *Control of Government Action: Text, Cases and Commentary* (LexisNexis Butterworths, 3rd ed, 2012) 389.

41 *Re Reference under Section 11 of Ombudsman Act 1976 for an Advisory Opinion; Ex Parte Director-General of Social Services* (1979) 2 ALD 86 [458] (Brennan J).

42 See John Willis, 'Delegatus Non Potest Delegare' (1943) 21 *Canadian Business Review* 257.

43 *Administrative Decisions (Judicial Review) Act 1977* (Cth).

not have jurisdiction to make the decision'. However, in practice, as Robin Creyke observes, the modern reality in large government agencies is that it is 'seldom practical for a principal nominated in legislation to personally make all decisions'.⁴⁴ As a result, a number of principles have emerged in statute and in case law that enable the nominated authority to delegate their decision-making power.

The case law that governs permissible delegations of administrative power recognise a distinction between *express* delegations and an *implied* authority to delegate. In relation to express delegations, the validity of an administrative decision turns upon 'the identity of the authority and the doer of the act'.⁴⁵ Typically, a statute will expressly provide that the principal decision-maker is authorised to delegate the power to another person,⁴⁶ and this is executed by way of a written delegation instrument.⁴⁷

Implied authorities to delegate may arise where the power is not expressly delegable, but the nominated person 'could not have been expected by the Parliament to have exercised it personally in the multitude of instances where its exercise would be required'.⁴⁸ Immigration matters come to mind here — particularly instances where a visa is to be revoked.⁴⁹ The key question in relation to implied delegations is whether a power consigned by statute to an authorised person requires the power to be personally exercised by its designated repository or a delegate.⁵⁰ This is typically a matter of statutory construction⁵¹ and depends on the nature of the power itself and other circumstances of the case.⁵² Such circumstances may include that Parliament could have assumed knowledge of a 'practical administrative necessity' that required the power to be delegated⁵³ or whether the power in question would adversely affect the rights of individuals.⁵⁴

The current approach to authorising automated government decisions

A key problem exists with the application of these delegation principles to automated systems, particularly where an automated system has been deployed in a way that it makes decisions with human input. Specifically, there is no human 'decision-maker' to whom an express delegation may take place.⁵⁵ To navigate this state of affairs, in the early 2000s the legislature sought to pre-empt this issue⁵⁶ by deeming any decision made by a computer

44 Creyke and McMillan, above n 40, 456.

45 *Re Reference under Section 11 of Ombudsman Act 1976 for an Advisory Opinion; Ex Parte Director-General of Social Services* (1979) 2 ALD 86, 93 (Brennan J).

46 See, eg, *Migration Act 1958* (Cth) s 496.

47 Creyke and McMillan, above n 40, 456.

48 *Carltona Ltd v Commissioners of Works* [1943] 2 All ER 506, 563 (Lord Greene MR); *Re Reference under Section 11 of Ombudsman Act 1976 for an Advisory Opinion; Ex Parte Director-General of Social Services* (1979) 2 ALD 86, 93 (Brennan J).

49 *Migration Act 1958* (Cth) s 501.

50 *Pattenden v Commissioner of Taxation* [2008] FCA 1590 [42] (Logan J).

51 *O'Reilly v Commissioner of State Bank of Victoria* (1982) 153 CLR 1 [10] (Gibbs CJ).

52 *Ibid*; *Re Reference under Section 11 of Ombudsman Act 1976 for an Advisory Opinion; Ex parte Director-General of Social Services* (1979) 2 ALD 86, 93 (Brennan J).

53 *Pattenden v Commissioner of Taxation* [2008] FCA 1590 [42] (Logan J).

54 *O'Reilly v Commissioner of State Bank of Victoria* (1982) 153 CLR 1 [12] (Gibbs CJ).

55 Perry, above n 3, 31.

56 See, eg, *Migration Legislation Amendment (Electronic Transactions and Methods of Notification) Act 2001* (Cth).

to be one made by a senior member of the executive, typically either a minister or the departmental secretary.⁵⁷ This is in accordance with the principle of responsible government and places the risk of any decision on the relevant minister. There are an increasing number of legislative provisions that delegate authority for decision-making to computers.⁵⁸ For example, s 495A of the *Migration Act 1956* (Cth) provides that:⁵⁹

(1) The Minister may arrange for the use, under the Minister's control, of computer programs for any purposes for which the Minister may, or must, under the designated migration law:

(a) make a decision; or

(b) exercise any power, or comply with any obligation; or

(c) do anything else related to making a decision, exercising a power, or complying with an obligation.

(2) The Minister is taken to have:

(a) made a decision; or

(b) exercised a power, or complied with an obligation; or

(c) done something else related to the making of a decision, the exercise of a power, or the compliance with an obligation;

that was made, exercised, complied with, or done (as the case requires) by the operation of a computer program ...

The *Migration Act 1958* (Cth) further provides that, where the computer malfunctions, the minister may substitute a computer-based decision with one that is more favourable to the applicant.⁶⁰ Ng and O'Sullivan observe that these provisions constitute an 'accountability structure that involves an individual', which might be interpreted as an intent on behalf of Parliament to preserve review rights in relation to automated decisions.⁶¹ Furthermore, the Administrative Review Council found that, where a computerised system is simply used as a decision support tool for a human officer who makes the actual decision, it would appear that such a legislative authority is not necessary.⁶²

Provisions such as s 495A of the Migration Act provide a legal authorisation⁶³ for the use of automated computer systems in the decision-making process. Specifically, the provisions grant the minister with a power to arrange for the use of computer programs for the making of a decision, where that takes place under the minister's control. This is despite the aforementioned golden rule.

57 See, eg, *Social Security (Administration) Act 1999* (Cth) s 6A, which provides that '[t]he Secretary may arrange for the use, under the Secretary's control, of computer programs for any purposes for which the Secretary may make decisions under the social security law' and that a 'decision made by the operation of a computer program under an arrangement made under subsection (1) is taken to be a decision made by the Secretary'.

58 See *Social Security (Administration) Act 1999* (Cth) s 6A, (deems such computer-generated decisions to be made by the Secretary); or the *Therapeutic Goods Act 1989* (Cth) s 7C.

59 *Migration Act 1956* (Cth) s 495A.

60 *Ibid* s 495B.

61 Ng and O'Sullivan, above n 2, 31.

62 Administrative Review Council, above n 35, pt 5.1.

63 *Administrative Decisions (Judicial Review) Act 1977* (Cth) s 5(1)(d).

However, it may be uncertain whether these provisions provide an authorisation for the use of automated systems where a minister or secretary does not have ‘control’⁶⁴ over a *specific* decision — particularly where there is no human input before the decision is communicated to the individual and where the system relies on adaptive machine-learned logic. The term ‘under the minister’s control’ could be interpreted as requiring that the minister have either *organisational* control over the department that implements the automated system or control over how an automated system approaches a *particular decision*, and exercise of power or an obligation. This would be a matter of statutory interpretation for the courts.

When undertaking an exercise of statutory interpretation, a court will seek to give the words of a statutory provision the meaning that the legislature is taken to have intended them to have.⁶⁵ This will principally involve a consideration of the literal meaning of the text but may also require consideration of the context and purpose of the provision.⁶⁶ Unfortunately, the Explanatory Memorandum relevant to s 495A of the Migration Act does not indicate what level of control Parliament may have intended a minister to hold over computer-based decision-making systems.⁶⁷ Furthermore, as the Bill was passed in 2000, it is unlikely that the legislature contemplated that automated systems could make decisions on the basis of an adaptive logic, as is the case with the current generation of machine learning based systems that have become widespread over the past five years. In light of the above, there is perhaps a risk that the implementation of a machine learning system to make administrative decisions, without any human input, could be invalidated on the basis that it is not used ‘under the minister’s control’.

In interpreting the requirement for computerised decision-making systems to be ‘under the minister’s control’, a possible approach the courts could adopt would be to require a level of control that is *proportionate* to potential harms of the decision being made. In particular, a higher level of control and oversight should be due for decisions that are likely to have a major impact on the rights and interests of individuals (for example, a criminal sentencing decision). This higher level of control could constitute a human approval of a decision by a government officer or a specific amount of testing to validate the systems’ results and legal alignment with statute.

Equally, as found within the *Therapeutic Goods Act 1989* (Cth), it might be that the oversight of the relevant human — in this case, the secretary — is statutorily barred from correcting a decision made by a computer:

7C Secretary may arrange for use of computer programs to make decisions

- (1) The Secretary may arrange for the use, under the Secretary’s control, of computer programs for any purposes for which the Secretary may make decisions under this Act or the regulations.
- (2) A decision made by the operation of a computer program under such an arrangement is taken to be a decision made by the Secretary.

64 Ng and O’Sullivan, above n 2, 31.

65 *Project Blue Sky Inc v Australian Broadcasting Authority* (1998) 194 CLR 355, 384.

66 *Ibid.*

67 Explanatory Memorandum, Migration Legislation Amendment (Electronic Transaction and Methods of Notification) Bill, 14 [64]–[69].

(3) The Secretary may substitute a decision (the *substituted decision*) for a decision (the *initial decision*) made by the operation of a computer program under such an arrangement if the Secretary is satisfied that the initial decision is incorrect.

(4) However, the substituted decision may only be made before the end of the period of 60 days beginning on the day the initial decision is made.

Sixty days is perhaps sufficient time currently for a workload managed by humans, at a set rate of possible decisions a day. However, when automated, this workload can infinitely expand, and there is significant risk in maintaining a human-centric time frame to supplementing decisions, when the statute clearly envisages automated decision-making capability.

Current approach: risks and fixes

The current ‘deemed decision’ approach poses a number of risks, including:

- (a) that there is lack of clear accountability for automated decisions, which may erode the rule of law in Australia; and
- (b) that this lack of accountability may lead to organisational and national risks for improperly supervised automated decisions.

A drawback of the current approach is that it offers a limited level of accountability for automated decisions, as well as poor transparency as to how an automated decision has been reached. Both limbs have the potential to erode fundamental aspects of the rule of law.⁶⁸ To the first limb: if an automated government decision is deemed to be made by the secretary or minister, who may or may not be aware of a particular decision, it is not clear who holds ownership of that decision. A lack of clear ownership can have multiple socio-legal effects. First, the decision-maker may be incorrectly identified, making review and appeals difficult for the affected individual (thereby reducing confidence in the system and raising discontent with government). Secondly, even when the decision-maker is identified, if the reasons given for the decision are not in a manner that is intelligible, finding grounds of appeal or jurisdictional error might also be impracticable — indeed, it might take machine learning to identify jurisdictional errors in the future. Even though a decision made by an algorithm has the potential to, on its face, be *more* transparent (in that the actual reasons, relevant and irrelevant considerations included, will be visible) it still remains a large point of risk. One solution which seems rather common sense is to ensure that reasons given by algorithms are in a manner that is comprehensible — ‘plain English’ in its most basic meaning.

Furthermore, leaving the review or overturning of automated decisions to a high-ranking government official such as the secretary or minister is unlikely to be a scalable approach in the long term, particularly as an increasing number of decisions in government departments are processed by automated means. It might that, just as implied delegation has come to be accepted, implied review by algorithms may have to occur. Such a process is known as adversarial machine learning — although it is similar to the idea of ‘let us make a super super AI to watch this super AI’. An easier process is to create a ‘dumb AI’ system that randomly

68 Zalnierute, Moses and Williams, above n 11.

audits the decisions made. It does not seek to improve the process but simply, like a random number generator, takes decisions made and offers them for human approval.

A preferable approach could be to implement, by statute, a requirement that certain government officials hold a delegable obligation of ‘independently justifying automated decisions’,⁶⁹ where a provision requires discretionary decision-making and is likely to have a high impact on the rights and interests of individuals. Such an approach could have a number of benefits. First, it would ensure that an accountable person reviews the output of an automated system and may provide a check and balance against systems that may be malfunctioning or may be operating with improper training data (that is, data bias). More broadly, it would confine and structure the use of automated systems that can encroach on fundamental rights and interests⁷⁰ and promote the ongoing consideration of whether a provision is amenable to be automated and whether the system is of a suitable quality to be used in that case. Finally, it would have an additional benefit of providing an avenue for transparency, in that the human reviewer would have the opportunity to provide reasons for the decision being made.

Without this, there is potential that the current approach cannot assure an appropriate level of oversight over automated decisions. There are both organisational and national risks to a lack of accountability and oversight. At an organisational level, there is the potential for a department to suffer reputational harm, as well as financial harm, if automated decisions are executed on a widespread scale without sufficient oversight in a department. At a national level, there is the possibility that there is a loss in trust in government in a broader sense.

Conclusion

Under such an esoteric maze of uncertainties, it is not unreasonable that governmental accountability and transparency will be lost. Without clear authorisations and structure in the legislation, it is likely that individuals will, in the face of an apathetic Leviathan (as Hobbes would know it), simply accept a decision rather than believe they can contest it. This was by all accounts the logic of many individuals affected by Robodebt.

What is the solution then? There are many, and looking to other countries is perhaps the best option. In Sweden, it is clear that ‘a human (must) confirm and take responsibility for each decision’. This reflects not only the European Union’s recognised right to not be subject to automated decision-making but would also appear an approach most in line with the Australian Government’s ‘golden rule’. Leaving legislation open-ended in a bid to preserve flexibility has some benefits in areas that require it due to their quick evolution — such as pandemic responses. Administrative decision-making, with its emphasis on the individual, is not one.

69 Ibid 445.

70 Creyke and McMillan, above n 40, 466. The comments of Professors Creyke and McMillan were made in relation to the arguments for public law delegation principles; however, the author contends that these arguments would also be relevant for the development of clearer accountability principles for automated decisions.