The use of algorithms in the New Zealand public sector

Joy Liddicoat, Colin Gavaghan, Alistair Knott, James Maclaurin and John Zerilli, University of Otago, with a preliminary assessment

The past few years have seen a growing awareness of — and concern about — the role of computer algorithms. Best-selling books such as Cathy O'Neill's *Weapons of Math Destruction* (Crown Publishing Group, New York, 2016) have shed light on some of their more problematic uses and effects, in contexts ranging from finance to criminal justice to employment. Recently, attention has been drawn to their use in New Zealand, with examples in ACC, Immigration NZ and other contexts, drawing concern and criticism from the media and academic commentators (see, for example, Kirsty Johnston "Privacy and profiling fears over secret ACC software" *The New Zealand Herald* (online ed, Auckland, 15 September 2017); Lincoln Tan "Immigration NZ's data profiling 'illegal' critics say" *The New Zealand Herald* (online ed, Auckland, 5 April 2018)).

Aside from these individual cases which have come to media attention, though, not much has been known about how — and how widely — algorithms are used in New Zealand. In October, Internal Affairs and StatsNZ took a first step in answering these questions, with the publication of their review of algorithm use across government agencies. The *Algorithm Assessment Report* (the Report) (<www.data.govt.nz/use-data/analyse-data/government-algorithm-tranparency>) documented 32 algorithms being used for a variety of purposes across 14 agencies. The Report also considered the extent to which agencies have "safeguards and assurance processes" (at 4) in place around algorithm use.

As a first step in forming an accurate picture of the extent and form of algorithm use across New Zealand government, the Report is timely and welcome. To what extent, though, are its findings likely to allay concerns about the use of algorithms by government? Are such concerns well-founded in any event, and if so, what sort of measures could be taken to address them? Funded by the New Zealand Law Foundation, our multi-disciplinary team at Otago University has been looking at these sorts of issues from a range of perspectives. In this article, we explore these questions and considers some implications for practitioners and policy-makers.

WHAT IS AN ALGORITHM?

The Report defines a computer algorithm as "a procedure or formula for solving a problem or carrying out a task" (at 5) It notes that the term can be taken to apply to very simple computerised processes, some of which have been in use for decades (the Report uses the term "business rules" to refer to simple algorithms that "make determinations about individuals or groups without a significant element of discretion"

(at 7)). The Report's focus, however, is primarily on more complex processes, which "can now model complex outcomes" and "use statistical methods and predict likely outcomes" (at 5).

The Report is primarily concerned with what it calls Operational Algorithms; those that "result in, or materially inform, decisions that impact significantly on individuals or groups" (at 7). It notes that algorithms are also used "for policy development and research", but treats those cases as out of scope, on the basis that "they have no direct or significant impact on individuals or groups" (at 7). Though we can see why they were not the main focus of this particular review, we would not want to give the impression that that latter cohort are immune from ethical or other concerns. While it may well be true that these have no *direct* impact at the level of individual decisions, it is less clear that algorithms which inform important policy decisions have no *significant* impacts at group levels. Certainly, some of our discussions with data scientists have suggested that data use for these purposes is worthy of closer consideration, for example, where algorithms are used to guide operational decisions on how to address social inequities where those inequities arise from previous policy decisions.

RISKS AND BENEFITS

The Report notes several advantages to algorithm use, ranging from (at 4):

> … the immediate, such as reducing costs to the taxpayer and speeding up the delivery of services, to the indirect, such as increasing New Zealand's productivity and improving the lives of people by reducing social harm.

Speed and efficiency are not the only advantages that have been claimed for algorithms. The possibility also exists that they could in many cases produce more accurate decisions and predictions, informed by quantities of data vastly greater than a human decision-maker could make use of. As the AI Forum recently noted in their own report (AI Forum NZ *Artificial Intelligence: Shaping a Future New Zealand* (May 2018) at 26):

> We live in a data rich world and human brains are not equipped to analyse today's vast quantities of structured and unstructured data, make connections, identify relationships and patterns across datasets.

But disadvantages and dangers have also been well documented, particularly those related to accuracy, transparency, bias and control.

Concerns about accuracy can arise where the algorithm is assumed to work without that assumption being thoroughly tested. A recent report from the AI Now Institute referred to

a troubling example used in the Washington DC juvenile criminal system (AI Now Institute *Litigating Algorithms: Challenging Government use of Algorithmic Decision Systems*. (September 2018) at 14):

> … the system in the DC case had been in use in the juvenile criminal system since 2004, but had not been challenged until 2018 … When [defence attorneys] dug into the validity behind the system, they found only two studies of its efficacy, neither of which made the case for the system's validity; one was 20 years old and the other was an unreviewed, unpublished Master's thesis. The long-held assumption that the system had been rigorously validated turned out to be untrue, even though many lives were shaped due to its unproven determination of 'risk'.

Regarding transparency, the concern is that an algorithm whose workings are not open to scrutiny is less open to being peer reviewed, its outputs less likely to challenged, and decisions made based upon them will be more difficult to appeal. Difficult need not mean impossible, and a considerable amount can be discerned about the accuracy of certain algorithms with sufficient information about inputs and outputs; seeing 'the working' isn't always necessary to check if the sums add up. However, as we discuss below, this form of scrutiny involves resources and expertise unlikely to available to most people affected by such decisions.

Perhaps the worst case scenario is where an algorithm operates as a 'black box' (a term popularised by Frank Pasquale in *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press, 2015)), its reasoning opaque even to those who designed or who rely on it. This would not have been a risk with previous generations of algorithms — the "very simple computerised processes" with which the Report is not concerned. But more advanced techniques such as machine learning are allowing for the development of "algorithms whose actions are difficult for humans to predict or whose decision-making logic is difficult to explain after the fact" (Brent Mittelstadt, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, Luciano Floridi "The ethics of algorithms: Mapping the debate" (2016) 3(2) Big Data and Society 1 at 3).

With regard to bias, the House of Lords Select Committee on Artificial Intelligence described the problem like this (House of Lords Select Committee on Artificial Intelligence *AI in the UK: ready, willing and able?* (April 2018), at 107):

> These systems are designed to spot patterns, and if the data is unrepresentative, or the patterns reflect historical patterns of prejudice, then the decisions which they make may be unrepresentative or discriminatory as well.

The sort of concerns that have been expressed about discretion are well captured in this quote from Cathy O'Neill (above, at 10):

> … you cannot appeal to a WMD [Weapon of Math Destruction]. That's part of their fearsome power. They do not listen. Nor do they bend. They're deaf not only to charm, threats, and cajoling but also to logic — even when there is good reason to question the data that feeds their conclusions.

Whether that is a bug or a feature of algorithms will depend on a range of considerations. Discretion, after all, is the same mechanism that allows for nepotism and cronyism. Both O'Neill's and other recent books, though, point to examples of algorithmic intransigence that seem like legitimate causes for concern when important decisions are wholly automated or ineffectively supervised.

How well are New Zealand government agencies addressing these risks? The Report paints a mixed picture. With regard to transparency, it notes that, while "[t]hree-quarters of participating agencies provide descriptions of their operational algorithms on their websites" (at 28), what these contain "vary significantly". While some provided "plain English descriptions of the rationale and use of algorithms including examples", others had "largely technical documents that may be difficult for anyone who is not familiar with data processing to understand" (at 28).

With regard to bias, the stocktake found "little consistency across government" in terms of "reviewing and assessing outcomes … to ensure there are no unfair, biased or discriminatory outcomes" (at 29). And regarding fears about algorithmic intransigence, it offers some reassurance that "[h]umans, rather than computers, review and decide on almost all significant decisions made by government agencies" (at 4).

The Report recommends improvements to current practice. To improve transparency, it recommends a two-tier approach, including both "simple summaries" and (at 34):

> … more detailed information about how data is collected and stored, the computer code used in the algorithms, and what role the algorithm plays in the decision-making process for those who are interested in more technical material.

The Report makes a number of other recommendations, including:

- that more attention should be paid to including stakeholder perspectives, and especially those of Māori people, during the development of algorithms;
- that consideration should be given to "establishing a centre of excellence to provide support and advice on best practice processes across government" (at 35); and
- that something along the lines of regulatory impact and privacy assessments could be employed by agencies wishing to employ algorithms, focusing on principles for safe and effective data use and legal obligations (at 33) (we discuss these principles later).

We strongly support all of those recommendations. With regard to another of its recommendations, however, we are somewhat more cautious. Throughout the Report, emphasis is placed on the importance in 'keeping a human in the loop', that is, ensuring that "a real person has exercised human judgement during the process and over the final decision" (at 31). While we recognise that this could be valuable for some reasons and in some contexts, as we explain below, we are somewhat wary of possible false reassurance in this regard.

It is interesting to observe that none of the Report's recommendations relate to creating new legal protections, safeguards or rights, nor to strengthening those that already exist. In fact, the Report barely refers to law at all; it is concerned with operational, procedural and practical matters. For our purposes, though, we are very interested in legal issues to which the use of operational algorithms in government may give rise. This is not at all to suggest that use of

algorithms outside of government raises no legal questions, merely to confine the scope of the present article to the scope of the stocktake.

In the next section, we consider some legal implications of algorithm use, including for decisions made or significantly informed by them, before considering some possible reforms.

LEGAL IMPLICATIONS OF ALGORITHM USE

At present, New Zealand has no primary legislation specifically dealing with algorithmic decision-making, nor any secondary legislation (such as codes, rules or regulations). This in itself need not be a cause for concern. Many new technologies have developed in New Zealand without the need for root and branch legal reform. Instead, for the most part, regulation of new technologies has occurred on an "as needed" basis, with specific laws in some areas (such as telecommunications and assisted reproductive technologies) and none in others (for example, despite its ubiquity, the statute book contains no definition of "the Internet"). Before we leap to the conclusion that new algorithm-specific rules or regulations are required, it makes sense to take stock of the adequacy of existing law for responding to the challenges presented by more advanced algorithms.

Operational algorithms that affect entitlements, such as some of those in the Report, may be subject to appeal or review depending on the regulatory context within which they are made. For example, decisions made under the Accident Compensation Act will be liable to the same processes for challenge and dispute resolution as those made by a human agent. Decisions relating to benefit entitlement or other forms of social security assistance may be subject to the Social Security Appeals process and automated tax assessments may be reviewed under the Inland Review or Tax Administration Acts.

Decisions that are made or assisted by algorithms may be subject to judicial review. Grounds may include, for example irrationality (taking into account irrelevant factors or failure to take into account relevant ones), bias or predetermination (for example, where the algorithmically generated decision has been accepted without clear evidence of human oversight) or whether the use of AI appears to have unduly fettered the decision-maker's discretion. Conversely, grounds for review may arise if a client is given a favourable predictive assessment which an official does not agree with and does not follow.

The capacity to challenge, appeal or review a decision relies significantly on being able to access information about how that decision was arrived at. In the absence of a specific statutory regime for review, the Official Information Act 1982 is likely to be relevant. Section 23 (1) provides that anyone who has been subject to a decision or recommendation by a department or Minister of the Crown is entitled to (inter alia) *the reasons* for that decision or recommendation (see also analogous provision under s 22 of the Local Government Official Information and Meetings Act 1987).

This raises the question of what, in the context of an algorithmic decision, would qualify as a 'reason'. In *Re Vixen Digital Limited* [2003] NZAR 418, the High Court has held that (at [43]):

> Where the legislature has specified that reasons must be given I should think those reasons must be sufficient to enable any body with a power of review to understand the process of thought whereby a conclusion was reached.

Equally the reasons must allow those with vested interests, like those of the appellant, to so understand the basis for decisions as to be better informed in predicting that which is or is not within the law. Further, in this case the public has a general interest in knowing and comprehending the standards that the Board sees as important.

In the context of algorithmic decisions, providing an explanation that is intelligible to affected persons and the general public could prove especially problematic. As Dr Janet Bastiman explained in her evidence to United Kingdom Parliament Science and Technology Committee (2017):

> The resulting systems can be explained mathematically, however the inputs for such systems are abstracted from the raw data to an extent where the numbers are practically meaningless to any outside observer.

What sorts of "reasons" can we expect from an intelligent machine? Deep learning involves multiple hidden layers of processing that are fiendishly intricate and virtually impossible to unsnarl.

An explanation meaningless to the person who requested it is unlikely to satisfy the requirements of the Official Information Act 1982 (OIA). The Act requires reasons in a written statement, including any findings on material issues of fact and (subject to a few exceptions) reference to the information on which the findings were based. Grounds for withholding such information are quite limited. While the question has not yet been judicially considered, it is difficult to see why s 23 of the Act should not apply to decisions made with the assistance of algorithmic tools, particularly in light of the Report indicating just how extensive the use of such tools in the public sector has been. Of course the OIA only applies to decisions in the public sector, so citizens affected by decision-making in the private sector are without a comparable right, whether the decisions are made by humans or by algorithms.

In the private sphere, at least, New Zealand's laws lag considerably behind those of other jurisdictions such as the European Union, whose new data protection laws arguably provide a more general right to explanations for fully automated decisions. Considerable attention is being paid to the challenge of rendering algorithmic systems transparent or explainable, and a number of suggested solutions have been canvassed. We briefly discuss some of these in the final section.

Agencies relying on algorithms to make or inform decisions must also be alert to their obligations around discrimination. The Human Rights Act 1993, for example, prohibits discrimination on a variety of grounds when offering services to the public, including in the areas of access to public goods and services, employment, housing, education and transport. The definition of discrimination includes both direct and indirect discrimination: an act or omission may be considered discriminatory whether the intention or effect of the act or omission results in disadvantage or other harm (Human Rights Act 1993 and see Human Rights Commission *Privacy, Data and Technology: Human Rights Challenges in the Digital Age* (Issues Paper 2018)).

Care should be taken by agencies using algorithms to assist decision-making to ensure that human rights violations do not arise unintentionally and unexpectedly from their use. This may involve more than simply excluding certain protected characteristics from the data. While it would be easy enough to programme an algorithm that does not include

data about race or gender, difficulties arise from the fact that other data may become a proxy for race or social status (such as an address or postcode), meaning that "unwanted discrimination can sneak back in" (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(1) Duke Law & Technology Review 18 at 29).

Particular difficulties may arise where characteristics such as race, sex or age statistically correlate with outcome variables that are relevant factors to take into account in a particular decision, for example, likelihood of victimisation, of having learning difficulties, of being convicted of property theft or failing to achieve seniority in certain jobs (at 29). Such correlations may be statistically accurate but may not be the desired outcome from a decision on access to an entitlement to a service. In these circumstances, the use of results without context or deeper critique may result in discrimination.

Another approach may be to specifically include these characteristics and undertake a human rights analysis of the results to better assess possible direct or indirect discrimination. The best approach is not yet clear and this is an area where more work is needed.

Where a decision on service entitlement is supported by an algorithm which has used the personal information of a specific individual, the Privacy Act 1993 will also apply. In these cases, the service provider must ensure that the information used has been collected lawfully and for a clear purpose (Privacy Act 1993, Privacy Principles 1–4). Before using it, they must be satisfied that the information is accurate, up to date and not misleading (Privacy Principle 8) and is being used for a purpose for which it is collected, unless a lawful exception applies (Privacy Principle 10). In these cases the customer will be entitled to ask for access to their personal information and, if they consider the personal information relied on is not correct, to seek correction (Privacy Principles 6 and 7). More specific personal information privacy codes may apply if the service relates to credit reporting or telecommunications (see Credit Information Privacy Code 2004 and Telecommunications Information Privacy Code 2003).

REGULATORY APPROACH — ARE CHANGES NEEDED?

In 2018 the House of Lords concluded that "[b]lanket AI-specific regulation, at this stage, would be inappropriate" (above at [16] and [386]). In New Zealand, regulators are taking a cautious approach, with each appearing to consider the specific issues within their various statutory mandates. For example, in 2018, the Government Chief Data Steward and the Privacy Commissioner issued six principles for effective data use (Government Chief Data Steward and Privacy Commissioner "Principles for safe and effective use of data and analytics" (2018)). The principles are intended to assist agencies "in data analytic activities, including algorithmic decision-making".

Other initatives are also being developed, for example, the Ministry of Social Development has developed a Privacy, Human Rights and Ethics (PHRaE) Framework (<www.msd.govt.nz/about-msd-and-our-work/work-programmes/initiatives/phrae/index.html>). This approach, along with that of the Financial Markets Authority, demonstrates that agencies in New Zealand are taking a pragmatic but principled approach. However, regulators do need to stay in touch with developments, particularly those overseas.

ALGORITHMIC TRANSPARENCY

The Council of Europe Data Protection Convention 108 and the General Data Protection Regulation (GDPR) contain specific provisions for algorithmic transparency. For example, art 13 of the GDPR imposes transparency obligations for automated decision-making, including when used for profiling. Convention 108 also provides for the right to request and obtain the reasoning underlying data processing. The United Kingdom House of Commons Science and Technology Committee recently noted that the right to an explanation is a key part of achieving accountability for government and recommended openness by default in certain circumstances, namely, explanations of decisions should be published when algorithms affect the rights or liberties of individuals. Several submissions on the Privacy Bill in New Zealand have also recommended new requirements for algorithmic transparency.

The Report does not address the question of whether specific legal provision should be made for algorithmic transparency, but it does recommend that agencies give consideration to making available detailed information about the algorithms they use, including the computer code on which they operate (at 34). Publishing technical details of the algorithm is potentially valuable; the potential for peer review by the wider community of experts is appealing from the perspective of ensuring accuracy. But it will not contribute much to the objective of providing a clear explanation of why a decision was made, and a basis for deciding whether it should be appealed or challenged. As the Report notes, clear explanations that are accessible to lay readers is also important.

RIGHT TO ERASE PERSONAL INFORMATION

Article 17 of the GDPR introduces a new right to erasure of personal information. Calls for this new right emerged in response to the concerns about weakening of the protections of obligations to delete information, protections that were designed to ensure that personal information is not kept any longer than is necessary. The use of personal information for algorithms and artificial intelligence technologies inevitably creates tension with these obligations as agencies will want to keep data for as long as possible. Article 17, which embodies the new right to erasure, enables individuals to assert the right to have data removed. However, it is unclear how the right to erasure and the closely related right of data portability would apply in the context of agencies using personal information for algorithmic service development or where an individual has an on-going relationship with the agency.

RIGHT OF REVIEW

The Courts Matters Bill contains what appears to be the first New Zealand statute based automated electronic decision-making system, including a right of review. New sections 86DA to 86DD provide for the Chief Executive to authorise an automated electronic decision-making system for setting fine payment arrangements, including a greater time for payment or payment by instalments. Where such a system is approved, the Chief Excutive must, under s 86A(4), also approve procedures for operating the system which must also include procedures for:

(a) setting the criteria for variations;

(b) identifying the information that will be sought from the individual;

(c) notifying the individual or their representative of the right to seek variation of the arrangement; and

(d) notifying of the right to review by a person of any automated decision.

New section 86DC provides that the Chief Executive may only approve an automated system if satisfied that "*each system has the capacity to do any actions required with reasonable reliability*" (emphasis added) and "there is a process available under which a person affected by an action performed by an electronic system can have that action reviewed by a person authorised by the chief executive to review those actions, without undue delay."

The Bill provides a useful model for automated decision-making, but its specificity does raise questions about the absence of such a provision in other legislative contexts where decisions are already being made. While the Report notes that the algorithms are only being used in cases where there are benefits, we note that any complaints about those decisions may give rise to questions about the legal basis for their use.

A 'HUMAN IN THE LOOP'

One message that emerges fairly strongly from the Algorithm Report is about the reassurance offered by retaining human oversight or decisional autonomy, at least over important decisions. While for example the new ACC system for approving claims will automatically approve all simple claims, "complex or sensitive claims will be reviewed by an ACC staff member" (at 5). Perhaps more importantly, only approvals will be automated; any claims that are declined must be processed by a human operator.

The Report's position on this is one likely to resonate with many people (at 31):

… where algorithms are material to decisions which affect people's lives in significant ways, it is reasonable to expect that a real person has exercised human judgement during the process and over the final decision.

This position also resonates with art 22 of the GDPR, which provides that a person:

… shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.

A degree of scepticism has, however, been expressed regarding to the extent that such a requirement offers genuine reassurance, at least in some situations. For some sorts of decisions, it may simply be that a well-designed algorithm would make more accurate decisions than any human, and inserting a human back into the system risks diluting that accuracy.

In other situations, there may be value in human oversight. Some sorts of 'errors', it has been suggested, may be difficult to predict or explain in a manner comprehensible to an algorithm, and may be more easily detected by humans – those which involve initiative in the face of unforeseen outcomes, perhaps, or an understanding of social and cultural factors.

Even where there is value in having human oversight, though, questions have arisen about how well humans will be able to discharge such a role. Human factors research has suggested that a range of attentional and attitudinal obstacles may stand in the way of effective human supervision of machine systems. As the accuracy of such systems increases, human ability to detect occasional errors deteriorates; "the operator assumes that the system is reliable and therefore failure detection deteriorates" (Kayvan Pazouki, Neil Forbes, Rosemary A Norman and Michael D Woodward "Investigation on the impact of human-automation interaction in maritime operations" (2018) 153 Ocean Engineering 297 at 299).

This is likely to be particularly pronounced in contexts such as driverless cars that require an inert but alert human supervisor behind the wheel. But in less urgent contexts too, more attention will need to be given to the question of how to ensure that any 'human in the loop' offers more than token reassurance.

THE LIMITATIONS OF RIGHTS-BASED MODELS

Rights to reasons, to privacy and to be free from discrimination undoubtedly have value in this context, and part of the challenge will be to ensure that they can be meaningfully applied in the context of algorithmic decision-making. Doubts have, however, been expressed as to the general adequacy of rights-based models for promoting accountability in algorithm use.

Some of these concerns relate to the capacity of ordinary people to assess algorithms for the sorts of factors we have identified, for example, accuracy and bias. As Edwards and Veale put it, "[i]ndividuals are mostly too time-poor, resource-poor, and lacking in the necessary expertise to meaningfully make use of these individual rights" ("Slave to the Algorithm", at 67). Virginia Eubanks has pointed out that in many cases we will be oblivious to the fact that we are subject to algorithmic decisions: "[m]any don't know that they are being targeted or don't have the energy or expertise to push back when they are" (Virginia Eubanks *Automating Inequality* (St Martin's Press, 2017) at 6).

One of our main concerns is that, even when individual rights work as intended, they are ultimately *individual* rights. Even if they grant us access to meaningful information about how decisions were made in our own cases, they will not reveal a 'wide angle' view necessary to evaluate whether these systems are exacerbating unfairness and inequality between sections of the population.

Such limitations have led an increasing number of commentators to look to supplement 'bottom up' rights-based models with some form of 'top-down' oversight agency, with the capacity to scrutinise the use of such algorithms in a broader and more specialised manner. The devil, of course, lies in the detail, and the form that such an agency should take — as well as details as to its remit and regulatory powers — are the subject of ongoing discussion. It seems, though, that broad agreement or consensus is coalescing around the perceived utility of some such agency.

CONCLUSION

There is a danger, when lawyers write about new technologies, to overstate the risks they present, and by implication, to overstate the merits of the status quo. We should be clear, then, that our aim here is neither to exaggerate the dangers of algorithmic decision tools, nor to adopt too rose-tinted a

in terms of s 52 of the Land Transfer Act 2017, let alone an order that purported to alter the title directly or permitted the plaintiff to enter property of which the defendant was still registered proprietor.

Although the concept of comity has an amorphous quality in private international law, it recognises the legitimate privilege that a sovereign nation enjoys over certain matters within its territorial sphere (*Hilton v Guyot* 159 US 113 at 163–4 (1895)). While it may sometimes be sufficient to defer on choice of law only (in other words, to assume jurisdiction but apply the law of the place where the property is situate), in other cases the courts properly recognise that a limitation on their subject-matter jurisdiction is also required (see for example *Ludgater Holdings Ltd v Gerling Australia Co Pty Ltd* [2010] NZSC 49, [2010] 3 NZLR 713). And where the plaintiff seeks a determination of legal title, the principle of effectiveness is directly engaged. It would create real difficulties, for example, if the Registrar-General of Land was required to determine whether to alter the register to give effect to a judgment from Ireland (or the United States or Venezuela) that purported to determine legal title to New Zealand land.

THE RESULT IN *FOSTER V CHRISTIE*

At first glance, the result in *Foster v Christie* seems unfortunate. The Associate Judge found that Ireland would have been the most appropriate forum for the trial of all the issues. Proceedings were already underway in that jurisdiction, and the consequence of the judge's decision was to split the dispute into proceedings on opposite sides of the world. The Irish court would presumably have the power to set aside a contract for the sale and purchase of New Zealand land for undue influence, even if it was governed by New Zealand law (compare *Attorney-General for England and Wales v R* [2002] 2 NZLR 91 (CA)).

Moreover, whatever the risk of impropriety in other cases there was little apparent risk of any actual conflict, in circumstances where the New Zealand court acknowledged merit in the case being heard in Ireland. Nor was there any evidence about whether the Irish court would regard *itself* as lacking jurisdiction to determine the undue influence claim.

Yet the Judge had good reason to conclude that the *Moçambique* rule was engaged, and that it compelled the court to retain jurisdiction. The crux of the plaintiff's case was that the severance of the joint tenancy was ineffective to affect her and her mother's title, because it was vitiated by undue influence. She sought to be recognised as the sole owner of the land. Expressed in terms of the *Griggs* case above, the plaintiff was not saying "I recognise my sister as the owner of this land, but you should compel her to perform her personal obligations towards me": the plaintiff asserted that she was entitled to be recognised as the owner of the land once the severance was reversed.

If the positions were reversed, and the plaintiff wanted to bring proceedings in Ireland, could she have formulated her claim to avoid the application of the *Moçambique* rule and bring it within the scope of the exception? In *Lord Cranstown v Johnston* (1796) 3 Ves Sen 170, a creditor used the processes of the West Indies courts to obtain ownership of the plaintiff's estate at a substantial undervalue. This amounted to equitable fraud, so Arden MR ordered the creditor to reconvey the land to the plaintiff.

The court will only exercise jurisdiction where there is privity of obligation between the plaintiff and the defendant (*Deschamps v Miller* [1908] 1 Ch 856). If the plaintiff could formulate a direct restitutionary claim against her sister that would be recognised by the Irish courts, there is a good argument that would be sufficient to avoid the application of the *Moçambique* rule. But the general rule is that a plaintiff is entitled to formulate her claim how she wishes, and in the form Ms Foster had formulated her claim in New Zealand, the *Moçambique* rule applies.   ❐

---

view of their human counterparts. In an increasingly data-saturated environment, it is entirely possible that algorithms will often out-perform humans in a range of tasks, including some of those covered in the most recent Report. We also entirely accept that human decision-makers far from perfect. As we have noted elsewhere (John Zerilli, Alistair Knott, James Maclaurin and Colin Gavaghan "Transparency in Algorithmic and Human Decision-Making: Is There a Double Standard?" *Philosophy & Technology* (published first online 5 September 2018)):

> It is true that human agents are able to furnish reasons for their decisions, but this is not the same as illuminating the cognitive processes leading to a conclusion. The cognitive processes underlying human choices, especially in areas in which a crucial element of intuition, personal impression, and unarticulated hunches are driving much of the deliberation, are in fact far from transparent.

Our aim in this paper — and in our project more generally — is to identify ways in which the benefits of algorithms can be optimised, and their risks and disadvantages minimised.

Towards this end, the Report is timely and a welcome first step in forming an accurate picture of the extent and form of algorithm use across New Zealand government. While the Report's findings are likely to allay some concerns about the use of algorithms by government, other concerns remain well-founded. We welcome the Report's recommendations for measures to address these concerns. However, more work is needed to provide legal clarity and to consider whether new rights, remedies or regulators are needed. The AI and Law Project will continue to explore these issues and publish research findings.   ❐