

Automated decision-making in (good) government

*Bernard McCabe**

The Snowy Mountains Hydro-electric Authority introduced government to the potential of computing in 1960. 'Snocom', the authority's first computer, was built with the assistance of academics from the University of Sydney. It was used to perform the complex mathematical calculations required in the design and engineering of the Snowy scheme.¹

Snocom was a transistor-based computer — one of the first of its kind in the world. Its magnetic drum memory held up to 2,048 words. It had the equivalent of 8 kilobytes of RAM.² It was essentially a bulky calculator. Yet it represented an important scientific advance. The Authority's computers (starting with Snocom) helped it construct the enormously complex scheme on time and, famously, under budget.

Modern computers are vastly more powerful than Snocom, and they are ubiquitous. Their capacity is turbocharged by networking, the internet and the cloud. Smart devices with user-friendly applications are a feature of daily life. Governments have taken to the new technology more slowly than some businesses, but computers and computerised processes have had profound effects on public administration. Public servants now use computerised information systems every day in their work.

The evolution of information technology and, more recently, the advent of artificial intelligence systems that 'learn' and adapt holds out the possibility of faster, cheaper, quicker and more accessible public services. That is promising, but administrative lawyers are yet to resolve all the potential obstacles to making these new decision-making processes accountable through Australia's system of administrative law.

Administrative lawyers have their work cut out, but the challenge goes further than merely ensuring accountability. The ultimate objective must be to promote good government. Policy-makers need to understand how automated processes can be accommodated within a framework of values concerning public administration. The task is not simply legal or technical. It requires a broader perspective on the way government is supposed to work, and — to some extent — what government is intended to do. That is an ambitious task. This article offers a modest contribution to the debate.

The article begins with a discussion of the concept of good government. The report of the Administrative Review Council (the ARC) titled *Automated Assistance in Administrative*

* Bernard McCabe is Deputy President and Head of the Taxation and Commercial and the Small Business Taxation Divisions of the Administrative Appeals Tribunal.

1 G Philipson, 'The 1950s and 60s Were Busy Times for Building Home-grown Computers' (online, 13 February 2017) <<https://ia.acs.org.au/article/2017/acs-heritage-project--chapter-9.html>>. During the 1950s, before Snocom was developed, the Authority had used more primitive devices, including the SILLIAC system, that were developed by the University of Sydney and the forerunner of the CSIRO.

2 T Wright, '40 Years Later, the Snowy Mountains Scheme is Back in Vogue', *Sydney Morning Herald* (online, 16 March 2017) <<https://www.smh.com.au/politics/federal/40-years-later-the-snowy-mountains-scheme-is-back-in-vogue-20170316-guzjis.html>>. For the purposes of comparison, the Samsung Galaxy S20 mobile phone has up to 512GB of memory and 12GB of RAM.

Decision Making is an important reference point.³ The report was published in 2004 — several lifetimes ago in cyber terms. Yet the insights in the report remain fresh because they are grounded on a clear conception of good government. The report identified five values as crucial elements of all good decision-making processes. The report also proposed a series of principles that should inform the design and implementation of automated processes. The principles were derived from the values.

I use the five values (and the related principles) identified by the ARC as a framework for discussing and evaluating a number of specific decision-making processes. In doing so, I consider some of the implications of automation — and where it can go wrong. The article concludes with a modest call to reinforce the central role of the Administrative Appeals Tribunal (AAT) in promoting good government in the Commonwealth's use of automated decision-making processes.

Good government

Section 51 of the *Constitution* provides that the Parliament has 'the power to make laws for the peace, order, and good government of the Commonwealth' with respect to a number of identified subjects. As a matter of constitutional law, the reference to 'peace, order, and good government' is not intended to create a limitation on the Parliament's legislative powers. There is no suggestion a court can evaluate a particular law to determine whether the law *actually* promotes the public interest according to some objective standard.⁴ And yet the expression reflects an aspiration for government that embodies concepts of sound public administration.

The Australian conception of good government has been driven by forces particular to this country. Human rights legislation has shaped the jurisprudence in relation to the role of government in many other modern western democracies. Australia does not have entrenched human rights legislation. Of course, Australian courts and the Parliament have not been blind to developments overseas. The Commonwealth is a signatory to a range of treaties that impose human rights obligations. But our courts and the Parliament do not have the lodestar of a bill of rights to guide the development of rules regulating public administration. They have had to make their own way in developing a system of administrative law (including administrative law doctrines) that ensures sound public administration.⁵

At the same time, the Australian constitutional settlement has required lawmakers to think more deeply about alternatives to judicial review in circumstances where a strict separation of powers provides for a more constrained role for the courts. That reflection prompted a burst of reform in administrative law in the 1970s and 1980s that commenced with the establishment of the Kerr committee⁶ and included the establishment of the Ombudsman,

3 Administrative Review Council, *Automated Assistance in Administrative Decision Making* (Report No 46, 2004).

4 *Union Steamship Company of Australia Pty Ltd v King* [1988] HCA 55 [16] (Mason CJ; Wilson, Brennan, Deane, Dawson, Toohey and Gaudron JJ).

5 See eg Yee-Fui Ng et al, 'Revitalising Public Law in a Technological Era: Rights, Transparency and Administrative Justice' (2020) 43 *University of New South Wales Law Journal* 1, 2.

6 Commonwealth, *Report of the Commonwealth Administrative Review Committee* (Parliamentary Paper No 144, 14 October 1971) 8.

the AAT, the ARC and the Federal Court, and the passage of the *Administrative Decisions (Judicial Review) Act 1978* (Cth) (ADJR Act) and the *Freedom of Information Act 1982* (Cth). These innovations created rights (such as the right in the ADJR Act to be provided with reasons for a decision) that were not generally recognised in the common law. They also created new rights of review with more flexible and accessible remedies.

The AAT was arguably the most striking of these innovations. The AAT was conceived as a general merits review tribunal that would permit a more wide-ranging review of government decision-making. While the courts were generally confined to addressing the lawfulness of government action given the separation of powers in the *Constitution*, the AAT was intended to reconsider the merits of individual reviewable decisions.

From the outset, the AAT was conceived as a mechanism for achieving administrative justice for individuals. But it also has a larger, systemic role. As Sir Gerard Brennan explained, the AAT was established on a judicial model so that it would have the capacity to generate clear, well-reasoned and objective decisions that established norms. Those norms could then be applied to decision-makers in the agency concerned, and elsewhere in the bureaucracy.⁷ In that sense, the AAT serves the cause of good government by promoting good decision-making across government. It does that by modelling good decision-making behaviour in the individual reviewable decisions that come before it.⁸

The AAT explicitly embraced its mission as a tool of good government from the time of its earliest decisions. In *Drake and Minister for Immigration and Ethnic Affairs (No 2)*,⁹ for example, Brennan J wrote perceptively about the proper place of policy in administrative decision-making. The President articulated a vision of orderly, objective and well-informed decision-making that was carefully attentive to the rights of individuals and conscious of the need for consistency and economy in a modern bureaucracy.

While the AAT was arguably the most interesting of the reforms enacted during the period, the role of the ARC was also important. The ARC served as a repository of expertise about good government and sound public administration. Its statutory remit envisaged it monitoring and reporting on government decision-makers and making recommendations for improvement.

This brings me to the ARC's 2004 report on automated assistance in decision-making processes. In the covering letter to the Attorney-General which accompanied the report, the President of the ARC referred to the 'potential for cost savings, efficiencies and greater accuracy in decision-making' that suggested automated computer systems that assisted administrators 'will become increasingly important tools of government'.¹⁰

7 Sir Gerard Brennan, 'The AAT — Twenty Years Forward' (Speech, Administrative Appeals Tribunal Twentieth Anniversary Conference, Canberra, 1 July 1996). See also G Brennan, 'The Future of Public Law — The Australian Administrative Appeals Tribunal' (1979) 4 *Otago Law Review* 286.

8 See eg *RBPk and Innovation And Science Australia* [2018] AATA 404 [11] (Thomas J; DP McCabe); see also B McCabe, 'Perspectives on Economy and Efficiency in Tribunal Decision-making' (2016) *AIAL Forum* 40, 45–6.

9 [1979] AATA 179.

10 The letter addressed to the Attorney-General which accompanied the report is dated 12 November 2004.

The report began with a discussion of the crucial elements of the administrative law system. The ARC identified five values that should be observed in the design and operation of all administrative decision-making processes. They are:

- lawfulness;
- fairness;
- rationality;
- openness or transparency; and
- efficiency.¹¹

Those values were themselves informed by concepts of administrative justice which include the ‘four basic requirements for just decision-making in a society governed by the rule of law’ identified by French J (as he then was) in 2001. These requirements were lawfulness, fairness, rationality and intelligibility.¹² His Honour said those values were evident in the various grounds of judicial review.

Justice French argued the requirements of lawfulness and fairness were covered by the ‘error of law’ and ‘procedural fairness’ grounds of review. The requirement of rationality was addressed by the obligation to take account of relevant considerations and ignore irrelevant considerations. The requirement of intelligibility was addressed in the requirement (arising under the ADJR Act, at least) that the decision-maker give proper reasons for the decision.¹³

The values identified by the ARC overlapped with those derived from the values evident in the decisions of the courts as they conducted judicial review. The two sets of values are not exactly the same. The view from the bench is different from the perspective of the ARC given the constitutional constraints on the courts. Justice French recognised as much. He pointed out other values or factors might make an important contribution in securing the larger objective of administrative justice. He noted that education of administrators on the precepts of administrative justice was crucial¹⁴ — although one might interpolate that sound public administration is more likely if administrators have been provided with *all* the training and education (along with the resources) they need to do their jobs. His Honour also referred to the importance of internal and external review mechanisms and the role of parliamentary scrutiny and the media.¹⁵ He also acknowledged the pressures facing a modern bureaucracy which deals with an enormous workload. With ‘the practical realities of administrative decision-making’ in mind, his Honour acknowledged the importance of attributes like accessibility, affordability and timeliness when evaluating the success of a system of administrative justice.¹⁶

11 Administrative Review Council, above n 3, 3.

12 RS French, ‘Judicial Review Rights’ (2001) 28 *AIAL Forum* 28, 30.

13 *Ibid* 33.

14 *Ibid* 35.

15 *Ibid*.

16 *Ibid* 34.

The ARC's perspective is necessarily wider than that. It is not simply concerned to promote administrative justice, although justice is of central importance. The ARC is concerned with good government and sound public administration. The five values it identified in the report embody a rich trove of thought on the subject that is applicable to all decision-making processes, not just those that are wholly or partly automated.

Corralling the technology to focus the debate

We will return to the values proposed by the ARC below. But first we must engage with some of the details of the ARC's report — in particular, the way in which the report approached the technological and conceptual issues involved in automated decision-making processes. That is important because the ARC report developed a series of principles that were to be used to guide the development and implementation of automated processes. The design of those principles was inevitably informed by the state of technology (and what was foreseeable) at the time.

The ARC distinguished between information technology systems that *support* decision-makers as they make decisions (including data storage and file management systems; word processing technology; calculators; and what are now known as smart forms and intelligent systems that included prompts to ask questions, seek information or apply rules) and 'expert systems'. The report explained an expert system is a computer that is programmed using a set of rules which enables it to mimic (or establish a proxy for) the thought processes of human experts as it processes data.¹⁷ A *legal* expert system was one that performed the role of a lawyer or other person with specific legal expertise.¹⁸

Most legal expert systems in contemplation at the time of the report were rule-based systems that assisted human decision-makers. They were programmed to follow an orderly process that reflected the legislative requirements. The expert system typically operated by proposing standard questions designed to elicit relevant information. The system might also include prompts on interpretation or policy issues to assist the decision-maker. The machine might also include links to explanatory material or suggest standard wording.

The report anticipated the power of modern apps and smart forms used in web portals. It pointed out expert systems in use at the time might increasingly work online. Importantly, the report also assumed some of these systems might draw conclusions without the intervention of a human decision-maker.¹⁹

The report acknowledged the potential for the development of expert systems which learned from data. Those systems isolate, assimilate and apply rules inferred from data points. This kind of expert system was not necessarily programmed with all the legislative rules: the systems were able to observe or infer the existence of the rules from the data, and learn from that experience. The ARC questioned whether these systems would be as useful as rule-based systems, and it doubted whether a 'neural network' or other form of artificial

¹⁷ Administrative Review Council, above n 3, 5.

¹⁸ *Ibid.*

¹⁹ *Ibid* 6.

intelligence could ever satisfy the requirement that decision-makers provide intelligible reasons for their decisions. But those systems might still support human decision-makers.²⁰

The ARC correctly anticipated developments in artificial intelligence and machine-based learning that have revolutionised commerce. Modern tech firms like Google, Amazon and Facebook use sophisticated algorithms and large datasets in search of patterns that might reveal useful information about trends and consumer preferences. They use that data to target advertising and other services. Pizza delivery chains are able to predict deliveries (and accurately devise staffing rosters in local stores) weeks into the future by analysing data from past sales that is cross-referenced with data about television programming, sporting events and other variables that might not be taken into account by human observers.²¹ In medicine, computerised processes play an increasingly important role in diagnosis. The technology is also useful in research: by quickly and rigorously sorting data, a computer can identify hitherto unseen connections between events or phenomena, predict events that were once thought to be random, or posit the existence of things that are yet to be discovered.

Those developments have had their analogue in the public sector. The Australian Taxation Office (ATO) is able to operate a vast self-assessment system which is kept in check by the ever-present threat of audit. The decision to audit is informed by sophisticated taxpayer profiles developed from vast amounts of data about the way taxpayers conduct their affairs. Services Australia, the entity that operates Centrelink, uses data analysis and data-matching processes to reduce the incidence of fraud and reporting. Data supplied to Centrelink can be verified against information supplied to other agencies. In each case, the computerised support system can prompt decisions (to commence an audit and potentially issue an amended assessment in tax matters, or to prompt a review and potentially raise a debt in social security matters) as well as provide information used in the decision-making process itself.

The distinction the ARC made between support and expert systems is becoming harder to maintain in the face of these technological developments. The evolution of ever more powerful computers and new ways to collect, store, process, analyse and share data has given administrators access to valuable tools. When combined with computerised techniques like decision trees that include prompts, templates and information from knowledge banks, 'support' systems are playing a larger, more direct role in many decision-making processes. That balance between automated processes and the humans that administer them might shift further as the humans are de-skilled. The ARC was aware of this risk: Principle 16 provided that officers using expert systems should continue to receive training so they could understand the relevant legislation and explain outcomes.²² But it is unclear whether that is occurring. With machines undertaking more administrative work, the humans who administer the system are likely to become less engaged in the specifics of the decision-making process. They may be rendered incapable of explaining (let alone

20 Ibid 9.

21 J Davidson, 'AI Tells Domino's When You Will Want a Pizza with Uncanny Accuracy', *Australian Financial Review* (online, 31 August 2020) <<http://www.afr.com/technology/ai-tells-domino-s-when-you-will-want-a-pizza-with-uncanny-accuracy-20200826-p55pl5?btis>>.

22 Administrative Review Council, above n 3, 40–2.

second-guessing) the outcomes of the automated process as they become dependent, even if they retain the formal ability to override the process.²³

The ARC's report must be read with all those caveats about technological developments in mind. Given all the change that has come to pass, it is likely to be more useful to focus on the values identified by the ARC. But the principles are still valuable. The principles were set out at the front of the report in a section titled 'Best practice principles for automated assistance in administrative decision-making'. I will refer to particular principles below as I discuss each of the five values and their application to particular decision-making processes.

Lawfulness

The use of an automated process in decision-making must be lawful. Where the automated process in question is clearly a support system that does not go to the heart of the decision-making, that is unlikely to be an issue — although the decision-maker must still ensure the system (such as it is) meets privacy and confidentiality rules.²⁴ More complicated issues arise where the automated process plays a more direct role in the decision so that the line becomes blurred between support system and decision-maker — or where the system itself becomes the de facto or de jure decision-maker.

Lawfulness might be an issue in two ways. The first relates to the conceptual question of whether it is legally possible for a computer to make a decision — and, if so, which ones? The second arises in relation to the operation of particular decision-making processes. Does the particular process apply the correct legal rules and do its job in a way that meets the standards imposed under administrative law?

Lawfulness in conception

To what extent can a computerised decision-making process be authorised formally to make a decision? As Perry J pointed out in one of a series of useful articles and addresses on this general topic:

It cannot be assumed that a statutory authority vested in a senior public servant which extends by implication to a properly authorised officer, will also extend to an automated system; nor that authority to delegate to a human decision-maker will permit 'delegation' to an automated system. Authority to use such systems should be transparent and express.²⁵

The ARC report in 2004 doubted whether it would be legally valid to delegate the formal decision-making power to a computer system. It pointed out both the ADJR Act and the *Administrative Appeals Tribunal Act 1975* (Cth) (AAT Act) assumed the decision-maker in

23 Ibid. Principle 17 dealt with the desirability of having officers available to step in if the automated process fails. While some organisations almost certainly retain that capacity, there is good reason for doubt whether the principle is universally observed in circumstances where automated decision-making processes have achieved such a level of penetration in many organisations.

24 Ibid 28, 30. Principle 8 provides: 'The people responsible for constructing an expert system must ensure that it is compatible with their agency's privacy obligations.'

25 M Perry, 'iDecide: Administrative Decision-making in the Digital World' (2017) 91 *Australian Law Journal* 29, 31.

each case was a person.²⁶ (In footnote 56, the report noted ss 5 and 6 of the ADJR Act and s 25 of the AAT Act ‘make reference to “his or her” when discussing the person exercising the decision-making power’, which was taken to indicate an expectation that the decision-maker would be human. The report also noted the relevant provisions of the *Acts Interpretation Act 1901* (Cth) referred to the power to delegate to a person holding, occupying or performing the duties of an office or position — the suggestion being that the person in question was a natural person.)

Section 6A of the *Social Security (Administration) Act 1999* (Cth) is an example of an authorising provision that deals with problem by using the device of a deeming provision. That device — at least notionally — leaves a natural person at the heart of the decision-making process. The section provides:

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 the social security law.
- (2) A decision made by the operation of a computer program under an arrangement made under sub-section (1) *is taken to be a decision made by the Secretary.* [Emphasis added.]

Provisions like this one are becoming more common, but they still raise questions. Perry J pointed out in her extrajudicial writing:

Nonetheless such deeming provisions require acceptance of highly artificial constructs of decision-making processes. More sophisticated approaches may need to be developed as these issues come to be litigated in the courts and these provisions fall to be considered.²⁷

Taking up this theme, Bateman and others have referred to a decision-making paradigm that contemplates the exercise of *human* judgment. Bateman, quoting Professor Lon Fuller, points to ‘human cognition’ as a key element in any system of administrative justice precisely because it ‘stopped law from producing gross injustice by the ridged [sic] application of unyielding logical stipulates’.²⁸ To put the argument differently, a decision made by a human (assuming, of course, the human is not a psychopath) is more likely to take proper account of the human consequences of that decision, even if subliminally. At any rate, Bateman and others argue the administrative law paradigm is constructed on the assumption a natural person will *at some point* take responsibility for the decision that has been made and be held accountable through the usual administrative law remedies.²⁹

Bateman acknowledges it is possible to treat inanimate objects (or things that have no physical existence at all, like a modern company) as legal actors with capacity. Those legal entities can incur obligations and be held responsible for their own actions or for actions taken in their names. In those circumstances, Bateman concedes there is nothing to prevent the Parliament from conferring legal personality on an algorithmic system. But he goes on to argue:

26 Administrative Review Council, above n 3, 19.

27 Perry, above n 25, 31.

28 W Bateman, ‘Algorithmic Decision-making and Legality: Public Law Dimensions’ (2020) 94 *Australian Law Journal* 520, 530.

29 *Ibid.*

Orthodox principles of public law would, by default, limit attempts to sever the exercise of statutory powers from an immediate exercise of *human* judgment.³⁰

Bateman relies on the High Court's decision in *Graham v Minister for Immigration and Border Protection*³¹ (*Graham*) in this connection. In *Graham*, the majority (Kiefel CJ and Bell, Gageler, Keane, Nettle and Gordon JJ) observed:

all power of government is limited by law. Within the limits of its jurisdiction where regularly invoked, the function of the judicial branch of government is to declare and enforce the law that limits its own power and the power of other branches of government through the application of judicial process and through the grant, where appropriate, of judicial remedies.³²

Those remedies — the prerogative writs — are referred to in s 75(v) of the *Constitution*. That subsection of the *Constitution* refers to the original jurisdiction of the High Court in relation to writs 'sought against *an officer* of the Commonwealth' (emphasis added). The remedies are well adapted to enforcing the administrative law obligations of *individual* decision-makers. It is less clear how well the remedies are adapted to dealing with a computer program that does not (a) engage in an identifiable conventional reasoning process; or (b) provide intelligible reasons for its decision which explain what the automated decision-maker regarded as relevant and determinative. There must also be *some* doubt over whether deeming that decision to be the decision of a natural person — such as the Secretary or other designated officer — addresses the challenge when the deemed decision-maker has not *in fact* turned his or her mind to the question at hand.

Questions over the efficacy of non-human decision-making have increased after the decision of the Full Federal Court in *Pintarich v Deputy Commissioner of Taxation*³³ (*Pintarich*). In *Pintarich*, a taxpayer had negotiated with an authorised officer of the ATO over a tax debt that included primary tax liability and an amount in respect of the general interest charge (GIC) which continued to accrue while the taxpayer's debt remained unpaid. In the course of negotiations, the officer agreed to accept a lump sum from the taxpayer in settlement of the debt. It was apparent the taxpayer understood the deal to include GIC, whereas the officer's notes from the meeting confirmed the question of GIC had not been resolved. The officer intended that GIC would be considered separately. The officer repaired to his office after the negotiations and entered the settlement information into the relevant system, which prepared a form letter recording the terms of the deal. The letter dated 8 December 2014 included the following sentence:

This payout figure is inclusive of an estimated general interest charge (GIC) amount calculated to 30 January.

The officer did not realise the sentence was included in the letter. It was in the template on the computerised system. The trial judge found (and the Full Court accepted) the officer had not reached a view about remission of the GIC when the letter was despatched.

30 Ibid 529–30 (emphasis added).

31 (2017) 263 CLR 1.

32 Ibid [39].

33 (2018) 262 FCR 41.

The taxpayer paid the lump sum amount in response to the letter. His bank loaned him money to make the payment on the strength of assurances the moneys would be used in full and final settlement of the tax debt. The taxpayer was understandably surprised when he subsequently learned the Deputy Commissioner did not accept the settlement included GIC. The Deputy Commissioner argued the letter containing that suggestion was issued in error. Mr Pintarich continued to negotiate about remission of the GIC under protest. In a letter dated 13 May 2016, the Deputy Commissioner purported to make a decision allowing a partial remission of the GIC. That decision was less favourable to the taxpayer than the deal he thought he had struck. Mr Pintarich sought judicial review of the remittal decision. He argued the Deputy Commissioner's decision of 13 May 2016 was ultra vires because the discretion had already been exercised on 8 December 2014.

The case was not, strictly speaking, concerned with a decision made by (or with the assistance of) a computer. The purported decision of 8 December 2014 was the product of a clerical error that resulted in a letter which did not accurately describe the terms of the settlement that had been reached following the negotiations between the taxpayer and the authorised officer, who was a natural person. A computerised system was certainly used by the decision-maker to document and communicate the decision, but there was no suggestion that any part of the decision-making process itself was delegated to, or undertaken by, an expert system. The issue before the Court was whether the *authorised officer* had made a decision on 8 December 2014 to exercise the discretion in s 8AAG of the *Taxation Administration Act 1953* (Cth) to remit the GIC.

In making its decision, the majority relied on the reasoning of Finn J in *Semunigus v Minister for Immigration and Multicultural Affairs*³⁴ (*Semunigus*). In that case, a member of the Refugee Review Tribunal had prepared written reasons for a decision. The member provided those reasons with the decision to registry staff for publication to the applicant. After the registry staff received the decision from the member but before it was published, the applicant lodged additional submissions with the Tribunal. The decision as written was subsequently published without reference to the new submissions. A question arose as to whether the Tribunal was already *functus officio* when it received the submissions (and therefore relieved from any obligation to take them into account). Finn J explained:

For present purposes I am prepared to hold that the making of a decision involves both reaching a conclusion on a matter as a result of a mental process having been engaged in and translating that conclusion into a decision by an overt act of such character as, in the circumstances, gives finality to the conclusion — as precludes the conclusion being revisited by the decision-maker at his or her option before the decision is to be regarded as final.³⁵

The majority in *Pintarich* accepted that the statement of Finn J 'capture[d] the elements that are generally involved in the making of a decision'.³⁶ The majority also noted the trial judge's uncontested finding of fact that the authorised officer had not turned his mind to the question at issue under s 8AAG or reached a conclusion. In those circumstances, the majority concluded, the decision-making process in relation to the remittal of GIC had not been completed on 8 December 2014 notwithstanding the representations contained

34 [1999] FCA 422.

35 *Ibid* [19] (emphasis added).

36 (2018) 262 FCR 41 [143].

in the letter of the same date because there had been no mental engagement with the statutory question.³⁷

The references to mental engagement by the decision-maker in the decision-making process have fed into the debate over automated decision-making. The potential significance of *Pintarich* in this regard was apparent in the reasoning of Kerr J in dissent. Justice Kerr, a former president of the AAT, warned against adopting too narrow a view of what constitutes a decision. His Honour explained:

while I agree that Finn J's statement still represents what is usually involved in a (valid) decision, I would respectfully observe that that may be rapidly becoming an artefact of the past.

The hitherto expectation that a 'decision' will usually involve human mental processes of reaching a conclusion prior to an outcome being expressed by an overt act is being challenged by automated 'intelligent' decision making systems that rely on algorithms to process applications and make decisions.

What was once inconceivable, that a complex decision might be made without any requirement of human mental processes is, for better or worse, rapidly becoming unexceptional. Automated systems are already routinely relied upon by a number of Australian government departments for bulk decision-making. Only on administrative (internal or external) and judicial review are humans involved.³⁸

In the course of his reasons, his Honour referred to the decision of the Western Australian Court of Appeal in *Polo Enterprises Pty Ltd v Shire of Broome*³⁹ (*Polo*). In that case, Martin CJ (with whom Newnes and Murphy JJA agreed) discussed the decision in *Semunigus* and pointed out the issue in *Polo* was different to a question over whether a decision-maker was *functus officio*. That difference was significant in the circumstances. Martin CJ observed:

it would not be appropriate to construe the passage [ie the passage in the judgment of Finn J in *Semunigus* containing the statement of what was generally required for a valid decision] as asserting that in every case and in every context there cannot be a 'decision' unless there has been a process of mental engagement by the decision-maker.⁴⁰

The majority in *Pintarich* also referred to the reasoning in *Polo* which suggested the statement about decision-making in *Semunigus* should not be applied reflexively — as well they might, given the statement of Finn J opened with the qualifier 'For present purposes'. Justices Moshinsky and Derington noted:

We have treated the statement of Finn J [in *Semunigus*] as a general statement of what is involved in the making of a decision, and we accept that it may not be applicable in relation to all issues.⁴¹

That qualification suggests the reasoning in *Pintarich* must be read carefully. The decision under review in that case involved the exercise of a discretion. Statutory provisions presumably invest a decision-maker with a genuine discretion in circumstances where it is too hard to anticipate and devise a rule that is capable of precise and consistent application. The discretion is available to provide a measure of flexibility to avoid arbitrary outcomes in particular cases. The Parliament (and ultimately the community) entrusts the final outcome

37 Ibid [153].

38 Ibid [45]–[47].

39 [2015] WASCA 201.

40 Ibid [78].

41 (2018) 262 FCR 41 [148].

of the process to the decision-maker's good judgment. Judgment is a concept ordinarily associated with individuals, not machines — even intelligent ones. It is easy to see why mental engagement might be regarded as an essential feature of most genuinely discretionary decisions. But not all administrative decisions involve the exercise of discretion, and not all decisions involving the exercise of discretion require fine judgment. Some discretionary provisions have such tightly drawn criteria that they do not involve the exercise of genuine judgment. The reasoning in *Pintarich* and *Semunigus* may be distinguishable in those cases.

The distinction between discretionary and non-discretionary decisions was anticipated by the ARC in its report in 2004. The report pointed out decision-makers must exercise discretionary powers personally and may not be fettered in doing so.⁴² The ARC's report made clear that some decisions were best left to humans — most obviously those decisions which involved the exercise of a genuine discretion. To that end, Principle 1 in the ARC's 2004 report stated:

Expert systems that make a decision — as opposed to helping a decision-maker make a decision — would generally be suitable only for decisions involving non-discretionary elements.⁴³

Principle 2 was even more explicit. It provided, 'Expert systems should not automate the exercise of discretion'.⁴⁴ Principles 3 and 4 went on to explain how expert systems might be used to *assist* the (human) decision-maker to make discretionary decisions but emphasised the discretion must not be fettered or directed in such a way that it was rendered nugatory.

Where does that leave us? At the level of theory, there may yet be questions over the lawfulness of automated decision-making, but those questions are more likely to arise in relation to discretionary decision-making. For the rest, we should not lose sight of the reality of decision-making. Administrative decisions which involve the exercise of statutory powers ultimately require the decision-maker to answer a question derived from a statute. That question may be hard to divine, but it will often be straightforward: if X (and assuming X is not in doubt), then Y. That style of decision-making is well suited to automated processes. In many such processes, there is unlikely to be any doubt about the reasons why a decision was made. There will be no difficulty in obtaining redress under established rules if an error does occur. Given that reality, it would be a pity to deprive the community of the benefits of automated decision-making in simple, high-volume decision-making. In these straightforward cases, accuracy, speed and efficiency count for a great deal.

There may yet be a technical fix to residual concerns over the conceptual and paradigmatic challenges posed by automated decision-making. Those engaged in a field of research and study known as 'ethical information technology' aim to adapt automated decision-making processes so they can be made accountable according to established administrative law principles. The ethical IT movement aims to address the so-called 'Black Box' problem

42 Administrative Review Council, above n 3, 12, 15–16.

43 Ibid viii.

44 Ibid.

in which the automated decision-making process is opaque and poorly understood.⁴⁵ Proponents of ethical IT call for automated decision-making processes to be programmed with the ability to provide an intelligible statement of reasons for the decisions they make in their own right, or on behalf of a deemed decision-maker.⁴⁶ Such a statement of reasons would be capable of identifying the relevant facts and considerations that were taken into account. That facility would make correcting errors easier — at least as easy as it would be in the case of a human decision-maker. With human decision-makers, it might be hard to tell if the statement of reasons accurately reflects the reasoning of the individual. (Computers are not disingenuous, after all.) These developments would also assist more complex decision-making processes gradually to conform to the administrative law values discussed in the ARC's report.

The caution evident in the ARC's report in relation to discretionary decision-making was prescient. Difficult questions remain over the extent to which it is possible (or even desirable) to automate any decision-making process involving the exercise of discretion. Other, more routine decision-making processes may present fewer legal issues — and to the extent they do raise questions, those questions may yet be amenable to technological fixes that bring them into conformity with the law.

Lawfulness in operation

It is incumbent on any decision-maker to get the law right and to apply it in accordance with established principles of administrative law. That is easier said than done.

Almost every administrative decision involves the decision-maker in answering a question derived from a rule set out in a statute. The trick lies in the precise formulation of the question in a given case. Every judge and tribunal member will be (sometimes painfully) aware that the hardest part of a case may lie in formulating the question correctly. Once the question is accurately formulated, the answer is often obvious. But the quintessentially *intellectual* process involved in framing the question is much harder than it might appear to outsiders or to software programmers and coders — or to some of the managers in the upper reaches of the bureaucracy who commission automated systems. Those managers do not necessarily have hands-on experience of the complexity of a particular decision-making process or the real-life circumstances that may be encountered.

Even in those cases where the correct formulation of the question is obvious or settled, it may be fiendishly difficult accurately to translate the rule into programming language that is capable of predictable application in all circumstances. Human decision-makers, for all their faults, can reason from a rule to deal with new, unusual or nuanced circumstances. (The common law has successfully operated on that model for 1,000 years.) Rule-based automated processes may end up being confined to the most routine decisions because

45 See, for example, M Kearns and A Roth, 'Ethical Algorithm Design Should Guide Technology Regulation', Brookings Institution (online, 13 January 2020) <<https://www.brookings.edu/research/ethical-algorithm-design-should-guide-technology-regulation/>>; see also Information Commissioner's Office, UK, and The Alan Turing Institute, *Explaining Decisions Made with AI* (online, 20 May 2020) <<https://ico.org.uk/for-organisations/guide-to-data-protection/key-data-protection-themes/explaining-decisions-made-with-ai/>>.

46 See, for example, Yee-Fui Ng et al, above n 5, 14.

rule-based systems are, by their nature, good at making binary choices between clearly recognised data points.

The likelihood of legal error increases exponentially when the decision-making process involves more steps, or more complicated enquiries. The potential for error — and the cost and inconvenience that results — has been amply demonstrated by the Robodebt initiative.

The Online Compliance Initiative, or ‘Robodebt’, was an initiative of the Department of Human Services. Robodebt was launched in 2016. It was intended to identify instances where individual recipients of Centrelink benefits were overpaid because they failed accurately to declare their income. In the past, officers of the department responded to tip-offs and computer-generated profiles by undertaking a comparison with employer-reported data held by the ATO. The process was slow and resource-intensive, and the department suspected a good deal of overpayments were not detected. Robodebt was intended to automate the data-matching *and* decision-making processes. The computer took over the task of comparing data provided to the department with the employer-generated data held by the ATO. Where anomalies were detected, the system was programmed to generate a letter asking the individual to update the information they provided to Centrelink. The computer then automatically adjusted the payments which the individual received in response to the data. Where necessary, the system also generated a letter informing the individual that a debt had been raised, which led, in the ordinary course, to recovery processes. The system was an immediate success in that it identified many individuals who had received overpayments — certainly many more individuals than would be identified using the processes in place before Robodebt was launched.

The legal authority for the decision-making process was found in s 6A of the *Social Security Act 1991* (Cth). That section deemed the decisions to raise debts to be decisions of the Secretary. That authorisation suggested compliance with Principle 5 in the ARC report, which provided:

The use of an expert system to make a decision — as opposed to helping a decision-maker to make a decision — should be legislatively sanctioned to ensure that it is compatible with the legal principles of authorised decision-making.⁴⁷

Yet the lawfulness of the scheme remained an issue. A report from the Ombudsman confirmed there were problems with the design and implementation of the scheme.⁴⁸ One of the central difficulties arose out of the design of the algorithm, the set of instructions that drove the program’s data-matching process. It used the yearly income figures from the ATO and divided them into equal fortnightly instalments for the purposes of calculating and comparing the individual’s income for welfare purposes. But that assumption does not hold

47 Administrative Review Council, above n 3, viii.

48 Acting Commonwealth Ombudsman, *Centrelink’s Automated Debt Raising and Recovery System* (Report No 02/2017, April 2017) <https://www.ombudsman.gov.au/__data/assets/pdf_file/0022/43528/Report-Centrelinks-automated-debt-raising-and-recovery-system-April-2017.pdf>.

good for casual workers whose income fluctuated over the year but who were required to report fortnightly.⁴⁹

At least in that respect, the computer got the law wrong and asked the wrong question. The process fell short of Principles 7 and 10 in the ARC report. Principle 7 provided:

The construction of an expert system must comply with the administrative law standards if decisions made in accordance with the rule base are to be lawful.

Decisions made by or with the assistance of expert systems must comply with administrative law standards in order to be legally valid.⁵⁰

Principle 10 provided:

Expert systems should be designed, used and maintained in such a way that they accurately and consistently reflect the relevant law and policy.⁵¹

The Robodebt initiative was shut down earlier this year. The government has reportedly agreed to refund \$721 million collected from around 370,000 Centrelink customers.⁵² This costly exercise serves as a reminder of the importance of ensuring that, at a minimum, automated decision-making processes are lawful in operation.

Fairness

All decision-making processes should be fair. Fairness in this context has two components. The process must be procedurally fair to the subject; and it should (at least up to a point) be substantively fair.

There were a number of problems with the operation of the Robodebt project that raised procedural fairness issues. The Ombudsman pointed to problems with the way in which information was collected from individuals. There was a high error rate in the data being supplied because individuals may have been unclear on what was required.⁵³ The agency failed to comply with Principle 21 in the ARC report, which provided:

Agencies should take steps to ensure that the data collected and used by expert systems for administrative decision making remain accurate and complete.⁵⁴

Those issues with the data were compounded when the decision-making process acted on the wrong information. The Ombudsman pointed out the burden of proof may have shifted unfairly onto individuals once they were identified for administrative action as part

49 For a useful discussion of the Robodebt program, see D Hogan-Doran, 'Computer Says "No": Automation, Algorithms and Artificial Intelligence in Government Decision-making' (2017) 13 *The Judicial Review* 345, 358ff.

50 Administrative Review Council, above n 3, viii.

51 *Ibid* ix.

52 J Hayne and M Doran, 'Government to Pay Back \$721 Million as it Scraps Robodebt for Centrelink Welfare Recipients', *ABC Online*, 29 May 2020 <<https://www.abc.net.au/news/2020-05-29/federal-government-refund-robodebt-scheme-repay-debts/12299410>>.

53 Acting Commonwealth Ombudsman, above n 48, 23.

54 Administrative Review Council, above n 3, x.

of the program.⁵⁵ Many vulnerable individuals were forced to undertake time-consuming and stressful appeal processes to correct the problem. Worryingly, a large number of disempowered individuals may have just given up and accepted the computer's erroneous decision.

There is another dimension to the requirement of fairness that has, to some extent, been overlooked in the furore over Robodebt. While the Robodebt project had an unacceptably high rate of error, the process did identify a large number of individuals who received overpayments as a consequence of under-reporting their income. They enjoyed the benefit of moneys they were not entitled to receive. Many of these individuals might have escaped notice had the agency persisted with its resource-intensive and time-consuming manual comparison and decision-making processes. Those processes were plainly inefficient. In the absence of Robodebt, at least some recipients of benefits would have continued to receive moneys to which they were not entitled. That is not fair to those who accurately account for their income. It may be that a properly designed automated data-matching and decision-making process would reduce overpayments and fraud by enabling the agency to undertake compliance investigations more thoroughly, more quickly and more often.

The decision-making process involved in Robodebt did not meet standards of fairness, but it must be acknowledged there are concerns over the fairness of the slow and inevitably arbitrary process it was intended to replace. There is no reason in principle why data-matching processes should not be extensively automated — and, to a significant extent, they are. Where the data-matching process identifies a problem, there may be a role for an automated process in calling the beneficiary to account. But the process must be properly designed with the law *and* fairness concerns in mind.

Rationality

A good decision-maker takes into account relevant and probative evidence and ignores irrelevant considerations in the course of reaching a reasoned and logical decision about how the law should be applied in a particular case. The requirement of rationality also contemplates striving for the holy grail of administrative decision-making: objectivity.

On the face of it, automated decision-making processes should excel when measured against this value. A rule-based expert system should be a model of rationality *assuming it is programmed correctly*. It should also be completely objective in the sense that it will not be swayed by human bias or emotion.

The practical reality is more complicated. We have already observed in the Robodebt case what can go wrong when the program is incorrectly programmed. The decisions made in that case were not, strictly speaking, irrational: there is no doubt the automated process was rigorously logical, as it applied the (perhaps misstated) rules to the (misunderstood) data it was instructed to interrogate. Yet that is unlikely to be comforting to the individual on the wrong end of the decision. The fact the wrong rules were rigorously applied to the wrong data lends the whole process a Kafka-esque quality.

55 Ibid 23.

If Robodebt involved unintended consequences, a case involving the *Business Names Registration Act 2011* (Cth) (Registration Act) can help illuminate the problems with intended and rational consequences that nonetheless bemuse or irritate the public. The Registration Act says individuals must register the business name they intend to use in carrying on a business. The Registration Act establishes the modern incarnation of a registration process that was formerly carried out manually by clerks in the office of the Registrar of Business Names or equivalent in each state and territory.

The computerised decision-making process that replaced the clerks of old is authorised under s 66 of the Registration Act. That section uses essentially the same language as that employed in s 6A of the *Social Security (Administration) Act 1999* (Cth) referred to earlier. In other words, the statute deems the determination reached by the automated decision-making process to be that of the regulator, the Australian Securities and Investments Commission (ASIC).

The registration process starts when an application is lodged using the appropriate ‘smart’ form on a website. The applicant must provide an Australian Business Number and answer a number of questions. The applicant must also enter the precise name it wishes to register. Section 24 of the Registration Act says ASIC must register the business name if the name is available and the other requirements are satisfied.

It follows that the question of whether a business name should be registered is essentially binary in nature. The name is either available to the applicant or it is not. There is no discretion involved in the decision. The decision-making process does not *appear* to require the decision-maker to evaluate or weigh complex facts or arguments.

A string of cases considered by the AAT suggest the decision in question is more complicated than it first appears. The problem arises because the register must be precise in its recording of existing business names to facilitate the computerised decision-making process. What of the situation where a business name is similar to — but not precisely the same as — an existing name? That is a challenge contemplated by the statute. Section 16 of the Registration Act sets out the objects of the legislation, and s 16(3) specifically refers to the purpose of ‘avoid[ing] confusion by ensuring that business names that are identical or nearly identical are not registered’. That was the problem in *Smith and Australian Securities and Investments Commission and Anee Donohoe trading as Central Coast Surf Academy*⁵⁶ (*Smith*). The proprietor of the existing business name ‘Central Coast Surf School’ was upset at a likely rival’s attempt to register the name ‘Central Coast Surf Academy’. The owner of the existing name objected to the application on the basis that ‘Central Coast Surf Academy’ was identical, or nearly identical, to the existing name.

It is easy to see why the owner of the existing name protested. There are many cases where businesses have been found to have engaged in misleading or deceptive conduct after using business names that are similar to a commercial rival or which might suggest an association that did not exist. The applicant in *Smith* argued it was common sense that the name ‘Central Coast Surf Academy’ should not be available for registration under s 25.

56 [2014] AATA 192 [1] (SM McCabe).

The words ‘identical’ and ‘nearly identical’ were defined in a ministerial determination. Section 6 of the *Business Names Registration (Availability of Names) Determination 2012* (the Determination) explained:

When comparing a business name with another name (other than a company name) to determine whether the names are identical or nearly identical, a word or expression in an item in Schedule 1 is to be taken to be the same as each other word or expression in the item.

ASIC insisted — and the AAT ultimately accepted — that the particular words in Schedule 1 were a precise and exhaustive list of what might be regarded as identical or nearly identical. The names and words on the list were not examples from which the decision-maker could extrapolate or analogise. As I explained in my reasons:

It seems one of the motivations for constructing the law in this way is to make it easier to automate the process of applying for a business name. Computers do not do well with examples: while they work more quickly (and presumably more cheaply) than the registry clerks of fond memory, they do not yet have the capacity to analogise from a list of examples. They prefer clear instructions capable of certain application. Computers don’t ‘do’ nuance, and they are not open to persuasion.⁵⁷

The list of names contained in the schedule was, at least at the time, a forest of single instances. There was a list of names that might be used by dance instructors which were taken to be identical or nearly identical to each other, including ‘academy of dance, dance academy, dance centre, dance studio, school of dance’. But the list in the schedule did not refer to an equivalent form of words when used by surfing instructors. In those circumstances, the expressions ‘Central Coast Surf School’ and ‘Central Coast Surf Academy’ were not identical or nearly identical to each other.

I acknowledged in my reasons in *Smith* that the decision-making process operated exactly as the legislation intended.⁵⁸ Yet, while the outcome was correct and rational in the sense it proceeded from fact to rule to conclusion in a straightforward way, the applicant’s experience of the process was hardly preferable if one is concerned about good government or public confidence in the rationality and wisdom of administrative decision-making.

Having said that, the decision in *Smith* may also be the exception that proves a rule. The case was decided during the relatively early stages of the system’s operation. The minister subsequently issued the *Business Names Registration (Availability of Names) Determination 2015*, which says names including the words ‘institute, academy, school, college’ should be treated as nearly identical if the only relevant difference between the names is one of those words. The schedule includes many other examples of common substitutes that are not linked to specific occupations. The outcome in *Smith* would likely be different under those revised rules. It follows the shortcomings of the automated decision-making process in that case were easily fixed.

I mentioned previously that automated decision-making processes should be more objective than decisions made manually. Objectivity should be less of an issue where the rule being applied is clearly drawn and the facts are easily ascertained. But the decision in *Smith*

57 Ibid [14].

58 Ibid [13].

demonstrates how objectivity (or at least the perception of objectivity) might be undermined in surprising ways. In *Smith*, ASIC pointed out that surf instructors or other businesses could lobby the minister (as dance instructors may have done) to have variations of their business names included on the list in the schedule. The need to do so evaporated once the Determination was changed in 2015, but the potential for special pleading is a worry.

The risk of bias, inconsistency and arbitrariness — the antithesis of rationality — is greater where the decision-making process in question requires the decision-maker to evaluate more complex facts or form a state of mind. A good example of such a process can be found in the provisions of the *Social Security Act 1991* (Cth) which require the decision-maker to identify whether an individual is a 'member of a couple'. Members of a couple may be paid social security benefits at a lower rate than individuals who do not meet the criteria set out in s 4(3) of the Act. The legislation assumes members of a couple require less generous assistance than single persons because couples have the opportunity to pool their resources. The policy objective which underlies the provision is clear and laudable. Social security benefits should be targeted and tailored to some extent according to need. That tailoring comes at a cost. To target and tailor, more information is required, and — particularly given the subtlety and variety of human relationships — more complex and nuanced distinctions may be necessary.

The legislation adopts a fudge in its definition of the enormously fluid concept of 'a couple'. Rather than attempting an exhaustive definition,⁵⁹ s 4(3) says a decision-maker must have regard to all the circumstances of a relationship between two people to determine whether they are members of a couple. In making that assessment, the test requires the decision-maker to consider a number of indicia which are thought to be characteristic of these relationships.

The definition, such as it is, demands an assessment that is as complex as it is subtle. That is probably inevitable given the richness and diversity of human relationships, but that truism does rather point to the practical difficulty of delivering on the policy objectives of the legislation. Any casual reader of decisions in the AAT on the topic will see that reasonable people may disagree on whether a particular relationship satisfies the definition. There is a real risk of subjectivity, bias and idiosyncrasy in decision-making if this sort of assessment is undertaken manually. But how would a rule-based automated decision-making process apply the criteria systematically without creating the risk of arbitrariness given the very subjectivity of some of the observations required?

The decision required in the 'member of a couple' case is not discretionary, but it does require a careful evaluation of facts before reaching a binary conclusion that is freighted with meaning and values. A rule-based automated decision is not likely to be especially useful in that endeavour beyond helping a decision-maker to assemble appropriate data and generating prompts to ask questions and provide commentary.

59 One is reminded of the words of Potter Stewart J in *Jacobeliss v Ohio* 378 US 184 (1964) in the United States Supreme Court. In that case, his Honour declined to exhaustively define the word 'obscenity' but pointed out: 'I know it when I see it': 197.

Might a more advanced system powered by artificial intelligence offer a solution to these more complex cases? The ARC was sceptical of systems that derived the rules from observed data, but the principal objection appeared to lie in technical concerns about how the decision-making process would be held accountable. I shall discuss that issue below, but for now I want to focus on the possibility of more objective decision-making on the basis of vast datasets using systems that can learn and adapt by observation. If the automated process was provided with enough information about couples, might it more reliably identify couples when they apply for benefits?

It turns out systems like this may yet incorporate biases inherent in the data, even if they are not themselves prone to bias. The data may be subject to distortions that are difficult to detect and correct. For example, same-sex couples were not regarded as being in marriage-like relationships under the social security legislation until 2009. It follows there is less historical data about relationships of this kind. Even now, there is likely to be under-reporting of these relationships in some sections of the community. Some couples are reluctant to identify themselves as such. They may not even think of themselves in those terms. The consequences for the integrity of the data are unclear. That inevitably calls into question the reliability of decision-making based on that data.

It is entirely possible that automated processes might assist decision-makers to identify and analyse data more conveniently and objectively. To the extent those processes do so, they promote rationality by ensuring relevant information is taken into consideration. But it is important to test the assumption that more information is inevitably better in decision-making. The quality of the data and the integrity of the datasets should never be taken as a given.

Openness or transparency

A decision-making process is effectively unaccountable if the decision-maker's reasoning is not apparent. Individual decision-makers can generate reasons which explain their reasoning, although there is always the danger of artifice or misstatement. The problem is more complicated in some automated decision-making processes.

In straightforward binary decisions like registering a business name, the reasoning process is rarely in doubt. The computer is programmed to apply a rule — in that case, the rule laid down in the Determination. If the computer rejects the application to register a particular name, the reasons for that decision should be perfectly clear. More complicated issues arise where the decision is made by a human with substantial input from a support system which the human does not understand. That may occur, for example, where an agency calculates the quantum of a debt using an obscure method that is not understood by the decision-maker responsible for making the decision to recover or waive the debt. Any statement of reasons provided by the human decision-maker in those circumstances will be incomplete.

Even more difficult issues arise in relation to a computerised decision-making process involving the application of artificial intelligence. As the ARC report pointed out, the outcome of such a process is dictated by the computer's application of a rule derived by data. But the operation

of the algorithm may not lend itself to an intelligible explanation because it does not engage in a recognisable reasoning process. The challenge is made more difficult by the fact the algorithm might be regarded as valuable intellectual property by the software developers. The developers may insist on the algorithm being kept confidential.

I have already mentioned that the so-called 'Black Box' problem might be addressed through 'ethical IT' initiatives. This field of research is developing innovations like software that creates virtual windows into the automated decision-making process. The software enables the computer to provide an explanation for what it did at each important point in the decision-making process. Further thought will need to be given to whether software developers should be able to restrict access to their algorithms and other key features of their products. It may be costly to require governments to rely on open-source software or pay for one-off products so they can be exposed to scrutiny.

Efficiency

Manual decision-making processes can be costly. They can be slow and prone to error. The processes can also be expensive for users who must line up in queues and wait for outcomes. Automated processes can often be accessed remotely, and at any time of the day or night. Some of the processes can provide immediate responses. Other processes that can quickly analyse vast datasets offer the promise of better targeted, more responsive and better informed decision-making. Data-matching programs can reduce fraud and waste.

The automated business name registration process established under the Registration Act was undoubtedly more efficient than the manual processes it replaced. With commerce occurring on a national scale, the old state-based systems were no longer viable. A central register was required 'to remove the inconvenience caused by the registration of business names under the law of more than one jurisdiction within Australia'.⁶⁰ Registry clerks were no longer required under the new arrangements, which presumably led to budgetary savings. Applicants might also have benefited from using a process that was quick and simple to use. Yet an applicant that is concerned about a business name being nearly identical may have less protection than the process involving the business name clerks. In the past, an applicant could negotiate with a business clerk over a particular name if there were concerns it might be similar. While that rough-and-ready process was hardly foolproof, it did offer some protection for the intellectual property in a business name. That meant there was less need for small businesses to resort to more expensive legal processes like trademarks or claims for misleading or deceptive conduct.

Efficiency appears to be a given in circumstances where armies of public servants are relieved of the obligation to make manual decisions. But the efficiency gains might not be as great as they first appear. As an agency's potential to process information increases, there may be a temptation to ask for ever *more* information — on the assumption that more and better data produces more accurate decisions. As the power of the system increases, the ambitions of those who use the data can be expected to grow. The net effect is that individuals are required to provide ever more data — and surrender ever more of their

⁶⁰ *Business Names Registration Act 2011* (Cth) s 16(1)(b).

privacy. While that might not be costly from the viewpoint of the decision-maker, the costs to the citizen and the community might be high.

It follows that increased capability might not yield the net efficiency gains that are promised. Budgetary savings might be dissipated by mission creep within the agency as it moves to exploit the increased capability — which potentially increases the regulatory burden on members of the community.

Claims that particular decision-making processes promote efficiency must be scrutinised carefully. The efficiencies may be achieved by simply shifting costs. The efficiencies may also be illusory. Some automated processes are expensive boondoggles that deliver none of the promised gains. There is also a real danger that agencies expend large amounts on poorly designed automated processes without making proper allowance for the costs in the event of failure. Those costs might be visited on vulnerable members of the community, as in the Robodebt case, where individuals were left to face stressful and time-consuming review processes. More research is required into the direct and indirect costs to the individual of using these systems and accessing review processes when they fail.

A way ahead?

It is important to understand how decisions made using automated processes differ from manual decision-making — but it is also important to see the ways in which they are the same. The values the ARC identified in its 2004 report should inform every decision-making process. It is therefore important to reflect on how decision-making processes that are wholly or partly automated might be accommodated by existing administrative law mechanisms of accountability and review.

The ARC referred to the importance of ensuring that automated decision-making processes conform to the precepts of administrative law. The report pointed out decisions should, at a minimum, be reviewable against the grounds identified in s 5 of the ADJR Act. But it also referred to the potential for merits review of administrative decisions made by expert systems.⁶¹

The decision-making processes in each agency tends to be carried on separately. Most agencies use computer systems that have been adapted to meet the needs of that organisation. Computers in some agencies are able to interact with computerised systems in other agencies subject to certain protocols, including privacy rules. The interaction between the systems operated by the Department of Human Services and the ATO is a good example of this process. While there are some common standards (especially technical standards) that apply across the government, the decision-making processes are typically shaped by the enactment which authorises them.

Given that diversity, the AAT plays an important unifying role in Australian public administration. The AAT reviews decisions made under more than 450 enactments. In each of those cases, the AAT forms part of the ‘continuum of administrative decision-making’ carried on by the executive government through the agency in question.

61 Administrative Review Council, above n 3, 23–7.

From its central vantage point, the AAT exerts a moderating influence on each decision-maker. That independent moderating influence is likely to be particularly important where the decision-making process within the agency is wholly or partly automated. The AAT can provide the all-important human control that provides direction and correction for machine learning processes. It can also provide a repository of decision-making expertise as an antidote to the de-skilling that may occur in some agencies as automated decision-making processes become more common.

To play its role effectively, the AAT will need to continue conducting its own decision-making processes manually. The ARC anticipated as much in its report. Principle 23 provided:

External reviews of administrative decisions should be done manually, in accordance with the procedures and practices of the particular tribunal or court conducting the review.⁶²

Technology will play a part in the AAT's operations, of course. There is a role for smart forms and other online tools to assist users to access the AAT's services, and modern technology (including hearings conducted in person and online using documents in digital form, and modern case management systems) will be a feature. But the essence of the model envisaged for the AAT by Sir Gerard Brennan — of a relatively small organisation in which experienced decision-makers used court-like processes and forensic tools to conduct rigorous reviews that serve as a model — is more relevant than ever in a world of automated decision-making. As Sir Gerard realised, that model of the AAT, while relatively costly, was the key to much larger efficiency gains in decision-making processes carried on elsewhere.

While AAT review will be an important mechanism for ensuring the integrity and efficacy of many expert systems, not all automated decision-making processes are likely to benefit from merits review. Many decision-making processes — as it happens, the ones that are amongst the easiest to automate — are simple and binary in nature. The factual enquiries are straightforward, and the law is clear. In those cases, at least, there is no particular advantage to be gained from the AAT's court-like processes. Making those decisions reviewable by the AAT risks overwhelming the AAT without adding to the quality of the decisions under review. A better system of internal review reinforced by prompt judicial review is likely to be more efficient. But the AAT already plays an important role in the review of discretionary decisions and more complex binary decisions where the facts or the law are such that the AAT's more court-like processes would assist. That function will become even more significant if the decision-making process below, or part of that process, is automated.

Perfecting the AAT's role as a form of institutionalised human control of automated decision-making processes requires a good deal of coordination between administrative lawyers and the individual agencies as they develop the systems in question. The AAT will need to learn more about how the decision-making process works in each case. The AAT will also be concerned to know how its decisions are dealt with once they are made. How will those decisions be taken into account by the original decision-maker in future cases? In the case of expert systems which learn from data, how precisely will AAT decisions be incorporated into the data or used to modify the algorithm which drives the process? Further thought is required on how the AAT can better play its role in the continuum of administrative

62 Ibid xi.

decision-making. Recognising the centrality of that role is a good start. That may yet require a change in culture from those agencies that regard AAT review as an annoyance or an optional extra.

Conclusion

‘Good government’ is not an empty slogan. The reality and perception of good government is a key to civic order and prosperity. Automated decision-making processes have a role to play in enhancing good government, but they need to be watched carefully to ensure they are wrought for the public good. The stakes are high: every serious misadventure in the implementation of automated decision-making processes will diminish the credit society extends to government.