

GOVERNMENTS AS REGULATORS AND CONSUMERS OF ETHICAL AI

Algorithmic Decision-Making (ADM) and Digital Technologies have created complex and multifaceted challenges for governments and the delivery of public services. Governments must operate simultaneously as the regulator and as a consumer, while maintaining public trust. The roll out of ADM systems has provided significant failures in public policy. For example, the Federal Government of Australia and the New South Wales State Government both deploying systems with significant flaws which led to serious outcomes for both welfare recipients and governments just in the last 10 years. Solutions to confront the failure of ADM systems can be addressed more effectively at a procurement and contract level as well as through the accountability and understanding within the public service. In order to ensure the fair use of digital technologies, governments must be willing to scrutinize their own purchase of technology, ideally through the lens of AI ethical frameworks, and address their own governance and procurement policies, thereby fulfilling their responsibility as service providers to citizens.

Gabby Bush* & Jeannie Paterson**



* Gabby Bush is the Program Manager at the Centre for AI and Digital Ethics at the University of Melbourne.

** Jeannie Paterson is a Professor of Law and the Co-Director of the Centre for AI and Digital Ethics at the University of Melbourne.

Increasingly sophisticated digital technologies provide significant opportunities for governments in making consistent, informed and efficient decisions. Thus, big data, algorithms, expert systems and machine learning are increasingly used to make predictions, recommendations and even decisions about access to government services, a process sometimes known as Algorithmic Decision-Making (ADM). While the safety, equity and fairness of these systems and tools are often criticized, governments continue to pursue the use of ADM in public service delivery. In this process, governments have consistently contracted out the development of ADM systems to private companies, while continuing to face significant backlash when these systems fail.¹ Simultaneously, as digital technologies increasingly present the risk of harm to many parts of society, governments are endeavoring to find solutions for protecting citizens when using technology, including through regulatory and legal responses.² Existing in this new technological landscape as regulator and legislator, consumer and contractor, is a unique challenge for governments.

Our contention is that to pursue safer and more ethical uses of digital technologies and AI across society, governments must look inwards and scrutinize their own processes in both contracting for and relying on ADM. ADM as an alternative to human processing and in augmenting human decision-making is not neutral: it has the potential to replicate and amplify existing biases, while also making the basis for decisions harder to interrogate. Significantly, ADM algorithms relied on by governments are built by private sector entities. In using outsourced ADM systems, governments risk letting decisions of the democratic state becoming decisions of the private sector.

In the last five years, Australia has utilized ADM algorithms to recover debt from welfare recipients that have perpetuated errors, including removing large amounts of money from the bank accounts of vulnerable people, and ultimately have been found to be outside the scope of the discretionary decision making power available to the cognate agencies. Governments in the U.S. and in the EU have used ADM systems for other purposes, with similar outcomes.³ There are, however, better options for procurement and use of ethical AI in governments. Tighter

¹ Jon Porter, "UK ditches exam results generated by biased algorithm after student protests," *The Verge*, 17 August 2020. <https://www.theverge.com/2020/8/17/21372045/uk-a-level-results-algorithm-biased-coronavirus-covid-19-pandemic-university-applications>; Rebecca Turner, "Robo-Debt condemned as a 'shameful chapter' in withering assessment by federal court judge," *ABC*, 11 June 2021. <https://www.abc.net.au/news/2021-06-11/robodebt-condemned-by-federal-court-judge-as-shameful-chapter/100207674>

² See eg JM Paterson, S. Chang, M. Cheong, C. Culnane, S. Dreyfus and D. McKay, "The Hidden Harms of Targeted Advertising by Algorithm and Interventions from the Consumer Protection Toolkit." *International Journal on Consumer Law and Practice*, Vol. 9 (2021), p. 1-24.

³ "Proposal for a Regulation of the European Parliament and of the Council (EU) COM (2021) 206 final laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts [2021]," (EU Draft AI Act.) Australian Human Rights Commission, *Human Rights and Technology Final Report* (2021).

regulation and more robust governance systems, teamed with ethical AI standards and scrutiny in AI procurement contracts, offer practical ways for governments to use new digital technologies to support rather than erode democratic values. Once governments have enacted fairer, more effective processes for the contracting and deployment of technology within their own services, these insights should flow through to governments' regulatory strategies for technology in society. Governments gain regulatory credibility for having their own house in order.

“While there are potential benefits in using new technologies, and particularly the insights from data, to government functions, those benefits lie in specific discrete field, perhaps automating clerical task or providing empirical evidence to guide policy.”

Algorithmic Decision-Making and Streamlined Services

Governments are often interested in streamlining processes and optimizing public service delivery. Ongoing cost cutting is a common and necessary process as funds must be divided amongst many portfolios. A body of literature points to the opportunities that AI technologies provide governments for more efficient processes.⁴ There are significant benefits and incentives for using automation – the utilization of available data allows for evidence based decision-making. Automation allows governments to fast track routine tasks that do not require discretion or judgement. Moreover, newer digital or algorithmic technologies, such as those premised on machine learning, computer vision, natural language processing or ‘artificial intelligence’, not only allow processing large amounts of data but allow opportunities to use that data to inform decision making in an adaptive and dynamic way.⁵ In the Australian Government, the Digital Transformation Agency and several statutes encourage the use of technology throughout government including for the automation of decisions and processes.⁶ Digital technologies can be innovative and novel in the functioning of the state, causing net benefit. For example, computer vision has been used to pick up illegal mobile phone use by drivers in New South Wales with the aim of reducing road accidents.⁷

⁴ Rajan Gupta and Saibal Kumar, *Introduction to Algorithmic Government* (Palgrave Macmillan, 2021), p. 5; Michael Veale and Irina Brass, “Administration by Algorithm?,” in *Algorithmic Regulation* (Oxford University Press, 2019), edited by Karen Yeung and Martin Lodge, p. 121 – 142.

⁵ Ben Dor, Lavi M and Cary Coglianese, “Procurement as AI Governance,” *IEEE Transactions on Technology and Society, Technology and Society*, Vol. 2, No. 4 (2021): p. 192.

⁶ Anna Huggins, “Addressing Disconnection: Automated Decision-Making, Administrative Law and Regulatory Reform,” *University of New South Wales Law Journal*, Vol. 44, No. 3 (2021): p. 1059.

⁷ New South Wales Ombudsman, *The New Machinery of Government: Using Maching Technology in Administrative Decision-Making*, Special report under section 31 of the *Ombudsman Act 1974*, 29 November 2021.

While there are potential benefits in using new technologies, and particularly the insights from data, to government functions, those benefits lie in a specific discrete field, perhaps automating clerical task or providing empirical evidence to guide policy. Once the ADM systems are used to make *discretionary* decisions, such as the allocation of welfare, the focus of policing, whether children are at risk or the likelihood of adolescent recidivism, we need to ask whether this kind of use of technology is consistent with the aspirations of good governance.

Flaws and irregularities in the performance of government functions may be amplified by ADM.⁸ Algorithmic decision-making systems can perpetuate a toxic status quo, entrench bias, discriminate against the marginalized, and lack human-based judgment and reason which cannot be fully quantified or written into code.⁹ Rubenstein wrote on this that “AI systems are social artifacts that embed and project human choices, biases and values.”¹⁰ Typically the justification for using these technologies is that using algorithmic processes removes the human bias, creating more efficient, neutral and replaceable decisions. Yet, critics argue that “Individual administrators can, and do, make errors but not at the scale and speed of ADM.”¹¹

The Risks of ADM in Government

Using ADM systems in government services carries considerable risks of harm for citizens and also the government’s social license. The primary risk is to citizens who may be the recipient of poor or biased decision-making at the hands of an algorithm – they may be the recipient of an incorrect debt notice, considered by an algorithm to be at risk of recidivism due to their race or denied a loan due to their previous post-code.¹² The deployment of flawed and biased systems undermines a core function of democratic governments, which is to treat its citizens fairly and respectfully, as well as investing public funds prudently. As seen in the Australian experience, poorly designed and deployed algorithmic systems are ultimately vulnerable to challenge by those institutions that oversee government exercises of power, such as tribunals, courts and ombudsmen. This may mean that government investment in algorithmic processes is ultimately rendered futile, as well as causing considerable harm to the typically already vulnerable individuals affected by the process.

The reasoning for increasing government use of ADM is complex – at best, the algorithms are used in good faith, to conserve public funds and streamline services;

⁸ New South Wales Ombudsman, 29 November 2021.

⁹ Veale and Brass, (2019), p. 122; Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (St. Martin’s Press, 2018); NSW Ombudsman, “The new machinery of Government.”

¹⁰ David S. Rubenstein, “Acquiring Ethical AI,” *Florida Law Review*, Vol. 73, No. 4 (2021): p. 761.

¹¹ Huggins, (2021), p. 1058.

¹² Daniel Ziffer, “Threat of ‘Postcode Discrimination’ as credit scores skewed by where you live,” *ABC*, 7 February 2022. <https://www.abc.net.au/news/2022-02-07/threat-of-postcode-discrimination-in-credit-scores/100723574>

at worst, the algorithms are used by governments looking to restrict access to social services in a manner which removes direct accountability from the public service in their deployment. Virginia Eubanks most famously referred to ADM in her book *Automating Inequality* as ‘systems which punish the poor.’¹³

“Ethical AI for procurement contracts would indeed address the issues of AI for government systems. However, this approach needs to be multifaceted.”

Regulating Digital Technologies

Governments are increasingly under pressure to increase regulation surrounding the use, consumption and creation of emerging digital technologies and to address ongoing issues of misinformation, deep fakes and the erosion of privacy online. The regulatory landscape is complex, with the digital platforms responsible for developing many of the new uses of AI transcending traditional nation-state boundaries. However, significant concerns highlight the need for this regulation, including a focus on uses of AI in ADM. Currently, many frameworks for the ethical use of AI have been written at a national and multinational level. These frameworks are typically not legally binding: in most cases they are at most a form of ‘soft law’ influencing behaviour without mandatory force. Increasingly, however, efforts are being made to formalize laws about ADM and AI.¹⁴

Governments are responsible for these regulatory systems. They are also effectively consumers in their own right. Governments have always held large amounts of data on their citizens and are now using ADM to process that data and deliver public services. In doing so, governments are faced with the challenge of not only trying to understand how to regulate the industry effectively for the safety of its citizens, but also how to operate and procure these systems as a purchaser. In this capacity, governments in some broad ways face challenges usually associated with traditional consumers, namely of considerable information asymmetries vis-à-vis the companies offering the technology. These asymmetries impinge on governments’ ability to contract for optional outcomes in technological tools. The added factor is that governments’ decisions about their consumption of technology impact on the individuals or citizens they are supposed to represent.

¹³ Eubanks, (2018).

¹⁴ Australian Human Rights Commission, (2021).

Private Companies and the State

Governments are usually not equipped to develop their own software, so the development of this software is outsourced to private companies. These companies are driven by commercial objectives, including, significantly, safeguarding their intellectual property and trade secrets.¹⁵

With ADM, programmers need to decide the outcomes, weight decisions and prioritize information. The concern for automating government services is that the decisions once made by politicians or bureaucrats become decisions made by private companies. As Rubenstein writes “the government’s decision to *outsource* AI development or deployment can be highly consequential. Especially if private vendors become de facto deciders.”¹⁶ Where governments have systems in place for transparency and accountability, private companies are required to do no such thing. Companies moreover tend to try to protect their intellectual property and operate with less than full transparency. This is cause for serious concern as private companies have increasing input into the decisions made and enacted by the state. The further concern is that the private companies are operating on a for-profit model.¹⁷ Are decisions made by a private, for-profit entity the best decisions for citizens of a democratic state?

Failures of ADM Systems

There have been well-documented examples of ADM in government services which have resulted in problematic outcomes, and ultimately the withdrawal of the system, including in Australia, England, the EU¹⁸ and the U.S.¹⁹ In 2020 the UK Government, from the Office of Qualifications and Examinations Regulation used a grading algorithm to decide the outcomes of students' A-Levels which had been disrupted by the COVID-19 Pandemic.²⁰ This algorithm used factors to weight the score, including the historical scores of the school each student attended. This meant that high-performing students from low-performing schools were given lower grades, while those who attended private schools were graded on average at a much higher level by the algorithm. This ADM system received significant backlash and with no processes for contesting the outcome the algorithm was pulled and the withdrawal

¹⁵ Dor and Coglianese, (2021), p. 193.

¹⁶ Rubenstein, (2021), p. 768.

¹⁷ Sarah Valentine, “Impoverished Algorithms: Misguided Governments, Flawed Technologies, and Social Control,” *Fordham Urban Law Journal*, Vol. 46, No. 2 (2019): p. 364–427.

¹⁸ Australian Human Rights Commission, (2021).

¹⁹ See generally Eubanks, (2018).

²⁰ Gabby Bush, Henrietta Lyons and Tim Miller, “Data isn’t neutral and neither are decision algorithms,” *Pursuit*, 15 September 2020. <https://pursuit.unimelb.edu.au/articles/data-isn-t-neutral-and-neither-are-decision-algorithms>

of the algorithmic grades completely.²¹

The Federal Government of Australia deployed an algorithm now known as Robo-Debt, to issue debt notices to welfare recipients, from 2016.²² The algorithm calculated the amount of welfare received and averaged this sum against the reported income of that year for each recipient. This algorithm found that thousands of welfare recipients had been overpaid and issued debt notices to recover the overpayments. However, the calculations made by the algorithm were in many cases, incorrect. For example, the algorithm calculated the average of a yearly income – not taking into account fluctuations of income and time spent on Centrelink. Nor did the algorithm account for “working credits” a scheme which allows welfare recipients to work for an allotted amount of hours without losing their welfare.²³ The result was that some 400,000 Centrelink recipients received false debt notices, totaling \$720 million and \$400 million in unlawful demands.²⁴ There was then no avenue for contesting this decision and in some cases, these debt notices were attributed as the cause of suicide in recipients.

This was not the last time that a government in Australia has used an algorithm to recoup debt from welfare recipients. Between 2016 and 2019 the New South Wales Government used a similar algorithm to garnish overpayments from welfare recipients.²⁵ While the algorithm was similar to the robot-debt algorithm, this one went one step further – instead of issuing a debt notice to the recipient in question, the algorithm gained access to the bank account of the recipient and removed the supposed debt directly from their account.

In 2021 the NSW Ombudsman released a report detailing the failure of this policy and the steps that should be taken to avoid a repeat of this issue. Pertinently, the report noted that in Australia there is no clear and available information on the extent of use of ADM in government decision-making. The report detailed a set of standards for the authorization of machine technology at a parliamentary level. These standards included a series of properties:

²¹ Catherine Carroll-Meehan, “A-levels governments u-turn has left universities in the lurch,” *The Conversation*, 20 August 2020. <https://theconversation.com/a-levels-governments-u-turn-has-left-universities-in-the-lurch-144763>

²² Luke Henriques-Gomes, “Robodebt: government to refund 470,000 unlawful Centrelink debts worth \$721m.” *The Guardian*, 29 May 2020. <https://www.theguardian.com/australia-news/2020/may/29/robodebt-government-to-repay-470000-unlawful-centrelink-debts-worth-721m>

²³ NSW Ombudsman, (2021), p. 27.

²⁴ Luke Henriques-Gomes, “Robodebt class action coalition agrees to pay 12bn to settle law suit,” *The Guardian*, 16 November 2020. <https://www.theguardian.com/australia-news/2020/nov/16/robodebt-class-action-coalition-agrees-to-pay-12bn-to-settle-lawsuit>

²⁵ NSW Ombudsman, (2021), p. 11.

- Is it visible?
- Is it avoidable?
- Is it subject to testing?
- Is it explainable?
- Is it accurate?
- Is it subject to audit?
- Is it replicable?
- Is it internally reviewable?
- Is it externally reviewable?
- Is it compensable?
- Is its privacy protected and data secure?

The NSW Ombudsman found that while it was permitted by legislation to issue garnishee orders, Revenue NSW had adopted new standards for garnishing debt, including a ‘minimum protected balance,’ meaning that the garnisheeing of debt was not allowed when the debtor had under a certain amount of funds.²⁶

These failures were ultimately addressed through litigation in court and review by the Ombudsman. Systems that were adopted to streamline processes had to be remedied through the legal system and independent bodies. The result was not more efficient systems for government. In other words, without proper procurement and governance processes, there is a false efficiency. ADM systems do not result in better outcomes for both citizens and government.

Building Accountability

The Australian cases raise a multitude of questions about government use of ADM, and also about the outsourcing the build of technology that impacts directly on core government functions, and, inevitably, vulnerable populations.²⁷ They signal the need for serious consideration of the governance of ADM systems in government, and also the processes by which they are contracted, deployed, scrutinized and monitored by governments.

Procurement Contracts for ADM in Government

As already noted, ADM systems are commonly not developed by governments themselves. Instead they are often, although not always, procured by government contracts to be delivered by private sector entities. Ben Dor and Coglianese in 2021, along with David S. Rubenstein have advocated for an approach that would see

²⁶ NSW Ombudsman, (2021).

²⁷ Valentine, (2019), p. 368.

ethical AI standards for accountability and transparency built into procurement contracts for government AI systems.²⁸ In doing so, the contracts would be reviewed and re-written for each time new technology is procured. The research claims that through this procurement process, ethical standards for ADM systems would be constantly updated through the process of contracting and procurement as new technology and AI systems are developed.²⁹ In principle, this approach would provide immediate solutions for ethical AI in government, however it falls heavily on those in the public service procurement departments to ensure all procurement contracts contain up to date ethical frameworks to ensure the ethical AI requirements are standardized and contracts are delivered adequately.

This approach has merits, as most government services are not developing their own AI systems it makes logical sense that the technology provided through procurement, be subject to ethical requirements in the contract.

Dickenson-Delaporte et al, in their work on self-regulation for the advertising sector noted that self-regulation would require action on behalf of the firm in a global market where the government will not step in.³⁰ The accountability placed on the private sector through procurement contracts, again removes the responsibility of the government and instead requires the private sector to prove their ethical practices but not vice versa. Governments need contracts to ensure they can scrutinize the ethical practices. As Rubenstein aptly wrote “some AI models *cannot* be explained because of their complexity, other AI models *will not* be explained because they don’t have to be.”³¹ The role of the contract is to ensure that the ADM System has to be explained both to and by the government.

Towards Ethical AI Procurement

Ethical AI for procurement contracts would indeed address the issues of AI for government systems. However, this approach needs to be multifaceted. First there must be regulation *requiring* this practice and oversight into these systems. Including ethical clauses in procurement contracts should not be undertaken at the discretion for the contracting agency or department.

Secondly, the procurement contracts should provide the framework for government scrutiny over ADM systems. Rather than the onus being placed on the contractor,

²⁸ Ben Dor et al, (2021), p. 192.

²⁹ Rubenstein, (2021), p. 779; Ben Dor et al, (2021), p. 192.

³⁰ Sonia Dickinson-Delaporte, Kathleen Mortimer, Gayle Kerr, David S Waller and Alice Kendrick, “Power and responsibility: Advertising self-regulation and consumer protection in a digital world,” *Journal of Consumer Affairs*, Vol. 54, No. 2, (2020): p. 675–700.

³¹ Rubenstein, (2021), p. 779.

by the contract; it must instead be agreed that the ethical clauses give government oversight – they are able to scrutinize the design and development of the algorithm, they are able to revisit the system to understand the outcomes that are produced and the government has oversight on any changes made to the algorithm from inception to delivery and maintenance.

Thirdly, in order for this to generate successfully enabling environments, it must be created within the government service to ensure the ethical delivery of ADM. It is not enough to outsource the build of an ADM system to a private company – departments looking to do so must also have computer science and machine learning expertise within their department. In order to properly scrutinize the system, governments must have the personnel to understand that system in its entirety.

Governments must also ensure there is a team able to respond to complaints about the system. It is not enough to simply deploy an algorithm nor is it enough to have one technical expert responsible for all queries and concerns. Government agencies or departments relying on ADM systems should ensure there are systems and *processes* for citizens and affected individuals to contest the operation of those built into the very way in which they use ADM.³² Humans must work in *collaboration* with the ADM system, acknowledging that the machine is not a flawless decision-maker.

Concluding Remarks

ADM systems in government provide a complex set of challenges that require detailed and multifaceted approaches. The notable failures we have highlighted primarily in Australia, require a significant overhaul and cultural shift in governance around ADM systems. In this piece we have focused on the procurement question. Where the decision is made to use an ADM system, any contracting out to a private company of responsibility for the design, construction and deployment of an ADM system must itself be subject to scrutiny and robust governance mechanisms by the government.

As both regulators and consumers, governments exist in a conflicting and unique position. It is not enough to explore regulation for the private sector, governments must also scrutinize their own processes. To fulfill their responsibility to the public, and their role as regulator, governments regulator and legislator, consumer and contractor, must consider the structural need for ethical ADM contracting and fairer ADM systems as a matter of urgency. The fair and robust function of public institutions is essential to a democratic state and one that must not be compromised by a superficial response to the perceived opportunities that ADM provides.

³² See further, Lyons Henrietta, Eduardo Velloso and Tim Miller, “Conceptualising Contestability,” *Proceedings of the ACM on Human-Computer Interaction*, Vol. 5, No. CSCW1 (2021): 1-25. doi:10.1145/3449180