

The death of law? Computationally personalised norms and the rule of law

Timothy, Endicott; Yeung, Karen

DOI:

[10.3138/utlj-2021-0011](https://doi.org/10.3138/utlj-2021-0011)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Timothy, E & Yeung, K 2021, 'The death of law? Computationally personalised norms and the rule of law', *University of Toronto Law Journal*. <https://doi.org/10.3138/utlj-2021-0011>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This is an accepted author manuscript of the following article published in University of Toronto Law Journal. The death of law? Computationally personalized norms and the rule of law
Timothy Endicott, Karen Yeung <https://doi.org/10.3138/utlj-2021-0011>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

The Death of Law? Computationally Personalised Norms and the Rule of Law

Timothy Endicott* and Karen Yeung†

Abstract:

The emergent power of big data analytics makes it possible to replace impersonal general legal rules with personalised, particular norms. We consider arguments that such a move would be generally beneficial, replacing crude, general laws with more efficiently targeted ways of meeting public policy goals and satisfying personal preferences. Those proposals pose a radical, new challenge to the rule of law. Data-driven legal personalisation offers some benefits that are worth pursuing, but we argue that the benefits can only legitimately be pursued where doing so is consistent with the agency that the state ought to accord to individuals, and with the agency that the state ought to accord to itself. These two principles –the principle of private agency and the principle of public agency– are prerequisites for the rule of law. Each is incompatible with unrestrained computational personalisation of law.

Keywords: machine learning and law, personalisation, algorithmic regulation, rule of law, agency, responsibility

Outline:

1. Introduction
2. The rule of law contrasted with technological management
3. The principles of private agency and public agency
4. Computationally personalised micro-directives
 - a. Why might personalisation be desirable?
 - b. Why might personalisation be undesirable?
5. Adjudicative micro-directives
6. Private agency and public agency: what the two principles require
 - a. The principle of private agency
 - b. The principle of public agency
7. The asymmetry between the two principles, and the connection
8. The value of uniformity
9. Conclusion

* Vinerian Professor of English Law, University of Oxford. Email: timothy.endicott@law.ox.ac.uk.

† Interdisciplinary Professorial Fellow in Law, Ethics and Informatics, Birmingham Law School. Email: k.yeung@bham.ac.uk.

The authors wish to thank the participants at an Oxford Internet Institute Colloquium titled 'What's wrong with using machine learning to personalise law?' in June 2018, and participants in a workshop on 'Responsibility and Control' organised by Maximilian Kiener in June 2021 under the Roots of Responsibility ERC Project. Special thanks to Emma Ahmed-Rengers, David Enoch, Joseph Raz, and Claire Field for comments. Karen Yeung acknowledges financial support for her research from the VW Stiftung, the Australian Research Council, and the University of Birmingham.

1. Introduction

A new prospect for the ordering of political communities in the 21st century is coming into view. It is emerging through the conjunction of massive data storage capacity, population-wide data harvesting from the continuous tracking of individual behaviour by networked digital devices, real-time data processing, and the development and deployment of powerful machine learning algorithms. This conjunction is enabling governance by various forms of ‘algorithmic regulation’,¹ a form of governance that relies on big data analytics: the use of information-processing hardware combined with software capable of sifting, sorting and interrogating vast quantities of data very quickly to search for data patterns, distilling those patterns into predictive analytics and applying the analytics to new data. The resulting techniques apply machine learning algorithms to data sets that may be vast and unstructured. The algorithms yield correlations between new data and patterns identified in the data flows to yield a particular, highly data-intensive form of knowledge.²

One of the most popular contemporary applications of big data analytics involves mining the digital traces of an individual’s online interactions. The resulting behavioural profile can support computational inferences as to a person’s dispositions, preferences and circumstances, to generate predictions about their future behaviour.³ Such predictions enable digital service providers, for example, to offer services to users that seek to predict each user’s preferences, automatically personalising their shopping recommendations or the results of web search queries.⁴ Automated, data-driven personalisation techniques of this kind could be put to use for many purposes that states currently seek to achieve through law.

The putative promise of this technological conjunction is that it could take social ordering beyond the crude, impersonal techniques of law, with its clumsy dependence on general rules. Several legal scholars have observed that big data analysis could be used to satisfy individuals’ preferences and to achieve policy objectives more directly, precisely and responsively through computationally personalised legal rights and duties.⁵ These

¹ K. Yeung, ‘Algorithmic Regulation: A Critical Interrogation’ (2018) 12 Reg & Gov 502. Yeung defines algorithmic regulation as ‘decision-making systems that regulate a domain of activity in order to manage risk or alter behaviour through continual computational generation of knowledge from data emitted and directly collected (in real time on a continuous basis) from numerous dynamic components pertaining to the regulated environment in order to identify, and if necessary, automatically refine (or prompt refinement of) the system’s operations in order to attain a pre-specified goal’: at p 507. See generally K. Yeung and M. Lodge eds., *Algorithmic Regulation* (Oxford: Oxford University Press, 2019). The automated, data-driven personalisation techniques that we discuss in this article are one distinctive form of algorithmic regulation, entailing the application of machine learning techniques to large behavioural data sets automatically to generate legally binding personalised norms applicable to individual persons or to determine automatically the legal rights and duties of a person for the purposes of ordering human affairs.

² J.E. Cohen, *Configuring the Networked Self* (New Haven: Yale University Press, 2012), 1919.

³ On the use of data-driven behavioural profiling, see M. Hildebrandt and S. Gutwirth (eds) *Profiling the European Citizen – Cross-disciplinary Perspectives* (Dordrecht: Springer, 2008).

⁴ K. Yeung, ‘“Hypernudge”: Big Data as a mode of regulation by design’ [2016] *Information, Communication & Society* 1-19.

⁵ B. Alarie, A. Niblett and A.H. Yoon, ‘Law in the Future’ (2016) 66 *University of Toronto Law Journal* 423-28; B. Alarie, ‘The Path of the Law: Towards Legal Singularity’ (2016) 66 *University of Toronto Law Journal* 443-55; A. J. Casey and A. Niblett, ‘Self-Driving Laws’ (2016) 66 *University of Toronto Law Journal* 429-442; P. Hacker (2017) ‘Personalizing EU Private Law: From Disclosures to Nudges and

automated techniques would also have incidental benefits. First, they would enable persons subject to the law to know their legal position with unprecedented clarity. We could escape at last from the costly and uncertain processes of learning (or getting advice as to) the general rules, and of predicting how they will be applied by human officials and institutions in a particular case. Computationally personalised norms could be communicated to each individual in real time: we could receive notifications through an app on a smartphone (we will call it a 'legal satnav'). Second, the particularised nature of the indications from the machine would do away with much of the costly and often unwieldy apparatus currently required for legal dispute resolution: a scheme of particularised directives would eliminate one major element in the stock in trade of adjudication –disputes over the application of general rules to particular cases.

The possibilities for computational personalisation are often presented as techniques for improving law, tailoring norms more sensitively to fit the circumstances and characteristics of each individual.⁶ We argue that these possibilities actually offer an *alternative* to law. That is because, as Ariel Porat and Lior Jacob Strahilevitz observe in an article advocating widespread computational personalisation, 'Law is impersonal'.⁷ Law orders social life through general rules. With automated, data-driven personalisation, society would be ordered by what Anthony J. Casey and Anthony Niblett call 'micro-directives': directives applicable to a particular person in a particular set of circumstances. Casey and Niblett predict that micro-directives 'will become the dominant form of law, culminating in the death of rules and standards'.⁸ But the death of rules and standards would not result in a newly dominant form of law. It would be the death of law. Computational personalisation would replace the rule of law with the technological management of private and public affairs⁹. Management is the particularised control and direction of affairs. If big data analytics became a standard technique of social ordering, society would no longer be ruled by law. Society could, for the first time, be ruled by technological management.

Section 2 explains the contrast between the rule of law and technological management. Section 3 offers an initial statement of the principles of private and public agency, and of our core argument: that computerised personalisation can only legitimately be used in governance where **its** use is compatible with the two principles and, therefore, with the rule of law. In sections 4 and 5 we address two examples of computationally personalised techniques of governance (computationally personalised micro-directives and adjudicative micro-directives). Section 6 unpacks the principles of private and public agency, and explains why they rule out indiscriminate use of the techniques outlined in sections 4 and 5. Section 7 discusses relations between the two principles, and section 8 explains the usefulness of uniform rules (an antithesis of personalisation) in regulation. The conclusion, section 9, points out collateral objections that might be made (and have been made) against replacing

Mandates', Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2914393; A. Porat and L.J. Strahilevitz, 'Personalizing Default Rules and Disclosure with Big Data' (2014) 112 *Michigan Law Review* 1417; Ben-Shahar and Porat, 'Personalizing Negligence Law' (2016) 91 *New York University Law Review* 627-8; C. Busch 'Implementing Personalized Law: Personalized Disclosures in Consumer Law and Data Privacy Law' (2019) 86 *University of Chicago Law Review*, Article 9. Available at: <https://chicagounbound.uchicago.edu/uclrev/vol86/iss2/9>

⁶ Alarie, *ibid*; Casey & Niblett, *ibid*; Porat and Strahilevitz, *ibid*.

⁷ Porat and Strahilevitz, n 5 above.

⁸ *ibid* 1404.

⁹ R.Brownsword, 'Technological management and the rule of law' (2016) 8 *Law Innov Technol* 100–140.

the rule of law with algorithmic regulation, and explains why our objection is more basic than those objections.

2. The rule of law contrasted with technological management

The rule of law involves governance by general rules, and by the regulated exercise of general legal powers by legal persons. The classic antitheses of the rule of law are dictatorship (in which the community is not ruled by law, but by the say-so of the dictator), and anarchy (in which the community is not ruled at all). Now we can envision a third antithesis of the rule of law: governance by particular norms generated by big data analytics.

The use of big data analytics in some aspects of governance would be compatible with the rule of law, and might even support it. And no one should turn the rule of law into a fetish. Its value has to be understood in light of a realistic awareness of the drawbacks of its central technique: governance by general rules. Those drawbacks are, we will argue, the inevitable result of any attempt by a state to accomplish two interrelated essentials: responsible government, and respect for the ability of individuals to take responsibility for their own lives. So although the rule of law is inevitably imperfect, even some of its drawbacks reflect its value. And we will argue that government through computational personalisation can only legitimately be pursued within the framework of the rule of law.

The value of governance by general rules lies in its opposition to arbitrariness, which is the antithesis of the rule of law. Government is arbitrary if it gives effect to the mere will of the rulers, without responding to the reasons on which they ought to act. It is the antithesis of responsible government, which is government that responds to those reasons. Applying general rules is a way in which public agencies and officials can act responsibly: they can point to the rule as a reason that distinguishes their action from their arbitrary say-so.

Yet, ironically, generality itself implicates forms of arbitrariness. Arbitrariness is unresponsiveness to reasons.¹⁰ The first form of arbitrariness in legal ordering arises purely from the generality of rules: they are over- and under-inclusive, so that their application is to some extent unresponsive to the reason for the rule.¹¹ They draw lines that could just as well have been drawn elsewhere, with the result that similar cases are treated very differently. And a form of arbitrariness arises from the need for general rules to be applied by human officials. To the extent that they do so capriciously (a risk that can be reduced but not eliminated), the law's effect is unresponsive to reason.

So there is a tension in the ideal of the rule of law. There are various ways in which a legal system can and should deal with the tension: the essential discretionary powers of executive agencies ought to be structured and constrained; rights to due process (including rights to contestation) must be respected in executive and judicial decision making; judicial appointment, training, and culture ought to support fidelity to the general rules; and the judges ought to hold to that fidelity where it most matters. These are all safeguards that the rule of law adopts against arbitrariness. The safeguards mitigate the arbitrariness, and support the fundamental aspect of governance by law that is *not* arbitrary: the agencies and

¹⁰ As Joseph Raz has put it, 'What is arbitrary action generally? It is action indifferent to the reasons for or against taking it. Arbitrary government is the use of power that is indifferent to the proper reasons for which power should be used.' J. Raz, 'The Law's Own Virtue' (2019) 39 *Oxford Journal of Legal Studies* 1–15, 5.

¹¹ See T. Endicott, 'The Generality of Law' in L.D. Almeida, A. Dolcetti, J. Edwards (eds), *Reading The Concept of Law* (Oxford: Hart Publishing, 2013) 15-36 at 26.

officials that make and that give effect to laws are responsible for doing so on the basis of reasons.

Computational personalisation promises to remove the need for safeguards against arbitrariness. It would dissolve the tension within the rule of law by abandoning the rule of law. It is an anti-law project. Recent proposals advocating governance by computationally personalised norms can be understood as a form of ‘technological solutionism’.¹² They pose a new theoretical challenge to the rule of law just as radical as earlier challenges from anarchism, authoritarianism, and Marxism. In this article, we begin to address that challenge. We will argue that the positive potential of computationally personalised norms to advance just and intelligent governance in contemporary societies depends on identifying specific aspects of governance that do not need to be conducted through general rules, or through the exercise by human beings of legal powers, so that those aspects of governance could justly (and more effectively) be accomplished through computational personalisation, instead.

The positive potential of computationally personalised norms is, therefore, strictly limited. Regardless of how our networked digital technologies evolve, and notwithstanding the inescapable elements of arbitrariness in legal ordering, there will still be value in the rule of law. Central aspects of the life of a political community will still need to be governed by clear, open, stable, prospective, general rules conferring rights and powers and imposing obligations and liabilities, promulgated by legislative institutions, and applied by tribunals comprising human judges exercising judicial authority in the process of adjudication.¹³

The rule of law supports and fosters the responsibility of persons in a political community for their own lives, and it fosters and supports responsible government. If persons are to have responsibility for their own lives, the state must treat each as an agent. Law does so both when it imposes duties, and when it confers powers. Computationally personalised governance (by which we mean the imposition by the state of binding particular norms generated by big data analytics) treats the person, instead, as an object (an object of assessment, or an object of gratification...).

And if the political community is to achieve responsible government, the state itself needs institutionalised ways of acting. It, too, needs agency. The state and public agencies within the state must demonstrably engage in reasoned action on behalf of the community. Then it becomes possible for them to be responsible and accountable to the community of individuals whom they govern, and on whose behalf they purport to act, for what they do and for what they fail to do.

If, instead, governance is conducted by particular directives resulting from pattern recognition based on machine learning, then the action triggered by those computationally generated outcomes entails a very different and much more limited kind of agency. The state’s only form of agency will lie in the sovereign act of requiring the individual to comply with whatever particular directive issues from the machine learning process. The state’s responsibility for the substance of the law will diminish. Members of the community (both in

¹² E.Morozov, *To Save Everything, Click Here* (London: Penguin Group, 2013).

¹³ For similar lists of the requirements of the rule of law see Lon Fuller, *The Morality of Law*, 2nd ed (Yale University Press, 1969), chapter 2, John Finnis, *Natural Law and Natural Rights* (Oxford University Press, 1980), pp 270-1, Joseph Raz, *The Authority of Law* (Oxford University Press, 1979) pp 214-8, 219. For Jeremy Bentham’s list, compiled a century and a half before Fuller, see T.Endicott, ‘Arbitrariness’ (2014) *Canadian Journal of Law and Jurisprudence* 49-72 at 53-6.

their capacity as persons subject to the law, and in their capacity as those in whose name the community is governed) will be unable to understand the substance of the requirements produced by these new techniques as an exercise of the state's responsibility for governance and for the resolution of legal disputes among members of the community.

3. The principles of private agency and public agency

Two basic principles must be respected, if a community is to attain the rule of law. They underlie its value.

1. **the principle of private agency:** the state must treat persons as capable of deciding and acting for themselves, by empowering them to order their own affairs in certain crucial respects.
2. **the principle of public agency:** the community must make itself capable of deciding and acting responsibly, by empowering and requiring officials and institutions to undertake demonstrably reasoned action on its behalf in certain crucial respects.

The principles of private and public agency are basic prerequisites for personal responsibility and for responsible government. They underwrite much of political morality. But they have scarcely come into focus in the long history of political theorising about individual liberty and political community.¹⁴ Never until the present day have these principles been susceptible to the wholesale challenge now arising from powerful digital transformations sweeping through contemporary society. There has never before been any alternative to empowering individuals to act for themselves, and empowering officials and institutions to govern.

The principles of private and public agency do not guarantee the rule of law. But no community can be ruled by clear, open, stable, prospective, general standards unless individuals are empowered by law to order certain aspects of their own affairs, and public agencies are empowered by law to take reasoned decisions to make and to apply the law.

What is more, these prerequisites for the rule of law help to explain its value. The rule of law does not guarantee democratic government, or good government. But the rule of law is a prerequisite for democratic government (which treats adult citizens as agents who share in the exercise of political power) and good government (which is government that makes well-reasoned responses to the right considerations.)

Our formulation of both principles carries the rider, 'in certain crucial respects'. The rule of law does not demand that *everything* should be regulated by clear, open, stable, prospective, general standards. In particular, it does not prohibit particular norms. Far from it. Every legal system needs arrangements for the making of particular norms as a technique for giving effect to general rules in particular cases (and for other purposes such as resolving

¹⁴ As we will show below, however, these principles are implicit or partly expressed in much theorising concerning the role of law in society, including Lon Fuller's thinking about the rule of law (n 13). And in her work on the rule of law and algorithmic regulation in general, Mireille Hildebrandt has taken a related approach to ours, emphasising that the rule of law 'gives effective standing to those subject to law'. 'Law as computation in the era of artificial legal intelligence' (2018) 68 *University of Toronto Law Journal* 12-35.

legal disputes, and for giving remedies and imposing sanctions for breach of general rules, and so on).¹⁵

But in some respects, public agencies and officials must govern by making general rules and applying them to particular cases, and must confer powers on private persons to regulate their own affairs. Different states in different circumstances may legitimately leave it to the discretion of officials or to machine learning systems to make particular directives with legal force in various aspects of governance. But no tolerably humane state could leave everything to the particular discretion of officials, or to computational personalisation. It may be impossible to give a complete answer to the crucial question ('which aspects of life should not be regulated through computational personalisation?'), and we certainly cannot do so here. Instead, we will illustrate those aspects with paradigmatic instances, and hope to explain their importance by illustrating them.

Since our purpose is to identify principles for distinguishing the legitimate uses of personalised norms based on big data analytics from the illegitimate, we will not survey the large, diverse and expanding array of possibilities and proposals for replacing the rule of law with management through new computational techniques.¹⁶ We will select examples that are useful for identifying and articulating the principles. We will focus on two proposed applications of computationally personalised norms, namely:

- **computationally personalised micro-directives**¹⁷ (section 2): that is, replacement of general norms with automatically generated particular norms tailored to a person's behavioural profile as gleaned from the traces of her networked digital interactions, and
- **adjudicative micro-directives** (section 3): the application of machine learning algorithms to identify patterns in the past decisions of a court or tribunal, to replace the court's application of general rules with a determination of a case based on a predictive inference as to how the court would have applied general norms to a particular new case.

The beneficial potential – tailoring the law to achieve socially desirable outcomes, and to satisfy private persons' preferences, while improving certainty as to our legal position, and reducing the need for dispute resolution– can only legitimately be pursued if its limits are understood. Our core argument is as follows:

Core argument: The state ought to respect and preserve personal responsibility, and ought to achieve responsible government. For those purposes it needs to treat persons as agents, and it needs to institute public

¹⁵ Endicott, above n 11; on the role of particular orders in a legal system see Stephen Smith, *Rights, Wrongs, and Injustices: The Structure of Remedial Law* (Oxford: OUP 2020).

¹⁶ For surveys of such proposals, see 'Symposium on Personalized Law' (2019) 86 *University of Chicago Law Rev* 217 and J.M. Barry, J.W.Hatfield, S. D. Kominers 'To Thine Own Self Be True? Incentive Problems in Personalized Law' *San Diego Legal Studies Paper No. 20-439*, 6 Mar 2020, footnote 6. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3536174.

¹⁷ Anthony J Casey and Anthony Niblett imagine or foresee a world in which, 'when a citizen faces a legal decision, she is informed of exactly how to comply with every relevant law before she acts. The citizen does not have to weigh the reasonableness of her actions nor does she have to search for the content of a law. She follows a simple directive that is optimized for her situation. We call these refined laws "micro-directives."' (2016) 66 *University of Toronto Law Journal* 429-442 at 430. See also A.J. Casey and A. Niblett, 'The Death of Rules and Standards' (2017) 92 *Indiana Law Journal* 1401-1447. Available at: <https://www.repository.law.indiana.edu/ilj/vol92/iss4/3>.

agencies to engage in demonstrably reasoned action on behalf of the community. The principle of private agency is a prerequisite for human autonomy, which is the capacity of persons to take responsibility for their lives. The principle of public agency is a prerequisite for responsible government. Computational personalisation should only be used as a technique of governance where doing so is compatible with both principles.

The core argument is an argument against indiscriminate use of computationally personalised norms. The use of such norms can only be justified in circumstances in which these principles –prerequisites for the rule of law– do not require regulation of the relevant conduct through reasoned acts of human agency.

Indiscriminate replacement of law by technological management through computationally personalised norms should be stubbornly resisted, despite the benefits on offer in terms of precision, cost reduction, and the targeted attainment of policy goals. Applications of this kind offer enhanced availability of real-time guidance to individuals, as to what treatment they can expect from the state. That facility cannot make up for the loss of autonomy and the loss of responsible government that will result, if private persons and public actors are not treated as agents capable of responsibility.

4. Computationally personalised micro-directives

Norms can be personalised without a machine. Personalisation is a technique for tailoring some kind of benefit, service, or opportunity, or a burden or obligation, to the individual preferences, dispositions, tastes, capacities, competences or circumstances of the particular individual to whom the benefit or burden applies. Personalisation is the antithesis of a general rule. Parents personalise norms when they tailor their children's household chores to their capacities, traits, likes and dislikes, and to everything else that the children have going on at any given moment.

We will point out some particular respects in which the state can legitimately personalise liabilities and entitlements, but we will argue that the state, unlike a parent, should generally adhere to the rule of law, which governs persons' conduct through general duty-imposing rules, and by conferring legal powers for people to personalise their own affairs. The law does not personalise the distribution of property when a property owner dies, for example. The law of wills empowers the testator to personalise his or her inheritance. Contrast the impersonal, general default rules that apply if a person dies without having made a will.

For *some* forms of social ordering that ought to depend on personal capacities of adults, the law may legitimately impose a form of personalisation. Consider the personalisation of driving licences, awarded after a driving test that is designed to evaluate the particular candidate's ability to drive safely. Contrast the *impersonal* right to vote, awarded by the application of a general rule conferring the franchise on a person when she reaches the voting age.

A voting age applies an extremely general rule with only a very minimal form of personalisation (tailoring the right to the vote to the particular person's age, but not to any other characteristics).¹⁸ Driving tests administered by human examiners are only *partly*

¹⁸ It is important to note that norms can be more or less personalised, and more or less general. On the variable generality of norms, see the well-established body of literature concerned with

personalised. The examiner applies a set of general rules in the form of criteria for passing the test. The question of whether to award a driving licence is highly personalised by comparison with the right to vote, because it is tailored to the applicant's own personal abilities by the examiner when evaluating whether the applicant's driving skills meet the general standard. But it is not entirely personalised. Through big data analytics, a machine could enable much more far-reaching personalisation, replacing the human examination process (through which an individual's driving skills are assessed against relatively crude general rules) with machine recognition of patterns in a vast set of data on a learner's driving behaviour, and awarding a driver's licence on the basis of a comparison of data as to the applicant's driving with the driving behaviour of the population. We should emphasise that in this article, when we use the word 'personalisation', we refer to personalisation in this unprecedented form, facilitated by new data-driven information technologies and practices. Its dramatic, new feature is that it offers a radical tailoring of norms to personal characteristics made possible by machine learning techniques applied to large data sets that enable individualised predictive inferences in real time. In the case of the award of driving licences, this computational, data-driven alternative to the current driving licence regimes would lack the human foibles of driving test examiners. Replacing a human driving examiner with a machine-generated assessment might be one early success for the use of big data analytics.

But we now have technologies with the potential to take personalisation even further. Imagine personalising the right to vote through big data analysis designed to use sophisticated statistical inferences from vast sets of behavioural data collected about individuals, to assess each individual's intellectual and emotional capacity to exercise the vote competently. With computational personalisation, the state could adopt a very powerful alternative to a crude general voting age.

Our core argument is an argument against computationally personalised eligibility to vote, but it is not necessarily an argument against computationally personalised driving tests. This distinction illustrates the fact that the principles of private and public agency apply to state governance in certain crucial respects. Computationally personalised evaluation of ability to drive need involve no failure to respect the agency of the applicant for a licence, and no failure of responsible government. But eligibility to vote is best governed by a general voting age because the state needs to confer a general right to vote simply on the basis of membership of the community.

The common thread in proposals for computationally personalised norms is the potential for achieving smart outcomes by tailoring norms (rights and powers, obligations and liabilities) to the predicted capacities, dispositions, circumstances and preferences of each individual. The outcomes would be smart in the sense that they would be generated through the automated application of big data analytics to individual profiles generated from

understanding the relative strengths and shortcomings of detailed rules in contrast to broadly framed principles. See for example J. Braithwaite, 'Rules and Principles: A Theory of Legal Certainty' (2001) 27 *Australian Journal of Legal Philosophy* 47-82; P. Schlag, 'Rules and Standards' (1985) 33 *UCLA L. Rev.* 379. Norms can be general as to the class of conduct and the class of persons they apply to, and as to the time and as to the place at which they apply, and they can be more or less general in each respect: T. Endicott, 'The Generality of Law' in Luís Duarte Almeida, Andrea Dolcetti, James Edwards eds, *Reading The Concept of Law* (Hart Publishing 2013) 15-36 at 17-19, 22.

behavioural data gleaned from each individual, and tailored to (a) their inferred preferences or dispositions or (b) their inferred capacities and capabilities.¹⁹

a. Why might personalisation be desirable?

Computational norm personalisation holds out the potential to overcome two common drawbacks of legal rules: vagueness (which renders their application indeterminate in some cases), and generality (which renders them over- and under-inclusive). Vagueness reduces the capacity of rules to provide clear guidance to those to whom the rules are addressed, and to those who enforce and give effect to the rules.

The precision of a speed limit has guidance value for drivers: it tells them what they are required to do.²⁰ It also has process value for officials: it gives a ground for effective intervention in the interests of road safety, that can be applied with relatively little scope for abuse of discretionary power and with relatively little scope for dispute.²¹ The precision of a Blood Alcohol Content ('BAC') level as a metric for driving impairment has an important process value (although it is not very useful to drivers as a guide to their conduct): a police officer with a breathalyser can use the precise standard as a guide in deciding whether to restrain or to prosecute a driver whose breathalyser test indicates that the driver has exceeded the BAC limit. If accurate breathalysers are used with integrity, the result is to protect defendants from a form of arbitrariness: capricious or merely misguided application of a vague standard by officials.

Precise general rules, however, result in their own form of arbitrariness: a mismatch in some cases between the purposes of the law and its requirements.²² Suppose that the policy purpose of restrictions on drunk driving is to protect road users (perhaps including the driver) from the dangerous consequences of cognitive impairment, while not preventing people from driving when alcohol consumption results in no material impairment.

A rule making it an offence to drive when an impairment makes it 'unduly dangerous' to do so would in principle correspond *perfectly* to the reason for the rule; it would have *fidelity value*.²³ But such a regime would lack the process value of a BAC limit. By comparison with the use of a BAC limit, it would incur a different and worse form of arbitrariness: it would subject drivers and the public to the rule of the officials who are responsible for applying the rule. It would be worse because of the great difficulty of determining the seriousness of impairment in adjudication without the BAC metric, which carries an imperfect but very significant correlation with impairment. The arbitrariness of human assessment of impairment is avoided, and the purpose of the law is achieved to some significant degree, by

¹⁹ Applying Yeung's taxonomy of algorithmic regulatory systems (in Yeung, 'Algorithmic Regulation: A Critical Interrogation', n 1 above at 507-509), computationally personalised norm techniques are adaptive, pre-emptive and can be either configured as recommender systems (as in the case of predictive legal guidance) or automated (as in the case of computationally personalised microdirectives and adjudicative microdirectives).

²⁰ J. Black, *Rules and Regulators* (Oxford: Clarendon Press, 1997) 5-45.

²¹ T. Endicott, 'The Value of Vagueness', in V. Bhatia, J. Engberg, M. Gotti and D. Heller (eds), *Vagueness in Normative Texts* (Bern: Peter Lang, 2005) 27-48.

²² For a classic discussion of overinclusivity and underinclusivity, see F. Schauer, *Playing by the Rules* (Oxford: OUP 1991).

²³ See the discussion of fidelity value in Endicott, n 23 above at 39, and the discussions in the regulatory literature of 'policy congruence'. See for example C. Diver, 'The Optimal Precision of Legal Rules' (1983) 93 *Yale Law Journal* 65-109; J. Braithwaite, 'Rules and Principles: A Theory of Legal Certainty' (2001) 27 *Australian Journal of Legal Philosophy* 47-82.

a law specifying the same maximum blood alcohol level for all drivers. But it is achieved imperfectly. Some drivers, being more significantly affected by alcohol consumption than others, will be free to drive while more significantly impaired than others. For them, the law may be insufficiently restrictive, whilst for others (who are less significantly affected by alcohol consumption), the law may be too restrictive. The results are somewhat arbitrary, in the sense that the application of the rule is not fully responsive to reason (i.e., to the reason for having the rule). In virtue of its precision and generality, the law lacks fidelity value. That form of arbitrariness may be well worth incurring. Law-makers must often decide whether fidelity value or guidance value or process value is more important in a particular context. But what if we could abolish general rules with their inevitable forms of arbitrariness, and replace them with particular norms?

Imagine personalised limits on alcohol use by drivers. We would not have to choose between the arbitrariness of a precise general rule, and the arbitrariness that arises from capricious application of a vague general rule. Computational personalisation could enhance the fidelity value of the scheme of regulation (that is, the congruence between the effect that the scheme has in particular cases, and its underlying policy goal) by more precisely evaluating whether each individual's driving pattern indicates 'unduly dangerous' driving behaviour. It is not hard to imagine machine assessment of driving performance in real time through the application of big data analysis of population-wide driving data and of an individual's driving, to identify an unsafe pattern in the individual's driving. The BAC, a general rule that uses a proxy for impairment, could be replaced with a smart algorithm assessing performance directly, and the legal satnav could communicate a resulting particular norm to each particular driver at any particular moment.²⁴

Personalising this aspect of the regulation of driving would do away with the arbitrariness of overinclusivity and underinclusivity associated with a single, general, impersonal BAC standard applicable to all drivers, without incurring the arbitrariness that arises from the application of vague legal rules. A computationally personalised assessment of impairment could provide tailored, precise guidance to those to whom the legal rules are addressed (provided they have a legal sat nav), and to police.²⁵ It might involve some of the other, collateral drawbacks of personalisation, such as privacy concerns. But it would not necessarily violate the principle of private agency or the principle of public agency. Our core argument is not an argument against this use of computational personalisation.

b. Why might personalisation be undesirable?

When a property owner dies without having made a will, how should her property be dealt with? Legal systems tend to use crude general default rules.²⁶ Porat and Strahilevitz propose a solution that relies on computational norm personalisation: although the owner did not write a will, we can apply big data analytics to the owner's digital footprint in order to *infer*

²⁴ Compare the thought experiment drawn vividly by Mireille Hildebrandt, *Smart Technologies and the End(s) of Law* (Cheltenham: Edward Elgar 2015), in which a car's computer detects that the driver is mentally fatigued and issues an automated warning – and when she is assessed as too tired to concentrate properly, it switches to fully automated mode (at 4).

²⁵ Similar forms of personalisation could offer medical professions, banks and vendors of goods more precise and detailed guidance concerning the information they are legally obliged to disclose to patients, clients and consumers. See for example P.Hacker, 'Personalizing EU Private Law: From Disclosures to Nudges and Mandates' (2017) 25 *European Review of Private Law* 651.

²⁶ For example, in England and Wales, the residue goes first to a surviving spouse, then to surviving issue, then to surviving parents, and so on: Administration of Estates Act 1925, Part IV.

her preferences about the allocation of her property on death.²⁷ Significant elements of such a system are already available in automated on-line product recommender systems. These systems apply machine learning techniques to the digital traces of an individual's online behaviour to infer customers' tastes, preferences and interests. Porat and Strahilevitz see an advantage in applying similar techniques to the problem of allocation of property on intestacy: 'more estates would be allocated in a way that better approximates the true preferences of the decedent'.²⁸

The suggestion is that the purpose of default rules in private law is to identify preferences – those that an individual *would presumably have expressed* through the exercise of a legal power, if he or she had exercised it. Porat and Strahilevitz presuppose that, on the death of an intestate property holder, the relevant question for the community is: what were the decedent's preferences as to the disposition of the property? On that view, if the decedent had made a will expressing those preferences, then the law of wills would have given the community a good way of answering the question as to what his or her preferences were. But if the deceased property holder has made no will, then the general default rules are, at best, a regrettably blunt proxy for the preferences of the deceased concerning the distribution of her property. Computational personalisation could replace that proxy with algorithmic identification of preferences:

'...their intestacy rules will be personalized and hence will more closely approximate what they would have wanted than will the status quo's one-size-fits-all approach.'²⁹

Moreover, they conclude that 'many individuals who would have otherwise needed to incur the expenses of drafting wills now may no longer need to do so',³⁰ if they know that the automated, personalised distribution of their property will align with their preferences. And we might go further than Porat and Strahilevitz go, to point out the potential of personalisation: when a testator does make a will, the provisions of the will are a blunt proxy for the preferences of the deceased, insofar as circumstances, influences or distractions may prevent her from expressing her actual preferences through the terms of a will, and insofar as transaction costs tend to prevent a testator from constantly modifying her will in real time as her preferences evolve. If big data analysis of preferences is highly effective, then computerized norm personalisation might better achieve allocation of property in accord with the preferences of the deceased person if we used it to replace not only the default rules on intestacy, *but also* the legal power to make a will.

But that approach (like Porat and Strahilevitz's own more muted version) would misconceive the purpose of the law of succession, *which is not to satisfy the preferences of the property holder*. To adopt preference satisfaction as the law's purpose would be to violate the principle of private agency. We need to explain this central point, which informs our fundamental objection to widespread governance of private relations by computational norm personalisation.

²⁷ Porat and Strahilevitz, n 5 above. They propose that this could be done by asking 'human guinea pigs' their preferences, analyzing the guinea pigs' big data profiles, and drawing inferences from similarities between the intestate person's profiles and the guinea pigs' profiles; see 1420.

²⁸ *ibid* 1420.

²⁹ *ibid* 1420.

³⁰ *ibid*.

The power to make a will, of course, enables you to secure an outcome that you prefer. But when you make a will, you are not providing information that the community then uses in order to dispose of the property in accordance with your preferences; you are acting, in the exercise of a normative power, to dispose of your property.³¹

To use that power for your own purposes, you cannot simply ask yourself what you prefer. Your preference is your disposition; it is up to you to decide what it is to be. If you use the power to make a will responsibly, your preferences will be your response to the reasons you have to do this or that with the property. If you use the power arbitrarily, you act in a way that is unresponsive to the reasons you have. The law leaves it to you to exercise the responsibility that it gives you. That responsibility for action is necessary if you are to have 'basic responsibility', as John Gardner called it.³²

[Basic responsibility] is the ability to explain oneself, to give an intelligible account of oneself, to answer for oneself, as a rational being. In short, it is exactly what it sounds like: response-ability, an ability to respond.³³

The law can only respect your basic responsibility if it treats you as an agent. The community gives you responsibility for choosing whether to exercise the power to make a will, and for the results of your choice. This crucial point – the importance of treating the testator as having power to *act* to determine the disposition of property – explains the legitimate force of the law's formal requirements for wills. By determining what counts as exercising the power, they enable an agent to *act* to dispose of her property for her own reasons. Satisfying the testator's preferences is not the state's purpose. The state gives legal effect not to her preference but to what she does, in a way that demonstrates respect for her agency and for her capacity to determine those matters that she ought to control.³⁴ What counts, for the purposes of the law of wills, is not what she *prefers*, but what she *does*.

If she does not exercise the power to make a will, whether deliberately or through sheer procrastination or even accident, the question for the community (reflected in the law of inheritance) is not what her preferences were. The general rules of intestacy may justly pursue distributive aims that have nothing to do with the deceased's preferences about the distribution of her estate. Those rules may give effect (for example) to good purposes such as provision for children and stability of family housing, without needing to find any rationale in actual or inferred preferences of the deceased. Therefore, congruence with the underlying policy purposes of the law would not be secured by computational norm personalisation.

If we are right about this, then the computationally personalised approach to succession management proposed by Porat and Strahilevitz is a paradigm case of the sort of

³¹ See Joseph Raz's account of normative powers, in which he argues that if people did not have powers such as the power to make a will, 'much would be lost beyond the value accruing from their use of these powers. The further loss is of the value of having these powers in expanding the range of free choices ...the value ...is, in part, in the ability to use them, in the ability to choose to use them or to choose not to use them, and not only in the consequences of those choices.' J. Raz, 'Normative Powers' (2019) <https://ssrn.com/abstract=3379368> -p 6-7.

³² J. Gardner, 'The Mark of Responsibility', Chapter 9 in *Offences and Defences* (Oxford: OUP 2007) at 182.

³³ *ibid.*

³⁴ Barry, et al, n 16 above, point out that while it is straightforward for Big Tech to use data-driven techniques to personalise its offerings of things we might 'want' and 'like', this is not what the law is for.

replacement of law by big data analytic techniques that would be an illegitimate departure from the rule of law.

If you are doubtful about this, consider a natural extension of their proposal that Porat and Strahilevitz do not discuss: the general management of a person's property by big data analysis –including the making of gifts *inter vivos*, charitable donations, payments of tax, payment of debts, investment decisions, making of improvements to real property, and the purchase of goods and payment for services. The state does not use information about the property owner's preferences to arrange those affairs; it empowers the property holder to arrange them. The crucial element in every property transaction is the individual's *action* in exercise of a legal power. By treating her action as necessary and sufficient for the making of a gift or a payment or a contract, the law treats her as an agent, so that she can have basic responsibility.

The conferral of powers in private law is a paradigm instance of the impersonality of law: the law is blind to the agent's preferences, desires, behavioural dispositions, or other personal or behavioural characteristics. Where the law is *not* blind to such characteristics –e.g., if it restricts the ability of minors to dispose of property– it limits the powers that it confers.³⁵ The law does not personalize the management of your property; it empowers *you* to personalize it. If we replaced your legal capacity to control your property with management based on big data analysis of your preferences, your property would no longer be *property* which you own: your goods would become goods that the state manages for the purpose of satisfying you. By empowering property owners to manage their own affairs, the law treats them as subjects with responsibility, rather than as objects whose preferences the state is to satisfy.

The Porat and Strahilevitz proposal overlooks the importance of intentional action in human affairs, and overlooks the crucial resulting significance of the legal conferral on private persons of legal power to act. Management of succession to property by computational norm personalisation would abandon the principle of private agency. If that technique were extended and applied to the management of a person's property more generally, it would be the end of private property. Property presupposes action by a person to deal with what he or she owns, through the exercise of normative powers. Generalised property management through the application of computational norm personalisation would fail to treat the owner as a normative agent. It would not even be a regrettable property regime; it would abolish property ownership altogether. If all goods and chattels and real estate were administered by a next-generation data analysis of everyone's preferences, no one would own property: they would merely benefit from it for so long as the state's socio-technical management system allowed. Insofar as a big-data-analysis property management scheme would deprive human beings of agency in a crucial respect, it would breach the principle of private agency.

And if you are inclined to think that it might be a good thing to replace private property with a form of collective computational management of goods that would more efficiently satisfy the preferences of persons who are currently subject to the rule of law, consider the potential for the acme of what Porat and Strahilevitz call 'personalized law *in general*':³⁶ marriage through big data analysis. Currently, computer dating sites use algorithms to identify potentially compatible couples. But the law leaves it to those persons to pursue the possibilities, and to order their own lives. The law of marriage is impersonal. With no

³⁵ We are grateful to Emma Rengers for pointing this out.

³⁶ Porat and Strahilevitz n 5 above at 1478.

analysis of the two personalities, it treats a marriage as brought about by the action of the two persons in going through a ceremony that the law treats as the exercise of a power conferred by a general rule. It enables them to personalize their marriage. If, instead, *the law* personalized their marriage, the machine could prevent a marriage that would not satisfy the preferences of the parties. Or it could go all the way and register two persons as legally married, where the arrangement is predicted to satisfy the preferences of the parties. They could receive an alert on their legal satnav. Even machines are not perfect, but a machine might be able to make more reliable predictions as to the sustainability of a relationship than the parties can, and more reliable predictions as to the resulting satisfaction of the parties' preferences. The fallible parties to an emotionally charged relationship doubtless act in non-ideal epistemic conditions, and there is no reason to expect that they could identify their own preferences, or work out which marriage (if any) can reliably be predicted to result in the satisfaction of those preferences, as accurately as a smart algorithm applying machine learning techniques to massive datasets to detect data patterns that can be found in successful, sustainable marital relationships.

Machine marriage, of course, is a *reductio ad absurdum* of the idea of personalized law in general. It is worth spelling out why it would be absurd. The reason is that preference satisfaction is not the law's purpose. The law's many purposes include the focal purpose of enabling persons to order their commercial and family relationships. The law has the best of reasons to confer power on persons to act (e.g., to deal with property or to marry) instead of using big data analytics to identify and satisfy their preferences. The reason for empowering persons to personalize these crucial aspects of their lives for themselves lies in the significance of the act. Each party to a marriage, for instance, signifies something to the other by their action, and their joint act signifies something to their community. They could not achieve that through machine marriage, which would reflect a computerized assessment, but would signify nothing. By marrying, two people take responsibility for the formation of the legal arrangement, and for its sustainability.

Note that this fundamental objection to machine marriage is not an objection to big data driven recommender systems that offer you computationally personalized suggestions about potential marriage partners whom you might be interested in dating. Similarly, our argument provides no objection to the application of big data personalization techniques to generate a draft will which you can then decide to execute, or to revise before you execute it, or to ignore. In both these cases, your agency is preserved.³⁷

5. Adjudicative micro-directives

Just as the next-generation successors of product recommender systems could be used to replace the making of wills through inferences as to the preferences of testators, they could open up the possibility of data-driven, machine-generated predictive inferences as to how a human court or tribunal would decide a dispute. Then private legal services firms or a public agency could offer clients a prediction as to how a court would apply the law to the client's case, utilising natural language processing techniques powered by machine learning which are already used to enable automated journalism. Why not, then, an automated legal advisor? The system would subject the texts of the court's previous decisions to big data analytic techniques combined with an analysis of the profile and circumstances of the client's case. Because the underlying machine learning techniques are based on identifying

³⁷ See J Danaher. 'The Ethics of Algorithmic Outsourcing in Everyday Life' in Yeung & Lodge, *supra* n.1 at 98.

data patterns and correlations, rather than conventional ‘if X then Y’ programming rules, the outputs generated by the machine would not be generated by rule-based analysis of previous cases. Instead, the system would undertake a sophisticated pattern-based comparison of the client’s new case to the corpus of past judicial decisions, yielding a predictive inference as to the decision likely to be given.³⁸ The outcome of the analysis could offer clients the sort of indication of their chances of success that a good solicitor gives, but with predictive accuracy that solicitors can seldom offer. If these outputs were made available via a legal satnav, each of us could learn what determination a court or tribunal would likely reach if our legal position were in issue.

But with the same algorithmic analysis, the machine could *replace* the human court. In other words, why not an automated judge? The output of the predictive analysis could be given legal effect as a micro-directive³⁹ –a legally binding personalised norm.

Consider, for example, how a legal system decides whether a person (let’s say, a driver for Uber or Lyft) is an employee or an independent contractor. The general approach in various jurisdictions is to use the distinction to determine questions of tax law and employment law. The law generally bases the legal status of the individual engaged in the provision of services on a range of considerations (particularly concerning the control that the putative employer exercises over their activities). If a legal dispute arises, it will be determined by a tribunal following an assessment of the overall relationship.

Benjamin Alarie has pointed out that big data analytic tools could replace the tribunal in making determinations about the employment status of an individual.⁴⁰ The relevant data set would primarily consist of the body of previous decisions made by a human tribunal in evaluating the employment status of individuals in particular cases. Let’s imagine a really good human employment tribunal and assume that, with big data analytics and an extensive body of previous decisions in which employment status has been determined, the machine could predict what that tribunal would do, with great accuracy.⁴¹ Alarie says,

‘In the context of worker classification, more data and better inference tools make possible much sharper predictions about the content of standards in the law. If one is unclear about how to classify a worker, a taxpayer merely needs to consult with the applicable tool and an extremely reliable answer can be provided in minutes.’⁴²

But then, we could move beyond treating the outcome generated by the inference tools as a prediction, and give it legal effect as a determination of the legal position. In other words, big data analytics could be used to automate tribunal decision-making, replacing the tribunal. We could save the cost of human dispute resolution –not only the infrastructure

³⁸ For an example of an attempt to build such a predictive system, see N. Aletras et al, ‘Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective’, (2016) 2 Peer J Comput Sci e93; DOI 10.7717/peerj-cs.93. For a critique, see M. Hildebrandt (2018) ‘Algorithmic regulation and the rule of law’ *Phil. Trans. R. Soc. A* 376: 20170355. <http://dx.doi.org/10.1098/rsta.2017.0355> p 25.

³⁹ Casey and Niblett n 17 above.

⁴⁰ Alarie, ‘The Path of the Law: Towards Legal Singularity’, n 5 above.

⁴¹ Alarie (ibid at 448) says that ‘computationally intensive machine learning algorithms’ can predict whether the Tax Court of Canada would classify a worker as an employee or an independent contractor with greater than 98% confidence. We do not know why he says that, or quite what it means, but we will assume that the predictions would be very reliable.

⁴² Alarie, n 5.

cost and the cost of paying judges, but the cost of advocacy for the parties, and the costs involved in delay. The resulting legal position could be available in real time on a person's legal satnav on any particular matter (contentious or not).

We will call the data, systems and processes that need to operate together to produce such an adjudicative micro-directive a 'machine tribunal', even though that term is actually an oxymoron: the machine would not do what tribunals do, which is *to engage in adjudication* (a matter we discuss further below). Rather, the legal system would be giving binding effect to the outputs of the machine's pattern analysis.

A machine tribunal would fail to give effect to a fundamental principle that underlies responsible government: the principle of public agency. We argued that, if the law gave the same legal force to a computer analysis of a deceased person's preferences as it gives to a will, it would violate the principle of private agency. By the same token, if a computer analysis of a tribunal's dispositions was given the binding legal force currently attached to a human tribunal's determination, it would violate the principle of public agency.

6. Private agency and public agency: what the two principles require

We have shown that there are circumstances in which predictive personalisation by big data analytics might justly be used, with socially valuable results. Assessment of driving ability and of alcohol impairment are examples. Big data analytic techniques could provide precise, tailored governance via personalised norms that would avoid the arbitrary over- and under-inclusiveness of general rules, and the arbitrariness that arises from capricious application of the law by human courts and tribunals.

We have also offered paradigmatic examples of applications that would be illegitimate: big data property management and big data management of the right to vote. These applications can be understood as forms of 'technological management'⁴³ of social behaviour, rather than governance through law. They are anti-law projects.

Proponents of such applications will, of course, be in favour of replacing legal regulation. The principles of private and public agency provide a framework for determining whether doing so is legitimate. In this framework, that legitimacy does not depend on whether the new applications can satisfy private and public preferences. Those techniques of technological management are illegitimate if they would deprive persons of the agency that they need in order to take responsibility for their own lives, or if they would deprive the state of the agency that it needs in order to take responsibility for making mandatory and power-conferring norms, for applying them to particular cases, and for resolving legal disputes.

a. The principle of private agency

A political community must treat persons as capable of, and responsible for, acts intended by them to order their own affairs.⁴⁴ There are aspects of a person's legal position that should not be determined by the community's predictions or inferences concerning that person's preferences or dispositions, but by that person's own actions in exercise of legal

⁴³ Brownsword, n 9 above.

⁴⁴ This claim depends on a doctrine of legal capacity for decision making; a good legal system will provide for assisted decision making when a person needs assistance, and for protection of the rights of persons who cannot make decisions.

powers (as with the making of a will, and dispositions of property in general, and marriage), or by general rules (as with eligibility to vote).

To articulate the requirements of the principle of private agency fully, it would be necessary to specify the aspects of private persons' affairs over which they ought to have active control. We cannot offer a full account of those aspects of a person's relationships with other persons and with the community. It would be an account of the private legal powers that persons ought to have. It would depend on cultural, social, political and economic characteristics of a jurisdiction, and there is not necessarily any injustice in the fact that different legal systems confer different private powers.

But some aspects of the question of what should and should not be determined by a person's actions are easy to answer, and do not depend on particularities of jurisdictions. For example, the principle of private agency does *not* demand that all adults should be lawfully entitled to drive if they decide to drive; the principle *does* demand that all adults should be lawfully entitled to marry or to vote if they decide to marry or to vote. The relationship of marriage is radically personal; therefore *the law* of marriage is radically impersonal. That is enough to demonstrate that –for all the complexity and variability of its requirements– the principle of private agency imposes fundamental limits on legal personalisation, whether predictively automated or otherwise.

The distribution of a person's property on death in accordance with her preferences as inferred through computationally personalised norms, rather than in accordance with a will, would also violate the principle of private agency. By treating the *act* of attesting to the will (in the form prescribed by law) as a uniquely effective ordering device, the traditional law of wills gives effect to the principle of private agency. If your exercise of a legal power to make a will is replaced by the output of big data analysis applied to behavioural data gleaned from tracking your actions and activities, you cannot be responsible for the outcome. Suppose that you can consult the latest advisory output from your legal satnav, to learn how your property would be allocated if you were to die today. You might be glad that things would turn out that way (we are assuming that the machine is good at inferring what you prefer). But you did not choose that allocation. You will have no cause to regret what you have done, and no cause to be glad about what you have done. You will not have done anything.

We have argued above that general management of property by big data analysis would, in fact, abolish property ownership. That is the case because property ownership depends on the principle of private agency. If a computer disposes of a chattel or a piece of real estate in accordance with your computationally inferred preferences, you will not be responsible for the transaction. Nor will you be responsible for any other computationally personalised disposition of property. Suppose the machine sends a gift to your uncle. He will be glad to receive it, and it may be nice for him to know that the data patterns identified by the computer presumably reflect a positive attitude towards him on your part. But you will have done nothing that he could thank you for.

This radical decline in personal responsibility would be a devastating loss for human beings.⁴⁵ That is the fundamental objection to a machine-learning property management system, and generally to any indiscriminate replacement of *actions* that attract legal consequences with machine-generated predictive inferences concerning what you and I

⁴⁵ For a discussion of loss of human autonomy from the turn to action-forcing forms of 'design-based' social control, see See K. Yeung, 'Can We Employ Design-Based Regulation While Avoiding *Brave New World*?' (2011) 3 *Law, Innovation and Technology* 1-29.

would prefer. The community holds you responsible for your actions: responsible for adhering to duties imposed by law, and responsible for the exercise of powers conferred by law. By imposing sanctions and remedies if you fail to discharge your duties, it *holds you responsible* for your wrongs. By giving legal effect to your exercise of legal powers, it gives you responsibility for the ordering of certain aspects of your own life. Machine property management would treat us as beings who are capable of gratification but not of responsibility.

Property management by big data analysis would violate the principle of private agency.⁴⁶ Its impact would be a setback for individual autonomy. For me to be autonomous, it is not enough that things that gratify me will happen to my goods; the law needs to empower me to deal with them intentionally and accountably.

b. The principle of public agency

Governments, like private agents, must have basic responsibility – the ability to respond to reasons.⁴⁷ And to have this capacity, the government must be able to engage in reasoned action and its actions must demonstrably count as the community’s actions.⁴⁸ Responsibility and accountability are not even possible unless the government has the ability to engage in reasoned action. Public officials and agencies must also exercise those capacities to deliver responsible government – that is, they must take reasoned action for purposes that are the community’s proper purposes, and must do so in a demonstrable manner, which typically requires mechanisms and institutions through which the government is rendered accountable for its words and deeds. The requirement that the government’s reasoned action be *demonstrable* is grounded on the government’s obligation of accountability or ‘answerability’, which is also rooted in the requirements of basic responsibility: it will not suffice for rulers merely to inform persons subject to their power as to what is going to happen to them; rulers must also be capable of rendering an account through which their actions can be explained, recognised and understood by members of the community *as* reasoned actions. If governments have the capacity for basic responsibility, and they discharge their responsibilities, the result is responsible government.

We argue that if action on behalf of a community (eg the UK) could not be demonstrably identified as reasoned action, then the community could not attain responsible government. Governments must act, and there are actions which the state has a moral duty to undertake on the ground of reasons concerning the substance of the decision to act. Such actions are necessary for responsible government, and they include action to establish and to maintain a legal system that gives effect to the rule of law. For example, the government has a duty to take executive and legislative action to respond to public needs, and a judicial duty to provide authoritative determinations of the law’s application, imposing appropriate legal

⁴⁶ This claim presupposes that ownership of property is one of the crucial respects in which persons should be able to order their own affairs; we cannot defend that presupposition here, but we can at least point out that it is not quite as ideologically loaded as it may appear; it is compatible with a very wide array of more or less socialist economic orders.

⁴⁷ See n 32 above.

⁴⁸ Will Bateman has argued that the principles of public law require that ‘statutory powers be exercised by agents who: have certain *cognitive capacities*; can give intelligible *reasons for the exercise of power*; and to whom social and political responsibility can be *attributed*’: ‘Algorithmic Decision-Making and Legality: Public Law Dimensions’ (2019) 93 *Australian Law Journal* (forthcoming). That argument defends a conception of responsible government that presupposes the principle of public agency.

remedies, sanctions and other orders on the basis of those determinations. The point of all techniques of public accountability is to hold those agencies to that responsibility. On this basis, we posit the principle of public agency as a requirement of responsible government.

The principle of public agency underwrites responsible government. But it is only the starting point for *good* government. The principle of public agency is not a principle of justice, or of democracy: by conferring power on a public agency, the law does not guarantee that the agency will use it responsibly; the law only makes it conceivable that the agency will do so. A dictator can adhere to the principle of public agency by making himself the sole public agent. But public agency is a *prerequisite* for every aspect of good government, and for democracy. Adhering to the principle of public agency creates the possibility of responsible government. Responsible government entails demonstrably reasoned action by public agencies on behalf of the community that responds to the right sorts of considerations. Responsible government, in turn, is a prerequisite for good government, which is government that responds well to those considerations.⁴⁹

There are intentional acts for which public agents and agencies ought to have responsibility – acts that they should be empowered to undertake, and for which they ought to be held responsible. A government agency must have capacity for such reasoned acts in order for it to be regarded as responsible for acting, and for deciding to use or not to use its powers, and for any failures to discharge the duties of government.

The principle of public agency applies at every level of government: the state itself must have public agency. It must be a person –an agent– in international law. And domestic governance must be carried out by agencies of the state. Local authorities, every executive official and agency, courts, legislatures, and voters must all be agents. The life of the community must be governed by their actions, and not merely in a way that is aligned with their preferences, dispositions or behaviours.

As with the principle of private agency, it is impossible for us to specify fully what aspects of governance must be carried out through the actions of public agencies. But as with the principle of private agency, it is possible to point out paradigms to illustrate those crucial aspects of governance. Replacing the orders of a court with the adjudicative micro-directives of a machine tribunal as the standard technique of dispute resolution would violate the principle of public agency. It would be a failure of responsible government. It would replace the rule of law with technological management.

The principle of public agency also applies to legislatures. One central prerequisite for responsible government (little noticed because it has never been put in question until now) is the requirement that a legislature should be empowered to make law, and should engage in a *reasoned act* in **the** exercise of its power (that is, an act that is based on reasons relevant to its exercise). The law treats the legislative act as the source of a binding general norm. The law treats the adjudicative act as the source of a binding particular norm that is an application of the general norm. The preferences and the dispositions of the legislature and the court are of no legal effect; it is what the legislature *does* and what the court *does* that counts. This is true of electoral, legislative, judicial, and executive decisions. It makes sense to hold a public agency to account (through the various techniques for accountability of public agencies) only insofar as the public agency in question is capable of acting on the basis of reasons.

⁴⁹ See T. Endicott, *Administrative Law* 4th ed (Oxford: OUP 2018) 15. Public agencies have basic responsibility in John Gardner's sense (n 31 above): 'response-ability, an ability to respond'.

To understand why an adjudicative microdirective that is given binding legal force would not constitute a reasoned action, imagine that a senior human tribunal, struggling under a huge caseload and backlog, is considering strategies to reduce its workload. One possibility might be to divert low value, high volume, relatively straightforward cases away from the tribunal. Two alternative possibilities are under consideration. First, it could convene a junior tribunal comprised of newly qualified solicitors to preside over and decide those cases. The senior human tribunal would delegate its judicial authority to the junior tribunal, which would then determine these cases, producing reasoned judgments which are deemed to be the judgments of the senior human tribunal. Alternatively, a team of software engineers and natural language processing experts could be commissioned to build a system that is capable of parsing the existing corpus of case law produced by the senior tribunal to generate predictive inferences as to how the senior tribunal would decide the case. The outcomes generated by this machine tribunal would then be deemed to be the decisions of the senior human tribunal.

Even if the machine tribunal was capable of producing perfectly accurate predictions as to how the senior human tribunal would decide the case, adopting this strategy would not entail reasoned action by the machine tribunal. It would thus fail to satisfy the principle of public agency, and would effectively abandon the rule of law. In contrast, if the senior tribunal delegated the task of adjudication to the junior human tribunal, there would be no abandonment of the rule of law unless decisions by the junior tribunal were merely arbitrary. Even if the junior tribunal's determinations were not always those that the senior tribunal would have arrived at, this proposal would respect the principle of public agency and the rule of law. Unlike the machine tribunal, the junior human tribunal would arrive at its decisions by hearing reasoned argument and evaluating those arguments by reference to the governing law. Unlike the junior human tribunal, the machine would not be engaging in an act of adjudication when generating binding outcomes.⁵⁰ That is because, as Lon Fuller famously put it, the essence and distinguishing characteristic of adjudication as a means for resolving disputes lies in participation through the presentation of proofs and reasoned argument. For Fuller, it is through adjudication that formal and institutional expression is given to the influence of reasoned argument in human affairs. By this he meant that adjudication is institutionally committed to a decision based on principle. It depends critically on governance through rules that are authoritatively applied by a neutral and independent arbitrator.⁵¹

Although it might be possible to provide a functional 'explanation' of how the machine tribunal generated a given outcome based on the inputs, mechanisms of pattern recognition and calculative mechanisms which its algorithmic model employs to produce a given outcome, these are not reasoned evaluations.⁵² When computer scientists working in the

⁵⁰ We are assuming no right of appeal from the decisions of either the junior human tribunal or the machine tribunal; such a compound process involving either scheme would be an intermediate case that might be compatible with the rule of law.

⁵¹ L Fuller, 'The Forms and Limits of Adjudication' (1978) 92 *Harvard Law Review* 353, 364.

⁵² Some forms of machine learning algorithms, such as decision trees, Bayesian classifiers, additive models, and sparse linear models generate interpretable models in that the model components (e.g., weight of a feature in a linear model, a path in a decision tree, or a specific rule) can be directly inspected to understand the model's predictions. These algorithms use a reasonably restricted number of internal components (i.e. paths, rules, or features) but provide traceability and transparency in their decision making. As long as the model is accurate for the prediction task, these approaches provide the visibility to understand the how the AI system arrived at a given output. However, other forms of machine learning techniques, particularly those that rely on deep learning

field of 'explainable AI' (XAI) refer to 'explanations', these should not be confused with what we call reasons. The aim of XAI is to provide visibility into how deep learning systems that are largely uninterpretable to humans (sometimes referred to as 'black box' models) produce outcomes and predictions.⁵³ These functional 'explanations' are not arrived at through the evaluation of reasoned argument necessary for adjudication and the proper exercise of judicial power. Accordingly, even if the machine could generate functional explanations as to how it arrived at any given outcome, it would not be engaging in reasoned action, nor would the process by which it generated determinations be a process of adjudication. Hence, if the senior human tribunal conferred binding legal force on the microdirectives generated by the machine, it would abdicate its judicial duty and abandon public agency and, therefore, the rule of law.

Generalised governance by computationally personalised norms would make responsible government impossible, because of the abandonment of public agency. In answer to this objection, it might be pointed out that our current modes of governance also fall short of responsible government, insofar as they are unresponsive to the interests at stake. Indeed, that is why machine-personalised governance appears attractive. If a legal dispute, for example, is determined by a tribunal of human beings who assess the disputing parties' positions and apply a vague general legal rule in light of various incommensurable considerations, then the parties' legal liabilities and entitlements are to some extent ruled by the mere say-so of those persons, and not by law. There is a significant element of arbitrariness in any scheme of adjudication by a human tribunal.

The element of arbitrariness in general rules is built into the rule of law; the element of arbitrariness in human adjudication is *at odds with* the rule of law, and yet it is an inevitable result of the rule of law (this is the tension in the ideal of the rule of law that we pointed out above). Every political community with the rule of law recognises, and has become accustomed to, the foibles and fallibility of human tribunals, and the under- and over-inclusiveness of general rules. Contemporary legal systems typically institute procedural mechanisms (including hearings before independent decision makers, the giving of reasons, and appeals) as safeguards against unjustified outcomes that may result from these vulnerabilities. Those safeguards do not eliminate the potential for arbitrariness when your legal position is determined by a human tribunal. But they reflect a commitment to governing relationships (between private parties, and between each of the parties and the community) in a way that is reasoned. The rule of law operates as a safeguard against the abuse of the various discretions that lie in the hands of the tribunal. Even taken together, the remaining forms of arbitrariness that accompany the rule of law –inevitable in any system of governance– are a price worth paying.

Governance by general rules applied by an independent tribunal through a process of adjudication is an imperfect attempt at responsible government. Machine tribunals would abandon the attempt. If we replaced the judgments of human courts with the

and convolutional neural networks, generate models that are inherently uninterpretable to human users. This class of machine learning algorithms sacrifice transparency and interpretability for prediction accuracy. See Hall, P., & Gill, N. (2019). *An introduction to machine learning interpretability*. Second edition. Sebastopol, CA: O'Reilly Media, Incorporated. Du, M., Liu, N., & Hu, X. (2018). Techniques for interpretable machine learning. *arXiv preprint arXiv:1808.00033*.

⁵³ C Rudin, 'Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead' (2019) 1.5 *Nature Machine Intelligence* 206-215.

output of predictive analytics, we would be replacing judgments with simulations of judgment.⁵⁴ Legally binding machine predictions of what a human tribunal would decide might be very accurate. They might well satisfy a form of Turing test for a machine tribunal. That is, it might be impossible to distinguish the machine tribunal from a human tribunal, merely by reference to its outputs. But the outputs would not be judgments. Judgments are decisions for which the decision maker has basic responsibility, and for which the decision maker could conceivably have reasons.

With the use of machine tribunals, the legal system would fail, for the first time, to adhere to the principle of public agency. In the long history of failures of responsible decision making by public authorities, their outputs have never been something for which no one can be responsible.

7. The asymmetry between the two principles, and the connection between them

There is a striking difference between our two principles, and we want to point out its implications. While we argue that both private persons and public bodies must have agency, it is also a fundamental feature of good government that it should not generally prevent individuals from acting arbitrarily. The empowerment of private persons in the law of succession, in the use of property in general, and even in marriage, empowers them to make arbitrary decisions with legal effect. It would be a massive injustice if the state passed judgment on the reasonableness of an individual's decision to marry, or to leave a bequest, or to make a purchase.⁵⁵ By contrast, the purpose for which the state and public bodies ought to have agency is not so that they can enjoy personal freedom and autonomy, but so that they can engage in responsible governmental decision making.

The asymmetry between the two principles reflects another asymmetry between state and individual: the state can impose legal duties on the individual, but the individual cannot impose legal duties on the state. That explains why the state should institute reflexive techniques for ensuring that it acts responsibly, but should only hold individuals to responsibilities where it has imposed them as legal duties through clear, open, general, prospective rules. We would gain nothing but a form of tyranny, if the state could require individuals to act in accordance with its assessment as to which individual decisions would be responsible. We would gain another form of tyranny, if the state did not require its agencies to act responsibly.

The state must respect the principle of private agency in order for people to have responsibility. The state must live up to the principle of public agency in order to act responsibly. In the traditional rule-of-law framework, general prospective laws purport to impose obligations and liabilities and to confer rights and powers on private individuals impersonally, and the state is responsible for applying those general norms retrospectively

⁵⁴ Compare Mireille Hildebrandt's argument that in the issuance of adjudicative micro-directives, computers would be 'merely simulating mathematically what human reasoning has come up with' (Hildebrandt, n 38 above at 21).

⁵⁵ We do not mean that these private actions cannot legitimately be regulated. Marriage law regulates who can marry whom, succession law may require that a surviving partner can continue to live in the family home or that children should be provided for by an estate, and the purchase of certain weapons may be prohibited; all of these are legitimate forms of regulation. The injustice we have in mind would arise if the state set out to supervise marriages, wills, and purchases, requiring the relevant legal powers to be exercised reasonably.

through adjudication by independent courts. Individuals in turn have responsibility for exercising their powers and for complying with their obligations.

The asymmetry reflects the reciprocity that, we have argued, results from the combination of respect for private agency, and commitment to public agency. The state exists for individuals, in the sense that it is its citizens' way of acting as an organised community. The individual does not exist for the state. Irresponsible private behaviour is, of course, unreasonable, and there are very many forms of irresponsible private behaviour that are the legitimate concern of the law and of public agencies; but there is no general principle that the state should pass judgment on whether a marriage or a bequest or a purchase is unreasonable. The complex governmental technique that underlies the rule of law involves (1) conferral of power on individuals to enable them to have responsibility (and therefore, autonomy), without requiring that they use their power responsibly, and (2) conferral of power on public agencies precisely for the purpose that they should exercise it responsibly. With private individuals, conferring power is necessary to support their autonomy; with public agencies, conferring power is necessary as a prerequisite for them to be held accountable for responsible government.

8. The value of uniformity

The logic underpinning the use of big data analytics to tailor legal norms to fit the behavioural traits, preferences, capacities and competences of individuals rests on the bluntness of the one-size-fits-all approach of general rules. But even though general rules can be over- and under-inclusive, there will be contexts in which uniformity is itself valuable, so that there are overwhelming practical reasons against personalising norms of behaviour.

In fact, no one could advocate complete replacement of general rules by big data analytics.⁵⁶ It would be absurd to personalise the rule of the road, using big data analysis to decide whether each particular driver should drive on the right or on the left, through computational analysis of their individual dispositions, capacities, or preferences. The absurdity illustrates a powerful reason not to personalise governance in general: where the community needs uniformity of action, good governance requires a general rule. There is no better example of the value of incurring some arbitrariness (specifically, the mismatch in the case of some particular drivers between the purposes of traffic regulation, and the application of the rule of the road) in the interest of crucial policy purposes that require uniformity (in this case, road safety and the value of vehicles being able to proceed quickly).

Uniform, general rules are often necessary for solving human coordination problems. Advocates of micro-directives as the dominant technique of regulation need not deny that. They could advocate the use of that technique only when it is optimal. They consider that it is optimal when its consequences align with pertinent private or public preferences; our argument is that those consequences can only legitimately be optimised when doing so is compatible with the principles of private and public agency.

Uniformity is not uniformly valuable; it depends on the particular problem of social ordering. Uniformity may not be required when coordination is required. One thing that machine learning can do very well is to optimise coordination between many moving parts in a

⁵⁶ Enthusiasts sometimes talk as if it would be progress. Casey and Niblett say with evident relish, 'Our long-run prediction is that micro-directives will become the dominant form of law, culminating in the death of rules and standards' (n 17 above at 1404), although they add that their analysis 'is positive rather than normative' (1405).

system. Uncoordinated human driving behaviour often leads to jams on popular routes, precisely because drivers are acting uniformly; in a smart traffic system, traffic could be coordinated by an arrangement requiring different vehicles to behave *differently* (taking a variety of routes, so that congestion decreases⁵⁷). Or suppose that the cycle of red and green traffic lights at an intersection has long been governed uniformly by a timer allotting the same amount of time to each direction; now the public authority replaces the timer with a computer operating a smart algorithm, using predictive big data analysis to give the east-west traffic longer when traffic in those lanes is likely to become heavy, or even to manage the overall traffic flow of the road network (perhaps on the basis of data from intersections around the city, or data from devices in vehicles, aided by machine learning concerning the complex patterns of traffic flow). The departure from uniformity might be a very useful improvement in traffic regulation. And the new scheme is no more contrary to the principle of private agency or to the principle of public agency than the old analogue timer. The reason is that the community does not need a reasoned decision, for which an agency can be held responsible, as to which lanes should have a green light at which times.

The alternation of red and green traffic lights is a solution to a pure coordination problem. It is a paradigm of a regulatory context in which a coordination solution is needed, but there is no issue of justice or good policy (which would call for an exercise of public agency) as to which direction of traffic should have a green light at which time. So we do not need a public agency to take responsibility for the fact that one direction has a green light at any moment, while the other has a red light. As a reminder of the reasons why computational personalisation can be legitimate, it is worth considering that there would be nothing wrong with having a police officer directing traffic at the intersection. Here, where particularised regulation of an aspect of the life of the community is legitimate, we might well achieve a cost saving and a worthwhile improvement in traffic engineering by managing traffic through computerised particular norms.

But the potential value of uniformity is important to the distinction between human pursuits that call for legal ordering, and human pursuits in which personalisation is legitimate. Ben-Shahar and Porat write:

‘In the same way that personalized medicine can save lives and avoid inefficient uniform treatments, personalized safety standards can reduce the social costs of accidents.’⁵⁸

We propose, to the contrary, that there is a deep difference between personalised medicine and personalised law: in medical care, the purpose is to treat the particular patient, and there is no inherent therapeutic value in uniformity of action.⁵⁹ It may be cost-efficient to locate a variety of medical services on the same site. A general policy of prescribing a uniform treatment (such as a standard dose of a medicine) may be an efficient way of achieving delivery of second-best treatments when medical personnel (or by patients

⁵⁷ This is already happening to some extent in some cities, through private uptake of Google’s GPS navigation app, Waze.

⁵⁸ Ben-Shahar and Porat n 5 above at 688.

⁵⁹ We hasten to add that uniformity in health care provision can have a social value (involving a form of solidarity), and can be a technique for respecting the needs of persons who would otherwise be victims of the economics of health care. And we do not mean to deny that certain forms of individualised medical care in the 21st century may have antitherapeutic consequences. See Donna Dickenson, *Me Medicine vs. We Medicine: Reclaiming Biotechnology for the Common Good* (Columbia University Press 2013).

administering their own treatment) are not able to make fine-grained, expert judgments in a hurry. But the focal criterion of success for medical care concerns the health of each individual patient, so that personalised medicine is *generally desirable*. A great deal of the prodigious advance of medicine has involved improvements in the personalisation of healthcare.

Legal regulation, by contrast, often has the achievement of uniformity *as its purpose*. The purpose of law is not to treat a particular person. Its primary purpose is to serve the public good of a community. Personalisation in law is only valuable when it is useful for that purpose. Very often, for a variety of reasons, that purpose rules out personalisation, and requires a uniform standard. Those reasons include a multitude of practical considerations, such as the reasons why it is useful to require vehicles to keep to the left or to the right on roads. But we can point to one very common value of uniformity in legal standards, which is to support fairness and solidarity in a community. The voting age provides a paradigm. It aims at restricting the vote to those who have the maturity that is needed for the responsible exercise of the right to vote. But maturity is a very personal matter, and adolescents grow up in different ways, at different rates. We could personalise the right to vote with an inference through big data analysis as to whether each individual has reached the required level of cognitive, emotional, and social maturity. Computer personalisation would avoid the arbitrariness of a uniform voting age, which does not precisely correspond to the developing maturity of each voter.

But the arbitrariness of a personalised voting age would be far worse. It would be arbitrary in a new sense: it would be based on considerations on which the community ought not to act. A voting age has a crucial expressive function.⁶⁰ That function depends on *not* personalising the franchise. A voting age signals the community's refusal to draw distinctions among persons concerning the maturity of their judgment, when allocating the political power of the vote. The general rule reflects the right of members of the political community in general to vote. So the use of a single uniform voting age –treating people uniformly in this respect– is a far better approach than personalising the franchise.

9. Conclusion

Unrestrained algorithmic personalisation of law would fail to give private persons and public agencies responsibility for things that they ought to be responsible for. That is our core argument. We have proceeded on the heroic assumption that it would be technically possible to build machine learning systems that can make the kind of inferences and predictions that would be necessary to generate forms of personalised law advocated by proponents of computational personalisation. In practice, it seems highly implausible that there would be an available source of suitable, high-quality data about every individual member of the population that would allow the making of accurate predictive inferences as to the entire gamut of preferences that have legal relevance.

While there might be more realistic prospects of configuring machine learning systems that can make accurate predictions about the outcomes of legal cases for narrowly defined questions, these models presuppose an existing corpus of judicial decisions as 'ground truth', providing a database that will support predictive inferences as to how other cases would be decided. But the proportion of legal complaints for which legal proceedings are initiated and which actually proceed to judgement is tiny, and may be an inadequate basis upon which to make predictions about how judges might decide novel cases that bear little

⁶⁰ See B. Morgan and K.Yeung, *An Introduction to Law and Regulation* (Cambridge: CUP, 2007) 6-7.

resemblance to those that form the database on which the proposed computational inferences would be based. Nothing in our argument should be taken to suggest that, apart from the rule-of-law concerns we have raised, computational personalisation could actually achieve the benefits held out by proponents.

And aside from questions of technical feasibility, there are sound objections to unrestrained personalisation of law, which impose stringent limits on the extent of personalisation that would be compatible with the purposes that a legal system must pursue:

- Inscrutability –that is, decision making that cannot be understood by the persons governed by it.
- The risk that biased, discriminatory or otherwise unjustified outputs may result if the data set lacks integrity, accuracy and reliability.
- Sclerosis: in spite of its promise of real-time response to information, big data analysis could trap public action in patterns identified in the existing dispositions and preferences reflected in the data set, preventing equitable decision making and the development of the law.
- Dangers of violating moral and legal rights to privacy and data protection.
- Risk of abuse of power by programmers or other officials.
- Risk of bias resulting from poor design of machine learning algorithms, even without any abuse of power.

Our core argument does not depend on any of these significant potential drawbacks of governance through computationally personalised norms. They manifest various collateral ways in which governance through big data analytic techniques might fail to respond to the relevant considerations and, therefore, could result in the unjust treatment of persons who are subject to micro-directives. The promise that has been posited by advocates of computational personalisation is that, if those collateral problems and risks can be addressed, personal and social preferences could be satisfied more effectively than they are through the rule of law. But our core argument is more basic than the collateral problems. If computational personalisation were treated as the dominant technique of ordering society, it would undermine responsible government by depriving the state of its basic responsibility. And it would undermine the autonomy of persons subject to the law by denying them basic responsibility.

The moral of the story has significance for the general theory of law: the purpose that a legal system ought to pursue is not the satisfaction of preferences. The purpose of a legal system is more complex than that, and includes the purpose of empowering private agents and public agencies to act.

It may be impossible to give a complete and general account of the crucial respects in which individuals ought to exercise responsibility for their lives, and the crucial respects in which agencies of the state ought to exercise responsibility for the substance of legal regulation. We certainly cannot do it here. What we have tried to do, instead, is to point out paradigm cases in which the principles of private agency and public agency require that the law empower private persons and public agents to make decisions for which they can be responsible.

Our paradigm cases of private matters that cannot legitimately be managed through computational personalisation are succession to property, the management of property in general, and marriage. Adjudication is our paradigm of an interaction between the state and

persons in which the content of the decision ought to be determined by reasoned action for which a public agent can be responsible (that is, for which the agent is capable of having reasons, and responding to them). Computational personalisation of adjudication in general would violate the principle of public agency. Computational personalisation of the right to vote is another paradigm in which the principle of public agency would be violated: in order to establish the right relationship between itself and members of the community (and among its members), the state has to take responsibility for a general, uniform standard for the franchise. Paradigms could, needless to say, be multiplied, to include machine criminal justice.

Finally, note the paradigm case in which computational personalisation would violate both the principle of public agency and the principle of private agency: it is the machine election. For if we replaced the vote with a big data analysis of the voters' preferences, they would lose agency at a point that connects the private and the public. The vote is the paradigm case of an action for which individuals must have responsibility, which they can only have if they act in exercise of a legal power conferred on them by the state. The two principles are connected. If (and only if) they both have effect, it is possible for there to be a relationship of moral community among citizens, and between them and the institutions of their state. Private agency and public agency, of course, do not guarantee a healthy democracy. But a healthy democracy is impossible without both.

What distinguishes law from technological management is that law is capable of recognising and treating individuals as responsible moral agents, and as members of a political community that is responsible for its decisions. In other words, the generality of law and the universality of its application bring us back to the principles of private agency and of public agency, which are referable to the importance of demonstrating respect for the autonomy of ordinary individuals, and to the community's need for public agencies that are capable of taking responsibility for making decisions on behalf of the community.

The principle of private agency and the principle of public agency have always been in the background of political and legal theory, generally taken for granted because the only effective ways of governing have been through empowering public agencies staffed by human beings to govern, and by empowering individual human beings to order crucial aspects of their own affairs. Today and for the future, the available techniques of governance include the use of algorithmic regulation. And so today and for the future, the principles of private and public agency provide critical compass points for identifying whether particular proposals to utilise big data analytics in the service of computational norm personalisation would be a damaging departure from the rule of law.

The principles of private agency and public agency will not be violated if we use big data analytics to manage things –such as the flow of traffic through an intersection– that can legitimately be subjected to technological management.⁶¹ The principles *will* be violated if we use big data analytics generally to abolish the exercise of legal powers by private persons and public agencies. We should steadfastly resist the wholesale application of such techniques, despite their potential to achieve new efficiencies.

⁶¹ And it is worth noting the obvious: that we have made no attempt to give a complete account of all the circumstances in which computational personalisation of norms would be legitimate. In some circumstances computational personalisation could actually increase human freedom and autonomy, if it enables lawmakers to replace general prohibitions with tailored restrictions, where a general restriction had previously been the only way to achieve a good public purpose. We are grateful to Joseph Raz for pointing this out.

Contemplating rule by computers reveals drawbacks in the rule of law. Legal regulation has clumsy and arbitrary aspects that computers could eliminate. And for all that, we are still better off with the rule of law.