

The new machinery of government:

using machine technology in administrative decision-making

A special report under section 31 of the *Ombudsman Act 1974*

29 November 2021

Acknowledgements

We thank Professor John McMillan AO, former Commonwealth and NSW Ombudsman, Dr Lachlan McCalman, Chief Practitioner, Gradient Institute, and Associate Professor Will Bateman, Associate Dean of Research at the ANU Law School for providing expert comments on early drafts of this report. We are also grateful to Bill Simpson-Young and Dr Tiberio Caetano of the Gradient Institute for the discussions we have had around some of the technical concepts covered in the report.

We also thank James Emmett SC and Myles Pulsford for allowing us to publish their legal opinion, and for their additional comments on a draft of the report. We also acknowledge and thank Revenue NSW for its co-operation particularly in the preparation of the Statement of Facts upon which that legal opinion was sought, and which is reproduced in annexure A.

All views expressed in this report are those of the NSW Ombudsman.

ISBN: 978-1-925885-92-7

NSW Ombudsman
Level 24, 580 George Street
Sydney NSW 2000

Phone: 02 9286 1000
Toll free (outside Sydney Metro Area): 1800 451 524
National Relay Service: 133 677

Website: www.ombo.nsw.gov.au
Email: nswombo@ombo.nsw.gov.au

© State of New South Wales, 29 November 2021
This publication is released under a Creative Commons License CC BY 4.0

The Hon Matthew Mason-Cox MLC
President
Legislative Council
Parliament House
SYDNEY NSW 2000

The Hon Jonathan O’Dea MP
Speaker
Legislative Assembly
Parliament House
SYDNEY NSW 2000

Dear Mr President and Mr Speaker

Pursuant to section 31 of the *Ombudsman Act 1974*, I am providing you with a report titled *The new machinery of government: using machine technology in administrative decision-making*.

I draw your attention to the provisions of section 31AA of the *Ombudsman Act 1974* in relation to tabling this report and request that you make the report public forthwith.

Yours sincerely



Paul Miller
NSW Ombudsman

29 November 2021

Contents

Foreword from the NSW Ombudsman	iv
1. Executive summary.....	1
1.1 Machine technology is on the rise, and offers many potential benefits	2
1.2 Why we have written this report	2
1.3 Administrative law and practice must be given central attention.....	3
1.4 Administrative law requirements for good decision-making.....	3
1.5 Good practice for designing and implementing machine technology	5
1.6 The role of Parliament in authorising machine technology	5
1.7 The way forward – starting with increased visibility	7
Part 1: Machine technology	8
2. Introduction	9
2.1 The rise of machine technology	9
2.2 Why is the Ombudsman interested in machine technology?	9
2.3 What has prompted this report?	11
2.4 The purpose and structure of this report	12
2.5 What we will do next	13
3. The new technologies	14
3.1 What we mean by ‘machine technology’	14
3.2 Machine technology is not just one thing.....	14
3.3 The need to design machine technology for particular applications.....	15
4. The promise of machine technology in government decision-making	16
4.1 Machine technology within a decision-making system	16
4.2 How are governments using machine technology?	18
4.3 The potential benefits of machine technology	19
Part 2: Administrative law and machine technology.....	21
5. Why administrative law matters for new technology	22
5.1 Public sector decision-making is different	22
5.2 Good government according to law	22
5.3 Existing laws apply to new technologies.....	23
5.4 How administrative law applies to new machine technologies.....	23
6. Key administrative law issues for machine technology	25
6.1 The essential requirements for good (and lawful) administrative decision-making	25
7. Proper authorisation.....	27
7.1 Legal power to make the decision	27
7.2 Legal authority of the person making the decision.....	28
7.3 The scope of the decision-making power – including the extent of any discretion	30
8. Appropriate procedures	34
8.1 The decision has followed a fair process	34
8.2 Other legal and ethical obligations	37
8.3 Reasons are given for the decision (particularly decisions that affect the rights or interests of individuals).....	39
9. Appropriate assessment	42
9.1 The decision answers the right question (which necessitates asking the right question).....	42
9.2 The decision is based on a proper analysis of relevant material	44
9.3 The decision is based on the merits and is reasonable in the circumstances.....	45

10. Adequate documentation	48
10.1 The circumstances surrounding the making of decisions is adequately documented and records kept	48
Part 3: Designing machine technology to comply with the law and fundamental principles of good government	49
11. Putting in place the right team	50
11.1 It's not an IT project	50
11.2 Having lawyers on the design team is essential.....	51
11.3 Ensuring policy choices are made by the right people, at the right level of authority	51
12. Determining the necessary degree of human involvement.....	53
12.1 The administrator must engage in an active mental process	53
12.2 The division of tasks between machine and human	54
12.3 The risk of technology complacency and 'creeping control'.....	55
12.4 Practical indications of active human involvement	56
13. Ensuring transparency	58
13.1 Reasons and the right to an explanation	58
13.2 Accountability and reviewability.....	60
13.3 Publishing source code	62
14. Verification, testing and ongoing monitoring.....	63
14.1 Testing before adoption.....	63
14.2 Undertake monitoring, review and periodic evaluation.....	64
15. Statutory provisions that authorise machine technology.....	70
15.1 Stating, in simple terms, that an administrator is authorised to use a machine	70
15.2 Attributing machine outputs to an administrator	71
15.3 More sophisticated authorisation provisions	73
15.4 Transforming the substantive statutory function	76
15.5 Mandatory properties of machine technology	77
16. Coda – new laws for new technology?.....	80
Endnotes	82
Annexure A – Revenue NSW case study	98

Machine technology in practice

Reviewing 'decisions' versus investigating 'conduct'	10
Mobile phone detection cameras.....	16
Automating procedural court rules	18
Some machine technology use by government will be 'legally unexceptional'	26
Services Australia Centrelink's automated income compliance program (robodebt).....	27
Machine technology and the 'rule against dictation'	31
Machine technology as a means of delivering 'guidance' to administrators	32
Algorithmic bias	35
Using machine technology to administer Commonwealth child support payments	40
Lost in translation – a simple error converting legislation into code	43
Machine technology in sentencing – COMPAS.....	46
Providing explanations for decisions that are 'instructive, informative and enlightening'	58
Use of the 'Structured Decision-Making' tool in the NSW child protection system.....	65
'Nearly identical' under the Commonwealth <i>Business Names Registration Act 2011</i>	72

Foreword from the NSW Ombudsman

We have entered a new digital age, and it is widely accepted that governments must transform themselves to be fit for this future.¹ The NSW Government's first Digital Strategy spoke of the need for government to be 'digital by design' and 'digital by default'.²

It is unsurprising then, that digital innovation has also begun to permeate the methods by which public officials and agencies exercise their roles as administrators – the ways they make decisions and exercise powers granted to them by Parliament through legislation.³

This report is about this shift toward machine technologies, a term we use for the broad range of digital and data enabled systems and processes that are, or might in future, be used to guide, assist or even determine when and in what ways administrative powers will be exercised.⁴

There is no doubt that machine technologies have the potential to bring significant benefits to government agencies and the public they serve, including in terms of speed, efficiency, accuracy and consistency.

However, the public sector has a unique constitutional role – it is that arm of government that administers laws, and as such it is uniquely subject to legal rules and standards of good conduct as to when and how it does so. This *administrative law* – the legal framework that controls government action – does not necessarily stand in the way of machine technology adoption, but it will significantly control the purposes to which it can be put and the ways in which it can operate in any particular context.

Failure to comply with the norms of administrative law risks *maladministration* – something at the forefront of the Ombudsman's jurisdiction. Some contraventions may also result in decisions or actions being held by a court to have been unlawful and/or invalid.

This is one reason why, as far back as 2004, the Administrative Review Council emphasised the need for lawyers to be actively involved in the design of machine technology for government⁵ – a key point we take up in **chapter 11** of this report.

Since that time, there is a small but growing body of legal academic literature, both in Australia⁶ and elsewhere,⁷ that seeks to examine public sector use of machine technology through the lens of administrative law.

It is not clear to us, however, that this body of work is always reaching the audience it needs to: law-makers, policy-makers, government lawyers and particularly those who are at front-lines of implementing and operationalising machine technologies.

A primary aim of this report is to help to bridge that gap. We also offer guidance on the important practical steps that agencies need to take when considering the adoption of machine technology to support the exercise of administrative functions.

At the end of this report we touch on the question of whether the rise of machine technologies may also warrant a reconsideration of the legal frameworks, institutional arrangements and rules that apply.

For example, it may be that traditional administrative law mechanisms of redress, such as judicial review or complaint to the Ombudsman, will be considered too slow or otherwise too individualised to provide an adequate response to forms of systemic maladministration that could arise from 'algorithmic bias'. Modified frameworks may be required – for example, to require proactive and ongoing external testing and auditing of systems, in addition to reactive individual review rights.

On the other hand, new or amended laws may also be needed to expressly facilitate the beneficial use of new technologies in some areas, where the operation or uncertainty of existing rules might otherwise unduly stand in the way.

Paul Miller
NSW Ombudsman

Executive summary



Our role at the NSW Ombudsman is to oversee government agencies and officials – helping to ensure they are conducting themselves lawfully, making decisions reasonably, and treating all individuals equitably and fairly (**chapter 2**).

When agencies and officials fail to do this they are said to have engaged in maladministration or, more formally, section 26 conduct (referring to section 26 of the *Ombudsman Act 1974* (NSW), which sets out the various categories of wrong conduct).

Clearly, the use by government agencies of machine technology – which might be referred to as *artificial intelligence* or *automated decision-making* (see **chapter 3**) – is not inherently a form of maladministration.

There are many situations in which government agencies could use appropriately-designed machine technologies to assist in the exercise of their functions, which would be compatible with lawful and appropriate conduct. Indeed, in some instances machine technology may improve aspects of good administrative conduct – such as accuracy and consistency in decision-making, as well as mitigating the risk of individual human bias.

However, if machine technology is designed and used in a way that does not accord with administrative law and associated principles of good administrative practice, then its use could constitute or involve maladministration. It could also result in legal challenges, including a risk that administrative decisions or actions may later be held by a court to have been unlawful or invalid.

1.1 Machine technology is on the rise, and offers many potential benefits

The use and sophistication of machine technology is increasing worldwide, and it has the potential to bring many potential benefits to government and the public (**chapter 4**).

These include:

- Efficiency and cost savings for government.
- Reduced red tape.
- Increased accuracy.
- Improved consistency.
- Increased productivity and re-focusing of staff to ‘higher value’ activities.
- Better customer service and experience.
- Insights and learning.

Of course, benefits cannot be assumed to follow as a matter of course, and it is important to be realistic about what benefits (and risks) particular technology will deliver in a particular context. Untested assumptions or utopian beliefs about technology should not drive automation strategies.

1.2 Why we have written this report

We were prompted to write this report after becoming aware of one agency (Revenue NSW) using machine technology for the performance of a discretionary statutory function (the garnisheeing of unpaid fine debts from individuals’ bank accounts), in a way that was having a significant impact on individuals, many of whom were already in situations of financial vulnerability.

Following a series of complaints to our office, Revenue NSW worked responsively with us over time to ensure that its garnishee system operated more fairly, by taking account of vulnerability and situations

of hardship. However, we still had questions as to whether Revenue NSW’s system of garnishee automation was legally consistent with its statutory functions.

We sought legal advice from Senior Counsel, which confirmed our doubts. The full Revenue NSW case study, including the legal advice, is set out in **annexure A**.

Currently, we do not know how many other NSW Government agencies are using, or developing, machine technology to assist them in the exercise of their statutory functions.

However, our experience with Revenue NSW and a scan of the Government’s published policies on the use of ‘AI’ and other digital technologies suggests that there may be inadequate attention being given to fundamental aspects of public law that are relevant to machine technology adoption.

1.3 Administrative law and practice must be given central attention

Some of the broader concerns about machine technology use by the private sector, in terms of privacy, human rights, ethics and so on, also apply (in some cases with greater force) to the public sector.

However, the powers, decisions and actions of government agencies and officials are constitutionally different from that of the general private sector.

This means that the public sector’s use of machine technology, particularly for the purposes of statutory decision-making, must also be assessed from an administrative law perspective (**chapter 5**). We believe that this assessment must be central to the use of this technology.

1.4 Administrative law requirements for good decision-making

For simplicity, we can broadly group the requirements for good decision-making in the following ways (**chapter 6**):

Proper authorisation – this means that there is legal power to make the relevant decision, that the person making the decision has the legal authority to do so, and that the decision is within the scope of decision-making power (including, in particular, within the bounds of any discretion conferred by the power) (**chapter 7**).

The requirement for proper authorisation means that statutory functions are not and cannot be directly given or delegated to a machine. It does not necessarily mean that the authorised person cannot be assisted by machine technology.

There is, however, no uniform answer as to what forms of machine technology can be used, and to what extent, in the performance of a particular statutory decision-making function. This must be carefully considered on a case-by-case basis by looking at the particular statute, its purpose, and the context in which it applies.

However, if the function is discretionary, machine technology must not be used in a way that would result in that discretion being fettered or effectively abandoned. In effect, this means that discretionary decision-making functions cannot be fully automated.

Appropriate procedures – this means that the decision has followed a fair process, that it has met other legal and ethical obligations, and that reasons are given for the decision (particularly where it significantly affects the rights or interests of individuals) (**chapter 8**).

Generally, a fair process requires decisions to be made without bias on the part of the decision maker ('no-bias rule') and following a fair hearing of the person affected ('hearing rule'). Machine technology can introduce the possibility of a different form of bias known as 'algorithmic bias'. Algorithmic bias arises when a machine produces results that are systemically prejudiced or unfair to certain groups of people. It is unclear whether the presence of algorithmic bias would necessarily constitute a breach of the no-bias rule (as that rule is traditionally concerned with actual or apprehended bias on the part of the particular decision maker). Even if it does not, however, algorithmic bias may still lead to unlawful decisions (because they are based on irrelevant consideration or contravene anti-discrimination laws) or other maladministration (because they involve or result in conduct that is unjust or improperly discriminatory).

Where machine technology is used in the exercise of a function under a particular statute it also needs to comply with other statutes and common law requirements. Privacy, freedom of information and anti-discrimination laws, in particular, will almost always be relevant.

Having appropriate procedures also means providing where required, or being able to provide where requested, reasons to those who are affected by a decision. In our view, this means also informing those affected if a machine has made (or contributed to the making of) a decision. Where reasons are required, they must be accurate, meaningful, and understandable, which can raise particular challenges when machine technology is used.

Appropriate assessment – this means that the decision answers the right question, that the decision is based on a proper analysis of relevant material, and that the decision is based on the merits and is reasonable in all the circumstances (**chapter 9**).

Using machine technology in the exercise of statutory functions means translating legislation and other guidance material (such as policy) into the form of machine-readable code. A key risk is the potential for errors in this translation process, and the consequent potential for errors and unlawful decisions being made at scale.

When designing and implementing machine technology, it is also essential to ensure that its use does not result in any obligatory considerations being overlooked or extraneous considerations coming into play. While the use of machine technology may enhance the consistency of outcomes, agencies with discretionary functions must be conscious of the duty to treat individual cases on their own merits.

Adequate documentation – agencies are required to properly document and keep records of decision-making (**chapter 10**).

In the context of machine technology, this means keeping sufficient records to enable comprehensive review and audit of decisions. Documentation relating to different 'versions' of the technology, and details of any updates or changes to the system, may be particularly important.

1.5 Good practice for designing and implementing machine technology

In light of the above, there are some key proactive steps that agencies should take when considering the design and adoption of machine technology that will help them to ensure they comply with principles of administrative law and good decision-making practice.

In particular, when setting out to design machine technology for use in the exercise of statutory functions, agencies should:

1. establish a multi-disciplinary design team that involves lawyers, policymakers, and operational experts, as well as technicians, with roles and responsibilities that are clearly defined (**chapter 11**)
2. assess the appropriate degree of human involvement in the decision-making processes, having regard to the nature of the particular function and the statute in question (**chapter 12**)
3. ensure appropriate transparency, including by deciding what can and should be disclosed about the use of machine technology to those whose interests may be affected (**chapter 13**)
4. test before operationalising, and establish ongoing monitoring, audit and review processes (**chapter 14**)
5. consider whether legislative amendment is necessary or prudent (**chapter 15**).

1.6 The role of Parliament in authorising machine technology

If legislation is introduced to enable the use of machine technology, then this provides an opportunity for public and Parliamentary debate on the properties that should be required of that technology.

Whether or not these are ultimately prescribed as mandatory requirements in the legislation itself, the kinds of questions that might be asked of government agencies that are seeking legislative authorisation of machine technology could include:

Properties	Example of qualities that could be prescribed
Is it visible ?	What information does the public, and especially those directly affected, need to be told regarding the involvement of the machine, how it works, its assessed accuracy, testing schedule etc? Are the design specifications and source code publicly available – for example as ‘open access information’ under the <i>Government Information (Public Access) Act 2009</i> ? Is an impact assessment required to be prepared and published? ⁸
Is it avoidable ?	Can an individual ‘opt out’ of the machine-led process and choose to have their case decided through a manual (human) process?
Is it subject to testing ?	What testing regime must be undertaken prior to operation, and at scheduled times thereafter? What are the purposes of testing (eg compliance with specifications, accuracy, identification of algorithmic bias)? Who is to undertake that testing? What standards are to apply (eg randomised control trials)? Are the results to be made public?
Is it explainable ?	What rights do those affected by the machine outputs have to be given reasons for those outcomes? Are reasons to be provided

	routinely or on request? In what form must those reasons be given and what information must they contain?
Is it accurate ?	To what extent must the predictions or inferences of the machine be demonstrated to be accurate? For example, is 'better than chance' sufficient, or is the tolerance for inaccuracy lower? How and when will accuracy be evaluated?
Is it subject to audit ?	What audit records must the machine maintain? What audits are to be conducted (internally and externally), by whom and for what purpose?
Is it replicable ?	Must the decision of the machine be replicable in the sense that, if exactly the same inputs were re-entered, the machine will consistently produce the same output, or can the machine improve or change over time? If the latter, must the machine be able to identify why the output now is different from what it was previously?
Is it internally reviewable ?	Are the outputs of the machine subject to internal review of a human decision maker? What is the nature of that review (eg full merits review)? Who has standing to seek such a review? Who has the ability to conduct that review and are they sufficiently senior and qualified to do so?
Is it externally reviewable ?	Are the outputs of the machine subject to external review or complaint to a human decision maker? What is the nature of that review (eg for example, merits review or review for error only)? Who has standing to seek such a review? If reviewable for error, what records are available to the review body to enable it to thoroughly inspect records and detect error?
Is it compensable ?	Are those who suffer detriment by an erroneous action of the machine entitled to compensation, and how is that determined?
Is it privacy protective and data secure ?	What privacy and data security measures and standards are required to be adhered to? Is a privacy impact assessment required to be undertaken and published? Are there particular rules limiting the collection, use and retention of personal information?

1.7 The way forward – starting with increased visibility

We are hopeful that this report will contribute to public and especially Parliamentary debate about the adoption of machine technology by government, and its proper limits and regulation.

In the final chapter of this report we identify avenues for future consideration, including a question around whether some forms or applications of machine technology might raise such significantly new issues and risks that consideration should be given to new forms of regulation – including mandatory requirements around transparency, pre-operation validation testing and routine auditing, and external review and oversight (**chapter 16**).

One risk, for example, may be that machine technology will be capable of producing new forms of extremely large-scale systemic injustices, to which the existing framework and institutions of administrative law are ill-equipped to respond.

However, a significant impediment to meaningful debate about the future governance of machine technology use by government is an almost complete lack of transparency about that use.

As mentioned above, we do not know how NSW Government agencies may currently be using machine technology to assist them in the exercise of statutory decision-making functions – and so we do not know how those systems have been designed, what they are being used for, and what (if any) assurance has been obtained that they are operating lawfully and in accordance with principles of good administrative practice.

This is a significant problem. Some technology use may be lawful and appropriately designed and used, but other technology may not.

While we do not consider that visibility is, of itself, a sufficient remedy to address potential concerns that might arise with the use of machine technology, it is an essential starting point.

Following this report, therefore, we will seek to work with relevant bodies, including Digital NSW (part of the Department of Customer Service) and the Office of Local Government, to comprehensively map current and proposed types and uses of machine technology (**chapter 2**). We will also look inward to consider what more we can do to support agencies and citizens, as well as our own staff, to understand the use of machine technology – and to ensure that administrative law and the enduring values of good public administration, including legality, transparency and fairness, are given central attention.

Part 1:

Machine technology



2. Introduction

2.1 The rise of machine technology

Use of machine technologies is increasing in the public sector, and their sophistication and use will only grow in the future.

Recent NSW Government announcements reveal an intention to increase work on – and investment in – machine technology. In September 2020, the NSW Artificial Intelligence (AI) Strategy was released which is ‘focused on improved service delivery and government decision-making’.⁹ At the same time, the Government also released the *Artificial Intelligence (AI) Ethics Policy*.¹⁰

In **chapter 3** we discuss what machine technology is, how it is currently being used by governments, and how it may be used in the future.

2.2 Why is the Ombudsman interested in machine technology?

We are always concerned to ensure that government agencies and officials conduct themselves lawfully, make decisions reasonably, and treat all individuals equitably and fairly.

One of our primary functions is to handle complaints about the conduct of government agencies and public officials. We can generally investigate these complaints if we think that conduct may fall within any of the following categories set out in section 26 of the *Ombudsman Act 1974*:

- (a) contrary to law,
- (b) unreasonable, unjust, oppressive or improperly discriminatory,
- (c) in accordance with any law or established practice but the law or practice is, or may be, unreasonable, unjust, oppressive or improperly discriminatory,
- (d) based wholly or partly on improper motives, irrelevant grounds or irrelevant consideration,
- (e) based wholly or partly on a mistake of law or fact,
- (f) conduct for which reasons should be given but are not given,
- (g) otherwise wrong.¹¹

Conduct of the kinds set out above may be said to constitute ‘maladministration’. Where we suspect maladministration, we can also make inquiries about and investigate conduct on an ‘own motion’ basis, without the need for someone to have made a complaint.

One way that an agency’s conduct may be constitute maladministration (ie be unlawful or unreasonable or unjust, etc) is if it is using machine technology in a way that is inconsistent with administrative law and principles of good administrative decision-making. This report highlights some of the ways this might happen.

Going forward, we will consider what further guidance we can provide to help agencies and public officials understand the matters that we will consider when handling complaints about the use of machine technology in the performance of their administrative functions.

Reviewing ‘decisions’ versus investigating ‘conduct’

It has been observed that, if a human decision maker is fully replaced by a machine to exercise administrative functions, one potentially adverse consequence may be that certain rights to challenge the exercise of those functions in court could be lost.¹²

This is because some rights may be premised on there being a ‘decision’ that can be the subject of challenge. As noted in **chapter 12**, the Federal Court has suggested that an essential element of a decision generally is that a relevant decision maker has engaged in a *subjective mental process* of reaching a conclusion. As an autonomous machine does not have a subjective mental capacity, a ‘decision’ of the machine may not be recognised by law as a decision.¹³

However, the automation of some or all of an agency’s activities should not limit the jurisdiction of the Ombudsman to receive complaints and undertake investigations about those activities. This is because the Ombudsman is concerned with *conduct*.

Under section 5(1) of the *Ombudsman Act*, conduct of a public authority is defined as follows: conduct means—

- (a) any action or inaction relating to a matter of administration, and
- (b) any alleged action or inaction relating to a matter of administration.

Conduct includes (but is much broader than) actions involved in making or implementing a decision.

For example, any or all of the following could be scrutinised by an Ombudsman to determine whether *conduct* has occurred that is unlawful, unreasonable, improperly discriminatory, unjust or otherwise wrong:

- the decision to adopt machine technology
- the way the machine has been designed
- the data used by the machine
- the policy and business rules underpinning the machine
- the people involved in designing and building the machine, what consultation occurred, and any external procurement activities
- whether and how the machine was validated, tested, audited and monitored
- whether and what safeguards were put in place to identify and address potential algorithmic bias
- whether and what information has been disclosed publicly about the machine
- the use of the machine for particular functions of the agency
- the consideration or effect that is given to the outputs of the machine, either generally or in a particular case.

More generally, if an agency uses machine technology then whatever that machine does will be ascribed to the agency itself – at least for the purposes of an Ombudsman investigation. Accordingly, if the processes or outputs of machine technology are unlawful, unreasonable, unjust or improperly discriminatory, then the agency’s conduct in using that machine will likely be considered by us to have been unlawful, unreasonable, unjust or improperly discriminatory under s 26 of the *Ombudsman Act*.

2.3 What has prompted this report?

The immediate impetus for this report was an investigation we commenced following complaints we received about Revenue NSW.

Revenue NSW's garnishee machine

Revenue NSW is the Government's debt collecting agency. The head of Revenue NSW, the Commissioner of Fines Administration, is permitted by legislation to issue garnishee orders to recover debts in certain circumstances. A garnishee order is an indirect method of recovering a debt from someone. The order allows a creditor, such as Revenue NSW, who is owed money by a debtor to recover the debt by obtaining payment from a third party who owes money to the debtor. Third parties who can be garnisheed include a person's employer (who owes the debtor their salary) or a person's bank (who owes the debtor what is held in the person's savings account). This power to issue garnishee orders is a debt recovery method that originated in (and is still available from) courts of law.

The complaints we handled were about Revenue NSW garnisheeing the bank accounts of people who had failed to pay fines they owed to the Government. Many of the complainants were financially vulnerable individuals who, in some cases, had their bank accounts emptied.

We engaged with Revenue NSW, and over time we became satisfied by the steps it was taking to address the issues raised in the complaints. For example, Revenue NSW adopted a 'minimum protected balance', meaning that its garnishee orders would not result in bank accounts being completely emptied and left with a nil balance. This 'minimum protected balance' protection was later put into legislation.

During our investigation, we became aware that Revenue NSW had been using machine technology in the exercise of those garnishee powers. However,

those who complained to us about Revenue NSW's activities were not complaining about the use of machine technology – they were not even aware of it.

They were just concerned that their money had been taken, in some cases leaving them with no money in their account to pay rent or buy groceries.

As we learned more about how Revenue NSW was using machine technology to issue garnishee orders, we became increasingly concerned about the lawfulness of its conduct. We used our power under section 31AC of the Ombudsman Act to make a number of formal suggestions including that Revenue NSW seek expert legal advice on the legality and design of its machine technology system.

Revenue NSW responded positively to our suggestions – for example, by publishing a new hardship policy. We decided to discontinue our investigation on the basis of actions Revenue NSW was already taking, as well as future actions it told us it would take.

Eleven months after we suggested that Revenue NSW seek legal advice, we followed up to check on any legal advice received and any action it had taken in response. We were advised that it had not sought that legal advice, either externally or from the legal branch of the Department of Customer Service (DCS), of which Revenue NSW is a part. Revenue NSW advised us that it did not consider it necessary to seek such advice, as it considered that changes it had made to its process for issuing garnishee orders had addressed any potential legal concerns.

We continued to have doubts and decided to seek our own legal advice about Revenue NSW's system – this also helped inform our broader understanding of the legal issues associated with public sector use of machine technologies. Revenue NSW cooperated throughout the process of seeking that legal advice, including by assisting in the preparation of a detailed statement of facts that we then provided to legal counsel for the purpose of obtaining their advice.

The legal advice, and the Revenue NSW case study, is set out in full in **annexure A** of this report. We understand Revenue NSW is currently considering further changes to its garnishee system. We will continue to monitor developments.

2.4 The purpose and structure of this report

This report provides a starting point for agencies and their officials to better understand why and when the Ombudsman (and other bodies, including courts) could hold concerns about their adoption and use of machine technology, and to identify some proactive steps they could take to ensure compliance with principles of administrative law and good practice.

It is not intended to be a comprehensive guide, either to the technology or to the legal and practice issues that its use might raise. Rather, we highlight some of the more important issues that we foresee will likely arise with the use of machine technology.

In doing so we hope to contribute to public debate about these technologies, and in particular their use by government, with a view to ensuring that fundamental and enduring 'public law values' are placed squarely at the centre of those discussions.¹⁴

We recognise that both machine technology and administrative law are, in their own different ways, highly technical fields that can be challenging for non-experts to understand. Indeed, one reason why machine technology use in the field of government administration may be particularly risky is because those who are expert in machine technology may lack experience in administrative law, and vice versa.¹⁵

However, we have sought as far as possible to write this report in non-technical language. Our hope is that it can be read and understood by any agency official likely to encounter machine technology, and by policymakers and the general public.

In this report we:

- outline what we mean by machine technology, its potential benefits and how we see our role in this context (**part 1**)
- examine and highlight the intersection between machine technology and administrative law and practice (**part 2**)
- offer some practical suggestions for machine technology design and implementation (**part 3**).

The report includes a number of short case studies as examples, as well as the more detailed case study of Revenue NSW's use of machine technology (**annexure A**).

We end the report with a question about whether there is a need for new laws – not to restrain innovation, but to ensure appropriate governance, transparency, accountability and oversight in government use of machine technology (**chapter 16**).

2.5 What we will do next

The NSW Parliamentary Research Service recently observed that, while there had been some international progress on transparency of automated decision-making,

*no Australian jurisdiction appeared to be working on creating a registry of automated decision-making systems.*¹⁶

Following the publication of this report, we will seek to work with relevant bodies – including Digital NSW (part of the Department of Customer Service) and the Office of Local Government – with a view to mapping in detail the types of machine technology currently in use, or under development, across NSW Government and Local Government. Following that work, we will explore whether there is a need for a centralised registry or other approaches to enhance transparency on an ongoing basis, such as by mandating that each individual agency make details of their machine technology use publicly available as ‘open access information’ under the *Government Information (Public Access) Act 2009* (as has been suggested by the NSW Information Commissioner).¹⁷

We will also begin work to develop more practical and comprehensive guidance to support agencies, recognising that the internal and external demands for them to adopt machine technology will inevitably continue to grow.

In particular, we will:

- Prepare a new edition of our publication, *Good Conduct and Administrative Practice: Guidelines for State and Local Government*,¹⁸ to include guidance around the use of machine technology
- Update our training services, including in particular our course on *Administrative Law in the Public Sector*,¹⁹ to specifically address the implications of machine technology on administrative law and practice.

3. The new technologies

3.1 What we mean by ‘machine technology’

The continual, rapid pace of technological change means that the terms used to name and describe various technologies are not settled and can differ depending on the context.

In this report we have chosen to use the term **machine technology** to refer to a broad cluster of current and future systems and processes that, once developed, run with limited or no human involvement, and whose output can be used to assist or even displace human decision-making (and specifically in the context of this report, within a public sector administrative context).²⁰ The complexity of this technology ranges from relatively rudimentary to extremely sophisticated.

A machine in this context does not necessarily mean a computer or other physically embodied device. Machine technology will often take the form of software code and, as we will see from the example in **chapter 14**, it may even involve a methodological tool that can be operationalised by simply using pen and paper.

We have sought where possible to avoid the use of terms such as ‘artificial intelligence’ (AI) or ‘automated decision-making’,²¹ although these would generally be covered by what we mean by machine technology.

Our focus is not on the technical aspects of machine technology, but on its outcomes and the risks involved in using it in the public sector.

3.2 Machine technology is not just one thing

While we have not attempted to define or classify the various types of machine technology that are currently in use and under development, one important distinction is between machine technology that adopts a ‘rule-based’ approach and those that involve adaptive ‘machine-learning’ techniques:

- A **rule-based process** is one that simulates a human decision-making process by following a logical set of rules or formulae which could ultimately be reduced to an expression (or a series of expressions) in the form of: ‘If x, then do y; if not-x, then do z’.²² This is sometimes described as ‘human coding’ or ‘good old fashioned AI’ – in which human programmers construct explicit rules for intelligent behaviour.²³

A critical feature of a rule-based process is that, at least in theory, its rules could be written out in a way that would be comprehensible to a human, or at least one who also understands the language of computer coding. In practice, however, some rule-based systems may involve such density and complexity that no human could ever realistically grasp their full end-to-end process.

Rule-based processes are often used to perform functions at scale because of bulk (‘brute force’) processing capability such as data-matching, processing online forms, calculation of amounts, and issuing of system-generated notices and correspondence.

- A **machine learning process** is one that first uses historical data known as ‘training sets’ – which may include the machine’s own ‘experiences’ – to identify correlations and patterns in data. It can then be fed new, previously unseen ‘real world’ data and make inferences and predictions based on whether and how that new data matches the correlations or patterns previously recognised in the training sets.

It does this by determining ‘features’ of the data and assigning ‘parameters’ (that is, weights) to those features by identifying, typically through an iterative process of trial and error, which of all the possible features and parameters optimise the proportion of ‘right’ inferences and predictions that it makes over time.

These systems are said to ‘learn’ because they are ‘capable of changing their behaviour to enhance their performance on some task through experience’²⁴ and without being explicitly programmed.

Machine learning systems can be used for various functions, including grouping together cohorts of people based on characteristics or categorising images.

Of course, a decision-making system may combine elements of machine learning processes and rule-based processes (as is the case in the current Revenue NSW garnishee system – see **annexure A**).

3.3 The need to design machine technology for particular applications

Whichever type of machine technology is used, every particular application will be unique to the task it has been designed for.

Leaving aside the speculative possibility of some future ‘general AI’ (that is, an intelligence able to understand and learn intellectual tasks equalling or surpassing that of humans),²⁵ every application of machine technology to a particular administrative decision-making context requires human designers to make decisions about the design of that technology in that context.

Even machine technology that may have self-improving (learning) capabilities will require humans to make a multitude of design choices. For example, human designers will ‘collect, curate and label’ training sets from which the technology will learn.²⁶ Human designers will set the objectives – that is, what it is the technology is learning to optimise. Further, while machine learning technology learns its own parameters through complex and iterative processes of ‘trial-error-adjustment-retrial’, there are various deeper aspects of the technology (known as ‘hyperparameters’) that humans must set up or ‘tune’ before learning can begin.²⁷

Machine technology is not just used for an administrative decision-making task; it must first be designed and built for such use.

We return to this important observation in **part 3** below, when we consider what steps can be taken to better ensure that machine technology is designed and built so that it is not used in a way that may be unlawful or result in maladministration.

4. The promise of machine technology in government decision-making

4.1 Machine technology within a decision-making system

The extent to which humans might be involved in the implementation of a system that utilises machine technology can vary widely. Generally speaking, where humans play some active role, the system can be referred to as a ‘human-in-the-loop’ system. For the purposes of administrative decision-making, the most important type of human-in-the-loop system is a ‘human-on-top’ system. In these instances, the final step in the system – say, to grant a permit, approve an application or provide a benefit – is ultimately made by a human with the outputs of the machine technology being used to inform or support their decision.

At the extreme other end are fully automated systems, in which the outputs of the machine technology (for example, to issue or cancel a licence of some kind) are both generated and actioned (that is, given effect as an administrative decision) without any intervening human decision-making or approval.

As we explain further in **chapter 7**, understanding the extent to which decision-making is automated and what role, if any, humans have in the performance of a function will often be critical to determining whether the use of the machine technology has been lawful and appropriate.

Mobile phone detection cameras

Since 1 March 2020, Transport for NSW (**TfNSW**) has been using mobile phone detection cameras, including fixed and transportable cameras, to identify drivers illegally using mobile phones.

In New South Wales it is generally illegal for a driver to use a mobile phone except in limited circumstances, such as to answer a phone call using hands-free Bluetooth or voice activation, or where the phone is in a fixed cradle that does not obscure vision.²⁸

Images taken by TfNSW’s mobile phone detection cameras are reviewed using machine learning (AI) systems to filter those images that may show potentially illegal mobile phone use while driving. The machine technology ‘automatically reviews images and detects potential offending drivers, and excludes images of non-offending drivers from further action.’²⁹

Images are then reviewed by an authorised officer and penalty notices are issued by Revenue NSW if illegal mobile phone use is determined.³⁰ This human verification process is an example of a human-on-top system, and is similar to the checks performed before a penalty notice is issued in relation to camera-detected speeding and red-light offences.³¹

In September 2019, to address concerns that the courts may be inundated with spurious challenges to infringement notices based on images caught by the cameras, the NSW Government introduced legislation that would ‘reverse the onus of proof’ by deeming that an object being held by a driver and shown in a photograph from a device approved for mobile phone use offences is a mobile phone unless the accused driver can establish that it is not.³²

The Bill passed the Legislative Assembly on 15 October 2019 and has been introduced into the Legislative Council. Debate on the Bill was adjourned on 20 November 2019, and has not resumed.

The Bill contains no provisions concerning the publication of information about, or the testing or auditing of, the technology. A Parliamentary Committee report on the Bill stated that:

While it is acknowledged that there will be human intervention and review prior to any infringement notices being issued, the task of winnowing down millions of images to identify prima facie criminal conduct will still be handled by artificial intelligence. Given this, there should be transparency in how the artificial intelligence identifies potential offenders including the ability to test whether or not the algorithms contain any inadvertent or inherent biases.³³

The Bill also contains no provisions relating to the use or destruction of captured images. The mobile phone detection scheme relies on broad permissive provisions in the *Privacy and Personal Information Protection Act 1998*, which allow the collection and use of personal information for law enforcement purposes.³⁴ In the second reading speech for the Bill, the Minister stated:

In relation to privacy, information relating to drivers and passengers is captured for law enforcement and road safety purposes only. As committed to during the introduction of the Road Transport Legislation Amendment (Road Safety) Act 2018, Transport for NSW undertook detailed consultation with the NSW Privacy Commissioner during the pilot of the program, and will continue to engage with both the Privacy Commissioner and the Information Commissioner on the rollout of the program.³⁵

On 19 November 2019, the NSW Privacy Commissioner issued a media release noting that:

The Privacy Commissioner provided advice and assistance to the agency to ensure that privacy rights were considered, and appropriate risk mitigation strategies put in place to minimise privacy harm to the public such as:

- minimising the retention of images
- cropping or pixilation of images when viewed for verification purposes
- the use of strong encryption and other security measures
- need for strong contractual requirements on any provider to comply with the PPIP Act.³⁶

There is otherwise limited publicly available information concerning the privacy protection measures that have been put in place.³⁷

More than 260,000 penalties have been issued since the mobile phone detection cameras became operational on 1 March 2020.³⁸

With the learning capability of the system, TfNSW expects that the machine technology will improve over time, meaning that it will become more accurate at detecting potential mobile phone use.³⁹ The technology is also currently being tested for use in detecting seatbelt offences.⁴⁰

4.2 How are governments using machine technology?

The use of simpler forms of machine technology in public sector decision-making is not new. However,

what is changing is the power, complexity, scale, and prevalence of machine technologies, and the extent to which they are increasingly replacing processes that have, up to now, been the exclusive domain of human decision-making.

One of the first such machine technology systems used by the NSW Government was the (then) Department of Fair Trading's automated business-name registration process in 1999. That system 'supported the registration of business names and the incorporation of associations'.⁴¹

Today, use of machine technology in the NSW public sector is likely to be extensive, and it is growing rapidly. We say 'likely', because, as already noted, there is currently no mandatory reporting or means of comprehensively tracking technology use by the NSW Government.

It is, however, clear that over the 2 decades since the adoption of an automated business names registration process (and as described in the examples throughout this report) machine technology is becoming a significant component of much government service delivery. Across NSW Government agencies, machine technology is a critical tool in a wide variety of areas, from traffic and fines enforcement⁴² to assessment of child protection risk,⁴³ from determining grants of legal aid⁴⁴ to triaging claims from injured workers.⁴⁵

Machine technology is also heavily relied on to deliver Australian Government services, including social services, immigration, and taxation. For example, the Australian Tax Office (ATO) has said that it is delivering greater automation and digital services and is using machine learning to accelerate decision-making.⁴⁶

The rise of machine technology is a global phenomenon and is increasingly being used by governments around the world for delivery of core government business.⁴⁷ Internationally, it has been frequently observed that machine technology is disproportionately used in areas that affect 'the most vulnerable in society' – in areas such as policing, healthcare, welfare eligibility, risk scoring and fraud detection.⁴⁸

Automating procedural court rules

In *Hemmett v Market Direct Group Pty Ltd [No 2]* [2018] WASC 310, a claimant had had his proceedings to recover a debt dismissed by an automated court case management system. The system was programmed to 'dismiss' claims, without human intervention, when claimants had not taken any action for a prescribed period of time. The claimant in this case was unable to bring a new claim, as the limitation period had expired.

The Supreme Court of Western Australia set aside the dismissal of the claimant's case, but on a technical point. Under the court rules, a case could only be dismissed from an 'Inactive Cases List', and the automated case management system did not keep such a list. Although the claim may have been registered as inactive in the system, that did not mean it was on an *actual* list of inactive cases, and so it could not be dismissed.

The Court did not need to consider whether the court rules around dismissing inactive cases could be automated. However, the Court gave a ‘provisional view’ that ‘a degree of automated decision-making – better described as ‘technology-assisted decision-making’ – may be permissible’ provided it ‘preserves accepted accountability structures’.⁴⁹

4.3 The potential benefits of machine technology

*There are many potential positive benefits of machine technology for public sector agencies and citizens.*⁵⁰

Frequently cited benefits include:

- **Efficiency and cost savings for government:** there are clear efficiency benefits and cost savings for agencies using technology to streamline processes and perform repetitive tasks at scale. In addition to making faster decisions, machine technology may also enable agencies to reach more citizens than is possible using staff alone. Increased efficiency also benefits citizens interacting with the government.
- **Reduced red tape:** citizens can benefit from machine technology as decisions are made faster and require less direct and unnecessary engagement with government. This is premised on a view that ‘citizens have limited time and energy to engage with government.’⁵¹ This is certainly the case where government activities impose red tape – resulting in time and financial costs for agencies and citizens as they identify and meet regulatory requirements. The NSW Government has said that it is seeking to use technology to ‘make compliance easy’⁵² for the citizen.
- **Increased accuracy:** machine technology is less prone to *certain types* of possible errors arising from inherently human frailties such as distraction, fatigue, negligence or lack of training. In this context, machine technology promises to be more accurate than human decision makers. Tools that use machine technology to support (rather than fully replace) staff through complex legislative rules could also support greater accuracy in the performance of functions.
- **Improved consistency:** as the output of machine technology is limited to what it has been designed to process, it will produce consistent outcomes. Some machine technology will produce more consistent outcomes than multiple human decision makers (although there may be exceptions, such as some machine learning processes that ‘learn’ over time, which prioritise improved accuracy over consistency across time: see **chapter 9**).

A related benefit is the ability to produce an audit trail of the steps taken to reach an outcome, which might in some instances be a more comprehensive form of transparency than a human decision maker’s account of how they arrived at an outcome.

- **Increased productivity and re-focusing of staff to ‘higher value’ activities:** the 2019 review of the Australian Public Service (APS) found that about 40% of APS employee time is currently spent on ‘highly automatable tasks.’⁵³ There is potential for machine technology to free up staff to focus on other functions that perhaps cannot (and arguably should not) be automated, such as complex individual case management.

- **Better customer service and experience:** there may be an expectation that government should keep pace with private sector service standards, providing instant, seamless and increasingly digital service. Reliability, speed and simplicity, as well as fewer mistakes, can contribute to a better experience for customers of public services. Automation of large-volume routine tasks could also allow client-facing staff to devote more of their time to providing more complex and caring services, giving more attention to those who need it.⁵⁴
- **Insights and learning:** an indirect benefit of machine technology use in decision-making is that it inherently involves the creation of a rich mine of data about both inputs and decision outputs, which can inform improvements in future public sector policies and practices.

Of course, benefits cannot be assumed to follow as a matter of course. For example, machine technology that has been coded with errors will not only result in inaccurate outcomes, it will also likely result in inaccuracies at much greater scale than would otherwise be possible (see **chapter 9**).

It is important to be realistic about what benefits (and risks) particular technology will deliver in a particular context, and not to allow untested assumptions or utopian beliefs about technology to drive automation strategies.⁵⁵

Part 2:

Administrative law
and machine technology



5. Why administrative law matters for new technology

5.1 Public sector decision-making is different

The use of machine technology by the private sector obviously raises technical, legal and ethical issues. Many of these issues also arise in the context of public sector use of machine technology.⁵⁶

Questions about the ethics of permitting risk allocation decisions to be made by autonomous devices, the collection and use of facial recognition and other personal information, and issues of potential bias and discrimination, are equally important to private and public sector use of machine technology.

However, the use of machine technology in the exercise of the government's administrative functions both heightens the impact of some of those considerations, and raises new ones. Public authorities exercise powers that impact virtually all aspects of an individual's life; there is 'scarcely any field of human activity which is not in some way open to aid or hindrance by the exercise of power by some public authority'.⁵⁷

The inherently 'public' nature of such functions (such as health, education, and transport) and the specific focus of some government service provision on vulnerable groups means that the government's use of machine technology will necessarily, and often significantly, impact most of society. Recipients of government services – unlike customers of private sector businesses – are also typically a captive market, unable to access alternative providers or to opt out entirely if they do not like the way decisions are made and services are provided.

Most importantly, governments do not just provide services; they also regulate the activity of citizens and exercise a monopoly over the use of public power and coercive force – taxation, licensing, law enforcement, punishment, forms of detention, and so on. It is in the exercise of functions like these, which can affect people's legal status, rights and interests, that administrative decision-making principles raise particular issues that are unique to the public sector.

5.2 Good government according to law

The government has a monopoly over public administrative power, but this means that the exercise of that power is controlled through public administrative law.

Any use of machine technology by government agencies must therefore be considered from an administrative law perspective (this is not to disregard or diminish other perspectives, such as broader ethical⁵⁸ and human rights⁵⁹ perspectives).

Ultimately, all administrative law principles may be seen to support a single underlying principle: while citizens may generally do whatever they please unless it is *prohibited* by law,

those exercising public or governmental power must not only avoid what is prohibited by law, they must also do only what they have been authorised by law to do.

That is, a government agency or public official needs express or implied legal authority to make and give effect to an administrative decision. This means that agencies and their administrators may exercise only those functions that have been granted to them – which today is usually done through legislation – along with any ancillary or incidental powers that are necessarily implied to facilitate the exercise of those functions.⁶⁰

The need for legal authority also means that functions can only be exercised ‘by reference to correct legal principles, correctly applied’.⁶¹ Those correct legal principles are concerned with upholding the values of good government decision-making, and include openness, fairness, accountability, consistency, rationality, legality and impartiality.

The ultimate aim of administrative law is good government *according to law*.⁶²

5.3 Existing laws apply to new technologies

When new technology is introduced, it is introduced into an existing legal environment.

No technology is ever introduced into a complete legal vacuum.

The technology may be more or less adapted to that legal environment, and the law itself may be more or less hospitable to the technology. It may also (at least initially) be unclear exactly how the legal environment will accommodate and respond to the technology.

Moreover, the combination of new technology and existing law can generate gaps, inconsistencies or other undesirable outcomes. Where this is the case it may be necessary or desirable to make conscious changes to the law, from minor tweaks to radical overhauls, to meet the challenges of the new technology – a possibility we explore further in **chapter 16**.

Of particular relevance to this report, the legal environment into which machine technologies are now being introduced is one that is governed by public administrative law – the law which controls government decision-making.⁶³

That legal environment includes courts, which undertake judicial review of administrative decisions, and administrative tribunals, which can have a role in undertaking merit reviews of some decisions. It also includes a range of integrity bodies such as ombudsman institutions, which – while they may not make legally binding determinations – have a broad remit in terms of investigating and making findings about wrong administrative conduct.

5.4 How administrative law applies to new machine technologies

Administrative law has developed over many centuries, although many of its modern features have developed in the last half century.⁶⁴ However, it is essentially principles-based and can therefore be considered, conceptually at least, to be ‘technology-agnostic’. This means that,

while the technology used in government decision-making may change, the underlying norms that underpin administrative law remain constant.

There is no reason to expect that administrative law will not evolve in response to the challenges raised by new technology. Indeed, the growth and importance of administrative law over the past half century or so is a by-product of its application and refinement to meet other challenges of modern government such as the rise of the welfare state, large scale bureaucracy, and privatisation.⁶⁵

We can be confident that these laws will continue to be interpreted and applied as the technological context continues to change. A recent survey of Australian academics and legal practitioners about the

impact of information technology on the teaching of administrative law found that ‘many interviewees expressed the view that technological change would not impact fundamental administrative law principles, but instead would be relevant to the interpretation and application of those principles in practice.’⁶⁶ Generally, that is the perspective we have also taken in this report.

In the next chapters we will consider some of the important elements of current administrative law and practice, and how they will likely affect and control the adoption and use of machine technology by government decision makers.

However, we also note in these chapters some potential gaps, or at least uncertainties, in the capacity of existing administrative law rules and associated frameworks to respond to novel issues that may arise with the use of new technologies. In the final chapter of this report, we will ask whether new or additional legal approaches should also be considered.

6. Key administrative law issues for machine technology

In the next 4 chapters we look at some of the key issues raised by administrative law that will likely be most important to machine-aided administrative decision-making.

6.1 The essential requirements for good (and lawful) administrative decision-making

When we provide training to non-lawyers in the public sector on administrative law,⁶⁷ we group the essential requirements of administrative law for good decision-making as follows:

A. Proper authorisation

1. There is a legal power to make the decision.
2. The person making the decision has the legal authority to do so.
3. The decision is within the scope of the decision-making power (including, in particular, within the bounds of any discretion that is a component of the power).

B. Appropriate procedures

4. The decision has followed a fair process.
5. The procedure meets other legal and ethical obligations.
6. Reasons are given for the decision (particularly decisions that affect the rights or interests of individuals).

C. Appropriate assessment

7. The decision answers the right question (which necessitates *asking* the right question).
8. The decision is based on a proper analysis of relevant material.
9. The decision is based on the merits and is reasonable in the circumstances.

D. Adequate documentation

10. The circumstances surrounding the making of decisions are adequately documented and records kept.

Administrative law is obviously more complex than this simple list may suggest, and there are more technically rigorous ways of classifying its requirements.⁶⁸ For simplicity, however, we will stick with the familiar and simple list above.

In the next chapters we will examine the use of machine technology in the context of administrative decision-making by looking at each of the above elements in turn. These chapters are not intended to be exhaustive or definitive. There are myriad ways in which administrative decision-making can go wrong – we are highlighting just some of the more obvious ways things can go wrong when machines are used.

In so doing, we aim to demonstrate why and how

the well-established elements of good administrative decision-making listed above must continue to be given central focus even – or perhaps especially – when new technologies are being used.

Some machine technology use by government will be 'legally unexceptional'

It is important to acknowledge that there will be many governmental uses of machine technology that will likely be considered legally unexceptional, in the sense that it will probably raise few or no significant concerns from an administrative law perspective.

This does not mean that the principles of administrative law summarised above and considered further below (and especially the underlying requirement of government agencies to act only within their legal authority) do not apply. However, the risks of automating some administrative tasks of government will obviously be much lower in some cases.

There are numerous straightforward administrative decisions that are non-discretionary – where 'if X' is the case then the decision must be 'Y' and where the question of whether X is the case will be obviously and incontrovertibly true or false. These would seem to be the kinds of functions that may be suited to automated processes, particularly as the reasons for any decision will be clear and there should be no difficulty in identifying if the decision was wrong, and obtaining redress if it was. In such very simple (and especially high-volume) decision-making, community benefits of automation in terms of accuracy, speed and efficiency may 'count for a great deal'.⁶⁹

7. Proper authorisation

7.1 Legal power to make the decision

We are primarily concerned with the legal principles that govern the performance of *statutory* functions – that is, activities where the source of power is legislation (including Acts, Regulations and other instruments). There are also ‘non-statutory’ sources of government power, such as the powers the Crown can exercise in common with other legal persons (sometimes referred to as ‘executive power’).⁷⁰ These include establishing and running workplaces and other enterprises, entering into contracts, procuring goods and services, and bringing and defending proceedings.⁷¹ Of course, machine technology can support the exercise of non-statutory powers too, and their use there has the potential to raise various legal and ethical concerns as well as potential benefits (see **chapter 4** above).

However, our main concern here is with the exercise of statutory powers. As is the case with statutory powers exercised exclusively by humans, those exercised by, or with the assistance of, machines will only be lawful if they are consistent with the statute that provides the source of the relevant power.

A decision cannot be made to do something that is not within the power given by the relevant statute. This will obviously continue to be true when a machine may be involved in the decision-making process. That point may seem obvious enough. However, its central importance cannot be overstated (see ‘Services Australia Centrelink’s automated compliance program (Robodebt)’ below).

Services Australia Centrelink’s automated income compliance program (Robodebt)

Much has been written about the issues associated with Centrelink’s automated compliance program, and it offers a cautionary tale for government use of machine technology.

Under its compliance program, Centrelink sought to use machine technology to raise and collect debts arising from overpayment of social security benefits and in some cases, apply a discretionary 10% recovery fee. The automated system used by Centrelink was flawed and generated erroneous debt notices. It did this by matching data from the Australian Tax Office with Centrelink data and averaging the income of social security recipients over a period of time. This failed to account for periods of fluctuating income, which were important for the correct calculation of social security benefits.⁷² The erroneous debt notices were sent directly to Centrelink customers following notification of a possible discrepancy in their payment.

The automated process replaced the previous manual fact-finding process. While the manual process may also have used averaged income data to identify and question *possible* overpayment, the automated process now treated this data as *evidence of a debt* under social security legislation. This triggered a shift in responsibility for proof of debt – the alleged debtor was required to prove that a debt was *not* owed.⁷³

Failure to pay could also result in garnishee action by Centrelink, as in the case of *Amato v The Commonwealth of Australia*.⁷⁴ In that matter, declarations and orders were made by consent acknowledging that the Commonwealth could not be satisfied, based on the income averaging method, that a debt was owed by the applicant, and that there was no foundation for imposing a penalty or taking garnishee action.⁷⁵ In November 2019, Centrelink stopped raising debts on the basis of income averaging and a class action lawsuit was filed against the Commonwealth of Australia. Centrelink is currently refunding eligible people who paid debts raised using averaged income information.⁷⁶

Centrelink's automated income compliance program has been subject to two Senate inquiries in addition to scrutiny through the Senate Estimates process. The Senate Community Affairs References Committee made several recommendations including a review of the legal requirements for all Services Australia compliance activities relating to overpayment.

This case study demonstrates the capacity for errors to impact on a large scale when machine technology is involved, and the importance of agencies obtaining and thoroughly considering legal advice when designing machine technologies to ensure they are applying the correct interpretation of legislation.

7.2 Legal authority of the person making the decision

When Parliament creates a statutory function, it gives someone (or more than one person) power to exercise that function. This person must be a 'legal person'. A legal person can be a natural person (a human being) or a legally-recognised entity, such as a statutory corporation. In other words, statutory functions are granted to *someone* who is legally capable of exercising powers and being held accountable for obligations.⁷⁷

Commonly, when Parliament gives someone power to exercise a function, it will also permit that person to formally delegate the function to a delegate.⁷⁸ Those delegates can then perform the function, as long as they comply with any conditions set out in the statute or the instrument of delegation.⁷⁹

Just as a statutory function can only be given by Parliament to a legal person, the function can only be delegated to a legal person.

When a function is delegated, the delegate can independently exercise the function in the same way as the person on whom Parliament conferred the function.⁸⁰ This is the way in which many statutory functions are performed.

At law, if a person purports to perform a function:

- without Parliament having given them the power to do so, and
- without a proper delegation

their exercise of the function may be invalid.

Statutory functions are not, and cannot be, granted to or delegated to a machine.⁸¹ The authority and responsibility for exercising a statutory function can only be conferred on or delegated to a legal subject (a someone) and not a legal object (a something).⁸²

Administrative assistance (the Carltona principle)

Even when a function has not been formally delegated, the person who has been conferred the function may be able to obtain assistance in the exercise of the function. Bodies corporate can only act through human agents, but even human administrators may be assisted in performing their statutory functions, at least to some extent.⁸³

This principle, sometimes referred to as the Carltona principle,⁸⁴ recognises that, in conferring a statutory function on an administrator, Parliament does not necessarily intend that the administrator personally undertake every detailed component or step of the function. As a matter of 'administrative necessity', some elements of a function might need to be shared with others who are taken to be acting on the administrator's behalf. The extent to which performance of functions can be shared under the Carltona principle will depend on the particular statutory function.

The reasoning underlying the Carltona principle appears to be sufficiently general that it could extend to permit at least some uses of machine technology. That is, if the holder of a statutory function, having regard to ‘practical necessity’,⁸⁵ cannot be expected to personally perform every step of a function in every case, there seems no reason why they should be limited to assistance only from human agents. Instead, they may be able share performance of components of the function with a machine.

However, whether using human or machine assistants, the Carltona principle only permits assistance that is consistent with the administrator remaining, at all times, the one who ultimately retains control of the function and is accountable for its performance. There may also be doubt as to whether assistance can extend to activities that are not routine or that involve the exercise of a statutory discretion.⁸⁶

Further, the principle is based on a *necessity* imperative. The holder of a statutory function cannot rely on it to authorise sharing performance of a function merely on the basis that it might be more *efficient* or otherwise desirable to do so.⁸⁷ While it is possible the Carltona principle may be extended in the future,⁸⁸ whether and how that might happen is not clear.

To date, the Carltona principle has been concerned only with the ability of administrators to rely on human agents. The reasoning that underpins that principle means it has the potential also to support some uses of machine technology.⁸⁹

Relevant inputs in decision-making

The Carltona principle is not the only means by which administrators may obtain assistance, whether from other people or other things, to help them better perform their functions.

For example, depending on the particular function, administrators can (and in some cases should, or even must) draw upon the scientific, medical and other technical expertise of others whose input will be relevant contributions to their decisions. Sometimes, this input can even be adopted as a component of the decision of an administrator for certain purposes. For example, an expert medical assessment that provides a report on a person’s level of impairment may be adopted by an administrator for the purposes of then making a compensation decision.⁹⁰

Of course, apart from these expert human inputs, administrators also use traditional forms of technology to provide inputs into their performance of statutory functions: to record, test, calculate, detect, measure, model, and so on.

Inevitably questions will arise as to the extent to which new machine technologies might be recognised as merely an example, or an extension, of these situations.

More simple machine technologies might be viewed as legally comparable to existing tools that administrators are permitted to use in the exercise of their functions.⁹¹ On the other hand, as machine technologies become more sophisticated, might their outputs be recognised as equivalent to the advice of a human expert, where an administrator may take such expertise into account in their decision-making?

We expect that, like the obtaining of expert advice and the use of traditional forms of technology, there will be at least some forms and uses of sophisticated machine technology that will come to be recognised as legitimate tools administrators can use to assist them to perform their functions, within the implicit authority conferred on them by the statute.

However, whether and the extent to which this is so will need to be carefully considered on a case-by-case basis, taking into account the particular statutory function, the proposed technology and the broader decision-making context in which the technology will be used.

7.3 The scope of the decision-making power – including the extent of any discretion

Most administrative powers given to decision makers by legislation involve at least some element of discretion.⁹²

This raises particular challenges when it comes to the automation of decision-making through machine technology, as any automation will need to be consistent with the discretionary nature of the power.

What is discretion?

We are using the term ‘discretion’ in a broad sense.⁹³ A function is ‘discretionary’ where there is no one right outcome, or where no one consideration or combination of considerations is ‘necessarily determinative’ of the outcome in all cases.⁹⁴ This means that there is some element of ‘decisional freedom’ – the outcome will be one about which there is ‘room for reasonable differences of opinion’⁹⁵ or a ‘choice of legally available alternatives’.⁹⁶

A discretionary function is therefore one in which an administrator has some freedom as to any one or more of the following:

- whether to exercise the function
- how to exercise the function
- the output of the function (i.e. what is ultimately done or not done).

That freedom does not have to be absolute. Indeed, no discretion is ever completely unfettered. There will always be some constraint on how the administrator acts. There is also a statutory presumption which, if it is not displaced, requires the discretion to be exercised in a manner that is consistent with standards of ‘legal reasonableness’.⁹⁷

The types of functions that fall within this broad concept of ‘discretion’ include those where an administrator:

- (1) is able to decide *whether or not* to exercise the function on any given occasion (typically identifiable by the use of the term ‘may’ in the relevant legislative provision)⁹⁸
- (2) can exercise the function in more than one way (for example, may grant a licence with or without conditions)
- (3) is called upon, when exercising the function, to take into account a range of factors, some of which may ‘pull in different directions’⁹⁹ without a fixed formula (or ‘recipe’) for applying those factors
- (4) is to evaluate whether someone falls within a category, class or definition contained in legislation (for example a ‘fit and proper person’ or ‘de facto relationship’) that involves an evaluative judgment¹⁰⁰
- (5) is permitted or required to exercise a function only if they possess a particular state of mind (such as ‘reasonable suspicion’, ‘reasonable belief’ or ‘satisfaction’).¹⁰¹

A technical note about ‘discretionary powers’

It should be noted that the concept of discretion we are using here may be broader than what a court or a lawyer will be referring to when they speak of a ‘discretionary power’.

For example, an administrator who is required to do Y if (and only if) satisfied of X may be said to have a legal duty (and not merely a discretionary power) to do Y. This is because, provided they are satisfied of X, they *must* do Y. Although being satisfied of X may involve some element of subjective evaluation or

choice, that does not, under this more technical sense of discretion, change the function from a duty to a discretionary power.

Where it is necessary in this report to distinguish between these different concepts of discretion, we refer to the narrower and more technical concept as a ‘formal discretionary power’. Otherwise, when we just refer to ‘discretion’ we just mean it in its more general sense as explained above.¹⁰²

The imperative to preserve discretion

By giving an administrator a discretion, Parliament has relinquished some element of control over individual outcomes, recognising that those outcomes cannot be prescribed or pre-ordained in advance by fixed rules.

But at the same time, Parliament is also prohibiting the administrator from setting and resorting to fixed rules that Parliament itself did not fix. If Parliament had intended to lay down fixed, pre-determinative rules for the exercise of these functions, it would have done so.¹⁰³ Where it has chosen not to do so, that decision must be respected.

This means that

exercising the discretions that Parliament has given to an administrator is just as important as administrators complying with any fixed rules Parliament has prescribed.

Potential issues with using machines in the exercise of discretionary decisions

Over time, administrative law has developed specific rules concerning the exercise of statutory discretions. These include the so-called ‘rule against dictation’ – see ‘Machine technology and the ‘rule against dictation’ below.

Machine technology and the ‘rule against dictation’

The so-called ‘rule against dictation’ requires an administrator who has been tasked with a discretionary function to exercise that discretion themselves, and not in automatic obedience to the directions or instructions of another.¹⁰⁴

This rule prohibits not just circumstances where the administrator is given an express order or direction from another to act in a particular manner. It also prohibits circumstances where the administrator feels obliged to exercise their discretion in a particular way based on the conclusions or wishes of another – even in circumstances that fall short of a direction or command.¹⁰⁵

This does not mean an administrator cannot *consider* the views of others. However, there is sometimes a fine line between taking into account what is said by another person (which is permitted) and acting at another’s behest (which is not).¹⁰⁶

We are not aware of any judicial consideration to suggest that the rule against dictation might be applied directly to machine technology. However, it is easy to see how the principle that underpins that rule – the need for the nominated administrator to exercise their *own* discretion – could have implications for the lawful use of machine technology by administrators. In particular, an administrator who is under a statutory duty to exercise

discretion may be acting unlawfully if they automatically or unreflectively adopt the instructions, recommendations, advice or output of someone *or something* else, as to do so would not constitute a genuine exercise of their own discretion.¹⁰⁷

There are also rules governing (and limiting) the use of policies and other guidance material to regulate the exercise of discretion – see ‘Machine technology as a means of delivering ‘guidance’ to administrators’ below.

Machine technology as a means of delivering ‘guidance’ to administrators

Administrators exercising discretion can be aided in that task by policies, guidelines and other similar resources (‘guidance material’). Might it be appropriate to consider at least some forms of machine technology as effectively just guidance material delivered a new way? Rule-based systems, in particular, that guide administrators through a series of questions and decision-trees may ultimately contain the same information in substance as what was previously provided to them in a written policy or manual.

The development and use of such guidance material to aid the exercise of discretionary functions is generally recognised, often encouraged,¹⁰⁸ and may sometimes even be required by statute.¹⁰⁹ Guidance material assists in addressing the tension that may exist between the flexibility and individualisation that discretion permits, and the consistency that public law and good administrative practice requires.¹¹⁰ Policies can improve consistency and certainty, mitigating the extent to which outcomes are affected by ‘individual predilection’ (that is, individual preference).¹¹¹

Guidance material (some of which can be relatively detailed)¹¹² is seen as particularly desirable in the exercise of what are sometimes described as ‘high volume’ functions.¹¹³ Use of this material in these contexts may be necessary to avoid ‘substantial injustice’,¹¹⁴ or ‘blinkered and individualised decision-making [which] would be a recipe for maladministration’.¹¹⁵

If machine technology were employed to guide the exercise of discretionary functions, established principles about the lawful use of existing forms of guidance material might be helpful. For example, courts have generally recognised that guidance material to assist in the exercise of discretionary functions will be lawful provided it:

- does not give effect to purposes inconsistent with the purposes of the legislation that created the function¹¹⁶
- leaves the range of discretion intact,¹¹⁷ and does not inappropriately exclude or narrow the interpretation of criteria prescribed by legislation¹¹⁸
- does not create inflexible rules that an administrator cannot depart from in the individual case¹¹⁹ and
- is not *treated* by administrators as giving rise to fixed determinative rules to be adhered to regardless of the merits of the individual case.¹²⁰

Even if a particular machine could, in principle, be considered as conceptually equivalent to some traditional guidance materials, there may be an important distinction in practice: the tendency humans have to uncritically accept and inappropriately rely upon the output of

machine technology systems. If that happens, then there is a risk that the machine moves beyond merely guiding the exercise of discretion (permissible) and instead operates in a way that means that discretion has effectively been lost (impermissible). This concern was noted by the Commonwealth Ombudsman in a 2007 report about a series investigations conducted into immigration decisions.¹²¹

Such rules are best viewed as applications of a more general principle that seeks to preserve the discretion that Parliament has incorporated into the function and conferred on a particular person or persons: where a statute gives discretion to an administrator, the administrator must remain capable of exercising, *and must in practice exercise*, that discretion. Those given a discretionary statutory function must ‘keep their minds open for the exceptional case’.¹²²

Given this principle, there may be risks in using some forms of machine technology in the exercise of statutory functions that have discretionary elements.

This was the view of the Administrative Review Council in 2004. It concluded that, while ‘expert systems’ might be used to *assist* an administrator to exercise a discretionary function, the exercise of the discretion should not be automated and any expert systems that are designed to assist in the exercise of discretionary functions should not fetter the exercise of that function by the administrator.¹²³

In summary, and at least on the current authorities, it should be assumed that:

Machine technology cannot be used in the exercise of discretionary functions if (and to the extent that) it would result in the discretion being effectively disregarded or fettered.

If a discretion has been given to an administrator, the discretion must remain the preserve of that administrator and there must, in law and in practice, continue to be a genuine exercise of that discretion by that administrator.¹²⁴

If the introduction of machine technology into a discretionary decision-making system has the effect that the administrator is no longer able to – or does not in practice – continue to exercise genuine discretion, that system will be inconsistent with the statute that granted the discretion, and its outputs will be unlawful.¹²⁵

In practice, this likely means that discretionary decisions cannot be fully automated by a machine.

8. Appropriate procedures

8.1 The decision has followed a fair process

Good administrative decision-making requires a fair process. The process will only be fair if it is also reasonably perceived to be fair by those affected. There are many elements that make up a fair process – those elements include transparency, accountability, and proper management of expectations.

In administrative law, a core requirement for a fair process is known as ‘procedural fairness’ (sometimes also termed ‘natural justice’). Procedural fairness requires:

- that the decision maker acts free of bias (the ‘no bias’ rule), and
- that those directly affected by decisions be given a genuine opportunity to be heard (the ‘hearing rule’).¹²⁶

Unless clearly excluded by legislation, administrators must apply procedural fairness principles when exercising statutory functions which could affect the rights or interests of individuals.¹²⁷ However, what procedural fairness requires in a given circumstances will vary depending on the legislative context¹²⁸ – there is no one-size-fits-all prescription that will answer all procedural fairness requirements.

The importance of affording procedural fairness may pose several challenges for machine technology.¹²⁹

Bias

One of the most commonly cited risks of machine technology is the introduction or amplification of inherent biases in the input data. This may be because of the way it has been coded, or it may be as a result of the way in which it has ‘learned’ from data sets that were themselves affected by historical or systemic (and perhaps hidden) biases.

However, the no-bias rule is traditionally concerned with the requirement that the administrator bring an ‘impartial mind’ to the making of the decision.¹³⁰ Bias may be ‘actual’ or ‘apprehended’.

Apprehended bias occurs where a fair-minded observer might reasonably perceive that a person might not bring an impartial mind to their task.¹³¹ Actual bias will be established when the administrator is ‘so committed to a conclusion already formed as to be incapable of alteration, whatever evidence or arguments may be presented’.¹³²

It is unclear if and how the no-bias rule will be applied to address other kinds of systemic bias that may be introduced by machines.¹³³ Similarly, it is unclear whether algorithmic bias – systematic and repetitive errors that result in unfair outcomes – may come to be accepted as a ground of ‘bias’ in judicial review proceedings.

Whether or not that happens, algorithmic bias may result in other kinds of unlawful conduct or maladministration – for example, if the effect is that the decision has been based on extraneous considerations’ (see chapter 9), is unreasonable, has involved unlawful discrimination, or otherwise generates systemically unjust or improperly discriminatory outcomes.¹³⁴

Algorithmic bias

One of the most significant concerns about machine technology, and in particular those that apply machine learning techniques, is its potential to reflect and amplify bias against minority and vulnerable groups. These concerns are heightened by the challenges in detecting such biases given the complexity and inherent opaqueness of the technology (see **chapter 13**).

Bias most commonly reflects and amplifies historical biases and inequalities contained, sometimes hidden, in the data sets from which machines learn.

Bias can also arise where that data set is unrepresentative or incomplete. For example, if training data¹³⁵ is more representative of some groups than others, the accuracy of the model's outputs tend to be systematically worse for under-represented groups. This has been observed in the case of facial recognition technology trained disproportionately on lighter skin tones and therefore significantly less accurate for darker skinned individuals.¹³⁶ On the other hand, bias can also result where more data fields are available for some groups than others. For example, in the United States a child welfare screening tool that is able to use data from means-tested programs (such as mental health counselling or drug treatment services) will have that data for lower-income families without having corresponding data on the use of similar services by wealthier families, with the result that child welfare risks may be disproportionately rated higher for poorer families.¹³⁷

Importantly, even training data that does not explicitly include sensitive attributes like race or gender may be susceptible to bias, because a learning algorithm can develop proxies for sensitive attributes. For example, in the United States zip codes can often be a proxy for race. Height or weight may be proxies for gender. Training data that excludes gender fields but includes names might give rise to gender proxies – for example, an algorithm may learn to generate different results if the name on record is 'Tony' or 'Toni'.¹³⁸

This means that an algorithm that is blind to a sensitive attribute may produce a similarly biased outcome as one that overtly uses the attribute in a discriminatory manner.¹³⁹ Indeed, in some cases to simply omit any sensitive attributes may be counter-productive – both because it may lead to complacency and a failure to recognise and address proxies hidden elsewhere in the data, and because it may prevent the designers from building in processes that attempt to 'correct' for historical bias (a so-called 'fairness-through-awareness' approach).¹⁴⁰

The key point is that algorithmic bias may arise without any intention to discriminate, without any awareness that it is occurring, and despite the best intentions of designers to exclude data fields that record any sensitive attributes or any obvious (to humans) proxies. This is one reason why formal evaluation, auditing and ongoing monitoring processes are essential.

Some of the growing number of examples from around the world where machine technology has been shown to generate or amplify bias include:

- Recently in the United Kingdom, government use of machine technology came under scrutiny when the Department for Education used automation to grade school leavers who could not sit exams due to the coronavirus pandemic. The system downgraded the results of large numbers of students in the state school system based on a model which was found to be favourably biased toward private schools. Ultimately, the grades awarded using automation were withdrawn in favour of predictions made by teachers.¹⁴¹
- A team of researchers were recently able to demonstrate hidden racial bias in an algorithm used widely in the US health care system to prioritise referral to a program aimed at improving outcomes for patients with complex medical needs. The algorithm used the cost

of healthcare as a proxy for illness, which resulted in bias against African American patients as they had systemically unequal access to care, meaning that less money is spent on their health care. By failing to take those systemic differences into account, the system effectively assumed that African American people were healthier than they were, and correspondingly required African American people to be sicker in order to be assessed as eligible for additional assistance. The algorithm developer later confirmed the results found by the researchers.¹⁴²

- Amazon’s experimental hiring tool used machine technology to review and score job applicants’ resumes from 1 to 5 stars. The experiment was shown to be biased towards men, because the machine had been trained using the resumes submitted to Amazon over the previous decade, and most of those were from men (in a technology industry still dominated by male employees).¹⁴³ The algorithm had learnt to model predicted employment outcomes based on word patterns in the resumes, rather than relevant skill sets. Consequently, it penalised resumes that included the word ‘women’s’ or referred to women’s colleges. Although Amazon ‘scrubbed’ the data to prevent it from overtly discriminating based on those parameters, there was no way to ensure the algorithm would not learn an alternative model that would also unfairly sort and rank male candidates higher, and the algorithm was scrapped.¹⁴⁴
- Amazon also used an algorithm to decide which neighbourhoods would be eligible for, or excluded from, its same-day Prime delivery system. The decision relied on whether a neighbourhood had a sufficient number of existing Prime members, proximity to a warehouse, and availability of willing delivery couriers. The purpose was to exclude unprofitable neighbourhoods. However, the result was unintentionally discriminatory, as the model resulted in the exclusion of poor and predominantly African American neighbourhoods.¹⁴⁵
- A Georgetown Law School study found significant overrepresentation of African American people in ‘mug-shot’ data bases. This meant that facial recognition networks used by law enforcement produced a biased effect, as the faces of African American people were more likely to be falsely matched.¹⁴⁶
- In a well-reported case, the Allegheny County Department of Human Services purchased a decision tool (the Allegheny Family Screening Tool) to generate scores as to which children are most likely to be removed from their homes within two years or to be re-referred to the child welfare office for suspected abuse. The tool was rebuilt after the County undertook an independent evaluation, which identified statistical unfairness, including racial bias.¹⁴⁷

A right to be heard

The hearing rule might typically require an administrator to notify an individual of a possible or proposed decision or course of action and invite them to respond, with information or arguments, before the decision is finalised or course of action taken.¹⁴⁸

One of the important questions that arises here is whether it is necessary, before the affected person is invited to provide their views on the proposed decision or action, to inform them if a machine has proposed, or been involved in proposing, that decision or action?

We are not aware of any court decision that has directly addressed this question. We expect that, from a strictly legal perspective, the answer may be that it depends on the particular decision or action in question and the nature and extent of the particular machine process involved.

However, as a matter of good administrative practice, our view is that this information should always be disclosed to the person.

Even if the administrator does not consider that the machine's involvement could be relevant in any way to anything the person might possibly wish to put forward for consideration, the person may have a different view.

Perhaps even more importantly, the mere fact that a machine has made (or was substantially involved in making) the proposed decision may be important to the person in deciding whether or not to make any submission at all. Knowing that an adverse decision was merely the output of some machine process may impel the affected person to make a submission when they otherwise might have simply accepted the outcome. If a decision is going to be made that affects them adversely, and even if it turns out that the decision is correct and cannot be changed, they may reasonably want to ensure that their situation has at least been considered by someone with a genuine (human) capacity to understand the decision and its consequences and impact for them. The right to be heard is not just about ensuring the correct decision is made in light of all relevant considerations – it is also a right of 'respect' to the person affected by the decision.

In our view, as a matter of good administrative practice, a right to be heard before a decision is finalised generally requires the person also to be told if a machine has made, or has materially contributed to the making of, the proposed decision.

8.2 Other legal and ethical obligations

Compliance with other lawful and ethical obligations in an administrative decision-making process involves such things as acting honestly and avoiding conflicts of interest. To some extent, the use of machine technology could help to mitigate the risk of contravening these obligations where they might otherwise result from human failings.

However, the requirement to meet other obligations also means complying with laws beyond the statute that creates the relevant function. This reflects an important element of the 'rule of law': like citizens, government must also abide by the law.

Some of these other laws, such as those governing privacy, freedom of information and anti-discrimination, have general application to most administrative functions, but will have particular implications for processes that involve machine automation:

(a) Privacy

Administrative agencies are required to comply with general privacy obligations concerning the collection, storage, use and disclosure of personal information. In NSW, those obligations are imposed by the *Privacy and Personal Information Protection Act 1998 (PIPPA)* and the *Health Records and Information Privacy Act 2002 (HRIPA)*.

Concerns have been raised about machine technology within the broader context of data protection, as governments increasingly digitise their operations. The use of machines will, in many instances involve collecting, translating and reducing personal information to a form that is suitable for use by machine technology.¹⁴⁹ That process will also be coloured by assumptions and interpretations of the designers of the system in determining what personal information is relevant to code and the significance of any particular piece of information.¹⁵⁰ This may raise particular issues for agencies' obligations concerning the currency, accuracy and completeness of the personal information they hold, and whether that might be potentially misleading. There may also be issues

to consider relating to obligations concerning the permitted uses, retention, safe storage, and destruction of the personal information that is being held.

Agencies will need to have proper regard to these issues in the design and implementation of machine technology systems. The NSW Privacy Commissioner has noted that presently there is no mandatory legal requirement even to conduct a privacy impact assessment before adopting machine technology, even though such technologies ‘can give rise to unique and complex privacy issues’.¹⁵¹

While allegations of violations of a person’s privacy are excluded from the Ombudsman’s jurisdiction (as they are matters for the NSW Privacy Commissioner), we suggest that a privacy impact assessment should be included as an essential element of any machine technology design process, and it should be made public.

(b) Freedom of information

NSW’s primary freedom of information law – the *Government Information (Public Access) Act 2009* (NSW) (**GIPA Act**) – deals with information and the records that hold them in a way that is intentionally technology-agnostic.¹⁵²

The aims of the legislation are to open government information to the public to maintain and advance a system of responsible and representative democratic government.

The GIPA Act places various obligations on agencies within NSW in respect of the publication and release of the information that they create and hold. The GIPA Act also provides rights for people to apply for access to government information.

These rights remain applicable where government uses technology to provide services and inform decisions.¹⁵³

The NSW Information Commissioner has issued guidance noting that:

This technology [automated decision-making systems that involve a computerised process that either assists or replaces the judgement of human decision-maker] can perform many functions that previously could only be done by humans. As these systems are adopted by governments, citizens will increasingly be subject to actions and decisions taken by, or with the assistance of, automated decision-making systems. To fully exercise their rights, it is important that individuals are able to access information on how a decision is made and what information was used to reach that decision.¹⁵⁴

(c) Anti-discrimination

Under the *Anti-Discrimination Act 1977* (NSW), it is unlawful, in the provision of services and in a broad range of other contexts, to discriminate against a person because of any of the following characteristics: disability (including disease and illness), sex (including pregnancy and breastfeeding), race, age, marital or domestic status, homosexuality, transgender status and carer’s responsibilities.¹⁵⁵ The protections afforded by this Act are long established and similar laws apply in other Australian jurisdictions.

Discrimination may be direct or indirect.¹⁵⁶ The use of machine technology can result in outcomes that involve either direct or indirect discrimination.¹⁵⁷ Furthermore, the use of such technology may also make it more difficult to detect or to understand if unlawful discrimination is occurring.

Apart from the difficulties in obtaining accessible information about how the technology has generated its output, there may be a tendency to assume that these systems are neutral, free of human bias and therefore incapable of unlawful discrimination.¹⁵⁸ This may lead agencies to discount or minimise the need to ensure the operation of machines comply with anti-discrimination obligations.

However, in the case of pre-programmed rules, a machine will obviously reflect and potentially amplify any intended or unintended pre-existing bias or assumption which influenced its programming.¹⁵⁹

More recent machine learning technologies can present a more complex challenge. Current forms of machine learning develop effectively their own rules based on statistically significant correlations from available data. Such self-developed rules may be constructed from and promote proxies for protected characteristics within their decision-making matrices, whether as a result of assumptions in the training data or learnt correlations.¹⁶⁰ The learning capabilities of these systems also means these rules are not static and may change over time, for better or worse.

This means that it is not sufficient for agencies to design machines in ways that are not overtly discriminatory, and to not use training data sets that explicitly record protected characteristics. Agencies will need to test, and then regularly monitor and audit, the operation of their learning machine to ensure it continues to operate in accordance with anti-discrimination legislation.¹⁶¹

There may also be less obvious legal rights and obligations that interact with particular statutory functions, which must be considered when functions are to be handled by machine technology. Few statutory functions operate independently of any other written or unwritten laws.

In the Revenue NSW case study, for example, the function being exercised was the issuing of garnishee orders on banks and other financial institutions under the *Fines Act 1996* (NSW). These orders direct third parties – financial institutions – to deduct specific amounts from the accounts of account holders and transfer them to Revenue NSW. Failure by a financial institution to comply with a direction of this kind can constitute an offence.

However, any technological process to issue such orders, even if it adhered to all of the provisions of the *Fines Act*, might still need to consider insolvency and bankruptcy legislation provisions that govern the priority in which debtors of an insolvent company or bankrupt individual are to be paid.¹⁶² There are also principles of unwritten law that prevent garnishee orders being issued on joint accounts held by financial institutions, or accounts where the funds in the account are held in trust.¹⁶³

Few statutory functions operate in isolation. Many of them must comply with other written and unwritten laws. It will be necessary to take steps to identify these and ensure that any use of machine technology is compliant with them.

Some requirements – such as those imposed by privacy, freedom of information and anti-discrimination laws – will likely always be relevant to some extent, and need to be considered. Care is particularly needed to ensure that the risk of algorithmic bias does not result in conduct or decisions that would amount to indirect discrimination.

8.3 Reasons are given for the decision (particularly decisions that affect the rights or interests of individuals)

The giving of reasons¹⁶⁴ is a basic principle of good administration – a person, especially one whose individual interests or rights have been adversely affected, is generally entitled to an explanation as to why that has happened.

In some cases the requirements of procedural fairness will mean that there is also a legal duty to provide reasons.¹⁶⁵ Sometimes this requirement is expressly imposed by statute, either specifically in relation to a particular decision, or more generally under certain circumstances.¹⁶⁶ In other cases it is implied because of the nature of the function, the person exercising it and the impact it has on those affected by it.¹⁶⁷ If, for example, there is a right of review or appeal, a requirement to provide reasons

will usually be implied, as knowing the reasons for a decision is essential to a person's ability to decide whether and how to challenge it.

Even where there is no specific legal requirement, the NSW Ombudsman's Office and other ombuds have consistently taken the view that reasons are an essential requirement of good administrative practice and should be provided to any person whose interests are significantly affected by a decision, at least upon request.

The purposes of reasons include:

(a) Transparency

The person is better able to see:

- The facts and reasoning that were the basis of the decision.
- That the decision was not made arbitrarily or based on mere speculation.
- To what extent any arguments they put forward have been understood or considered.
- Whether they have been dealt with fairly.
- Whether or not they might choose to exercise any rights of objection, review, appeal or complaint, and the arguments they will have to respond to if they do.
- How they might need to adjust their position to achieve more favourable decisions in the future.

(b) Accountability

Decision makers who are required to explain their decision have a greater incentive to ensure those decisions are defensible and based on acknowledged facts. Supervisors, as well as those with an external review role, are also in a better position to assess the decision, including whether it was reached lawfully, based on relevant considerations, and on the merits of the case.

(c) Quality

Decision makers who are required to explain their decision have a greater incentive to carefully identify and assess the relevant issues and apply rigour in their reasoning. Other decision makers can use reasons as guidance for the assessment or determination of similar issues in future.¹⁶⁸

Where reasons are required, the degree of detail required in those reasons may also be prescribed by legislation but, where it is not, will depend on the nature of the administrator and the function, and the circumstances in which it is exercised.¹⁶⁹

The prevalence of complaints, proceedings and applications in which individuals claim they have not been given adequate reasons for decisions is testament to the importance (and the complexity involved) in ensuring that actions are adequately explained and capable of being understood.

Using machine technology to administer Commonwealth child support payments

Commonwealth child support legislation imposes an obligation on certain parents to make periodic payments through the Child Support Registrar to the other parent of their child or children. The legislation recognises that the liable parent may also make payments to third persons, for example when school fees are paid directly to the child's school.

When these payments are made, the registrar has the discretion, in certain circumstances, to deduct those other payments from the amount that would otherwise be required to be paid to the other parent.

In proceedings concerning the registrar's decisions to deduct certain payments of this kind, it became evident that the registrar had been using an 'automation' process. This had been done by automatically crediting the amount of school fees paid by the liable parent over several months. The registrar would then advise the other parent that the payment had been made and an equivalent amount had been deducted from the child support payment, giving them an opportunity to object.

The Tribunal criticised this practice, of issuing 'provisional' decisions and inviting the payee to object, as not in compliance with the legislation, describing it as 'well intentioned' but 'legally flawed'.¹⁷⁰ The Tribunal also added that '[i]t must be doubtful that the sort of automation contended for can be consistent with a discretionary administrative decision in any event'.¹⁷¹

Where reasons are required to be given, there is no uniform standard that will be adequate in all circumstances. Generally, reasons should, when read as a whole, show the 'actual path of reasoning' taken by the administrator,¹⁷² or provide 'an explanation connecting any findings of fact with the ultimate decision'.¹⁷³ Interpretation legislation in some Australian jurisdictions may also mandate certain content.¹⁷⁴

The giving of reasons is just as important when machine technology has been involved in the making of decisions.

In our view, if machine technology has been materially involved in the process of making a decision, and the decision is one for which reasons ought to be given, then it would be unreasonable for the statement of those reasons to fail to include some explanation of the fact, nature and extent of the machine technology's involvement.

Reasons why machine involvement should be disclosed when giving reasons include:¹⁷⁵

- Failure to disclose this information appears misleading by omission, as most people would otherwise assume that the statement reflects the human decision maker's own reasoning processes. If, however, someone or something else has taken much of the cognitive load of the decision-making, then the decision maker's reliance or consideration of their work is part of the reasons for the decision and should be disclosed as such.
- The involvement of a machine in the decision-making process may affect the ways in which the decision can be challenged or reviewed. Informing the person affected of that involvement is therefore important to give them a genuine opportunity to decide whether and how to exercise their rights of challenge or review.
- Disclosure provides the opportunity to build confidence in decisions and decision-making processes. People may be reassured by the use of properly-designed technology that has helped to ensure the impartiality, rigour, consistency and accuracy of decisions. On the other hand, secrecy about the involvement of technology is likely to undermine public confidence and raise suspicions about why that involvement was not disclosed.¹⁷⁶

In **chapter 13** we will consider in more detail how requirements of transparency, including the requirement to give reasons, should be considered when designing and implementing machine technology.

9. Appropriate assessment

9.1 The decision *answers* the right question (which necessitates *asking* the right question)

The development of machine technology for use in the exercise of statutory functions will require the translation of legislation, and related guidance material (such as directions or policies), into a form capable of being turned into code that machines can read.

Whether, and to what extent that can be done will depend on the type and style of legislation, the correct interpretation of the legislation,¹⁷⁷ and the capabilities and limitations of the particular technology employed. Importantly, it will also depend on the expertise of those who translate the legal text into code, and the processes developed to undertake that task.

Perhaps the most basic error that can be made when introducing automation into the exercise of statutory functions is to misinterpret or misapply the legislative scheme – effectively, to ask the machine the wrong question. Of course, this is not a risk that is confined to technology: human beings are quite capable of misinterpreting and misapplying legislation.

However, for a variety of reasons the risk may be magnified when technology is involved:

(a) The *likelihood* of error may be greater

Laws drafted and made by humans, to be read and implemented by humans, do not readily lend themselves to translation into code. Computer code generally is ‘more precise and has a narrower vocabulary than natural language’ used in legislation.¹⁷⁸ The need to translate law into code introduces an additional step in realising the intention of Parliament.

Those involved in designing technology to exercise statutory functions will typically not have expertise or experience in interpreting legislation or exercising administrative functions.

This translation process must also be repeated every time the relevant legislation is amended, and when judicial decisions or changes to other laws affect the way the legislation is interpreted or applied.

The risk of error with technologically embedded compliance processes is highlighted by a growing body of court cases where contraventions of statutory obligations by private entities have been attributed, at least in part, to ‘information technology system issues’.¹⁷⁹

(b) The *consequences* of error may be more significant

Error in code will almost inevitably affect more outcomes (and therefore more individuals) than an error committed by a particular administrator. One of the key advantages of machines – their potential to process high volumes of data at high speed – means that errors may be replicated at a rate exceeding that of any human administrator. Consequently, the number of people adversely affected by a single error may be substantial.

(c) The *detection* of error may be more difficult

While administrators have internal processes for detecting human errors in the exercising of functions, detection of errors in outcomes of machine processes will call for an interrogation in a manner beyond the capability of most administrators. Those affected by erroneous decisions, particularly if they are already vulnerable, may also be less able to identify or effectively challenge a machine error that arises from within the technology design and where the error is not immediately apparent in the output of any individual case.¹⁸⁰

(d) Rectifying an error may be more costly and take more time

If a human decision maker makes an error, then their conduct can easily be corrected for future decisions. Even a systemic error perpetuated by an error in policy or other guidance material can generally be remedied quickly. However, if an error is detected in machine technology, fixing the error may be difficult, costly and time consuming. This may be particularly so if the technology has been procured through an external vendor. An agency that is aware that machine technology contains an error, but is unable immediately to fix those errors, may be in a difficult position if the move to automation has left it with no other means of exercising the function.

Any errors in the translation process may mean that, even in circumstances where technology can otherwise be used consistently with principles of administrative law, doubts will arise about the legality and reliability of any decisions and actions of the public agency relying upon the machine process.

Lost in translation – a simple error converting legislation into code

A recent complaint handled by our office illustrates the challenges involved in automating even those statutory functions that on paper seem very simple.

If the holder of a NSW driver licence exceeds the permitted number of demerit points within a 3 year period, Transport for NSW (**TfNSW**) may suspend their licence and/or declare them ineligible to obtain a licence for a period of time (a **licence suspension**).¹⁸¹

If the driver does not wish to serve the licence suspension period, they can opt instead to enter into a 12-month good behaviour period.¹⁸² If the holder incurs any more than 2 demerit points within that period, they must incur a licence suspension for a period twice as long as the original licence suspension.¹⁸³ A licence suspension in those circumstances is not discretionary.

Licence suspensions are initiated through the issuing of suspension notices. These notices specify the date (a certain number of days after the date of the notice) when the suspension will begin and how long it will last.

TfNSW has automated its process for issuing notices of these licence suspensions through the use of machine technology (the **DRIVES** system). DRIVES has been programmed in such a way that a different process is followed depending on how long the driver's licence has until expiry at the time of the suspension:

- If there are 35 or more days left to expiry, the notice will be automatically issued.
- If there are fewer than 35 days left to expiry, no notice is issued. Instead, when the driver applies to renew their licence they will be denied a licence and given a licence suspension notice.

In the complaint we received, a driver had incurred more than 2 demerit points during the good behaviour period. At the relevant time, their licence had fewer than 35 days to expiry.

Aware that they faced a 6-month period of licence suspension given the recent demerits, the driver did not immediately seek to renew their licence. They assumed that their licence was or would promptly be suspended.

However, because the automated notice system was programmed not to issue a notice unless or until a new licence application was made, no licence suspension was triggered.

Some months later, the driver applied for a new licence. The application was refused and they were only then issued with notice of a 6 month licence suspension. As a notice can only set a licence suspension to commence in the future, this meant that the suspension period only began then and not months earlier when it should have been triggered.

The case suggests that those coding TfNSW's machine made certain assumptions, including that any driver whose licence had expired would apply promptly for a new licence.

The lengthy delay before the notice of licence suspension was issued meant that there was a lengthy delay before the suspension period commenced. TfNSW acknowledged the error but noted that it had no power to 'backdate' the suspension. It did, however, apologise to the complainant.

It seems possible that the machine would have been coded differently had the legislation explicitly set a specific time limit within which any notice must be sent – that is, if the legislation expressly stated that a notice of licence suspension must be sent within so many days of a bond breach occurring. However, although the legislation does not say this, a requirement to issue a notice within a *reasonable time* is implied by common law, taking into account the purpose of the statutory requirement and the surrounding legislative provisions.¹⁸⁴ Such an implied requirement may not have been obvious to those involved in designing the code for the machine if they were not experienced in statutory interpretation.

TfNSW has acknowledged to us that its code is incorrect in this respect, and that notices of licence suspension should always be issued promptly. However, while it is committed to fixing the error, it will not be possible to do so until the next scheduled system update. In the interim, it will consider whether there are any interim measures it can put in place until the system can be corrected.

9.2 The decision is based on a proper analysis of relevant material

No decision-making discretion is given to an administrator in absolute or unconditional terms. All functions are qualified to some extent by available, obligatory and extraneous considerations, each of which impact on the exercise of a function in different ways:

Available considerations¹⁸⁵ are facts or matters that *may* be taken into account but are not required to be taken into account.

Obligatory considerations¹⁸⁶ are facts or matters that Parliament has determined *must* be taken into account when exercising a discretionary power.

A failure to take an obligatory factor into account may render the exercise of power unlawful, and potentially invalid.¹⁸⁷

In many cases, obligatory considerations are expressly stated in the legislation that created the function.¹⁸⁸ Even where legislation giving a function to an administrator does not expressly set out any obligatory considerations, they may be implied having regard to the subject matter, scope and purpose of the statute.¹⁸⁹

Extraneous considerations¹⁹⁰ are facts or matters that *must not* be taken into account by an administrator when exercising a function.

Extraneous considerations, like obligatory considerations, may be set out in the statute¹⁹¹ or (more generally) they will be implied having regard to the subject matter, scope and purpose of the power being exercised.¹⁹² Where an administrator takes into account an extraneous

consideration, their conduct can be seen to reflect ‘an extraneous or improper purpose or to render the decision arbitrary or capricious’.¹⁹³ The result may be that the decision is invalid.

While an administrator may be given considerable freedom with regard to available considerations, they *must* have regard to obligatory considerations and *must not* have regard to extraneous considerations.¹⁹⁴

When designing and implementing machine technology, it is essential to ensure that doing so does not result in any obligatory considerations being overlooked or extraneous considerations coming into play.

9.3 The decision is based on the merits and is reasonable in the circumstances

As already noted, many administrative functions involve an element of discretion. This permits administrators to deliver appropriately individualised outcomes when exercising functions in contexts and circumstances that are inherently unique, and which cannot be perfectly foreseen or prescribed in advance. In other words, decisions must be based ‘on the merits’ of each particular case.

This requirement overlaps many of the matters already discussed – for example, a decision that is affected by bias, or that has been made on the basis of discriminatory or otherwise extraneous considerations, is not a decision made on its (legal) merits.

At the same time, any exercise of discretion must also be *reasonable*. In assessing what is reasonable in any particular case it may be appropriate to consider whether the decision is consistent with decisions made in other cases – that is, are like cases treated alike and are different cases treated differently?¹⁹⁵

While not a stand-alone ground of review in administrative law, a lack of consistency may indicate a decision has not been made on its merits, is arbitrary and not reasonable, or is ‘infected’ by some other specific error. In addition, there will also be ‘limits beyond which such inconsistency itself constitutes a form of injustice.’¹⁹⁶

Machine technology has the potential to enhance consistency of outcomes between like cases by controlling for the risk of human biases and other idiosyncrasies associated with multiple human decision makers.

However, there may be a tension between the attainment of this consistency and the requirement to treat each individual case on its merits.

Although a judicial rather than an administrative decision, criminal sentencing decisions provide a clear example where such tensions could arise. Under the laws of sentencing, the task of the sentencer is said to be to arrive at an ‘instinctive synthesis’.¹⁹⁷ Each individual to be sentenced and each case is inherently unique, and there are multiple factors the sentencer must take into account. These factors are ‘incommensurable, and indeed, in many respects, inconsistent’.¹⁹⁸ For this reason, the goal of ‘reasonable consistency’ between sentences is considered incapable of any ‘mathematical expression’.¹⁹⁹

In cases such as this, consistency is important, but the kind of consistency that is sought is a consistent application of principles and reasoning to each different decision, and not merely some formulaic or statistical consistency. The High Court has already raised caution about using statistics and ‘guideline judgments’ in exercising sentencing discretion:²⁰⁰

[R]ecording what sentences have been imposed in other cases is useful if, but only if, it is accompanied by an articulation of what are to be seen as the unifying principles which those disparate sentences may reveal. The production of bare statistics about sentences that have been passed tells the judge who is about to pass sentence on an offender very little that is useful if the sentencing judge is not also told why those sentences were fixed as they were.²⁰¹

...
 To focus on the result of the sentencing task, to the exclusion of the reasons which support the result, is to depart from fundamental principles of equal justice. Equal justice requires identity of outcome in cases that are relevantly identical. It requires different outcomes in cases that are different in some relevant respect. Publishing a table of predicted or intended outcomes masks the task of identifying what are relevant differences.²⁰²

This does not mean that machine technology could not come to play any role at all in such complex and inherently case-by-case decision-making. However, it does suggest that the role of machine technology will be limited.

For example, a machine that outputs a suggested sentence or sentence range would seem to be significantly less legally safe to the sentencing decision maker than a machine that provides, with some empirically validated degree of accuracy,²⁰³ a numerical rating of the risk of reoffending – that is, something the sentencer could feasibly consider within their overall ‘intuitive synthesis’.²⁰⁴ Of course, there are other obvious risks there, including the risk that inherent but hidden biases in the historic data sets – such as racial stereotyping – will be entrenched or even amplified in such machine-generated risk ratings (see **chapter 8**).

Unsurprisingly, in other jurisdictions where machine technology has been used for judicial decisions like sentencing, significant concerns have been raised about them – see below ‘Machine technology in sentencing – COMPAS’.

Machine technology in sentencing – COMPAS

The Correctional Offender Management Profiling for Alternative Sanctions (**COMPAS**) is a risk assessment tool used in the United States that aims to predict the likelihood of reoffending. Originally designed for use in post-sentencing supervision decisions, it is now used in sentencing and other criminal justice processes. COMPAS works by analysing an individual against certain criteria and historical data – an output is produced by way of a score ranking them from low risk to high risk of committing future crime. The specifics of how the system uses inference in performing its calculations are not known.

COMPAS has been widely criticised due to doubts about its accuracy and concerns about discrimination, with one researcher claiming that COMPAS ‘is no more accurate or fair than predictions made by people with little or no criminal justice expertise.’²⁰⁵

A review of the tool by ProPublica (a non-profit organisation) found that the system was unreliable in predicting violent crime, as well as being racially biased. ProPublica claims that it wrongly predicts African American defendants to be reoffenders at almost twice the rate as it does for white defendants.²⁰⁶ It is understood that ‘race’ itself is not a variable in COMPAS, but that bias appears to result from the relationships between race and other characteristics concerning social economic factors (which may operate as proxies for race), as well as because of historical training data that reflects human biases present in policing decisions.

The creator of COMPAS has disputed ProPublica’s findings,²⁰⁷ and it has since been noted that the tool creator and ProPublica each measured fairness based on different and incompatible measures.²⁰⁸

Tools like COMPAS also raise again concerns about the degree of reliance and trust decision makers place in the outputs of machine technology. ProPublica provides an example of a US judge who overturned a plea deal more favourable to the defendant to impose an increased

prison sentence after considering the COMPAS prediction that the defendant had a high risk of future violent crime.²⁰⁹

In 2016, the Wisconsin Supreme Court upheld the use of COMPAS in sentencing.²¹⁰ The Court required, however, that a written warning (or 'advisement') be given to judges using COMPAS about its limitations. However, questions remain about how effective such a warning could be.²¹¹ The editors of the Harvard Law Review observe:

[The larger problem] is simply that most judges are unlikely to understand algorithmic risk assessments... [T]he court was mistaken to think that as long as judges are informed about COMPAS's potential inaccuracy, they can discount appropriately.

Additionally, the warning will likely be ineffectual in changing the way judges think about risk assessments given the pressure within the judicial system to use these assessments as well as the cognitive biases supporting data reliance.... Research suggests that it is challenging and unusual for individuals to defy algorithmic recommendations.²¹² [references omitted]

10. Adequate documentation

10.1 The circumstances surrounding the making of decisions is adequately documented and records kept

One of the easiest, and unfortunately most common, findings an Ombudsman can make is that an agency failed to properly document and keep records of its decision-making.

In general, the basic documentation required to be kept for any decision will include:

- details of the decision itself
- reasons for the decision
- the identity of the decision maker
- the date of the decision
- copies of any written notification, or file-notes of any non-written communication, of the decision, to the person affected or to any other person.

These requirements apply equally to decisions made with the assistance of machines. Under the *State Records Act 1988* (and related information legislation, including the GIPA Act) terms like ‘record’, ‘information’ and ‘document’ are defined in technology-neutral ways:

record means any document or other source of information compiled, recorded or stored in written form or on film, or by electronic process, or in any other manner or by any other means.²¹³

Machine technology, however, can raise particular issues for the maintenance of appropriate records. For that reason, it is important that agencies’ recordkeeping policies include provisions that directly address the records required to be kept in relation to any machine technology in use by the agency.

For example, agencies will need to ensure that they maintain a register of all versions of their systems, with their dates and a description of the changes made between each version.²¹⁴ The changes should specifically note any updates that have arisen because of a change in law or policy. Ideally, previous versions should be kept in full, so that past decisions made using that version can be replicated and reviewed, if necessary.

In **chapter 13**, we consider in more detail the importance, when designing machine technology, of identifying the information that will need to be kept and published to ensure transparency and promote accountability.

Part 3:

Designing machine
technology to comply
with the law and
fundamental principles
of good government



11. Putting in place the right team

In this and the next 4 chapters, we consider some of the practical steps that government agencies should take when designing and implementing machine technology to support the exercise of an existing statutory function.

These chapters proceed on the assumption that the relevant agency or official already has the relevant function, is currently exercising that function without the use of machine technology, and is contemplating the adoption of some form of machine technology to assist its exercise of that function in the future.

We focus on 5 critical steps that agencies should take:

1. establish a multi-disciplinary design team (lawyers, policymakers, operational experts, and technicians), with clearly allocated roles and responsibilities (**chapter 11**)
2. assess the appropriate degree of human involvement in the decision-making processes, having regard to the nature of the particular function and the statute in question (**chapter 12**)
3. ensure transparency, by deciding what can and should be disclosed about the use of machine technology to those whose interests may be affected by the relevant function (**chapter 13**)
4. test *before* operationalising, and establish *ongoing* monitoring, audit and review processes (**chapter 14**)
5. consider whether legislative amendment is necessary or prudent (**chapter 15**).

In this chapter we start with step one: establishing the right design team.

11.1 It's not an IT project

Adopting machine technology to support a government function should not be viewed as simply, or primarily, an information technology project.

It is, rather, a coordinated exercise of legal, policy and operational process development, aided by technology. As such, the team that is formed to design and implement the project needs to include – and indeed, be led by – those with relevant, and sufficiently senior, legal, policy and operational expertise, who can work with and guide the technology specialists.

Each of these individuals must have appropriate degrees of involvement and authority throughout the project. Legal, policy and operational experts should not be relegated to ‘consultation’ or ‘advisory’ status, and nor should they simply be given a near-finished product for review or endorsement.

It is clearly better for all parties (including for the efficiency and reputation of the agency itself) if machine technology is designed by those who are best placed to know whether it is delivering demonstrably lawful and fair decisions, rather than having to try to ‘retrofit’ that expertise into the system later when it is challenged in court proceedings or an Ombudsman investigation.²¹⁵

11.2 Having lawyers on the design team is essential

Our concerns with Revenue NSW (see **annexure A**) arose largely because it did not seek any expert internal or external legal advice on the design and operation of its machine-aided garnishee system, even after we told them they should.

Agencies that have been exercising functions under a statute for a long period of time no doubt develop a good understanding of how that legislation operates. However, when new technologies and new modes of exercising the function are being considered, it is essential that the source legislation be carefully considered afresh.

The task of interpreting a statute to arrive at its correct meaning can be a complex task at the best of times, and one that can challenge both highly experienced administrative officials and lawyers.²¹⁶ Even legal rules that on their face appear to be straightforward and ‘black and white’, and which may be the most appropriate candidates for machine technology development, can nonetheless have a nuanced scope and meaning. They may also be subject to administrative law principles – such as underlying assumptions (for example, the principle of legality)²¹⁷ and procedural fairness obligations – which would not be apparent on the face of the legislation.

This is not to overstate the mystery of the law. However, it is to say that legislative interpretation requires specialist skills, and the challenge involved is likely to be especially pronounced when seeking to translate law into what amounts to a different language – ie a form capable of being executed by a machine.²¹⁸

The structure, underlying assumptions, nuance and ambiguities of the English language, the internal logic of a particular statute and its relationship to the unwritten law, and the need to apply the text in the real world means that statutory provisions do not typically translate in any straightforward and definitive way into a form required by machine technology. While interpreting legislation starts with interpreting the ordinary and grammatical meaning of the words used, it also involves considering the context in which those words are used.²¹⁹ That context includes the surrounding legislative provisions and the statute taken as a whole. An analysis of that context may result in those words ultimately being given a different meaning than a purely literal interpretation would produce.²²⁰ Any coding of the relevant law (or parts of it) will almost always require the making of interpretative choices to enable it to fit it within the different language and logic of the machine form.²²¹ This must be done in a way that does not result in the meaning or effect of the law being impermissibly altered.

Most government agencies are well-resourced with highly-qualified legal professionals who are skilled in statutory interpretation. Agencies that wish to use machine technology should ensure that they utilise that expertise from the very outset of any design process.

11.3 Ensuring policy choices are made by the right people, at the right level of authority

Developing and implementing machine technology will rarely, if ever, be a simple mechanical process. Any design team, presented with a description of a function and an explanation of how it is currently being performed, will not come up with precisely the same device, with exactly the same specifications, functionality and performance, as any other design team. A multitude of decision-points and therefore choices will arise at various steps in the process.

Some of those choices may have profound impacts on the operation of the machine technology, and therefore on the ultimate exercise of the relevant functions and their impact in the real world. They are therefore important public policy choices.

Who is making those choices?

We have some concern if some of these choices are effectively outsourced to the technicians tasked with developing the technical design specifications for a machine and/or for the detailed coding, design and build of the machine itself. There are three reasons why this could be problematic:

1. First, these technicians will generally have neither the necessary legal and policy expertise, nor the administrative operational experience, to appreciate all of the legal, policy and operational issues that may come into play when developing machine technology for the exercise of statutory administrative decision-making and its impact in the real-world environment.
2. Secondly, coding and other highly technical skills are frequently obtained through private sector procurement. Even if those technicians did have knowledge and experience to make appropriate legal policy and administrative decisions affecting the public, it would be inappropriate for them to do so. They are not public servants, and are not subject to the same constitutional, employment, cultural and professional ethics frameworks that apply to those who work in public service.
3. Thirdly, the kinds of policy choices involved when designing machines can be thought of as elements relating to the 'quality' of the machine. Some of those choices may have cost implications. If there are such quality/cost trade-offs, the decision should not be made unilaterally by the (profit-driven) vendor as that decision may differ from the decision the government agency itself would make in the public interest. There may be other choices that also involve public policy trade-offs other than cost (e.g. whether the machine is designed to minimise 'false positives' or 'false negatives'). However, depending on the outsourcing process, it may be that the vendor will not even need to inform the government agency that these choices can be made.

Concerns of 'undue influence' might be raised if important decisions about the design and specification of a machine, including what data sets are or are not appropriate to be included, are left to private sector technology vendors (and especially of those vendors might later claim trade secrecy over those matters).²²²

Lessons from legislative drafting?

When Government prepares legislation, the process involves senior policymaker (such as Cabinet, Ministers and senior public officials) deciding the overall policy objectives and parameters, in-house legal experts within agencies drawing up highly-detailed 'drafting instructions', and expert drafters at the office of the NSW Parliamentary Counsel preparing the draft statute or other instrument.

This process is iterative, recognising that the drafting process may identify gaps or ambiguities, or other policy decisions that need to be resolved. When this happens, the drafters seek further instructions from the agency. Ultimately, the final Bill returns to the senior-policymaker, together with a report certifying that it has implemented their policy decisions.

A similar process could be considered for the development of decision-making machines – with design specifications taking the place of drafting instructions, and machine code taking the place of legislation. Such a process means that, as with the drafting of legislation, the process of designing and developing machine technology in respect of legislative functions:

- Is seen as an inherently iterative process between those making design/policy decisions and those implementing them.
- Clearly differentiates the roles and authorities of the policymakers, those who translate policy into instructions/specifications, and the technicians.
- Ensures that any significant decisions are only made by appropriate government members of the team, and are escalated where necessary.

12. Determining the necessary degree of human involvement

In **chapter 7** above, we noted that it is legally essential that a person given a discretionary decision-making function must genuinely exercise that discretion and make their own decision.

Discretionary decision-making requires some degree of human involvement – it can never be fully automated.

How far a discretionary decision-making process *can* be automated is not an easy question. It needs to be assessed in the context of the particular function and the statute in question.

As already discussed, the law has recognised that policies may play a legitimate role in guiding discretionary decision makers, provided decisions are not impermissibly fettered by the terms of the policy or the way it is used. The law likewise recognises that a discretionary decision maker may take into account, and where appropriate act on, the advice and recommendations of others, provided they are not impermissibly acting under the other's dictation and abdicating their own discretion.²²³ This reflects the practical reality that administrators often need to rely on others, such as their staff and other experts, when carrying out their functions.

It would seem a short and legally uncontroversial step to accept that an administrator exercising a discretionary function is also not precluded from *considering* the outputs of a relevant and well-designed machine.²²⁴ However, those outputs must not impermissibly *control* the administrator's exercise of the function.

12.1 The administrator must engage in an active mental process

Minimally, any statutory discretion requires there to be a person (the person to whom the discretion has been given or delegated) who makes a *decision* whether and how to exercise discretion in the particular case or cases before them.

However, merely placing a 'human-on-top' of a process will not, of itself, validate the use of machine technology in the exercise of a discretionary function. As the external legal advice we obtained noted:

Although the response of administrative law to the use of information technology may be nascent, ordinary administrative law principles require there to be a "process of reasoning" for the exercise of discretions. This can also be seen in our conceptions of what it means to make a "decision", with two members of the Full Federal Court ... accepting that one of the elements generally involved in a "decision" is "reaching a conclusion on a matter as a result of a mental process having been engaged in."²²⁵ [case references omitted]

This means that, even if a person officially 'signs off' at the end of a process, the decision-making process may still be unlawful if in reality that person is merely acting as a rubber stamp, accepting the outputs of a machine 'as a matter of course' and 'without engaging in a mental process to justify that conclusion'.²²⁶

The need for functions to be exercised by the person to whom it is given (or delegated) has also been emphasised in Federal Court decisions concerning the exercise of immigration discretions, which have referred to the need for there to be 'active intellectual consideration',²²⁷ an 'active intellectual process',²²⁸ or 'the reality of consideration'²²⁹ by an administrator when making a discretionary decision.

Among other things, these cases looked at the amount of time an administrator had between when they received relevant material and the time when they made their decision. In some cases, this time period was shown to have been too short for the administrator to have even read the material before them. The Court concluded that there could not have been any 'active intellectual consideration'

undertaken in the exercise of the function, and therefore overturned the decisions on the basis that there had been no valid exercise of discretion.²³⁰

Not all administrative functions have consequences as significant as those concerning immigration. The ‘reality of consideration’ may look different in different administrative contexts, in proportion to the nature of the function being exercised and the consequences it has for those it may affect. However, the principle remains relevant to the exercise of all statutory functions by administrators: in the exercise of a statutory discretion given to a person, some level of genuine and active decision-making by that person is required. As noted in **chapter 7**, where Parliament has chosen not to adopt fixed rules for the exercise of a statutory function, the discretion it has given to an administrator must be recognised and exercised.

12.2 The division of tasks between machine and human

In designing a machine technology supported decision-making process, thought needs to be given not only to ensuring that the human decision maker genuinely makes the final decision, but also to the division of tasks between human and machine throughout the decision-making process.

As we have already seen, most discretionary decisions will include a range of obligatory and available considerations. Some of those considerations may be more appropriate than others to be addressed by machine technology.

Consider, for example, a simple statutory payment scheme that requires an administrator to decide whether to make a discretionary payment to a person having regard to their:

- Age (the person must be above a certain age to be eligible).
- Place of residence (the person must live in a certain area to be eligible).
- ‘Need’ (the person’s need for the payment is to be taken into account).

There would seem to be no issue with the decision maker being assisted by a machine that can generate outputs about a person’s age and place of residence.²³¹ For example, a machine might sort or filter a list of all those who have applied for the payment by reference to those two fields in order to identify those eligible.

The decision maker would then be required to separately consider the question of need. Provided the decision maker does so and considers *both* the outputs of the machine (age and place of residence) and then additionally considers need, the decision maker will have met the legal requirement of having taken into account all obligatory considerations.²³²

A more sophisticated machine might go further and also generate an output that seeks to rank or score applicants by assessed need, having regard to parameters such as their income, assets, dependents, and so on.

Unlike ‘age’ or ‘place of residence’ (for which there will, generally, be an objectively right or wrong factual answer) assessing ‘need’ involves the exercise of a complex, evaluative judgment. Therefore, even if the machine generates some score or ranking of need based on pre-determined criteria, the decision maker will still need to apply their own ‘active intellectual consideration’ to the output, as well as take into account any other considerations that have not been addressed (or have not been addressed fully) by the machine’s outputs. They will also need to determine the relative weight given to each of these considerations in their overall decision: ‘What is required is a human judge exercising their discretion to decide which factors are the most important in a particular factual scenario.’²³³

12.3 The risk of technology complacency and ‘creeping control’

The simple example above may suggest that designing a machine system that meets a minimum required threshold for human involvement in discretionary decision-making will not be particularly challenging.

However, what matters is not just that there is the required degree of human involvement on paper; there must be that human involvement in practice.

Even if a decision-making system is appropriately designed with a human decision maker in the process who is to (lawfully) *consider* machine outputs, there are a number of reasons why, over time, the decision maker may tend toward ‘technology complacency’. This means that their decisions may tend to become increasingly (and potentially unlawfully) *controlled* by the machine’s outputs.

Reasons for this tendency may include:

(a) A bias toward uncritical acceptance

It is well-recognised that there can be a natural bias for administrators to uncritically accept information provided to them by technology, especially where the outputs generated are presented in a form that appears to constitute objectively quantifiable fact.²³⁴

In the Commonwealth Ombudsman’s review of immigration detention decisions, for example, the Ombudsman noted a tendency of government staff to accept the accuracy of the information they accessed through the use of technology, even in the face of conflicting or contradictory information from other sources.²³⁵

(b) Blame-avoidance and ‘path of least resistance’

Even if a human administrator is not completely certain that the machine has produced the ‘right’ answer, they know that accepting that output (even if it turns out to have been wrong) is unlikely to result in them being held personally responsible for any adverse outcome. If, on the other hand, they actively overrule the machine, then the risk of their being blamed for the outcome is likely to be very high.

Accepting the output becomes not just less work, but it is also the lowest risk option for the individual decision maker.

A decision maker may avoid even questioning the output generated by the machine for fear of causing unnecessary delay or being seen as causing problems for management. Questioning the output may be particularly challenging if the machine also required a significant investment of public funds, or if it was designed, launched and lauded as the ‘next great thing’ by those more senior than the decision maker.

(c) Practical and technical impediments to scrutiny

There are also practical and technical challenges to even the most conscientious administrator who seeks to make their own independent decision rather than being controlled by a machine’s output.

Typically, most machine technology will be designed and maintained in non-operational technical areas that are organisationally and functionally separated from the administrators who will use its outputs to exercise functions. Administrators may be insufficiently aware of the scope, capacities and limits of the machine to even know *when* they should be seeking further information or clarification about how it works. If they do wish to seek clarification or more information to inform their decision whether to adopt a machine output, the machine may not be configured to provide them with the particular information needed or they may not have access to it. Pursuing those questions may require engaging at a technical level with technology support personnel, which may not be feasible, either practically or culturally.

Agencies need to be wary of the risk that even a well-designed decision-making process involving machine technology could come to cross a line in practice which may render decisions made using it unlawful. The risk is likely to increase with the level of technical opacity of the machine.

12.4 Practical indications of active human involvement

When designing and implementing machine technology, government agencies must therefore also consider how the system will work in practice and over time, having regard to ‘soft’ issues like natural human biases and behaviour and organisational culture.

They must also recognise that those who in future will be making decisions supported by the machine will not necessarily be the people who were involved in its original conception, design and implementation. The controls and mitigations that are needed to avoid ‘creeping control’ by the machine technology will need to be fully documented so they can be rigorously applied going forward.

The following are some of the factors that are likely to be relevant to consider in determining whether there is an appropriate degree of human involvement in a machine-supported decision-making system:²³⁶

(a) Time

Does the process afford the administrator sufficient time to properly consider the outputs of the machine and any other relevant individual circumstances of the case(s) in respect of which the function is being exercised? Does the administrator take this time in practice?

(b) Access to source information

Is the administrator able to consider the source material used by the machine? Do they have access to other material and information that may be relevant to their decision?

(c) Seniority and experience

Does the administrator have the appropriate organisational seniority and level of experience that would be expected for the type of decision they are making (with or without the support of machine technology)?

(d) Decision-making ownership

Does the administrator always take ownership of their decisions, even when they are following the outputs of the machine? Organisationally, is the administrator considered responsible for the decisions they make?

(e) Cultural acceptance

Are there systems in place to overcome or mitigate automation-related complacency or technology bias, to scrutinise and raise queries about the output of the machine technology, to undertake their own further inquiries, and – if the administrator considers it appropriate to do so – to reject the output of the machine? Is the authority of the administrator to question and reject the machine’s outputs respected and encouraged? Does it happen in practice?

(f) Understanding of the reasoning process

Does the administrator have a thorough understanding of the operation of the machine technology as a whole, at least conceptually, in order to be able to form a view on a reasonable and rational basis about its outputs?²³⁷ Is the administrator able to provide comprehensible reasons for their decision?

(g) Input into decision-making process design

Can the administrator make or require changes to be made to the machine to better support their decision-making?

(h) Appreciation of decision-making impacts

Does the administrator have a genuine understanding of what their decision (and what a different decision) would mean in reality, including for the individuals who may be affected by the decision?²³⁸

The list above also highlights the importance of ensuring that those humans who will be involved in using the machine technology are given the appropriate training and skills to ensure they are assisted but not controlled by its outputs.

It is particularly important that the relevant administrator, and others responsible for analysing or working with the outputs of the technology, have a sufficient understanding of the technology and what its outputs actually *mean* in order to be able to use them appropriately. This is likely to mean that comprehensive training, both formal and on-the-job, will be required.

This training will need to be ongoing, as the technology is modified or updated, as staff may change, and as reinforcement may be required for existing staff to mitigate the risk of declining skills or creeping complacency over time.

There is also a need to ensure that the machine itself is designed so that its outputs will be presented in a manner – whether that be through dashboard designs or data visualisations – that will be most conducive to active mental engagement, human understanding and appropriate scepticism.²³⁹

13. Ensuring transparency

13.1 Reasons and the right to an explanation

In designing machine technology, agencies must ensure that meaningful reasons will still be able to be provided to those whose legal or other significant interests may be affected by decisions. Those reasons should also note whether machine technology was involved in that decision. In our view, the information to be provided in this regard should include, at a minimum:

- (a) the fact that machine technology was involved
- (b) the nature and extent of that involvement
- (c) what information about them is processed by the machine, including any assumptions, proxies or inferences
- (d) the particular version of the technology, program or application used, and the date of that version, and
- (e) an explanation of how the technology works in a way that is meaningful and intelligible to an ordinary person.

Of course, the statement should also include the usual requirements for decision notices, including details of how the decision may be challenged or reviewed, and by whom – see **chapter 8**.

Providing explanations for decisions that are ‘instructive, informative and enlightening’

There appear to be very few court or tribunal decisions that have grappled with the adequacy of explanations given when machine technology has been involved in decision-making. In *Schouten and Secretary, Department of Education, Employment and Workplace Relations*,²⁴⁰ the issue arose in an application to review the amount of social services benefit (Youth Allowance) payable by Centrelink to an individual.

While affirming that the amount being paid to the individual was correct, the Tribunal noted that it was not until a government employee gave evidence to the Tribunal about the process employed to calculate the rate of benefit payable that the individual, and the Tribunal itself, could understand the process. The Tribunal noted that the case highlighted ‘the difficulty where government agencies make “automated decisions” and the decision is complex.’ It noted that:

The citizen will not understand and therefore be unable to challenge a decision about which they feel aggrieved unless provided with a plain English explanation of the basis for the decision. As in this case, the initial decision-maker is sometimes unable to provide that explanation. The Administrative Review Council in its report to the Attorney-General, “Automated Assistance in Administrative Decision-Making Report No. 46” noted that care was needed to ensure that the values of transparency and external scrutiny are not compromised where automated decision-making is employed ... A major challenge for government agencies dealing with citizens is to ensure that their decisions are instructive, informative and enlightening. In this case, Centrelink has not met that challenge.²⁴¹

How can reasons be provided when machine technology is used?

The use of machines can create additional challenges when providing reasons. One challenge is the tension that may exist between providing reasons that are technically accurate (in terms of describing how the machine works, whether or not it is also possible to show exactly how it came to generate a certain output)²⁴² and providing reasons that serve as an explanation of, or justification for, the decision that will be intelligible and useful to the person affected.

It is clear that an explanation for an outcome that is technically accurate, but is otherwise unintelligible, cannot achieve its purpose and should therefore not be accepted as appropriate reasons at all.

When a human makes a decision, the reasons given do not refer to their brain chemistry or the intricate details of a process that commences with a particular set of synapses firing and culminates in a movement of the physical body giving rise to vocalised or written words. Likewise, explaining how a machine works, even if that explanation is fully comprehensive and accurate, will not necessarily satisfy the requirement to provide ‘reasons’ for its outputs.

*Reasons must be accurate; they must also be meaningful and intelligible to the person who is to receive them. They must provide an ‘explanation’.*²⁴³

While there has been much discussion about whether a person affected by a relevant machine-made decision should have a right to the underlying code used by the machine, a more immediate issue is ensuring that a statement of reasons is prepared with their purpose and audience in mind. Even where code is made available, this is unlikely to satisfy a requirement to provide ‘reasons’ even to the small number of individuals who could understand it and in the case of the very simplest of codes (the *Business Names Registration Act 2011* (Cth) may be an example). It is hard to see how the provision of source code could satisfy a requirement to give reasons – it would not, for example, set out findings on questions of fact, refer to the evidence on which those findings were based, or otherwise explain why the decision was made.

Generally, reason statements should be in plain English, and provide information that would be intelligible to a person with no legal or other relevant technical training. What is required is something approximating a ‘path of reasoning’, bridging the relevant findings of fact with the outcome.

In the case of machine-assisted decisions, such an explanation might include information about the machine’s objectives, what data has been utilised, its accuracy or success rate, and information about whether and what is measured. It would seem at least arguable that doing so would satisfy the requirement for reasons.

This does not mean that the more technical details of the design and operation of a machine should not also be provided. We note the NSW Information Commissioner has advised that such information should, at least presumptively, be treated and made available as ‘open access information’.²⁴⁴ We agree with that sentiment. However, merely publishing technical specifications or the underlying code will generally not satisfy a requirement to provide reasons.

The risk of automating reason-giving

Just as machine technology could be used to generate decisions or components of decisions, so too it is easy to imagine machine technology being used to generate statements of reasons for decisions or components of the reasons.

Already, template letters or standard paragraphs with formulaic expressions of reasons are not uncommon in use by government agencies. Nor are they impermissible, provided ‘the formula is used to guide the steps in making the decision and reveals no legal error’. However, a formula must not be used in a way that would ‘cloak the decision with the appearance of conformity with the law when the decision is infected’ by error. In such a case, ‘the use of the formula may even be evidence of an actionable abuse of power by the decision-maker’.²⁴⁵

The use of even more sophisticated machine technologies for the generation of statements of reasons raises greater concerns that they ‘will provide a façade of accuracy or objectivity that masks flawed decisions’. That is, machine-generated statements of reasons may ‘merely enhance the *appearance* of a lawfully made decision’.²⁴⁶

We suggest that it is safer, if a machine technology process is to be used in the decision-making process, that this process not also be tasked with generating a statement of reasons for the decision. Instead, the machine could produce the necessary audit records of its inputs, outputs, and processing, which can be taken into account by the human administrator as they develop a statement of reasons. Where practical, if the human administrator actually authors the statement of reasons (rather than simply adopting a statement that has been generated by a machine), this could provide important evidence to support a claim that the relevant administrator did in fact engage in the ‘process of mental reasoning’ necessary for them to be considered a genuine decision maker (see **chapter 12**).

13.2 Accountability and reviewability

In traditional administrative decision-making, a properly prepared statement of reasons will promote accountability in at least two ways:

- **explainability** – it enables the person who is affected by the decision to understand it, and provides a meaningful justification for the decision, and
- **reviewability** – it provides the primary basis upon which the decision and the process that led to that decision can be reviewed, whether by the person affected themselves, or by another person or body, such as an ombuds or a court, to verify that it was lawful, reasonable and otherwise complied with norms of good decision-making.

With machine-assisted decision-making, however, these two aspects of accountability tend to become more distinct.

In particular, a statement of reasons for a machine-assisted decision that provides an appropriate and readable explanation of the decision to the person affected (explainability) is less likely to provide a sufficient basis upon which the decision and its associated decision-making processes can be properly assessed and reviewed (reviewability).

Reviewability will generally necessitate both a broader and deeper (more introspective) examination of what has occurred in the process leading to the decision, including matters relating to the design, training and testing of the machine, the data used in decision-making, the context of its deployment, and all of the surrounding technical and organisational workflows.²⁴⁷

This means, when designing and deploying machine technology, that it will not be sufficient that a (traditional) statement of reasons can be generated for each decision.

Agencies must also consider what other information needs to be kept and published to ensure that their processes and decisions can be properly reviewed for compliance with legal and good decision-making requirements (including avoiding legal non-compliance because of non-obvious features such as algorithmic bias).

In particular, agencies must ensure that any decision-making process is designed so that full, and meaningful, records of the process will be available that can enable the Ombudsman, courts or other review bodies to be satisfied that there has been no unlawful, unjust or otherwise wrong conduct.

A failure to keep such records may itself lead to an inference that the agency has engaged in wrong conduct.

Table: Why transparency of machine technology is important²⁴⁸

Role of transparency	Purpose or benefit
Dignity and respect	Respecting a person's right to an explanation as to why, how and by whom decisions were made that affect their legal or other significant interests (especially when the decision has not gone their way)
Accountability	Enhancing accountability in the exercise of public power and exposing and deterring unethical, negligent or otherwise inappropriate conduct
Early warning system	Increasing opportunities for early identification and rectification of legal and other flaws
Stakeholder input and crowd-sourcing	Encouraging both expert and lay input to improve the technology or its 'fitness' to particular contexts
Informed choice	Enabling individual choice, including whether to 'opt out' of machine processes (if possible) and/or to seek a review of a machine outcome by a human
Informed public debate	Informing democratic deliberation about the relevant function and the associated use of the technology in the exercise of that function, and about machine technology generally
Review	Allowing identification of grounds for potential challenge and enabling proper inquiry into decisions and outcomes to be undertaken to identify any error or unfairness

Recordkeeping policies and practices

Agencies that use, or intend to use, machine technology should therefore also ensure that their recordkeeping policies and practices are reviewed and that they explicitly address the records that are needed to be generated and retained in respect of machines – see **chapter 10**.

This may also mean explicitly providing that previous versions of technology that can 'read' the relevant records are also properly kept and maintained, and that staff continue to be trained to know how to use them.

Again, if an agency does not have in place recordkeeping policies and practices that ensure proper records of decision-making processes are kept and can be comprehensively reviewed, that failure may justify a finding that the agency has engaged in wrong conduct.

13.3 Publishing source code

As we noted above, providing reasons does not necessarily require releasing the detailed specifications or source code for a machine. Indeed, doing so would rarely satisfy the requirement for ‘explainability’ – that is, the provision of reasons that the person can understand.

However, scrutiny of the underlying specifications and code may be necessary if decisions are to be properly reviewable. Accordingly, these records will need to be available for review and oversight bodies.

In any case, there should be (at least) a presumption in favour of proactively publishing the specifications and source code of machine technology used for government decision-making.

As well as enhancing the transparency and accountability of government decision-making, doing so has the added benefit that it exposes the technology to appraisal by outside experts.

Indeed, just as government may release policy white-papers or exposure draft bills to draw on the expertise of interested stakeholders, agencies should consider releasing draft specifications, code and even ‘beta’ versions of new machines to draw on external expertise and help to identify flaws or potential improvements *before* the technology is put into operation.²⁴⁹

Trade secrets and commercial-in-confidence arrangements

A key transparency issue arises when an agency engages an external provider for machine technology expertise. Trade secrets and commercial-in-confidence arrangements should not be more important than the value of transparency and the requirement, where it exists, to provide reasons. Contractual confidentiality obligations negotiated between parties must also be read as being subject to legislation that compels the production of information to a court, tribunal or regulatory or integrity body.²⁵⁰

Furthermore, even if courts are willing to protect algorithms as intellectual property, the tension can be avoided by good procurement practices that demand transparency from industry, and ensure ‘that trade secrets and copyright claims do not trump the values of good governance’.²⁵¹ We agree with the advice that ‘[o]fficials should refuse to work with vendors who are not willing to make their system sufficiently transparent for appropriate auditing and review.’²⁵²

The NSW Information Commissioner has noted that ‘...there is scope to strengthen existing information access laws to better facilitate access to AI-informed decision-making, particularly where governments partner with the private sector and NGOs in using these technologies.’²⁵³

Section 121 of the GIPA Act sets out requirements for inclusion of a contractual provision relating to an immediate right of access by the agency to certain information held by a contractor. Such a provision would mean that in effect, certain information held by the contractor would be government information for the purposes of the GIPA Act. However, there are exceptions to s 121 and it only applies in certain circumstances.

The Information and Privacy Commission’s guidance for agencies negotiating confidentiality clauses is to ask and consider the responses to three key questions:

1. Who holds the information?
2. In what form is it held?
3. How will access be provided?²⁵⁴

As a minimum, agencies should ensure that the terms of any commercial contracts they enter in respect of machine technology will not preclude them from providing comprehensive details (including the source code and data sets) to the Ombudsman, courts or other review bodies as required for them to review the agency’s conduct for legal compliance.

14. Verification, testing and ongoing monitoring

14.1 Testing before adoption

Agencies need to identify ways of testing that go beyond whether the machine technology is performing according to its programming to consider whether the outputs of the machine technology are legal, fair and reasonable.

Verification and validation testing of the outputs of the machine technology must be relevant to the specific functional area, including whether it is delivering effectively against the relevant legislative mandate and policy imperatives.

Legal audit of the correctness of legal interpretation

Given the inherent risk of interpretive errors being embedded in the code of automated systems, an initial verification process should involve a thorough legal audit of the system *before* it is implemented.

Ultimately, only a court can provide a conclusive determination of the meaning of a statute. However, as courts are generally unable to provide advisory opinions, legal advice on the correctness of the interpretation of a statute encoded in a machine will need to be sought from legal experts.²⁵⁵

Ideally, those tasked with undertaking a legal audit prior to launch should not be the same lawyers as those who were involved in the design of the technology. A risk-based assessment may be appropriate to guide the nature and scope of legal audit and who should do it (for example, whether it is appropriate to seek a formal opinion from senior counsel).

Validation and accuracy testing

There are various examples that demonstrate the need to verify and validate machine technology at the outset and periodically after implementation. The domestic and family violence risk assessment tools used by Police in NSW and the ACT are illustrative. Those tools have been found to perform poorly in terms of predictive validity.²⁵⁶ In 2018, the NSW Bureau of Crime Statistics and Research (**BOCSAR**) examined the Domestic Violence Safety Assessment Tool used by NSW Police to determine its ability to accurately predict a victim's risk of repeat intimate partner victimisation. BOCSAR concluded the tool performed poorly and found that the 'study highlights the importance of empirical validation when developing a risk assessment tool'.²⁵⁷

The Queensland Police are reportedly now trialling a new risk assessment tool to be used in the domestic and family violence context – incorporating lessons learned from other jurisdictions. The Queensland Police are aware of the potential for bias in the data model and will 'develop a framework about monitoring and managing models before they are rolled out' in addition to a 'model monitoring tool' to identify and address bias on an ongoing basis.²⁵⁸

Testing for algorithmic bias and other unintended consequences

Systems and processes need to be established up front to safeguard against inaccuracy and unintended consequences, such as algorithmic bias. It is important at the project planning stage and as part of the risk management strategy for the machine technology that agencies determine testing procedures enabling them to define:

- (a) What will be tested, including key components of the machine technology such as data, training data models, and business rules.
- (b) What testing methods will be used from the range of possible techniques available to test the robustness of the machine technology and identify vulnerabilities and other issues prior to operationalising.

- (c) The frequency of testing including what major system modifications would trigger additional unscheduled testing.
- (d) Who will be involved in the design and performance of the testing. We note that the European Commission Expert Group suggests that testing processes ‘should be designed and performed by an as diverse group of people as possible’.²⁵⁹

Establishment of quality assurance process and audit trail

Prior to implementation agencies also need to develop appropriate quality assurance processes and establish performance metrics for ongoing system monitoring. A key consideration of quality assurance is the ability of the machine technology to generate a comprehensive audit trail to support scrutinising the system and ensure transparency and accountability.

14.2 Undertake monitoring, review and periodic evaluation

It is essential that machine technology be subject to ongoing monitoring, review and periodic evaluation to ensure that the technology continues to support lawful decision-making consistent with principles of good public administration. Adopting machine technology to support the exercise of an administrative function should not involve a ‘set and forget’ approach.

Agencies need to assess whether the machine technology is working as expected, and must actively continue to monitor its accuracy and the fairness of outcomes. The use of enforcement cameras such as fixed-speed cameras and red-light cameras provides an example of existing NSW legislation requiring ongoing confirmation of the accuracy of tools used in a machine technology system. Enforcement cameras must be routinely tested for accuracy and calibrated every 12 months.²⁶⁰ Certification of the cameras is required at 90-day intervals with the testing and calibration performed by a TfNSW team, an accredited laboratory under the national scheme.²⁶¹

A monitoring and review regime recognises that changes in the external environment can’t be ‘known’ to the machine technology – for example, statutory amendment or judicial interpretation that shifts the basis upon which the machine technology has been designed; or even something like a natural disaster or other external event that might require adjustments to be made to policy settings. It is also important to ensure that any changes over time – especially through machine learning – operate to increase accuracy and fairness, and do not introduce any unintended consequences such as algorithmic bias.

Machine technology governance must be fit for purpose and keep pace with machine technology capabilities.²⁶² We noted above that agencies should establish early a monitoring and review cycle, including assigning responsibilities, the scope of information and data to be reviewed, and a mechanism for monitoring the progress of any recommended changes.²⁶³

Ongoing monitoring and review of machine technology as business-as-usual may include:

1. A sustainable schedule of review and internal reporting on outcomes aligned with existing governance arrangements for risk analysis and mitigation.
2. Routine certification, testing and auditing of machine technology undertaken by an appropriate independent expert.
3. Systematic review of identified errors, false positive and false negatives.
4. Audits of the machine technology outputs as part of the agency’s overall quality assurance processes. This might include consideration of information owned by the agency such as complaints data and feedback from staff that may provide insights into the operation of the machine technology.²⁶⁴

5. Random auditing or ‘benchmarking’ of individual cases, by holding out a sample of cases for human decision independently of the machine process.²⁶⁵

Comprehensive records of how an agency has undertaken monitoring, review and evaluation of the machine technology are not only an important part of transparency and accountability, but may also be required if the machine technology is subject to external review by an oversight body such as ours or by a court.

Use of the ‘Structured Decision-Making’ tool in the NSW child protection system

Since 2011²⁶⁶, the Department of Communities and Justice (DCJ) has been using a set of tools, known as ‘Structured Decision-Making’ (SDM) to assist in the performance of its functions under the *Children and Young Persons (Care and Protection) Act 1998*.

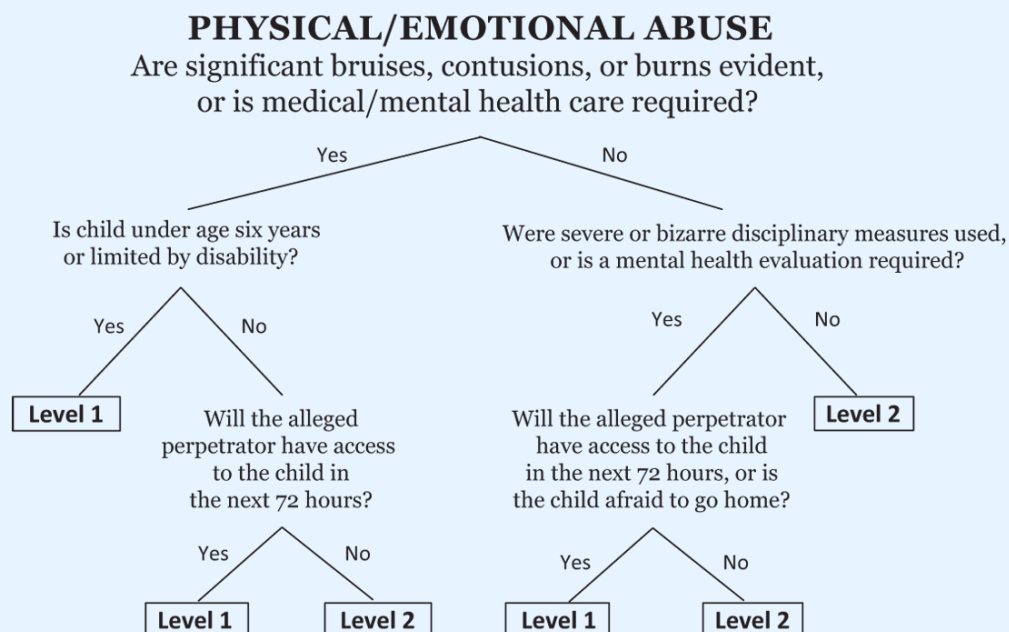
SDM was developed and is trade-marked in the United States by a non-government not-for-profit organisation called ‘Evident Change’. Evident Change was previously known as the National Council on Crime and Delinquency, which was ‘established in 1907 to assist private and public agencies serving delinquent youth’.²⁶⁷

DCJ’s website describes SDM as:

‘a process that ensures each key decision in child protection is informed by information known through research to be relevant to that decision. A number of decision-making tools underpin SDM and assist staff in making key decisions.’²⁶⁸

The core components of SDM typically comprise decision trees as well as scoring checklists. Additionally, a written narrative is required to be entered by the user to capture analysis and conclusions about particular items included in the score (or not).

An example element of a decision tree might look something like this:²⁶⁹



The reference in the above decision tree to ‘levels’ refers to the suggested response that may be warranted by a case worker; for example:

Assigned Response (select one level): Level 1 = within 24 hours
 Level 2 = within 72 hours

An example element of a scoring checklist might be:²⁷⁰

NEGLECT	Score
N1. Current complaint is for neglect	
a. No.....0	
b. Yes2	_____
N2. Prior investigations (assign highest score that applies)	
a. None -1	
b. One or more, <u>abuse</u> only..... 1	
c. One or two for <u>neglect</u>2	
d. Three or more for <u>neglect</u> 3	_____
N3. Household has previously received CPS (<i>voluntary/court-ordered</i>)	
a. No.....0	
b. Yes3	_____
N4. Number of children involved in the CA/N incident	
a. One, two, or three0	
b. Four or more2	_____
N5. Age of youngest child in the home (Age = ____)	
a. Two or older0	
b. Under two1	_____
N6. Primary caregiver provides physical care inconsistent with child needs	
a. No.....0	
b. Yes1	_____
N7. Primary caregiver has a history of abuse or neglect as a child	
a. No.....0	
b. Yes2	_____
N8. Primary caregiver has/had a mental health problem	
a. None/Not applicable0	
b. One or more apply1	_____
____ During the last 12 months AND/OR	
____ Prior to the last 12 months	

By tallying the various scores for individual items an overall score is obtained for some relevant multi-factor consideration – in this example ‘risk’:

SCORED RISK LEVEL. Assign the family’s scored risk level based on the highest score on either the neglect or abuse indices, using the following chart:

<u>Neglect Score</u>	<u>Abuse Score</u>	<u>Scored Risk Level</u>
_____ -1 - 0	_____ 0 - 1	_____ Low
_____ 1 - 3	_____ 2 - 4	_____ Moderate
_____ 4 - 8	_____ 5 - 8	_____ High
_____ 9 +	_____ 9 +	_____ Very High

These SDM tools are an example of assisted decision-making developed using machine technology but which may not necessarily be digital in their operation. Many of the SDM tools are capable of being used, and in practice have often been used by case workers in the field, in a paper form.

The US designed SDM tools are used in the NSW child protection system in relation to decision-making around reporting possible risk of significant harm, screening reports, determining response, assessing safety and risk and assessing when it is safe for a child to be restored home.²⁷¹ The publicly available information indicates that SDM is used by DCJ in the performance of statutory functions *alongside* professional judgement, as opposed to *replacing* human judgement.²⁷²

As noted above, the tools take the user through a series of definitions and questions considered relevant to the decision being made. The outputs of SDM in child protection is dependent on the specific tool. Some tools guide users to suggested actions, while others suggest an assessment outcome based on the data inputs.

For example, at the Child Protection Helpline, the SDM Screening and Response Priority Tool may generate an outcome of 'screened in', meaning the matter is to be referred to the local DCJ office for response. Based on the data inputs, the SDM will also generate a suggested response time, which may be from 'within 24 hours' to 'within 10 days'.²⁷³ We understand there is a limited discretionary override available to staff to change the response time based on the individual circumstances of a case and professional judgement.

DCJ's view is that SDM ensures consistency, accuracy and timeliness in decision-making.²⁷⁴ In a 2017 examination of the use of SDM in Los Angeles, the Office of Child Protection there also noted that:

[o]ne of the strongest identified benefits of using SDM is that, because it is a data-driven tool, it is more objective than professional judgment. When used correctly, it weighs its information uniformly and is not subject to human biases and stereotypes. It can help to guide a case worker's thinking about a case, particularly when the factors of that case are not clear cut. It may also help to address disproportionality by assessing case characteristics, risk factors, and family functioning equally across families of varying social backgrounds.²⁷⁵ [references omitted]

However, SDM tools are also subject to the same vulnerabilities as decision-making in other contexts – including user error, knowledge and training gaps, and non-compliance. The LA Office of Child Protection also noted the potential for information entered to be 'manipulated or skewed to support predetermined thinking'.²⁷⁶

The Australian Institute of Family Studies (AIFS) identified other potential weaknesses of consensus-based and actuarial risk assessment tools like those used in SDM. Actuarial tools like the one used by DCJ to assess risk may not consider unusual or context specific factors and be insufficiently flexible to incorporate professional judgment. For example, an SDM concerning risk to a child of harm may not give the user the option to input information about the strengths of a family unit, which could be relevant to the outcome. Although the user may retain an ability to exercise professional judgment to 'override' the results of the SDM, it may not be clear even to a highly expert user how such an additional factor should be weighed against the output (a composite risk score) generated by the SDM.

The AIFS also raised the potential for trust in SDM to impact whether a user rejects or accepts an output. In particular, bias may occur where a user assumes the SDM is always accurate.²⁷⁷

The Los Angeles Office of Child Safety noted that:

[o]ne of the most cited weaknesses of SDM is that, because the model is proprietary, there is a lack of transparency about how its algorithms are constructed and various

factors weighted (thus earning its classification as a “black box” model). This is concerning to users and evaluators alike, as no way exists to understand how the decision-making process is being influenced by these elements, and if any systemic biases are inherent in the tool.²⁷⁸

All tools used in the performance of administrative functions must be considered thoroughly before implementation and subject to ongoing monitoring and review to ensure that they support lawful decision-making consistent with principles of good public administration.

The US-based creator of SDM, Evident Change, states that it ‘works closely with each jurisdiction to ensure that assessments are constructed, validated, and customized for the population served. All risk assessments are tested to ensure racial equity...’²⁷⁹ and that ongoing evaluation of SDM tools is strongly encouraged.²⁸⁰ DCJ’s website states that a ‘preliminary risk calibration study’²⁸¹ was to be completed as part of the implementation process but it is not clear what ongoing validation of SDM was conducted by DCJ after 2011. There is little publicly available information about what jurisdiction-specific calibration and evaluation has taken place in relation to the use of SDM in the context of NSW families and children or in respect of different local populations within this State.

In 2017, the NSW Legislative Council General Purpose Standing Committee No 2 *Inquiry into Child Protection* questioned the effectiveness of DCJ’s SDM tools and recommended an independent review of them.²⁸² Later in 2019, issues with SDM including cultural bias were considered by the *Family Is Culture: Independent Review of Aboriginal Children and Young People in Out of Home Care (FIC)* review. The FIC review found that in practice there was little Aboriginal consultation in the application of the SDM which ‘considerably reduces the competency of the tool’.²⁸³ It found that the SDM could be manipulated by staff to result in a punitive approach to assessing Aboriginal families.²⁸⁴

The FIC review made a similar recommendation to the 2017 NSW Legislative Council General Purpose Standing Committee, that there be an independent review of SDM tools.²⁸⁵ The FIC review added that the independent review should occur in partnership with Aboriginal communities to examine adequacy from a cultural perspective.²⁸⁶ In November 2020, the NSW Government reported that DCJ was scoping a possible review of SDM tools in consultation with Aboriginal stakeholders – to be completed by July 2021.²⁸⁷

In June 2021, DCJ advised us that a Quality Services Review of its SDM tools would commence later that month. DCJ noted that:

The reviews will focus on co-designing updates with Aboriginal people, practitioners and researchers to improve racial equity, validity and accuracy to NSW population data, practice and legislative and policy settings. Implementation of the updated tools will focus on workforce and leadership development and bolstering systems to safeguard practice and decisions. These factors are pivotal requirements to ensure any assessment tool is used effectively, accurately and consistently.²⁸⁸

We have been told the review is expected to be completed over a period of 2 years including implementation of the updated SDM tools and any related changes to practice and process. Additionally, DCJ advised that it was developing an additional SDM Family Strengths and Needs Assessment tool which will be used ‘to develop a more fulsome understanding of the family’s experiences and characteristics and to support practitioners to case plan with families in an approach that targets their needs, and utilises their strengths, rather than just recognising danger or risk.’

Cost implications

One of the key benefits of machine technologies for government is its potential efficiencies and consequent cost savings. However,

when preparing a ‘business case’ for a proposed machine technology project, it is important that all costs are factored into the cost-benefit equation.

In particular the need that has been highlighted in this chapter for rigorous pre-deployment testing, as well as ongoing monitoring and auditing, is a significant cost that must be taken into account.²⁸⁹ So too is the cost of maintaining and updating the machine over time (including as legislation may change in the future), as well as the current and future training needs of operational staff.

It would also be prudent to also consider contingency costs that might be incurred in future if things go wrong – for example, if an error is detected in the machine design that means that it needs to be substantially re-coded or manual work-arounds put in place. Of course, errors can also have the potential to result in costly legal disputes and compensation claims.

Simplistically comparing a machine’s build and basic operating costs against the expenses (usually in wages) of existing manual processes will present an inaccurately inflated picture of the financial benefits of the technology.

15. Statutory provisions that authorise machine technology

As we have seen in **chapters 5 to 10**, there can be legal risks associated with the use of machine technology to support the exercise of a statutory function, especially one that requires a decision maker to exercise discretion.

It may be that, after applying the steps identified in **chapters 11 to 14**, the design team concludes that it would be unlawful or legally risky for the proposed new technology-assisted decision-making process it has designed to be used for the particular function as the law currently stands.

That raises the question: can and should the statute be amended to expressly authorise the use of machine technology?

15.1 Stating, in simple terms, that an administrator is authorised to use a machine

The simplest form of authorisation provision would be to merely add to the existing statutory function a statement that the person named is authorised to use a machine for the purpose of exercising the function.

As far as we are aware, no Australian court has ruled on the effect of such a provision.

Our preliminary view is that such a simple provision may be of limited effect. This is because, at least in most cases, we do not think that it is necessary to expressly authorise a decision maker to use technology when exercising a function (see **chapter 7**). Therefore, a provision that states that they are authorised to do this, may be doing little more than making explicit what is already implicit.

That said, adding such a provision may be useful, if only for the avoidance of doubt. However, there are some potential risks to be aware of with this approach:

1. Complacency

Amending the relevant legislation to ‘authorise machine technology use’ might have a tendency to lead an agency to falsely believe that the issues we have posed in **chapters 7-10** have been fully dealt with and can be safely disregarded. That is not the case.

The Commonwealth Ombudsman’s guidance refers to the authority for making automated decisions being put ‘beyond doubt’ if specifically enabled by legislation.²⁹⁰ This must be read with caution. Merely authorising in general terms that machine technology may be used does not necessarily mean that any specific use of that technology will be lawful. A general authority to use technology would not, for example, mean that the technology has been authorised to be used in a way that is biased, that results in a decision maker taking into account extraneous considerations, or that breaches privacy or anti-discrimination laws.

2. Unintended consequences for other statutory provisions

A second risk is a potential ambiguity affecting other statutory provisions in the same or other Acts.

For example, if an authorising provision is included in one Act but not in other Acts (or especially if it is included in one part of an Act, but not in other parts of that same Act), then questions might arise as to whether Parliament intended that machine technology is *not* authorised in those other Acts or parts of the Act that have not also expressly authorised it. That is, if the authority to use a machine is expressed in one place but not in another, was the omission in that other place deliberate, and what does that omission mean?

3. Potential uncertainty in interpreting the legislation

The third risk is that it may not be obvious how the new provision (that authorises machine technology use) can be interpreted in a way that is consistent with the function itself.

For example, if a statutory provision currently gives a named decision maker a very open discretion, and then a further provision is added simply to authorise that person to use machine technology in the exercise of that discretionary power, it may be unclear how the two provisions should be read together consistently.

Does the express authority to use a machine mean that Parliament intended that the decision maker could now do things that would otherwise involve an impermissible fetter of their discretion? This seems to be a possibility suggested by the Commonwealth Ombudsman in its *Better Practice Guide*. It suggests that, where legislation has expressly stated that the use of automation technology is authorised, future courts might decide that it is then acceptable for discretions to be automated in limited circumstances such as where the automatic output is set only to apply beneficially to the person affected.²⁹¹

Or did Parliament intend that the new provision, authorising the use of technology, is to be limited in standard ways by the discretionary nature of the function? That is, is the authority to use technology to be ‘read down’ so that a machine can be used, but only in ways that are consistent with the decision maker retaining and personally exercising (and not fettering) their discretion?

Ultimately the proper interpretation of the particular statute will only be able to be resolved by a court. Unless and until that happens, there may be a great deal of uncertainty (and therefore legal risk) about what the new provision actually authorises the decision maker to do.

Obviously, the only safe course in the interim is to assume that any authorising provision will be interpreted narrowly.

4. Lost opportunity to give proper consideration to legal and policy issues

The approach of simply authorising technology use is a simplistic approach that gives insufficient attention to the kinds of issues that we have raised in this report.

If it is thought necessary to expressly authorise by legislation the use of machine technology for a particular function, then in our view much more comprehensive consideration should be given as to what that legislation should include beyond simply stating that technology use is ‘authorised’ in general terms (see below).

15.2 Attributing machine outputs to an administrator

Some Commonwealth legislation has gone further than simply authorising the use of machine technology. It also provides that the output of machine technology is or may be ‘taken to be’ the decision of the administrator.²⁹²

In some cases, the provision does not provide for the administrator to over-rule or substitute their own decision for the machine’s output, which could be problematic in practice.

Other Commonwealth legislation seeks to address that problem. For example, s 495B of the *Migration Act 1958* (Cth) expressly authorises the human decision maker (the Minister) to over-rule the machine technology’s deemed decision, but only following certification that ‘the computer program was not functioning correctly’ and if the substituted decision is more ‘favourable’ to the relevant person.²⁹³

‘Nearly identical’ under the Commonwealth *Business Names Registration Act 2011*

The *Business Names Registration Act 2011* (**BNR Act**) established a scheme for business name registration specifically developed with the use of machine technology in mind.

One of the objects of the BNR Act is to avoid confusion by ensuring that business names that are identical or nearly identical are not registered: BNR Act s 16(3)(a). The Act requires the Australian Securities and Investments Commission (**ASIC**) to register the business name submitted by an entity if, among other things, the name is ‘available to an entity’. A name is available to an entity if it is ‘not identical or nearly identical’ to a range of names prescribed by BNR Act s 25. The terms ‘identical’ and ‘nearly identical’ are defined in accordance with s 26,²⁹⁴ which allows the Minister to make rules determining ‘whether a name is identical or nearly identical to another name’.

The current determination made under s 26 is the *Business Names Registration (Availability of Names) Determination 2015*. It sets out the rules that must be applied when determining whether a name is identical or nearly identical to company names or other names on the Register.

Section 66 of the BNR Act provides that **ASIC** can use computer programs ‘for any purposes for which ASIC may make decisions’ under the BNR Act, and that a ‘decision’ made by a computer program ‘is taken to be a decision made by ASIC’. There are some 14 programmed system rules that apply the business name availability rules set out in the BNR Act and the Determination.²⁹⁵

It appears, however, that the system’s rigid rule-based concept of what is nearly identical does not always square with human common sense. The names in the table below would seem to be confusingly similar.

Available name	Already registered name
Northern Beaches Tutoring Service	Northern Beaches Private Tutoring Services ²⁹⁶
Perth Martial Arts Centre	Perth Martial Arts Academy ²⁹⁷
Rainbow Beach Plumbing	Rainbow Beach Plumbing Services Pty Limited ²⁹⁸
Central Coast Surf Academy	Central Coast Surf School ²⁹⁹
Appaloosa Association of Australia	Australian Appaloosa Association Ltd ³⁰⁰
Cainscrete Plumbing	Cairns Concrete Plumbing ³⁰¹

However, in each of the above cases, the computer program determined that the first name was ‘available’ even though the second name was already registered. That is, the program did not see the two names as being identical or nearly identical.

In each case, the Administrative Appeals Tribunal (albeit in some cases with obvious reluctance) upheld the decisions on the basis that consideration of whether a business name is nearly identical is to be determined solely by applying the rules set out in the Determination, which are coded into the computer program.

These decisions were made even when the Tribunal acknowledged that the outcomes might ‘mislead and confuse’,³⁰² give rise to ‘anomaly’,³⁰³ be ‘inconsistent and arbitrary’,³⁰⁴ be ‘counter intuitive’ and ‘neither pleasing nor sensible’,³⁰⁵ or could be ‘quite absurd’.³⁰⁶

However, in three other cases, including a recent case in October 2021,³⁰⁷ the Tribunal has taken a different approach. In these cases the Tribunal has set aside ASIC’s decision to register business names that the computer program decided were ‘available’. The Tribunal considered (contrary to the decision of ASIC’s computer) that the following names were identical or nearly identical:

Already registered name	Name not available
Melbourne Children’s Psychology Clinic	Melbourne Child Psychology ³⁰⁸
Solar Repairs Perth	Solar Repairs Pty Ltd ³⁰⁹
Voices of Casey	Voices of Casey Choir ³¹⁰

In these decisions, the Tribunal refused to accept that the coded rules left no room for human discretion, and held that it was necessary for the concept of ‘nearly identical’ to be determined having regard to ‘the ordinary meaning of that term having regard to its legislative purpose’³¹¹ – something that the computer program was not able to do, and which required human intervention.

15.3 More sophisticated authorisation provisions

To date, the approach taken (primarily in the Commonwealth) to authorising in legislation the use of machine technology has involved very simple provisions of the kind described above.

While we are now seeing some legislation take a slightly more sophisticated approach – for example, the *Migration Act* provision referred to above – they remain high-level and focused on permitting, rather than regulating, the use of machine technology.

More refined approaches might include the following:

(a) Separating the discretionary and non-discretionary components of a decision

It may be possible for legislation to be amended to more clearly differentiate between those elements of a function that are authorised and expected to be done by a machine, and those that are reserved for the human administrator. This may require specifically identifying these different elements in the legislation for the first time.

For example, amendments might be made to more clearly differentiate between bright-line rules of eligibility and ineligibility (eg, a person is not eligible for a practising certificate as a legal practitioner unless they have completed a certified course of study) and discretionary issues on which judgment is required (eg, a practising certificate may only be granted if the person is ‘fit and proper’). Machine technology might be authorised to determine the first (eligibility according to rules) but not the second (eligibility according to broad discretionary or evaluative issues).

Consideration might also be given to whether even those individual elements might be amended or further broken down to facilitate processing by a machine. For example, if an

element is currently expressed in terms that confer a discretion, could the element be further sub-divided into non-discretionary and discretionary components?

This approach may be challenging because, as currently drafted, few statutory provisions expressly state whether, and in which respects, an administrator has ‘discretion’.

(b) Converting discretionary powers into non-discretionary rules

As we have seen, attempts to automate discretionary powers raise particular legal risks (see **chapter 7**). It may be tempting then, to simply amend the relevant statutory provision to remove discretion and thereby facilitate the adoption of machine technology.

That is, if the function currently involves a discretionary power to do something, the function might be redrafted so that it is expressed instead as a non-discretionary duty to do that thing whenever fixed and clear rules say that it must be done. These would be rules that a machine can process, without any suggestion that it has fettered the discretion of the decision maker – because the decision maker no longer has discretion.

This is essentially the kind of approach that has been taken in the Commonwealth *Business Names Registration Act* (see above). As that example shows, however, even in circumstances where it appears that clear and unbending rules would be appropriate, the removal of all possible discretion can lead to results that, to human intuition, might seem absurd or defy common sense.

We suggest that great care be taken before taking this approach. The prospect that machine technology will create an incentive toward legislation that eliminates all discretion in favour of fixed rules could raise concerns.

Discretion exists in the law for a reason, including to ensure that officials can provide appropriately individualised solutions that take into account the unique context of the unique human whose status, rights or interests may be affected by the exercise of functions on a particular occasion.³¹² Discretion also exists because it is frequently impossible to precisely and comprehensively detail in hard rules all possible situations that the law might need to deal with in practice. Even where that might in theory be possible, it may be undesirable. The modern trend toward overly complex and prescriptive legislative drafting has been criticised.³¹³

Removing discretion could also mean that any right to seek merits review of the decision may become ‘a meaningless and empty charade’ if the person conducting the review then lacks even a residual discretion (for example, in the event of absurdity) to make a different decision.³¹⁴

(c) Authorising the automation of discretion for ‘beneficial’ decisions only

An alternative is to retain the discretionary nature of the function, but for the legislation to be amended to authorise the exercise of that discretion to be automated in limited cases. If this authority is expressed clearly enough, it should override the usual presumption that when Parliament confers a discretion on a person it intends for the discretion to be exercised by that person and not to be fettered.

One suggestion sometimes made is that legislative authority could be given for discretion to be automated in this way only in circumstances where ‘it is to apply beneficially to the person affected’.³¹⁵ However, this approach may have two limitations:

- First, the approach would seem feasible only for a function that is clearly expressed to be binary in nature – that is, it involves a simple yes/no decision. In other cases, whether the outcome is beneficial or not may be perceived differently by the agency and the person affected. If, for example, a decision appears favourable but does not provide the person with everything they wanted, would that count as beneficial?

- Second, the implicit suggestion in this proposal that no one can be harmed if the automated process is only able to exercise powers beneficially may not be valid. Any automated process can make two types of errors – false positives (type 1) and false negatives (type 2). In this case, a false positive would mean that a machine has wrongly determined to exercise the power beneficially in respect of a person.

The possibility of these kinds of errors may involve systemic discrimination and injustice.³¹⁶ If, for example, one group of people (Group A) is systemically more likely to be the subject of false positives than another group (Group B), then even though the machine is only making beneficial determinations in any individual case, Group B may be said to be indirectly harmed – in the sense that those in Group B will be systemically subject to less favourable treatment.

(d) Authorising the automation of discretion with a right of review

Another similar approach is for legislation to expressly authorise the automation of a discretionary power but subject to certain rights of objection and review.

For example, legislation could provide that a machine can make an initial determination as to whether the discretionary power will be exercised, provided the person affected is provided with advance notice of that proposed exercise and can request instead that a human make a decision.³¹⁷

In practice, of course, this would likely have much the same effect as the approach in (c) above (as presumably people will only object to a determination made by the machine if it is unfavourable to them).

This approach may, however, avoid the first limitation identified in (c) above, as it will be up to the person affected to decide whether they consider the determination is beneficial or worth objecting to. The second limitation in (c) above might still apply – that is, a provision of this kind could still lead to systemic unfairness if the machine has a greater propensity to make favourable determinations (including ‘false positives’) in respect of some groups than others.

This approach may also raise other concerns, including that people who can and should object to the machine’s determination not doing so. This may be because of the known propensity for people to accept technology-generated outcomes as correct, or that some people may be less able, through vulnerability, to exercise their right to object and request a human-made decision.

(e) Beyond authorisation: regulation of machine technology

In the current legislative approaches we have seen, there has been little consideration given to including in the legislation not merely a general authorisation to use a machine, but also specific requirements to ensure that its use will be consistent with administrative law values of the kind discussed in this report.

We see this as a missed opportunity. However, before we take up this issue further (in **section 15.5** below), there is an important alternative approach to the legislative authorisation of machine technology use that could be considered.

15.4 Transforming the substantive statutory function

An alternative approach to authorising machine technology for use in the exercise of a statutory function is to replace the relevant statutory function itself.

We are, of course, not suggesting that it is appropriate to amend legislation as a way of sidestepping the principles and concerns of administrative law. However, where a machine is appropriately designed to generate good public policy outcomes, then it may be appropriate to consider reframing the entire statutory function itself rather than seeking to simply overlay machine authorisation onto the existing function.

This approach requires a coordinated exercise of legislative and machine design. A simple example is as follows:

- Assume a relevant transport agency has the function of deciding whether or not a person should be granted a driver licence. Currently this function is given to the Secretary of that agency, who can delegate it to any officer of the agency. Such a delegation has been made to hundreds of front line officers above a certain level of seniority.
- The legislation provides that a licence may only be granted if certain conditions apply (eg the person is over a certain age and has been certified as having passed a driving test) and must not be granted if certain other conditions apply (eg the person is subject to a current licence disqualification period).
- Now assume the transport agency wants to automate the process of issuing licences. Following the steps we set out in **chapters 11 to 14**, assume it designs a state-of-the-art machine that would perform the function flawlessly.
- One approach the agency could take is simply to roll out the machine, and hope that its use is not unlawful. However, this is a very high-risk approach, especially if the current function is expressed in the legislation as involving some element of discretion (see **chapter 7**). Amending the Act to include a simple provision stating that the Secretary (or delegate) is authorised to use a machine in the exercise of the function may not completely remove the risk (see above).
- An alternative approach may be to design a wholly new legislative scheme to replace the current driver licence issuing process. Under the new scheme, instead of the Secretary having a statutory function (personally or through delegates) of deciding whether to issue driver licences, the Secretary's function would be to approve and authorise the operation of a machine that issues such licences. (The legislation might also provide for the Secretary to retain a separate discretionary power to issue licences outside the automated process.)
- In this way, one statutory function (issuing driver licences) is replaced by a new statutory function (approving a machine that issues driver licences). And under the new legislative scheme, this relevant function (approving the machine) is appropriately performed by a legally responsible and accountable human administrator, the Secretary.

This approach could in some cases be preferable and produce a better public policy outcome than either alternative of attempting to automate without legislative authorisation or enacting an authority provision of the kind we discussed in the previous sections. Advantages to this approach may include:

1. First, and most obviously, there will be less doubt as to the legal efficacy of decisions made by the machine.
2. Secondly, there will be transparency, as people will know that decisions are being made by the machine, and the extent to which that is happening.

3. Thirdly, this approach does not circumvent administrative law, but it does change where it is focused. The decisions of the Secretary in relation to approving and operating the machine will be decisions that are subject to the usual requirements of administrative law and good administrative practice.

This approach does have a major drawback, which is the potential for fewer or weakened mechanisms for legal redress for those who may be wronged by the process, or harmed by the outcomes of the new system.³¹⁸

However, done right, the preparation, introduction and enactment of a new legislative scheme provides an opportunity for full public and Parliamentary debate about what legal redress avenues are required, and what other properties the machine-driven scheme must exhibit to ensure that it will uphold norms of good public administration.

In the hypothetical example above concerning proposed new legislation for a machine-operated driver licensing scheme, Parliament might also consider additional legislated elements of that scheme such as:

- mandating that the machine's specifications (and any updates) be made public
- requiring the machine to be subject, prior to deployment and at regular intervals during its operation, to external legal and technology audits, with findings to be made public
- introducing a clear right of full merits review to an appropriately senior (and *human*) officer of the agency for any determination made by the machine that a person wishes to challenge
- ensuring that any person aggrieved by the machine outputs has the ability to some form of external review or right to complain to an appropriate oversight body.

In other words, designing legislation *and* machines together may provide an opportunity for better control of machine use – including by emphasising the primacy of legislation and by ensuring that the machine is fully visible to lawmakers, to the public, and to review bodies.

15.5 Mandating properties of machine technology

Whichever approach is taken, if legislation is to be drafted and debated to authorise machine technology, this also presents the opportunity to ensure that the technology has all of the properties necessary for its use to meet legal, Parliamentary and community expectations of good administrative practice.

This does not appear to be a naïve hope. When, for example, legislation was introduced for the use of machine technology for the detection of mobile phone offences while driving (see **chapter 4**), the legislation as introduced would have done no more than facilitate its use by reversing the onus of proof on drivers who wished to dispute infringement notices in court.³¹⁹ However, the debate, both in the Parliament and in a Parliamentary Committee,³²⁰ raised broader issues including the privacy and security of the personal data collected and the potential for algorithmic bias. Although the legislation now appears to have stalled in its entirety, a number of amendments had been proposed, including to expressly legislate rules for the proper destruction of images and personal data.³²¹

The following is not intended as an exhaustive list, but provides an illustration of the kinds of properties that could be considered when legislating a new function for the approval of a machine technology system. The properties that are most important will differ depending on the context.

For example, in some contexts, having stronger properties in terms of *reviewability* may mean that weaker properties in terms of *explainability* could be acceptable. Where there is a possibility of algorithmic bias, having stronger properties relating to *testing* and *auditing* might be particularly important.

Properties	Example of qualities that could be prescribed
Is it visible ?	What information does the public, and especially those directly affected, need to be told regarding the involvement of the machine, how it works, its assessed accuracy, testing schedule etc? Are the design specifications and source code publicly available – for example as ‘open access information’ under the GIPA Act? Is an impact assessment required to be prepared and published? ³²²
Is it avoidable ?	Can an individual ‘opt out’ of the machine-led process and choose to have their case decided through a manual (human) process?
Is it subject to testing ?	What testing regime must be undertaken prior to operation, and at scheduled times thereafter? What are the purposes of testing (eg compliance with specifications, accuracy, identification of algorithmic bias)? Who is to undertake that testing? What standards are to apply (eg randomised control trials)? Are the results to be made public?
Is it explainable ?	What rights do those affected by the machine outputs have to be given reasons for those outcomes? Are reasons to be provided routinely or on request? In what form must those reasons be given and what information must they contain?
Is it accurate ?	To what extent must the predictions or inferences of the machine be demonstrated to be accurate? For example, is ‘better than chance’ sufficient, or is the tolerance for inaccuracy lower? How and when will accuracy be evaluated?
Is it subject to audit ?	What audit records must the machine maintain? What audits are to be conducted (internally and externally), by whom and for what purpose?
Is it replicable ?	Must the decision of the machine be replicable in the sense that, if exactly the same inputs were re-entered, the machine will consistently produce the same output, or can the machine improve or change over time? If the latter, must the machine be able to identify why the output now is different from what it was previously?
Is it internally reviewable ?	Are the outputs of the machine subject to internal review by a human decision maker? What is the nature of that review (eg full merits review)? Who has standing to seek such a review? Who has the ability to conduct that review and are they sufficiently senior and qualified to do so?

<p>Is it externally reviewable?</p>	<p>Are the outputs of the machine subject to external review or complaint to a human decision maker?</p> <p>What is the nature of that review (eg for example, merits review or review for error only)? Who has standing to seek such a review? If reviewable for error, what records are available to the review body to enable it to thoroughly inspect records and detect error?</p>
<p>Is it compensable?</p>	<p>Are those who suffer detriment by an erroneous action of the machine entitled to compensation, and how is that determined?</p>
<p>Is it privacy protective and data secure?</p>	<p>What privacy and data security measures and standards are required to be adhered to? Is a privacy impact assessment required to be undertaken and published? Are there particular rules limiting the collection, use and retention of personal information?</p>

16. Coda – new laws for new technology?

Throughout this report we have focused on how existing laws and norms of public sector administrative decision-making may control the use of machine technology when used in that context.

Uncertainties and gaps in the existing legal framework

However, we have also observed that there are likely to be, at least initially, significant uncertainties and potentially significant gaps in the existing legal framework given what are likely to be rapid and revolutionary changes to the way government conducts itself in coming years.

One risk, for example, may be that machine technology will be capable of producing extremely large-scale *systemic* injustices that are not possible or likely under current technologies. The existing framework of administrative law, which is typically concerned with the protection of individual rights and interests, may be ill-equipped or at least too slow to respond.³²³

Indeed, the fact that administrative law is primarily developed through the decisions of courts, tribunals and other review bodies is one of its strengths, as it provides flexibility – including to accommodate changing technologies. However, it also means that any consideration and determinative rulings are inherently ‘after the fact’. The pace at which legal certainty is provided may be substantially slower than is desirable.

Because oversight decisions first require a challenge to be brought, courts and others must generally wait for opportunities to arise when they can consider and offer certainty about the application or extension of legal and ethical norms to new situations and new technologies.

Those opportunities may arise even less rapidly or frequently in the case of the new machine technologies, given the following:

- (a) **invisibility** – currently, and despite the views we set out in **chapter 13**, government agencies are not routinely publishing or informing those affected about their use of machine technology
- (b) **technical opacity** – the complexity of the technology may make it harder for individuals wronged by decisions to recognise error or maladministration, even if intelligible reasons for the individual decision in their case are given³²⁴
- (c) **systematisation** – errors introduced by the technology are more likely to be systemic in nature, rather than just affecting a particular individual, which may make it less likely that any individual will challenge the decision
- (d) **vulnerability** – in the public sector context machine technology has more frequently been used in ways that affect people in lower socio-economic groups or who are otherwise more vulnerable, and who may accordingly have less capacity or resources to recognise and challenge potentially unlawful decisions.

We finish this short final chapter then, by simply asking the question of whether existing laws and associated institutional frameworks are adequate, and whether new laws should be considered.

Modernising administrative law for the new machinery of government

In the previous chapter, we noted that if a statute is to be amended to specifically authorise a particular use of machine technology, this creates an opportunity for Parliament to consider scaffolding a governance framework around that technology. That could include stipulating certain properties the system must exhibit in terms of transparency, accuracy, auditability, reviewability, and so on.

However, is there a need to consider more generally applicable legal or institutional reform, particularly to ensure that machine technology is subject to appropriate governance, oversight and review when used in a government context?³²⁵

There may be precedent for this approach. The machinery of Australia's modern administrative law – the administrative decisions tribunals, ombudsman institutions, privacy commissions, and (in some jurisdictions) codified judicial review legislation – was largely installed in a short period of intense legislative reform, responding to what was then the new technology of modern government at the time.³²⁶

The Government of Canada has also recently taken steps in this direction, with its 'Directive on Automated Decision-Making'. The Directive was issued in 2019 as part of the Government's commitment to using artificial intelligence 'in a manner that is compatible with core administrative law principles such as transparency, accountability, legality, and procedural fairness'. The Directive sets out requirements to increase transparency of such systems including public notice of the use of an automated decision system, the provision of reasons for decisions and release of source code. Quality assurance requirements include testing and monitoring, ensuring data quality and consultation with legal services to ensure the system is legally compliant.


The Directive aims to set core requirements for increased transparency and reduced risk for government use of machine technology. However, the Directive does not apply to all agencies and there are certain limitations of scope such as application only to automated decision systems developed or procured after 1 April 2020.³²⁷

As we come to understand better how machine technology will impact on government decision-making, consideration may need to be given to whether, and if so how, the legal and institutional framework might again need to be modernised to address the new challenges.

In the interim, we will also continue to consider the role and value the NSW Ombudsman can and should bring to this area, given our existing statutory functions and resources.

Ombudsman institutions³²⁸ have proven useful in many areas where traditional regulation and judicial enforcement is inadequate or inefficient. They seem particularly well-placed to also play an active role in the burgeoning fields of machine technology given their independence, ability to operate with greater agility and informality than judicial processes, and powers to require agency co-operation and access. Ombudsman institutions also have the ability to not only respond reactively to individual complaints but also to proactively inquire into potential systemic issues, and the ability to make public reports and recommendations to improve practices, policies and legislation.³²⁹ On the other hand, it must also be recognised that ombudsman institutions may be limited at present by a lack of the deep technical skills and resources needed for any sophisticated deconstruction and interrogation of data quality and modelling, which may, at least in some cases, be required for effective scrutiny and investigation of machine technology.³³⁰

Endnotes



-
- 1 See Australian Government, Digital Transformation Strategy 2018-2025, (Webpage) <<https://www.dta.gov.au/digital-transformation-strategy>>.
 - 2 NSW Government, *NSW Digital Government Strategy*, (Strategy, May 2017) <<https://www.digital.nsw.gov.au/sites/default/files/DigitalStrategy.pdf>>; that Strategy has been revised and replaced by NSW Government, *Beyond Digital* (Strategy, November 2019) <https://www.digital.nsw.gov.au/sites/default/files/Beyond_Digital.pdf>.
 - 3 See Andrew Le Sueur, 'Robot Government: Automated Decision-Making and its Implications for Parliament' in Alexander Horne & Andrew Le Sueur (eds), *Parliament: Legislation and Accountability*, Hart Publishing: Oxford, 2016, at 181.
 - 4 We have chosen to use this term – machine technology – as it allows us to sidestep technical debates, for example about what does and does not count as artificial intelligence. We also generally avoid terms like 'automated decision-making' where possible, as it may be confusing (commentators who use this term generally mean it to include not only (fully) automated decision-making but also assisted decision-making) and because it may tend toward question-begging argument (e.g., 'a decision that is automated is a *decision* because it is an 'automated decision'): cf Yee-Fui Ng & Maria O'Sullivan, 'Deliberation and Automation – When is a Decision a 'Decision'?' (2019) 26 *Australian Journal of Administrative Law* 21.
 - 5 Administrative Review Council, *Automated Assistance in Administrative Decision Making*, Report No 46, 1 January 2004 <<https://www.ag.gov.au/sites/default/files/2020-03/report-46.pdf>>.
 - 6 See eg Dominique Hogan-Doran, 'Computer Says "No": Automation, Algorithms and Artificial Intelligence in Government Decision-making' (2017) 13 *The Judicial Review* 345; Monika Zalnierute, Lyria Bennett Moses and George Williams, 'The Rule of Law and Automation of Government Decision-Making' (2019) 82(3) *Modern Law Review* 1; Will Bateman, 'Algorithmic Decision-making and Legality: Public Law Dimensions' (2020) 94 *Australian Law Journal* 520; Katie Miller, 'The Application of Administrative Law Principles to Technology-Assisted Decision-making' (2016) 86 *Australian Institute of Administrative Law Forum* 20, 23, 24, 26 ('Application of Administrative Law Principles'); Sarah Lim, 'Re-thinking Bias in the Age of Automation' (2019) 26 *Australian Journal of Administrative Law* 35; Melissa Perry & Alexander Smith, 'iDecide: The Legal Implications of Automated Decision-Making' (Speech, Cambridge Centre for Public Law Conference 2014, 15-17 September 2014); Anna Huggins, 'Addressing Disconnection: Automated Decision-making, Administrative Law and Regulatory Reform' (2021) 44 *University of New South Wales Law Journal* 1048 ('Addressing Disconnection'); Matthew Groves, 'Fairness in Automated Decision-Making', in Janina Boughey and Katie Miller, *The Automated State: implications, challenges and opportunities*, The Federation Press: Alexandria, 2021, at 31.
 - 7 See eg Jennifer Cobbe, 'Administrative Law and the Machines of Government: Judicial Review of Automated Public-Sector Decision making' (2019) 39(4) *Legal Studies* 636; Jennifer Cobbe et al, 'Centering the Rule of Law in the Digital State' (2020) 53(10) *IEEE Computer* 1 ('Centering the Rule of Law'); Marion Oswald, 'Algorithm-assisted Decision-making in the Public Sector: Framing the Issues Using Administrative Law Rules Governing Discretionary Powers' (2018) 376(2128) *Philosophical Transactions of the Royal Society A* 1; Madeleine Waller & Paul Waller, 'Why Predictive Algorithms are So Risky for Public Sector Bodies' (Article, October 2020), available at <https://ssrn.com/abstract=3716166>; Le Sueur (n 3); Jennifer Raso, 'AI and Administrative Law' in Florian Martin-Bariteau & Teresa Scassa (eds), *Artificial Intelligence and the Law in Canada*, Toronto: LexisNexis Canada, 2021.
 - 8 A number of commentators have proposed 'algorithmic impact assessment' processes be undertaken similar to environment or privacy impact assessments: see, eg, Michele Loi, *Automated Decision Making in the Public Sector: An Impact Assessment Tool for Public Authorities* (Report, Algorithm Watch, 2021); Nicol Turner Lee, Paul Resnick and Genie Barton, *Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms* (Report, Brookings, 22 May 2019).
 - 9 Department of Customer Service, *AI Strategy* (Strategy, 4 September 2020) <<https://www.digital.nsw.gov.au/policy/artificial-intelligence-ai/ai-strategy>>.
 - 10 Department of Customer Service, *NSW Government Artificial Intelligence (AI) Ethics Policy* (Policy, No. 2, September 4 2020) <<https://arp.nsw.gov.au/assets/ars/attachments/AI-Ethics-Policy.pdf>>.
 - 11 However, there is certain conduct that the Ombudsman is not able to investigate. We cannot investigate conduct of the NSW Police Force or conduct of a court. Complaints about the conduct of the NSW Police Force may be made to the Law Enforcement Conduct Commission: *Law Enforcement Conduct Commission Act 2016* (NSW) s 35.
Of particular relevance in the context of machine technology, we also cannot investigate "[c]onduct of a public authority relating to alleged violations of the privacy of persons": *Ombudsman Act 1974* (NSW) Sch 1 item 17. Complaints about alleged violations of, or interference with, the privacy of a person in NSW may be made the Privacy Commissioner: *Privacy and Personal Information Protection Act 1998* (NSW) s 45.
 - 12 This appears likely to be the case in respect of some statutory rights of review, for example under the Commonwealth *Administrative Decisions (Judicial Review) Act 1977* (ADJR Act), which are premised on there being an 'administrative decision' to be challenged: Huggins, 'Addressing Disconnection', (n 6), 1059-1064. It is less clear that common law judicial review rights (as apply in New South Wales) would be similarly affected. The Commonwealth ADJR Act expressly requires any challenge to be about a 'decision' - whether that be the decision itself (s 5), conduct in making the decision (s 6), or a

failure to make a decision (s 7). Although judicial review at common law is also, as a matter of practice, also typically concerned with decision-making (and the grounds of challenge are similar), its central concern is ultimately with lawful administrative action or inaction. The origins of those common law rights are the traditional orders (prerogative writs) that compel the performance of a public duty (mandamus), prohibit conduct outside of jurisdiction (prohibition) or quash past conduct for which there was no jurisdiction (certiorari).

There have separately been questions raised as to whether the constitutionally entrenched rights of judicial review (Commonwealth Constitution s 75(v)) may be affected by a move toward the automation of administrative decision-making, as those rights refer to relevant orders being ‘sought against an officer of the Commonwealth’: Yee-Fui Ng and Maria O’Sullivan, (n4), 31-32. On the other hand, it might be that this constitutional provision could ultimately come to limit the ability of Government to adopt fully-autonomous machines. In particular, might it be inconsistent with this provision – and therefore constitutionally impermissible – for an agency to put in place autonomous mechanisms in such a way that would result in there being no ‘officer of the Commonwealth’ against whom orders could be sought for legal (jurisdictional) errors? Cf Will Bateman & Julia Powles, ‘Response to the Commission’s Discussion Paper’, https://humanrights.gov.au/sites/default/files/2020-07/julia_powles_and_will_bateman.pdf (‘Any liability rules which sought to circumvent that constitutional rule (section 75(v)) would be invalid...’).

- 13 *Pintarich v Federal Commissioner of Taxation* [2018] FCAFC 79; (2018) 262 FCR 41. The situation is complicated where legislation purports to deem the output of a machine to be a decision by a relevant human administrator (see **chapter 15**).
- 14 See eg Cobbe et al, ‘Centering the Rule of Law’ (n 7).
- 15 There are attempts being made to bridge these fields, eg, The Australian Society for Law and Technology <<https://www.technolawsociety.org/>> and the Australian Society for Computers and the Law <<http://auscl.org/>>.
- 16 Daniel Montoya and Alice Rummery, *The use of artificial intelligence by government: parliamentary and legal issues* (e-brief, NSW Parliamentary Research Service, September 2020) 20.
- 17 Information and Privacy Commission, *Automated decision-making, digital government and preserving information access rights – for agencies* (Fact Sheet, September 2020) 1 <<https://www.ipc.nsw.gov.au/fact-sheet-automated-decision-making-digital-government-and-preserving-information-access-rights-agencies>>.
- 18 NSW Ombudsman, *Good Conduct and Administrative Practice: Guidelines for State and Local Government* (Guide, 3rd ed, March 2017) <https://www.ombo.nsw.gov.au/_data/assets/pdf_file/0016/3634/Good-conduct-and-administrative-practice-guidelines-for-state-and-local-government.pdf>.
- 19 ‘Administrative law in the public sector’, *NSW Ombudsman* (Web Page) <<https://www.ombo.nsw.gov.au/training-workshops-and-events/our-workshops/investigation-training/administrative-law-for-investigators>>.
- 20 Cf Commonwealth Ombudsman, *Automated Decision-making Better Practice Guide* (Guide, 2019) 5 <https://www.ombudsman.gov.au/_data/assets/pdf_file/0030/109596/OMB1188-Automated-Decision-Making-Report_Final-A1898885.pdf>.
- 21 The type of machine technology we reference in this report will involve an element of ‘automation’, in the sense that it will have the ability to carry out its particular task with no significant human participation during the performance of that function. And the output of that task has the potential to replace or supplement at least some aspects of what would otherwise be a human decision-making process, hence the device makes or at least contributes to ‘decisions’. However, we chose not to use the term ‘automated decision technologies’ because, as we will see, one of the most important issues in considering the legality of machine technology use is the question of: how *much* automation?
- 22 It is sometimes suggested that a key feature of a rule-based process (as distinct from a machine learning process) is that, for a given set of inputs, it will consistently generate the same output – ie if the input is x, then the output will always be the action y. However, this is not necessarily so, as a rule-based process can be designed to include elements of chance. For example, the input ‘x’ may be a “0” or “1” generated by a machine that is programmed to return one or other of those numbers with a certain probability (ie, the equivalent of a coin toss, albeit where the coin can be designed to have a pre-specified bias rather than 50:50).
- 23 Melanie Mitchell, *Artificial Intelligence: A Guide for Thinking Humans* (Pelican Books, 2019) 110.
- 24 Harry Surden, ‘Machine Learning and Law’ (2014) 89 *Washington Law Review* 87, 89.
- 25 Also known as ‘strong AI’: see Mitchell (n 23) 40-45.
- 26 *Ibid* 110.
- 27 *Ibid* 111.
- 28 *Road Rules 2014* (NSW) r 300.
- 29 ‘Mobile phone detection cameras’, *Transport for NSW* (Web Page) <[Mobile phone detection cameras - Mobile phone use - Staying safe - NSW Centre for Road Safety](#)>.
- 30 A penalty notice can be disputed by seeking review through Revenue NSW and/or electing to take the matter to court. It is noted that the *Road Transport Act 2013* (NSW) provides special evidentiary provisions covering a range of ‘detectable traffic offences’, which include ‘mobile phone use offences’.

- 31 'Mobile phone detection cameras' (n 29).
- 32 Road Transport Amendment (Mobile Phone Detection) Bill 2019 (NSW), introduced in the Legislative Assembly on 24 September 2019.
- 33 Legislative Council Portfolio Committee No. 5 – Legal Affairs, *Road Transport Amendment (Mobile Phone Detection) Bill 2019* (Report, No. 52, November 2019) 26.
- 34 See eg *Privacy and Personal Information Protection Act 1998* (NSW) s 23.
- 35 New South Wales, *Parliamentary Debates*, Legislative Assembly, 24 September 2019, 1507 (Mr Constance, Minister for Transport and Roads).
- 36 NSW Privacy Commissioner, 'Statement relating to mobile phone detection cameras' (Statement, 19 November 2019) <<https://www.ipc.nsw.gov.au/statements/statement-relating-mobile-phone-detection-cameras>>.
- 37 See 'FAQs: How will privacy be protected', *Transport for New South Wales* (Web Page) <<https://roadsafety.transport.nsw.gov.au/stayingsafe/mobilephones/technology.html#faq6>>.
- 38 Total number issued was 261,768 as at 1 October 2021. 'Data and statistics', *Revenue NSW* (Web Page) <<https://www.revenue.nsw.gov.au/help-centre/resources-library/statistics>>.
- 39 Amy Cheng, 'AI phone detection busts 500 drivers a day', *Government News* (online, 17 December 2019) <<https://www.governmentnews.com.au/ai-phone-detection-busts-500-drivers-a-day/>>.
- 40 'Mobile phone detection cameras', *Transport for NSW* (Web Page) <[Mobile phone detection cameras - Mobile phone use - Staying safe - NSW Centre for Road Safety](#)>.
- 41 Daniel Montoya and Alice Rummary, *The use of artificial intelligence by government: parliamentary and legal issues*' (e-brief, NSW Parliamentary Research Service, September 2020) 5.
- 42 See **annexure A**.
- 43 See **chapter 14**.
- 44 NSW Legal Aid, *Annual Report 2019-2020* (Report, November 2020) 60 <https://www.legalaid.nsw.gov.au/data/assets/pdf_file/0003/42546/201116_2020_LANSW_Annual_Report_WEB.pdf>.
- 45 'The 'human' side of claims triage', *icare* (Web Page) <<https://www.icare.nsw.gov.au/news-and-stories/the-human-side-of-claims-triage/#gref>>.
- 46 'How we use data and analytics', *Australian Taxation Office* (Web Page, 30 June 2021) <https://www.ato.gov.au/About-ATO/Commitments-and-reporting/Information-and-privacy/How-we-use-data-and-analytics/?=Redirected_URL#Deliveringgreaterautomationanddigitaliser>.
- 47 Including health, criminal justice and education settings. See for example, the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) risk assessment tool to predict the likelihood of reoffending – it is used in some jurisdictions in the United States as part of sentencing and other criminal justice processes. Machine technology in education can be found in Sweden in managing student loans and grants: Monika Zalnieriute, Lyria Bennett Moses and George Williams, 'The Rule of Law and Automation of Government Decision-Making' (2019) 82(3) *Modern Law Review* 1, 25. A 2019 survey of US federal agency use of AI found that many agencies have experimented with AI and machine learning: David Freeman Engstrom et al *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies* (Report, February 2020) <<https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf>>.
- 48 Cobbe et al, 'Centering the Rule of Law' (n 7) 4.
- 49 *Hemmett v Market Direct Group Pty Limited [No 2]* [2018] WASC 310, [69].
- 50 These benefits have been acknowledged by many academics and other legal commentators, even as they have also expressed concerns that some uses of the technology may not be compatible with administrative law principles: see Karen Lee and Ellen Rock, *The Impact of Information Technologies Upon the Teaching of Administrative Law* (Report, December 2019) 6, citing John Carroll and Amanda Ryan, *Artificial Intelligence & Automated and Computer Assisted Decision Making in Government* (Clayton Utz, October 2018); Hogan-Doran (n 6) 346- 7; Katie Miller, 'The Application of Administrative Law Principles to Technology-Assisted Decision-making' (2016) 86 *Australian Institute of Administrative Law Forum* 20, 23, 24, 26 ('Application of Administrative Law Principles'); Yee-Fui Ng & Maria O'Sullivan, (n4), 21; Lim (n 6) 35; Bateman (n 6).
- 51 Matthias Daub et al 'Automation in government: Harnessing technology to transform customer experience', *McKinsey & Company* (online, 28 September 2020) <<https://www.mckinsey.com/industries/public-and-social-sector/our-insights/automation-in-government-harnessing-technology-to-transform-customer-experience>>.
- 52 Minister for Digital, Minister for Customer Service, 'Digitising the rule of government to make compliance easy' (Media release, 1 October 2020) <<https://www.nsw.gov.au/media-releases/digitising-rules-of-government-to-make-compliance-easy>>.
- 53 Independent Panel of the APS Review, *Our Public Service Future* (Report, 20 September 2019) 25 <<https://pmc.gov.au/sites/default/files/publications/independent-review-aps.pdf>>.
- 54 Matthias Daub et al (n 51).

- 55 Cobbe et al, 'Centering the Rule of Law' (n 7) 3.
- 56 Many of these are discussed in Australian Human Rights Commission, *Human Rights and Technology* (Final Report, 1 March 2021) ('*Human Rights and Technology*').
- 57 New South Wales Law Reform Commission, *Appeals in Administration* (Report 16, December 1972) [6].
- 58 Cf Madeleine Waller and Paul Waller (n 7), who argue that consideration of ethics may be 'superfluous':
 'The understanding of "ethical behaviour" depends on social context: time, place and social norms. Hence we suggest that in the context of public administration, laws on human rights, statutory administrative functions, and data protection provide the basis for appraising the use of algorithms: maladministration is the primary concern rather than a breach of "ethics": at 4-5, 11.
- 59 Of course, although not explicitly couched in 'human rights' terms, a core pre-occupation of administrative law and good administrative practice is the protection of fundamental human rights: see Australian Human Rights Commission, *Human Rights and Technology* (n 56) 55.
- 60 *Jackson v Sterling Industries Limited* [1987] HCA 23; (1987) 162 CLR 612, 623-624.
- 61 *Plaintiff M61/2010E v The Commonwealth (Offshore Processing Case)* [2010] HCA 41; (2010) 243 CLR 319, 354 [78].
- 62 *Corporation of the City of Enfield v Development Assessment Commission* [2000] HCA 5; (2000) 199 CLR 135, 157 [56].
- 63 See also Australian Human Rights Commission, *Human Rights and Technology* (n 56) 55 ("Generally, the [administrative] law... will apply, regardless of how the relevant administrative decisions are made, or the technology used to make the decisions.").
- 64 See eg Robin Creyke, 'Administrative Justice – Towards Integrity in Government' (2007) 31(3) *Melbourne University Law Review* 705.
- 65 Oswald (n 7) 3. Other areas of law have also done this, for example Electronic Transactions legislation first enacted in 1999 and 2000 and now in effect in all Australian jurisdictions.
- 66 Lee and Rock (n 50) 10.
- 67 'Administrative law in the public sector', *NSW Ombudsman* (Web Page) <<https://www.ombo.nsw.gov.au/training-workshops-and-events/our-workshops/investigation-training/administrative-law-for-investigators>>.
- 68 For example, requirements can be grouped according to whether a failure to comply with them gives rise to a right to challenge the decision in the Courts by way of judicial review, and if they do the various individual 'grounds' of such review. They can also be grouped broadly by considering whether a failure to comply with them would mean: (a) the decision is invalid (jurisdictional error); (b) there has been some other breach of law (other legal error); or (c) the decision, or its processes, is otherwise wrong (for example, in a way that could result in an adverse finding under section 26 of the Ombudsman Act).
- 69 Bernard McCabe, 'Automated decision-making in (good) government' (2020) 100 *Australian Institute of Administrative Law Forum*, 106, 117.
- 70 Not to be confused with the broader use of the term 'executive power', when used in the context of distinguishing the three branches of government – legislature, executive and judiciary.
- 71 Robin Creyke et al, *Control of Government Action: Text Cases & Commentary* (LexisNexis Butterworths, 5th ed, 2019) 543. There are also a limited category of 'prerogative' powers, which are non-statutory functions that only the Crown can exercise as a residue of English history – such as pardoning criminals, conferring honours, and declaring war. Although these are not immune from aspects of administrative law (see eg *R (on the application of Miller) v The Prime Minister* [2019] UKSC 41 (Miller (No 2)) for simplicity we do not focus on them in this report. It seems, in any case, unlikely that matters of prerogative power will be (or should be) automated by the use of machines any time soon.
- 72 Commonwealth Ombudsman, Centrelink's automated debt raising and recovery system – A report about the Department of Human Services' online compliance intervention system for debt raising and recovery (Report, No 2, April 2017) <https://www.ombudsman.gov.au/data/assets/pdf_file/0022/43528/Report-Centrelinks-automated-debt-raising-and-recovery-system-April-2017.pdf>.
- 73 Louise Macleod, 'Lessons Learnt About Digital Transformation and Public Administration: Centrelink's Online Compliance Intervention' (2017) 89 *AIAL Forum* 59; Senate Community Affairs References Committee, *Centrelink's compliance program* (Second interim report, September 2020) 20-22 <https://www.aph.gov.au/Parliamentary_Business/Committees/Senate/Community_Affairs/Centrelinkcompliance/Second_Interim_Report/section?id=committees%2freportsen%2f024338%2f27760#footnote56ref>.
- 74 Federal Court of Australia, VID611/2019, 27 November 2019.
- 75 Ibid.
- 76 'Class action settlement', *Services Australia* (Web Page, 4 November 2021) <<https://www.servicesaustralia.gov.au/individuals/subjects/information-about-refunds-income-compliance-program/class-action-notices>>.
- 77 See the reference to Constitution Ch III, s 75(v) in n 12.

- 78 *Racecourse Co-operative Sugar Association Ltd v Attorney-General (Q)* [1979] HCA 50; (1979) 142 CLR 460, 481. As an example, the *Companion Animals Act 1998* (NSW) gives certain functions to the Chief Executive of the Office of Local Government, and also permits the Chief Executive to delegate any of his or her functions under the Act (other than their power of delegation) or regulations to any officer of the Office of Local Government: *Companion Animals Act 1998* (NSW) s 88. Some functions that are delegated may also be able to be sub-delegated by the delegate to a sub-delegate.
- 79 Legislation also sometimes provides that someone who is given a power may 'authorise' someone else to exercise that power. This means of empowering another person to exercise a power will, in many respects, resemble a delegation of power: In *Re Reference Under s 11 of the Ombudsman Act 1976 for an Advisory Opinion; Ex parte Director-General of Social Services* (1979) 2 ALD 86, 94, Brennan J referred to a 'confusing similarity' between the two. For our purpose it is not necessary to distinguish between authorisation and delegation.
- 80 Subject to any conditions that the instrument of delegation might impose on the delegate. The person to whom the function is given by Parliament will usually also retain the ability to exercise the function and will, at law, remain responsible for the exercise of the function.
- 81 Currently, the law recognises as 'legal persons' both individuals and certain artificial persons, such as companies and other legally incorporated bodies. Despite suggestions that AI may one day develop to such a degree that the law might recognise such a system as having legal personality, this is clearly not the case today. See Bateman (n 6) 529-30.
- 82 Of course, it is conceivable that legislation could be amended so that something that is now required or permitted to be done by a human administrator is instead being done in practice by a machine. However, depending on how the legislation is drafted, the proper legal characterisation of what has occurred may be that the statutory function has been repealed and replaced with a provision of an entirely different nature.
- Consider, for example, a statute that gives a function to a police officer to direct cars to go or to stop, with drivers then under a legal duty to obey those directions. It is, of course, possible to amend the legislation to replace the police officer with a set of traffic lights. We might even say (in lay terms) that the green light now performs the *function* of telling drivers when to go, and the red light has the *function* of telling them when to stop. However, those are not 'legal functions'. Indeed, in that scenario, there will likely be quite different legal functions now in place. Initially, there was a legal function (on the police officer) to direct traffic. Now the police officer's legal function may be simply to enforce obedience with the traffic light signals. There will likely be other, newly created legal functions as well (perhaps on the head of the relevant roads and traffic authority), for example, to design, place, maintain, etc the system of traffic lights. A *legal* function requires there to be a *legal* person, who is its repository.
- 83 *Re Reference Under s 11 of the Ombudsman Act 1976 for an Advisory Opinion; Ex parte Director-General of Social Services* (1979) 2 ALD 86; *O'Reilly v Commissioners of State Bank of Victoria* [1983] HCA 47; (1983) 153 CLR 1. However, an administrator cannot abdicate to others those elements of a function where the administrator must form their own opinion: see *New South Wales Aboriginal Land Council v Minister Administering the Crown Lands Act (the Nelson Bay Claim)* [2014] NSWCA 377.
- 84 *After Carltona Ltd v Commissioner of Works* [1943] 2 All ER 560.
- 85 *O'Reilly v Commissioners of State Bank of Victoria* [1983] HCA 47; (1983) 153 CLR 1, 12.
- 86 *New South Wales Aboriginal Land Council v Minister Administering the Crown Lands Act* [2014] NSWCA 377, [38].
- 87 *New South Wales Aboriginal Land Council v Minister Administering the Crown Lands Act* [2014] NSWCA 377, [38].
- 88 See Miller, 'Application of Administrative Law Principles' (n 50) 20, 22. Miller argues that '[t]he need to avoid administrative 'black boxes' which are immune from review or accountability may provide a basis for extending the *Carltona* principle to public servants in the context of technology-assisted decision-making to ensure that actions of technology assistants are attributable to a human decision-maker who can be held accountable.'
- 89 Given uncertainty around the application of the *Carltona* principle (which is based on an inference as to Parliament's intent), the Commonwealth Ombudsman has suggested that the authority to use machine technology 'will only be beyond doubt if specifically enabled by legislation': Commonwealth Ombudsman (n 20) 9. That is, rather than *inferring* that Parliament must have intended that administrators be able to seek the assistance of machines, Parliament could *expressly* state that intention.
- As discussed in **chapter 15**, there are already some rudimentary examples of such legislative provisions but, as we will see, they are not without their own problems.
- 90 See, for example, *Commissioner of Victims Rights v Dobbie* [2019] NSWCA 183, which involved legislation requiring a decision maker to obtain and have regard to a report written by a relevantly qualified person but not being legally bound to accept and act on that assessment.
- 91 See, in the US context, Cary Coglianese and David Lehr, 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era' (2017) 105 *The Georgetown Law Journal* 1147, 1181-1182 ('[I]n most instances automated artificial intelligence systems, once constructed by humans, will typically function as legally permissible measurement tools.').
- 92 *Burns v Australian National University* [1982] FCA 59; (1982) 40 ALR 707, 714.

- 93 While noting that 'discretionary' is sometimes used in a narrower sense in law, the broader concept of the term is the one which is most relevant to this report.
- 94 *Jago v District Court (NSW)* [1989] HCA 46; (1989) 168 CLR 23, 76, cited in *Coal and Allied Operations Pty Ltd v Australian Industrial Relations Commission* [2000] HCA 47; (2000) 203 CLR 194, [19]. Narrower concepts of discretion than the one used in this report have been used in other administrative law contexts: see for example the discussion in *DAO v R* (2011) 81 NSWLR 568; [2011] NSWCCA 63, [46]-[52].
- 95 *Norbis v Norbis* [1986] HCA 17; (1986) 161 CLR 513, 518.
- 96 Louis L Jaffe, *Judicial Control of Administrative Action* (Little, Brown & Company, abridged student edition, 1965) 586.
- 97 *Minister for Immigration and Citizenship v Li* [2013] HCA 18; (2013) 249 CLR 332, [66].
- 98 Subject to a contrary intention: *Interpretation Act 1987* (NSW) s 5(2). The use of the word 'may' indicates that 'the power may be exercised or not, at discretion': *Interpretation Act 1987* (NSW) s 9(1). For situations where 'may' does not confer a discretion, see the discussion in *Ward v Williams* [1955] HCA 4; (1955) 92 CLR 496, 505-508; and *Finance Facilities Pty Ltd v Federal Commissioner of Taxation* [1971] HCA 12; (1971) 127 CLR 106, 134-135, 138-139.
- 99 *Duffy v Da Rin* [2014] 87 NSWLR 495; (2014) 312 ALR 340, [53]. See also *Randren House Pty Ltd v Water Administration Ministerial Corporation* [2020] NSWCA 14, which discusses polycentricity in the exercise of administrative functions.
- 100 This example is given by Melissa Perry & Alexander Smith, 'iDecide: The Legal Implications of Automated Decision-Making' (Speech, Cambridge Centre for Public Law Conference 2014, 15-17 September 2014), cited in Michael Guihot & Lyria Bennett Moses, *Artificial Intelligence, Robots and the Law* (LexisNexis, 2020) 140.
- We have included this category although, it would not be characterised as a *discretionary* function in the narrower legal sense. For example, in discussing a legislative provision that requires a decision to be made if a person is a 'member of a couple', Deputy President McCabe of the Administrative Appeals Tribunal has said:
- '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 a generating prompts to ask questions and provide commentary.'
- McCabe (n 69), 124.
- 101 This fifth category of discretion involves a special legal concept known as 'subjective jurisdictional facts', where an administrator is required to be 'satisfied' of certain things, or to possess some other particular state of mind (such as 'reasonable belief') as a threshold before exercising a function. This category is also one that, in some contexts, is not considered 'discretionary'.
- 102 Note that, in the advice of Counsel concerning the Revenue NSW case study (**annexure A**), any references to 'discretion' will generally be referring to this narrower and more formal sense of a discretionary power rather than the broader concept to which we refer.
- 103 *NEAT Domestic Trading Pty Limited v AWB Limited* [2003] HCA 35; (2003) 216 CLR 277, [138].
- 104 *Bread Manufacturers of NSW v Evans* [1981] HCA 69; (1981) 180 CLR 404, 418; *Ex Parte Mitchell James Holdings Pty Ltd* [2001] WASCA 286, [23]-[33]; *Habib v Minister for Foreign Affairs* [2010] FCA 1203; *Nashua Australia Pty Ltd v Channon* (1981) 58 FLR 325, 341-3 (NSWSC).
- 105 *Evans v Donaldson* [1909] HCA 46; (1909) 9 CLR 140; *Rendell v Release on Licence Board* (1987) 10 NSWLR 499, 506-507; *Bread Manufacturers of NSW v Evans* [1981] HCA 69; (1981) 180 CLR 404, 418-419.
- 106 *Telstra Corporation Ltd v Kendall* [1995] FCA 1481; (1995) 55 FCR 221, 231.
- 107 *Rendell v Release on Licence Board* (1987) 10 NSWLR 499, 504.
- 108 *Re Drake and Minister for Immigration and Ethnic Affairs (No 2)* [1979] AATA 179; (1979) 2 ALD 634, 641; *British Oxygen Co Ltd v Minister of Technology* [1971] AC 610, 625; *Re Romato*; *Ex parte Mitchell James Holdings Pty Ltd* [2001] WASCA 286, [26]. When it has statutory sanction or recognition, guidance material is more likely to be considered a 'valid' fetter to the exercise of discretion than guidance material that does not: *Neat Domestic Trading Pty Ltd v AWB Ltd* [2003] HCA 35; (2003) 216 CLR 277, 286-287.
- 109 Although, when this occurs, that guidance material becomes a "substitute regime" that must be complied with: *Montenegro v Secretary, Department of Education* [2020] FCAFC 210, [25].
- 110 *Searle v Commonwealth of Australia* [2019] NSWCA 127; (2019) 100 NSWLR 55, [250], citing *Plaintiff M64/2015 v Minister for Immigration and Border Protection* [2015] HCA 50; (2015) 258 CLR 173, [54], [68]-[69] and *Rendell v Release on Licence Board* (1987) 10 NSWLR 499, 504A-B.
- 111 *Re Drake and Minister for Immigration and Ethnic Affairs (No 2)* [1979] AATA 179; (1979) 2 ALD 634.
- 112 *Minister for Home Affairs v G* [2019] FCAFC 79; (2019) 164 ALD 103.
- 113 *Plaintiff M64/2015 v Minister for Immigration and Border Protection* [2015] HCA 50; (2015) 258 CLR 173, [54] (citations omitted).

- 114 This point was made in *Plaintiff M64/2015 v Minister for Immigration and Border Protection* [2015] HCA 50; (2015) 258 CLR 173, 198 [68] (Gageler J), citing *Nevistic v Immigration and Ethnic Affairs* [1981] FCA 41; (1981) 34 ALR 639, 647, in which Deane J noted the role of policy in decision-making that involves ‘competition or correlativity between rights, advantages, obligations and disadvantages’.
- 115 *Plaintiff M64/2015 v Minister for Immigration and Border Protection* [2015] HCA 50; (2015) 258 CLR 173, 199 [69].
- 116 *Re Drake and Minister for Immigration and Ethnic Affairs (No 2)* [1979] AATA 179; (1979) 2 ALD 634; *Marks v Shire of Swanhill* [1974] VR 896; *Green v Daniels* [1977] HCA 18; (1977) 13 ALR 1; *Rendell v Release on Licence Board* (1987) 10 NSWLR 499, 504.
- 117 *Re Drake and Minister for Immigration and Ethnic Affairs (No 2)* [1979] AATA 179; (1979) 2 ALD 634, 14.
- 118 *Babar v Minister for Immigration, Citizenship, Migrant Services and Multicultural Affairs* [2020] FCAFC 38; (2020) 275 FCR 413.
- 119 *NEAT Domestic Trading Pty Ltd v AWB Ltd* [2003] HCA 35; (2003) 216 CLR 277, 320. (This rule does not apply if the statute has itself fettered that discretion by expressly mandating that the policy be implemented).
- 120 *British Oxygen Co Ltd v Minister of Technology* [1971] AC 610, 625; *Rendell v Release on Licence Board* (1987) 10 NSWLR 499; *Re Romato*; *Ex parte Mitchell James Holdings Pty Ltd* [2001] WASCA 286, [27].
- 121 Commonwealth Ombudsman, *Lessons for Public Administration: Ombudsman Investigation of Referred Immigration Cases*, August 2007 <https://www.ombudsman.gov.au/__data/assets/pdf_file/0018/26244/investigation_2007_11.pdf>.
- 122 *NEAT Domestic Trading Pty Limited v AWB Limited* [2003] HCA 35; (2003) 216 CLR 277, [150] citing, among other authorities *R v Port of London Authority*; *Ex parte Kynoch Ltd* [1919] 1 KB 176, 184; *Green v Daniels* [1977] HCA 18; (1977) 51 ALJR 463, 467 and *Kioa v West* [1985] HCA 81; (1985) 159 CLR 550, 632-633.
- 123 Administrative Review Council, *Automated Assistance in Administrative Decision Making* (Report no. 46, 2004) 15-16.
- 124 James Emmett SC and Myles Pulsford, ‘Legality of Automated decision-making procedures for the making of garnishee orders’ (Joint Opinion, 29 October 2020) 11 [35]: ‘Subject to consideration of issues like agency (see *Carltona Ltd v Commissioner of Works* [1943] 2 All ER 560) and delegation, to be validly exercised a discretionary power must be exercised by the repository of that power.’
- 125 Of course, machines themselves are inherently incapable of exercising discretion. Even if machines could exercise discretion, their doing so would not be consistent with the legislation, which has conferred the discretion on a particular (human) administrator.
- 126 *Plaintiff S157/2002 v Commonwealth* [2003] HCA 2; (2003) 211 CLR 476, [25].
- 127 *Kioa v West* [1985] HCA 81; (1985) 159 CLR 550, 584; *Annetts v McCann* [1990] HCA 57; (1990) 170 CLR 596, 598.
- 128 *Minister for Immigration and Border Protection v WZARH* [2015] HCA 40; (2015) 256 CLR 326. This includes consideration of the interests and purposes which the statutory power serves to protect: *Kioa v West* [1985] HCA 81; (1985) 159 CLR 550, 585.
- 129 See Matthew Groves in Janina Boughey & Katie Miller (n 6).
- 130 *Isbester v Knox City Council* [2015] HCA 20; (2015) 255 CLR 135, 149 [31], 153-4 [50].
- 131 *Ebner v Official Trustee in Bankruptcy* [2000] HCA 63; (2000) 205 CLR 337, 350 [33].
- 132 *Minister for Immigration and Multicultural Affairs v Jia* [2001] HCA 17; (2001) 205 CLR 507, 532 [72].
- 133 Lee and Rock (n 50) 8, quoting Interview with Janina Boughey, Senior Lecturer, UNSW (Ellen Rock, Email Interview, 22 October 2019); see Lim (n 6), concluding that ‘the rule against bias is altogether too narrow and human-focused to be engaged by decisions made by predictive systems’: at 44.
- 134 Lim (n 6), arguing that the relevancy and reasonableness grounds of review are available and amendable to modification to deal with issues of algorithmic bias: at 44.
- 135 See **chapter 3**.
- 136 See <<https://www.media.mit.edu/projects/gender-shades/overview/>>.
- 137 Virginia Eubanks, ‘A Child Abuse Prediction Model Fails Poor Families’ *Wired* 15 January 2018, cited in David Freeman Engstrom et al (n47).
- 138 Turner Lee, Resnick and Barton (n 8).
- 139 Engstrom et al (n 47) 80.
- 140 *Ibid*.
- 141 ‘A-levels and GCSEs: How did the exam algorithm work?’, *BBC* (online, 20 August 2020) <<https://www.bbc.com/news/explainers-53807730>>.
- 142 Ziad Obermeyer et al, ‘Dissecting racial bias in an algorithm used to manage the health of populations’ (2019) 366 (6464) *Science AAAS* 447, 7 <<https://www.science.org/doi/10.1126/science.aax2342>>.
- 143 Katie Miller, ‘A matter of perspective: Discrimination, bias and inequality in AI’ in C Bertram et al (eds), *Closer to the Machine: Technical, social, and legal aspects of AI* (Office of the Victorian Information Commissioner, 2019) 29.

- 144 Turner Lee, Resnick and Barton (n 8).
- 145 Ibid.
- 146 Ibid.
- 147 Alexandra Chouldechova et al, 'A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions' (2018) 81 Proceedings of Machine Learning Research 1 <<https://proceedings.mlr.press/v81/chouldechova18a/chouldechova18a.pdf>>.
- 148 *Kioa v West* [1985] HCA 81; (1985) 159 CLR 550, 587.
- 149 Anna Huggins, 'Executive Power in the Digital Age: Automation, Statutory Interpretation and Administrative Law' in J Boughy and L Burton Crawford (eds), *Interpreting Executive Power* (Federation Press, 2020) 111, 118 ('Executive Power in the Digital Age').
- 150 For example, in its initial implementation, Centrelink's automated debt raising and recovery system ("robodebt") relied on averaged yearly income to calculate and commence recovery of supposed overpayments in circumstances where recipients did not provided updated information when notified: Ibid 122.
- 151 The Information and Privacy Commission has published guidance for undertaking privacy impact assessments in NSW *Guide to Privacy Impact Assessments in NSW* (Fact Sheet, May 2020) <<https://www.ipc.nsw.gov.au/guide-privacy-impact-assessments-nsw>>.
- 152 See eg *Government Information (Public Access) Act 2009* (NSW) Sch 4 cl 10, which defines **record** to mean 'any document or other source of information compiled, recorded or stored in written form or by electronic process, or in any other manner or by any other means.'
- 153 Information and Privacy Commission, *Automated decision-making, digital government and preserving information access rights – for citizens* (Fact Sheet, September 2020) 1 <<https://www.ipc.nsw.gov.au/fact-sheet-automated-decision-making-digital-government-and-preserving-information-access-rights-citizens>>.
- 154 Ibid.
- 155 *Anti-Discrimination Act 1977*, ss 7(1), 24(1), 38B(1), 49B(1), 49T(1), 49ZYA(1), 49ZG(1), 99(1).
- 156 *Anti-Discrimination Act 1977*, ss 7(1)(c), 24(1)(b), 38B(1)(c), 49B(1)(b), 49T(1)(b), 49ZYA(1)(b), 49ZG(1)(b), 99(1)(b).
- 157 Cobbe et al, 'Centering the Rule of Law' (n 7) 7; Council of Europe, Frederik Borgesius, *Discrimination, Artificial Intelligence, and Algorithmic Decision-Making*, 2018, 15; Anya E. R. Prince and Daniel Schwarcz, 'Proxy Discrimination in the Age of Artificial Intelligence and Big Data' (2020) 105 *Iowa Law Review* 1257.
- 158 Le Sueur (n 3) 183, 191.
- 159 Cobbe et al, 'Centering the Rule of Law' (n 7) 7; Australian Human Rights Commission, 'Using artificial intelligence to make decisions: Addressing the problem of algorithmic bias' (Technical Paper, 2020) 15 ('Using artificial intelligence to make decisions').
- 160 Cobbe et al, 'Centering the Rule of Law' (n 7) 6.
- 161 Australian Human Rights Commission, 'Using artificial intelligence to make decisions' (n 159) 23.
- 162 *Fitz Jersey Pty Ltd v Atlas Construction Group Pty Ltd* [2017] NSWCA 53; (2017) 94 NSWLR 606, [70].
- 163 *Commissioner of State Revenue v Can Barz Pty Ltd* [2016] QCA 323.
- 164 In law, the term 'reasons' is often used 'to encompass a decision, the reasons for the decision and the findings of fact giving rise to the decision'. More precisely, the reasons set out the process of reasoning that has led the decision maker from the findings to the decision: *Minister for Immigration and Multicultural Affairs v W157/00A* [2002] FCAFC 281; (2002) 125 FCR 433, [37] (FCAFC).
- 165 *Public Service Board of New South Wales v Osmond* [1986] HCA 7; (1986) 159 CLR 656; *Wingfoot Australia Partners Pty Ltd v Kocak* [2013] HCA 43; (2013) 252 CLR 480.
- 166 Such as where judicial review proceedings have been commenced in the NSW Supreme Court (*Uniform Civil Procedure Rules 2005* (NSW) r 59.9) or in the Federal Court, Federal Circuit and Family Court of Australia (Division 2) (*Administrative Decisions (Judicial Review) Act 1977* (Cth) s 13).
- 167 *Public Service Board of New South Wales v Osmond* [1986] HCA 7; (1986) 159 CLR 656, 666-7; *Campbelltown City Council v Vegan* [2006] NSWCA 284; (2006) 67 NSWLR 372, [24]. See, for example, the list of factors identified in *Soliman v University of Technology, Sydney* [2012] FCAFC 146; (2012) 207 FCR 277, [46].
- 168 *Segal v Waverley Council* [2005] NSWCA 310; (2005) 64 NSWLR 177, [49]-[50]; *Re Minister for Immigration and Multicultural Affairs; Ex Parte Palme* [2003] HCA 56; (2003) 216 CLR 212, 242 [105]. Both cite de Smith, Woolf and Jowell, *Judicial Review of Administrative Action* (1995, 5th ed) 459 [9-042], who note that having to explain the basis on which a decision is made is 'a salutary discipline for those who have to decide anything that adversely affects others'.
- 169 *Wainohu v New South Wales* [2011] HCA 24; (2011) 243 CLR 181, [56]; *Wingfoot Australia Partners Pty Ltd v Kocak* [2013] HCA 43; (2013) 252 CLR 480, [44]-[45]; *Sydney Ferries v Morton* [2010] NSWCA 156, [78]-[79].
- 170 *Fuller and Brown (Child support)* [2016] AATA 2007, [17].
- 171 Ibid.

- 172 *Wingfoot Australia Partners Pty Ltd v Kocak* [2013] HCA 43; (2013) 252 CLR 480, [48]. See also *Minister for Immigration and Ethnic Affairs v Wu Shan Liang* [1996] HCA 6; (1996) 185 CLR 259, 271-272.
- 173 *Public Service Association and Professional Officers' Association Amalgamated Union of New South Wales v Secretary of the Treasury* [2014] NSWCA 112, [46].
- 174 See, for example, *Acts Interpretation Act 1954* (Cth) s25D.
- 175 See Miller, 'Application of Administrative Law Principles' (n 50) 25-26.
- 176 Katie Miller, past President of the Law Institute of Victoria has said '[m]y position that a statement of reasons *should* disclose technology assistance is strengthened by considering the contrary question: why *shouldn't* a statement of reasons disclose technology assistance?': Ibid 26.
- 177 Robert French, 'Rationality and Reason in Administrative Law - Would a Roll of the Dice be Just as Good?' (Australian Academy of Law Annual Lecture, Perth, 29 November 2017) 3.
- 178 Huggins, 'Addressing Disconnection' (n 6) 1053.
- 179 See *Australian Energy Regulator v AGL Sales Pty Limited* [2020] FCA 1623, [90], citing *Australian Competition and Consumer Commission v Medibank Private Limited* [2020] FCA 1030, [21]-[24], *Australian Securities and Investments Commission v National Australia Bank Limited* [2020] FCA 1494, [51], *Australian Securities and Investments Commission v Australia and New Zealand Banking Group Limited (No 3)* [2020] FCA 1421 and *Australian Securities and Investments Commission v BT Funds Management Limited* [2021] FCA 844, [43].
- 180 Cobbe et al, 'Centering the Rule of Law' (n 7) 13.
- 181 *Road Transport Act 2013* (NSW) ss 33, 35.
- 182 Ibid s 36.
- 183 Ibid s 36(4).
- 184 See *Minister for Immigration v Li* [2013] HCA 18; (2013) 249 CLR 332, [102]; *Koon Wing Lau v Calwell* [1949] HCA 65; (1949) 80 CLR 533, 573-4; *Giris Pty Ltd v Commissioner of Taxation* [1969] HCA 5; (1969) 119 CLR 365, 383; *Re Federal Commissioner of Taxation; Ex parte Australena Investments Pty Ltd* (1983) 50 ALR 577, 578.
- 185 Also known as 'permissible considerations'.
- 186 Also described as 'mandatory considerations': *Lo v Chief Commissioner of State Revenue* [2013] NSWCA 180; (2013) 85 NSWLR 86, [9] or 'relevant considerations': *Minister for Aboriginal Affairs v Peko-Wallsend Ltd* [1986] HCA 40; (1986) 162 CLR 24.
- 187 *Minister for Aboriginal Affairs v Peko-Wallsend Ltd* [1986] HCA 40; (1986) 162 CLR 24, 39.
- 188 For example, the *Residential Tenancies Act 2010* (NSW) provides that, in considering whether to terminate a social housing tenancy agreement, the authority '*must have regard to*' a range of factors, including the likelihood that neighbouring residents or other persons will suffer serious adverse effects in the future if the tenancy is not terminated: s 154E.
- 189 *Minister for Aboriginal Affairs v Peko-Wallsend Ltd* [1986] HCA 40; (1986) 162 CLR 24, 39-40; *Wattie v Industrial Relations Secretary on behalf of the Secretary of the Department of Justice (No 2)* [2018] NSWCA 124, [134].
- 190 *NEAT Domestic Trading Pty Ltd v AWB Ltd* [2003] HCA 35; (2003) 216 CLR 277, [20]. Extraneous considerations are sometimes also referred to as "irrelevant considerations". Relevant considerations, on the hand, include all obligatory considerations as well as any available considerations that are of relevance to the decision at hand.
- 191 For example, the *Government Information (Public Access) Act 2009* (NSW) provides in s 15(c) that, in deciding whether there is an overriding public interest against disclosure of information, the decision maker must not take into account the fact that 'disclosure of information might cause embarrassment to, or a loss of confidence in, the Government'.
- 192 *Lo v Chief Commissioner of State Revenue* [2013] NSWCA 180; (2013) 85 NSWLR 86, [9].
- 193 Ibid.
- 194 While Parliament determines those considerations that are obligatory or available, how those considerations are taken into account is generally a matter for the administrator. This means that, provided the administrator takes any obligatory considerations into account, the administrator may ultimately give decisive weight or little weight to the consideration or (after consideration) may dismiss it altogether: *Allianz Australia Insurance Ltd v Cervantes* [2012] NSWCA 244; (2012) 61 MVR 443, [16].
- 195 *Hili v The Queen* [2010] HCA 45; (2010) 242 CLR 520, [48]-[49].
- 196 *Wong v The Queen* [2001] HCA 64; (2001) 207 CLR 584, 591 [6]. The same principles apply to discretionary decisions made by administrators: see eg *Ibrahim v Minister for Immigration and Multicultural Affairs* [2000] FCA 1309; (2000) 63 ALD 37, [15]; *Segal v Waverley Council* [2005] NSWCA 310; (2005) 64 NSWLR 177.
- 197 *Wong v The Queen* [2001] HCA 64; (2001) 207 CLR 584, 611 [75].
- 198 *Director of Public Prosecutions v Dalgliesh (a pseudonym)* [2017] HCA 41; (2017) 262 CLR 428, [4].
- 199 *Hili v The Queen* [2010] HCA 45; (2010) 242 CLR 520, [48]-[49].

- 200 *Wong v The Queen* [2001] HCA 64; (2001) 207 CLR 584, 601-2 [45].
- 201 *Ibid* 606 [59].
- 202 *Ibid* 608 [65].
- 203 Empirical validation is important because, unless the machine’s predictions have been demonstrated as having a degree of accuracy at least better than chance, it is difficult to see how they could be legally relevant to the decision. Even if empirically validated, the degree of accuracy revealed in the testing may also affect the weight to be given to the predictions, including relative to other considerations. Taking into account machine outputs that are inaccurate (or whose accuracy has not been empirically validated as any better on average than chance) would seem to involve taking into account an irrelevant and therefore legally impermissible consideration. This is one reason why, in **chapter 14**, we emphasise the importance of testing, and ongoing testing, of machine outputs.
- 204 Of course, such technology may raise other concerns, including algorithmic bias and the risk of non-obvious discrimination based on protected factors (such as race).
- 205 Julia Dressel and Hany Farid, ‘The accuracy, fairness, and limits of predicting recidivism’ (2018) 4(1) *Science Advances* 1, 3.
- 206 Julia Angwin et al, ‘Machine Bias’, *ProPublica* (online, 23 May 2016) <<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>>.
- 207 William Dieterich, Christina Mendoza and Tim Brennan, *COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity* (Report, Northpoint Inc., 8 July 2016) <http://go.volarisgroup.com/rs/430-MBX-989/images/ProPublica_Commentary_Final_070616.pdf>.
- 208 The Washington Post, ‘A computer program used for bail and sentencing decisions was labelled biased against Blacks. It’s actually not that clear’ October 2016 <<https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/>>.
- 209 Julia Angwin et al, (n206).
- 210 *State v Loomis*, 881 N.W.2d 749 (Wis. 2016).
- 211 ‘State v Loomis: Wisconsin Supreme Court Requires Warning Before Use of Algorithmic Risk Assessments in Sentencing’ (2017) 130 *Harvard Law Review* 1530 <<https://harvardlawreview.org/2017/03/state-v-loomis/>>.
- 212 *Ibid*.
- 213 *State Records Act 1988* (NSW) s 3.
- 214 See Miller, ‘Application of Administrative Law Principles’ (n 50) 26.
- 215 *Ibid* 31.
- 216 Huggins, ‘Executive Power in the Digital Age’ (n 149) 117; McCabe (n 69) 118.
- 217 Cf the reversal of the onus of proof of the existence of a debt in the initial implementation of the Commonwealth “Robodebt” system: Huggins (n 149) 125.
- 218 Cf McCabe (n 69) 118.
- 219 *Alcan (NT) Alumina Pty Ltd v Commissioner of Territory Revenue* [2009] HCA 41; (2009) 239 CLR 27, [4], [47].
- 220 *R v A2* [2019] HCA 35; (2019) 93 ALJR 1106, [32].
- 221 Huggins, ‘Executive Power in the Digital Age’ (n 149) 118.
- 222 Elizabeth Joh, ‘The Undue Influence of Surveillance Technology Companies on Policing’ (2017) 92 *New York University Law Review* 19; see also Jake Goldenfein, ‘Algorithmic Transparency and Decision-making Accountability: Thoughts for buying machine learning algorithms’ in C Bertram et al (eds), *Closer to the Machine: Technical, social, and legal aspects of AI* (Office of the Victorian Information Commissioner, 2019).
- 223 *Telstra Corporation Ltd v Kendall* [1995] FCA 1481; (1995) 55 FCR 221, 231.
- 224 Of course, the method by which that ‘output’ was created would need to be otherwise consistent with the legislation – for example, a decision maker’s consideration of the ‘expert’ output of machine technology would be unlawful if it resulted in the decision-making taking into account (directly or indirectly) discriminatory or otherwise prohibited considerations, or failing to take into account mandatory relevant considerations (see **chapter 9**).
- 225 Emmett and Pulsford (n 124) 23 [82].
- 226 *Ibid* 23 [83].
- 227 *Navoto v Minister for Home Affairs* [2019] FCAFC 135, [89].
- 228 *Carrascalao v Minister for Immigration and Border Protection* [2017] FCAFC 107; (2017) 252 FCR 352, [46]; *Chetcuti v Minister for Immigration and Border Protection* [2019] FCAFC 112, [65].
- 229 *Minister for Immigration and Border Protection v Maioha* [2018] FCAFC 216; (2018) 267 FCR 643, [45]. In *Hands v Minister for Immigration and Border Protection* [2018] FCAFC 225, [3], Allsop CJ described this, in the context of decisions made under the *Migration Act 1958* (Cth), as the need for an ‘honest confrontation’ with the human consequences of administrative decision-making.

- 230 *Carrascalao v Minister for Immigration and Border Protection* [2017] FCAFC 107; (2017) 252 FCR 352; *Chetcuti v Minister for Immigration and Border Protection* [2019] FCAFC 112.
- 231 Subject, of course, to ensuring that the machine technology is applying the correct and full interpretation of those terms. Depending on the particular statute, 'Place of residence', for example, could refer to a person's one principal abode or to any place that the person resides from time to time (and there could be more than one of those). A machine technology that applies data from an existing data source may produce generally correct, but occasionally incorrect, outputs if the data source was generated for a different context.
- 232 Justice Perry in a speech provides a different example of s 4AA of the *Family Law Act 1975*, that relates to the decision whether a couple is in a de facto relationship, noting that some of the specified criteria may be amenable to analysis by machine technology (eg that the persons are not legally married and not related by family) but others not able to be automated (eg whether the persons are considered to be in a relationship as a couple living together on a genuine domestic basis, having regard to a number of prescribed sub-considerations): Perry and Smith cited in *Guihot & Bennett Moses* (n 100) 140.
- 233 *Guihot & Bennett Moses* (n 100) 141.
- 234 Raja Parasuraman and Dietrich Manzey, 'Complacency and Bias in Human Use of Automation: An Attentional Integration' (2010) 52(3) *Human Factors: The Journal of the Human Factors and Ergonomics Society* 381.
- 235 Commonwealth Ombudsman, Lessons for public administration: Ombudsman Investigation of Referred Immigration Cases (Report, August 2007) 10.
- 236 See also 'What does the GDPR say about automated decision-making and profiling?', *Information Commissioner's Office (UK)* (Web Page) <<https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/automated-decision-making-and-profiling/what-does-the-gdpr-say-about-automated-decision-making-and-profiling/#id2>>.
- 237 See further Counsel's advice at **annexure A** and refer to *Guihot & Bennett Moses* (n 100) 160.
- 238 *Hands v Minister for Immigration and Border Protection* [2018] FCAFC 225; (2018) 267 FCR 628, [3].
- 239 *Goldenfein* (n 222) 49-50.
- 240 [2011] AATA 365.
- 241 [2011] AATA 365, [39].
- 242 On the problem of inherent opacity (ie a lack of 'explainability' as to how certain inputs lead to certain outputs) of some forms of machine technology, see Marc Cheong and Kobi Leins, 'Who Oversees the Government's Automated Decision-Making? Modernising Regulation and Review of Australian Automated Administrative Decision-Making' in *Janina Boughey and Katie Miller* (n 6) 174.
- 243 *Guihot & Bennett Moses* (n 100) 151-159.
- 244 Information and Privacy Commission, *Automated decision-making, digital government and preserving information access rights – for agencies* (Fact Sheet, September 2020) 1 <<https://www.ipc.nsw.gov.au/fact-sheet-automated-decision-making-digital-government-and-preserving-information-access-rights-agencies>>.
- 245 *Minister for Immigration and Ethnic Affairs v Wu Shan Liang* [1996] HCA 6; (1996) 185 CLR 259, [16].
- 246 *Miller*, 'Application of Administrative Law Principles' (n 50) 26.
- 247 *Cobbe et al*, 'Centering the Rule of Law' (n 7) 10.
- 248 This table has been modified from *Loi* (n 8) 19.
- 249 *Miller*, 'Application of Administrative Law Principles' (n 50) 32.
- 250 *Smorgon v Australia and New Zealand Banking Group Limited* [1976] HCA 53; (1976) 134 CLR 475, 489.
- 251 *Goldenfein* (n 222) 46.
- 252 *Ibid*.
- 253 Information and Privacy Commission, *Human Rights and Technology Discussion Paper* (Submission, February 2020) <https://humanrights.gov.au/sites/default/files/2020-07/10_-_information_and_privacy_commission_nsw_1.pdf>.
- 254 Information and Privacy Commission, *Automated decision making and access to information under the GIPA Act* (Case Summary) <<https://www.ipc.nsw.gov.au/case-summary-automated-decision-making-and-access-information-under-gipa-act>>.
- 255 *Huggins*, 'Addressing Disconnection' (n 6) 1052, 1057 (suggesting that 'ideally the courts would offer an advisory jurisdiction in which pro-active judicial advice regarding the correctness of the interpretation of a statute encoded in an automated system is available before that system is implemented').
- 256 *Hayley Gleeson*, 'The 39-question tool transforming the way Victoria Police assesses family violence risk', *ABC* (online, 17 May 2021) <<https://www.abc.net.au/news/2021-05-17/victoria-police-39-question-actuarial-tool-family-violence-risk/100130532>>. See *Christopher Dowling and Anthony Morgan*, 'Predicting repeat domestic violence: Improving police risk assessment tools' (2019) 581 *Australian Institute of Criminology: Trends & Issues in crime and criminal justice* 1,1, finding that the Family Violence Assessment Tool (FVRAT) in use by the ACT police 'is not a strong predictor of repeat

- domestic violence’, but that a refined version of the tool that reduced the predictive items included in the tool (from 37 to 10) was more accurate; Clare Ringland, ‘[The Domestic Violence Safety Assessment Tool \(DVSAT\) and intimate partner repeat victimisation](#)’ (2018) 213 *Bureau of Crime Statistics and Research: Contemporary Issues in Crime and Justice* 1, 1, finding that the DV Safety Action Tool used in NSW performed only slightly better than chance in predicting repeat DV. Cf Melanie Millsted and Sarah Coghlan, ‘Predictors of recidivism amongst police recorded family violence perpetrators’, *Victorian Crime Statistics Agency* (In Brief, No. 4, May 2016), finding that Victoria’s L17 tool was a significant predictor of repeat domestic violence.
- 257 Clare Ringland, (n256).
- 258 Ben Smee, ‘Queensland police to trial AI tool designed to predict and prevent domestic violence incidents’, *The Guardian* (online, 13 September 2021) <<https://amp-theguardian-com.cdn.ampproject.org/c/s/amp.theguardian.com/australia-news/2021/sep/14/queensland-police-to-trial-ai-tool-designed-to-predict-and-prevent-domestic-violence-incident>>.
- 259 European Commission, High-Level Expert Group on Artificial Intelligence, *Ethics Guidelines for Trustworthy AI*, April 2019, 22.
- 260 In accordance with *Road Transport Act 2013* (NSW) s 137 and *Road Transport (General) Regulation 2021* (NSW) cl 29.
- 261 ‘Calibration and certification’, *Transport for NSW* (Web Page) <https://roadsafety.transport.nsw.gov.au/speeding/speedcameras/calibration_certification.html>.
- 262 Peter Leonard, ‘How to avoid the ethical pitfalls of artificial intelligence and machine learning’, *BusinessThink* (online, 20 May 2021) <<https://www.businessthink.unsw.edu.au/articles/ethical-pitfalls-artificial-intelligence-machine-learning>>.
- 263 Commonwealth Ombudsman (n 20) 27.
- 264 Ibid.
- 265 Engstrom et al (n 47) 7.
- 266 SDM was introduced by DCJ following a recommendation made in the 2008 *Report of the Special Commission of Inquiry into Child Protection Services in NSW* (the Wood inquiry). The Wood inquiry had considered the lack of written guidance for staff screening reports. Wood recommended that the then Department of Community Services test use of SDM tools for assessments and interventions relating to a child or young person’s care: *Special Commission of Inquiry into Child Protection Services in NSW* (Final Report, November 2008) xv (Recommendation 9.1).
- 267 Children’s Research Centre, *The Structured Decision Making Model: An Evidence-based Approach to Human Services* (Report, 1 March 2008) Preface <http://www.evidentchange.org/sites/default/files/publication_pdf/2008_sdm_book.pdf>.
- 268 ‘Structured Decision Making’, *Communities and Justice* (Web Page, 24 September 2019) <<https://www.facs.nsw.gov.au/providers/children-families/child-protection-services/making-decisions#:~:text=SDM%20is%20a%20process%20that,staff%20in%20making%20key%20decisions>>.
- 269 This example is from the Connecticut Department of Children and Families: Children’s Research Centre, *The Structured Decision Making Model: An Evidence-based Approach to Human Services* (Report, 1 March 2008) 6 <http://www.evidentchange.org/sites/default/files/publication_pdf/2008_sdm_book.pdf>.
- 270 This example is from the California Family Risk Assessment form: Children’s Research Centre, *The Structured Decision Making Model: An Evidence-based Approach to Human Services* (Report, 1 March 2008) 10 <http://www.evidentchange.org/sites/default/files/publication_pdf/2008_sdm_book.pdf>.
- 271 The Mandatory Reporter Guide helps mandatory reporters decide whether a child is suspected to be at Risk of Significant Harm (ROSH) and a report to the Child Protection Helpline should be made. Child Protection Helpline Caseworkers use the Screening and Response Priority Tools to determine whether reports meet the threshold of ROSH and if so, the timeframe for a response. DCJ Caseworkers use the Safety Assessment, Risk Assessment and Risk Reassessment to assess safety and risk for children and families at specific points in time.
- The Restoration Assessment tool is used in making decisions about restoration of children in out of home care. A more detailed description of the SDM tools can be found on the DCJ website: ‘Text alternative to NSW Structured Decision Making Framework diagram’, *Communities and Justice* (Web Page, 3 May 2021) <<https://www.facs.nsw.gov.au/providers/children-families/interagency-guidelines/assessing-wellbeing,-safety-and-risk/chapters/the-nsw-structured-decision-making-framework/text-alternative-to-nsw-structured-decision-making-framework-diagram>>.
- 272 DCJ’s SDM policy and procedures are not published. DCJ’s ‘Permanency Case Management Policy: Rules and Practice Guidance 2019’ states that, ‘Departmental casework practitioners assess reports using Structured Decision Making (SDM) or Alternative Assessment frameworks, alongside professional judgement’: Department of Communities and Justice, *Permanency Case Management Policy: Rules and Practice Guidance* (Guide, Ver. 4.4, 2021) 18 <https://www.facs.nsw.gov.au/data/assets/pdf_file/0010/595198/PCMP-Rules-and-Practice-Guidance-2019.pdf/_recache>.
- 273 ‘Assessing wellbeing, safety and risk’, *Communities and Justice* (Web Page, 3 May 2021) <<https://www.facs.nsw.gov.au/providers/children-families/interagency-guidelines/assessing-wellbeing,-safety-and-risk/chapters/child-protection-helpline-the-helpline>>.

- 274 'Glossary', *Communities and Justice* (Web Page, 11 May 2020)
<<https://www.facs.nsw.gov.au/resources/statistics/glossary#ChildProtectionHelpline>>.
- 275 Judge Michael Nash, *Examination of Using Structured Decision Making and Predictive Analytics in Assessing Safety and Risk in Child Welfare* (Report, 4 May 2017) 3
<http://file.lacounty.gov/SDSInter/bos/bc/1023048_05.04.17OCPReportonRiskAssessmentTools_SDMandPredictiveAnalytics.pdf>.
- 276 Ibid 4.
- 277 Child Family Community Australia, *Risk assessment instruments in child protection* (Resource Sheet, June 2016)
<<https://aifs.gov.au/cfca/publications/risk-assessment-child-protection>>.
- 278 Nash (n 275) 3-4
<http://file.lacounty.gov/SDSInter/bos/bc/1023048_05.04.17OCPReportonRiskAssessmentTools_SDMandPredictiveAnalytics.pdf>.
- 279 'Customized risk & Needs Assessment', *Evident Change* (Web Page, 2021)
<<https://www.evidentchange.org/assessment/customized-risk-needs-assessment>>.
- 280 Nicole Mickelson, Traci LaLiberte and Kristine Piescher, *Assessing Risk: A Comparison of Tools for Child Welfare Practice with Indigenous Families* (Report, Centre for Advanced Studies in Child Welfare, University of Minnesota, 2017) 7
<https://casw.umn.edu/wp-content/uploads/2018/01/Risk-Assessment_FinalReport.pdf>. Note the reference is to the 'National Council on Crime and Delinquency' now known as 'Evident Change'.
- 281 Raelene Freitag and Rick Wiebush, Children's Research Centre, *An Introduction to the Structured Decision Making (SDM) System* (PowerPoint Slide, March 2009)
<https://www.facs.nsw.gov.au/data/assets/pdf_file/0007/592783/childprotection_sdm.pdf>.
- 282 Megan Davis, *Family Is Culture: Independent Review of Aboriginal Children and Young People in Out of Home Care* (FIC) (Review Report, October 2019) 215
<https://www.familyisculture.nsw.gov.au/data/assets/pdf_file/0011/726329/Family-Is-Culture-Review-Report.pdf>.
- 283 Ibid 214.
- 284 Ibid 219.
- 285 Ibid 216, 220.
- 286 Ibid 220 (Recommendation 56): 'The Department of Communities and Justice should commission an independent review of its structured decision-making tools and processes to identify how they can be improved to enhance objectivity within child protection assessments. This review should be undertaken in partnership with Aboriginal community and stakeholders to ensure that it examines the cultural adequacy of current risk and safety paradigms and tools.'
- 287 Department of Communities and Justice, *Family is Culture* (Progress Report, 25 November 2020) 20.
- 288 Letter to the NSW Ombudsman from the Department of Communities and Justice, Office of the Senior Practitioner on behalf of Michael Coutts-Trotter, Secretary, 21 June 2021.
- 289 See Sarah Crossman and Rachel Dixon, 'Government Procurement and Project Management for Automated Decision-Making Systems', in Janina Boughey and Katie Miller (n 6) 154, 170.
- 290 Commonwealth Ombudsman (n 20) 9.
- 291 Ibid. The Commonwealth Ombudsman suggests that a future court might also read these kinds of provisions (that is, a discretionary power together with a general authorisation to use machine technology) in a way that allows 'discretions to be automated' where 'the person affected is provided advance notice of the decision to be made, permitted to make submissions, and able to ask for the decision to be made or reviewed by a human decision maker'.
Having regard to the lack of judicial consideration of these kinds of provisions, we think it would extremely unsafe for any agency to assume that a court will adopt either of the interpretations suggested in the Commonwealth *Better Practice Guide*. Certainly if Parliament had intended this to be the effect of authorising legislation, it could more clearly and directly express that intention.
- 292 The Commonwealth *Business Names Registration Act 2011* s 66. Similar provisions can be found in other Federal legislation such as the *A New Tax System (Family Assistance) (Administration) Act 1999* (Cth) s 223, *Child Support (Assessment) Act 1989* (Cth) s 12A and *Social Security (Administration) Act 1999* (Cth) s 6A.
- 293 See also s 62F of the *Business Names Registration Act 2011* (Cth), which provides that administrators may make use of 'processes to assist decision making (such as computer applications and systems)'. While a decision made with such assistance will be taken to be a decision of the administrator, the administrator is permitted to later substitute their own decision for a decision made with assistance, if the administrator is satisfied that the initial decision is incorrect.
- 294 *Business Names Registration Act 2011* (Cth) s 3.
- 295 *Australian Appaloosa Association Ltd and Australian Securities and Investments Commission* [2019] AATA 2195, [41].
- 296 *Hazeldine and Australian Securities and Investments Commission* [2019] AATA 366.
- 297 *Perth Martial Arts Academy and Australian Securities and Investments Commission* [2018] AATA 3664.

- 298 *Boyce and Australian Securities and Investments Commission* [2015] AATA 768.
- 299 *Smith and Australian Securities and Investments Commission* [2014] AATA 192.
- 300 *Australian Appaloosa Association Ltd and Australian Securities and Investments Commission* [2019] AATA 2195, [42].
- 301 *B & L Whittaker Pty Ltd and Australian Securities and Investments Commission* [2014] AATA 302, [2].
- 302 *Australian Appaloosa Association Ltd and Australian Securities and Investments Commission* [2019] AATA 2195, [42]; *Smith and Australian Securities and Investments Commission* [2014] AATA 192, [3].
- 303 *Perth Martial Arts Academy and Australian Securities and Investments Commission* [2018] AATA 3664, [38].
- 304 *Smith and Australian Securities and Investments Commission* [2014] AATA 192, [13].
- 305 *B & L Whittaker Pty Ltd and Australian Securities and Investments Commission* [2014] AATA 302, [2].
- 306 *B & L Whittaker Pty Ltd and Australian Securities and Investments Commission* [2014] AATA 302, [16].
- 307 *George and Australian Securities and Investments Commission* [2021] AATA 3615.
- 308 *G C Swinburne and F J McFarlane and Australian Securities and Investments Commission* [2014] AATA 602.
- 309 *Stasiw and Australian Securities and Investments Commission* [2015] AATA 328. (Both names were to be used by Perth-based businesses).
- 310 *George and Australian Securities and Investments Commission* [2021] AATA 3615. (Both names were to be used by choirs in the same location).
- 311 *Stasiw and Australian Securities and Investments Commission* [2015] AATA 328, [54].
- 312 See Le Sueur (n 3) 192, citing Lipsky: ‘Street-level bureaucracies have discretion because the nature of the service provision calls for human judgment that cannot be programmed and for which machines cannot substitute’.
- 313 See eg, Chris Reed. ‘How to Make Bad Law: Lessons from Cyberspace’ cited in Le Sueur, (n 3), 192. See also Australian Law Reform Commission, Unnecessary Complexity in Australia’s Financial Services Law – Fact Sheet, <<https://www.alrc.gov.au/wp-content/uploads/2021/02/Complexity-in-Aust-Financial-Services-Laws-Fact-Sheet.pdf>>.
- 314 Justice Duncan Kerr, ‘Foreword’, in Janina Boughey and Katie Miller, (n 6), at v, viii; see ‘Nearly identical’ under the *Commonwealth Business Names Registration Act 2011* (above, **section 15.2**).
- 315 Commonwealth Ombudsman (n 20) 9.
- 316 The machine could also make a wrong decision to decline to make a beneficial decision (a false negative), which would presumably shunt the person into an alternative decision-making process with a human decision-maker. Even if the correct decision is ultimately made by the human decision-maker, the process itself might be seen as involving inequity. Anyone who has experienced both automated passport control and lining-up for manual passport terminals will recognise that the latter involves less favourable treatment even if the ultimate decision (to allow entry) is the same.
- 317 Commonwealth Ombudsman (n 20) 9. Cf ‘Using machine technology to administer Commonwealth child support payments’ above in **chapter 8** (an example where this approach was attempted *without* express legislative authorisation, and found to be legally impermissible).
- 318 See n 12.
- 319 Road Transport Amendment (Mobile Phone Detection) Bill 2019 (NSW), introduced in the legislative Assembly on 24 September 2019 sch 1 item 1.1. See ‘Mobile phone detection cameras’ **chapter 4**.
- 320 Legislative Council Portfolio Committee No. 5 – Legal Affairs, *Road Transport Amendment (Mobile Phone Detection) Bill 2019* (Report, No. 52, November 2019).
- 321 See ‘Road Transport Amendment (Mobile Phone Detection) Bill 2019’, *Parliament of New South Wales* (Web Page) <<https://www.parliament.nsw.gov.au/bills/Pages/bill-details.aspx?pk=3692>>.
- 322 A number of commentators have proposed ‘algorithmic impact assessment’ processes be undertaken similar to environment or privacy impact assessments: see, eg Loi (n 8); Turner Lee, Resnick and Barton (n 8).
- 323 See Raso (n 7), Joel Townsend, ‘Better Decisions? Robodebt and the Failings of Merits Review’, in Janina Boughey and Katie Miller (n 6) 52, 56 (discussing the limits of existing merits review systems to address high volume, technology-assisted decision-making).
- 324 See Lilian Edwards and Michael Veale, ‘Slave to the Algorithm? Why a “Right to an Explanation” is Probably Not the Remedy You Are Looking For’ (2017) 16 *Duke Law & Technology Review* 18, 83-84.
- 325 See, eg, Cobbe et al, ‘Centering the Rule of Law’ (n 7) 15 (‘Given the limitations of existing laws and oversight mechanisms,...as well as the potential impact on vulnerable members of society, we argue for a comprehensive statutory framework to address public sector automation.’); Bateman (n 6) 530 (‘Attaining the efficiency gains promised by public sector automation in a way that minimizes legal risk is best achieved by developing a legislative framework that governs the exercise and review of automated statutory powers in a way which protects the substantive values of public law. Other jurisdictions have made steps in that direction, and there is no reason Australia could not follow suit.’); see also Terry Carney, ‘Robo-debt Illegality: The seven veils of failed guarantees of the rule of law?’ (2019) 44(1) *Alternative Law Journal* 4.

-
- 326 Robin Creyke, (n 64).
- 327 See 'Directive on Automated Decision-Making', *Government of Canada* (Web Page) <<https://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=32592>>.
- 328 This is true also of bodies that may not necessarily bear the title of Ombudsman, but which perform similar and in some cases more specialised roles, including for example Human Rights Commissions or Information and Privacy Commissions.
- 329 Cf Simon Chesterman, *We, the Robots? Regulating Artificial Intelligence and the Limits of the Law* (Cambridge University Press, 2021) 220-222 (suggesting the establishment of 'an AI Ombudsperson').
- 330 Cf Coglianesi and Lehr (n 91) 1190 (suggesting oversight approaches including 'the establishment of a body of neutral and independent statistical experts to provide oversight and review, or more likely a prior rule making process informed by an expert advisory committee or subjected to a peer review process').

