

1. Introduction: Life and the law in the era of data-driven agency

Mireille Hildebrandt and Kieron O'Hara¹

WHAT'S NEXT?

This volume targets the issue of how data-driven agency affects everyday life and everyday law, highlighting potential transformations of human agency and the emergence of new types of human-machine hybrid intelligence. Human agency will be considered as depending – in part – on the affordances of our technological environment, meaning that insofar as modern law assumes a specific type of human agency, we cannot take for granted that this assumption will survive the era of data-driven information and communication infrastructures. This invites a reflection on ‘the political’ in the realm of algorithmic systems that act on the feedback they gain from their sensors and/or other data input. Are we moving towards a driverless democracy that overcomes old school notions such as sovereignty, or will a driverless democracy result in a clueless form of *ungovernment*? Shall we lose track of the values and institutions that we need to preserve if we wish to survive the scaling of personalized manipulation in the realm of public debate? And what if a driverless law arrives to overdetermine the choice architecture of our machinic environments?

THE END(S)

Data-Driven Agency

In her monograph, *Smart Technologies and the End(s) of Law. Novel Entanglements between Law and Technology* (which will be referred to as *The End(s)* throughout this volume), Hildebrandt develops a theoretical framework to understand what data-driven agency means for law and the Rule of Law. She defines agency broadly as the ability to observe an environment and to act upon it based on such observation. Data-driven agency is a type of agency where observations are limited to digital data

and actions are informed by the computational processing of such data. This brings any deterministic data-driven decision-system under the concept of data-driven agency, for instance an application that determines when to start the central heating based on a decline in temperature, or one that determines social security benefits based on input of relevant data and the specified decision tree. Next to deterministic systems, which are in principle predictable, we now have systems that apply machine learning (ML), meaning that the system updates its own operational rules based on the feedback it receives. The thermostat tries to guess whether you are home, what time you will arrive, and which temperatures you prefer based on your past behaviours. The social security system may include an early warning system that flags people whose behaviours correlate with fraud, autonomically requiring them to provide additional information or autonomically instigating more detailed monitoring (and, potentially, autonomically deciding on punitive or other sanctions).

Issues of Privacy, Non-Discrimination and the Presumption of Innocence

Clearly both types of data-driven agency have their promise and drawbacks. *The End(s)* investigates how the surge of agential applications in our 'everyday' lifeworld reconfigures issues of privacy, non-discrimination, due process and the presumption of innocence, while remaining largely invisible, and hard to test and contest. These systems cross borders into what we thought was our private life and refine business and government abilities to apply fine-grained, mathematically-based discriminations. They determine who will be monitored more closely and allow human decision-makers to hide behind novel modulations of 'computer says no'. The aim of Hildebrandt's theoretical framework is to provide conceptual instruments to assess the implications of both deterministic and the more unpredictable data-driven applications, in relation to a more in-depth understanding of the grammar that informs law and the Rule of Law. She argues that this grammar cannot be taken for granted and will have to be reinvented in part, to ensure that individual human persons and groups can develop the capability to speak law to power, notably to the power generated by data-driven agents.

Theorizing the Implications of Data-Driven Agency

We believe that nothing is as practical as good theory. We witness an upsurge of initiatives around fair, accountable and transparent machine learning (FAT ML), the fake news challenge (FNC), and artificial intelligence now (AINOW), as well as concerns about singularity or

recursively-improving artificial intelligence (RIAI) leading to artificial superintelligence (ASI). Simultaneously, we see growing fears for a progressive erosion of privacy, non-discrimination, freedom of speech and freedom of information. To ensure the effectiveness and continuity of such initiatives in the long run it is crucial that scholarship is developed that faces and explores the assumptions that inform data-driven computing systems, detecting how and to what extent they undermine the familiar assumptions of previous information and communication infrastructures (ICIs) such as script, typescript, analogue telephone and telegraph and mass media. *The End(s)* has generated fruitful debate with a wide range of interlocutors, raising a number of pivotal questions. Reviews have been published (Diver 2018, McGee 2016, Keenan 2016, Jewell 2016, Van der Sloot 2015), including a special issue of *Critical Analysis of Law* featuring five reviews by leading thinkers in the domain, with Hildebrandt's reply.² However, the questions raised by *The End(s)* deserve further and deeper responses, developed into more precise arguments, not constrained by the format of a review. We have therefore invited a multi-disciplinary range of authors to contribute considered responses, developing their own positions on the core issues raised by *The End(s)*.

Pertinent Explorations of Core Themes of *The End(s)*

This edited volume can be read independently of *The End(s)*. It contains a wealth of new ideas on the interrelated topics explored in *The End(s)*. It targets two main research domains investigated by Hildebrandt, both of which concern the dramatic changes initiated by practices around data-driven innovation that have resulted in what has been called an onlife world (Floridi 2014).³ The first domain involves the shape of this new world and the ways in which it will, or may, affect human agency within it. Second, this book looks into the kind of politics this world does and does not afford and considers the challenges it contains to law and the Rule of Law.

This introductory overview by the editors is followed by a debate between the editors, and it ends with a response by Hildebrandt. In between the reader will find two parts: one on 'human agency in a data-driven environment' and one on 'the political and the law'.

PART I. HUMAN AGENCY IN A DATA-DRIVEN ENVIRONMENT

With regard to the rise of an onlife world, we note that smart technologies are not yet mature; we see hype next to impressive substance. Much of

the potential of smart technologies, for good or ill, is premised on a view of an integrated set of AI methods running over linked data, yet reality is (currently) far more fragmented and the results less rosy than some would have it. In fact, two of the founding fathers of machine learning emphasize the limited capacities of reinforcement learning, and deep learning (DL) (Brooks 2017, Le Vine 2017), warning against unrealistic expectations. However, we cannot rest our civil protections on the hopes that the technology will not deliver. Instead, we need to face the fact that many decision-making systems are based on a belief in the salience of ML and DL, with far reaching implications for those depending on them ('if machines define a situation as real, it is real in its consequences').⁴

Philosopher of robotics Paul Dumouchel considers the epistemology of big data, focusing in particular on two aspects which impact on its political and legal effects. Firstly, machine learning is a 'black box' process, where explaining why or how a particular output was reached is non-trivial – even designers of machine learning systems cannot explain a decision because there is no propositional content to their internal workings. Dumouchel marks the implications of the belief that data and correlations are all we need, warning against the rise of a new reductionism in the slipstream of a new 'Unity of Science' project, where all science becomes data science. The idea that causality does not matter and research on data is equivalent to empirical research, however, seems more prominent in the social domain than in the natural sciences. Dumouchel recounts that in the natural sciences data is often collected for a specific purpose, whereas the social domain tends to work with the data-exhaust produced as a side-effect of commercial applications or government services. This seems to make the social domain even more vulnerable to unfounded beliefs in the 'truth' of big data inferences. Secondly, when machine learning is embedded into real-world decision-making, it alters the world it models. For instance, smart signage on a motorway does not just describe traffic flow – it affects the flow itself, creating feedback loops that may be hard to predict, and which may result in the production of self-fulfilling prophecies. Dumouchel highlights the fact that in many applications of data science knowing and doing are conflated, and this is where human agency is at stake, as these systems act upon the behaviours they observe. This confirms a dire need to reflect on the potential repercussions of questionable beliefs in supposedly objective, neutral and omniscient algorithms.

Legal scholar Julie Cohen takes on Dumouchel's epistemological challenges, and their consequences for human agency. In parallel with Hildebrandt's notion of a 'digital unconscious', Cohen explores the metaphor of the human limbic system, a set of brain functions associated with emotion, motivation, learning and memory. Her analysis builds on

a rich account of how online platforms manage to subliminally influence the behaviour of their users by way of algorithmic systems that optimise for more behavioural data, more user engagement and thus for a kind of participatory consumer surveillance. Given these developments, the issue of the human response and adaptation to smart technologies is further explored, in particular, the contrast between the propositional, surface level of human and social reasoning, and the lower level, non-propositional, non-declarative depth of so-called deep learning. Cohen uses contrasting notions such as cognitive/precognitive and conscious/subconscious to explore the distinction between the big data ‘limbic’ constructs and the idea that technology would lead to human freedom, and that ‘knowledge is power’. In her new book, *Between Truth and Power. The Legal Constructions of Informational Capitalism* (forthcoming 2019), she further elaborates on the affordances of this new limbic system, with keen attention to the economic incentive structure that drives their exploitation. This implies that restricting one’s analysis to individual agency will not do, and meaningful consent must be considered from the perspective of what the digital unconscious affords.

Computer scientist and human-computer interaction researcher mc schraefel and her co-authors aim to turn meaningless into meaningful consent by way of ‘seamful’ interaction design. They discuss the interface issues between humans and the machines that surreptitiously adjust their new onlife world. This frame, tellingly dubbed the Internet of Things (IoT), in many ways resembles Cohen’s limbic system and Hildebrandt’s digital unconscious. schraefel et al. provide important insights on how to rethink human capabilities in an environment that thrives on subliminal influencing. They notably consider the role of human machine interaction (HMI) and user experience (UX) design, to argue for the importance of making unexpected usage of data apparent to data subjects using implicit signals at the point of contact between human and machine. Given the notion of what they call ‘apparency’, humans could at least be alerted to the use of their data, and use whatever legal (data protection) powers they may have to demand to see what goes on in the black box. Apparency, in this sense, must precede transparency if machines are to be accountable. This seems to bridge Hildebrandt’s focal interest in Gibson’s concept of affordances with Norman’s understanding of affordances, the latter being core to intuitive design. schraefel et al. note that apparency will create an onlife world that is not seamless (as IoT aficionados would have it) but seamful (Chalmers and Galani 2004), thus developing technical articulations of Cohen’s (2012) semantic discontinuity. A seamful environment will provide both semantic transparency (about what one consents to) and pragmatic transparency (about the actual consequences of clicking the

consent button). This connects with Dumouchel's emphasis on the new connection between knowing and doing on the side of the limbic system; algorithmic systems built on big data do not merely know or predict us, but act on this knowledge by e.g. pre-empting us. Apparency and pragmatic transparency, finally, are meant to integrate actionable transparency on the side of the user, who wishes to prevent or even pre-empt their IoT environments from making incorrect, unfair or otherwise undesirable decisions. This connects the need for apparency and transparency with the need for controllability and accountability, thus closing the loop that should inform the design of our new online world.

The main concern of *legal philosopher Sylvie Delacroix* and *policy scientist Michael Veale* is neither privacy or autonomy, but 'a fundamental commitment to equality' that aims to protect human agency against manipulability on the basis of invisible profiling and targeting (which subjects human agency to differential treatment based on statistical inferences). Instead of focusing on the fairness or correctness of algorithmic decisions, they target the implications of profiling for human agency and its relationship with the 'double contingency' that is inherent in human intercourse. They analyse Hildebrandt's suggestion of 'counter-profiling' to see how this may afford new ways to nourish and protect our 'double contingency' and the agency it enables. To explain how this relates to their commitment to equality they explore the example of 'social cruelty', notably that of genocide, which implies a fundamental intent to treat a specific group of people as unequal and not worthy of respect. The authors explain how this eradicates the capability of constructing one's own identity in the process of interacting with others. Their position is that the invisible, profile-based optimization of current data-driven environments entails a similar 'social cruelty', as it defines people based on their inferred characteristics, increasingly forcing them to behave as their dynamic profiles assume they will behave. They mark the use of affective computing, based on emotion-recognition, which enables precisely the kind of subliminal influencing that Julie Cohen refers to in the second chapter. To counter such 'social cruelty' the authors develop some ideas for passive and active interventions. They reject the attempt to reduce transparency requirements to a mere rectification of epistemic imbalances, referring to the 'transparency fallacy' as a mistaken attempt to put the burden of informed consent on the shoulders of individual agency. In fact, they seem to agree with Dumouchel and schraefel et al. on the integration of knowing and behaving that is core to smart systems, arguing that rebalancing cannot be restricted to 'mere' knowledge symmetry, requiring what schraefel et al. call pragmatic (actionable) transparency, next to semantic transparency. In terms of passive interventions, in alignment

with the previous chapter, Delacroix and Veale emphasize the need to turn the seamlessness of smart environments inside out, creating triggers for users to resist the way they are targeted, notably by showing them how others may 'read' and 'target' them. In terms of active interventions they develop two very interesting ideas to counter the 'social cruelty' of being defined by smart environments: surprises and appropriation. This is a rich contribution to the discourse, way beyond attempts to counter overdetermination by means of 'propositional, surface level of human and social reasoning', instead countering 'the lower level, non-propositional, non-declarative depth of so-called deep learning' (Cohen's limbic system) by way of intuitive counterings.

Philosopher of technology Charles Ess follows Hildebrandt's lead to explore the intriguing possibility of a 'post-digital world', invoking virtue ethics to explore a world in which the very medium of law itself is evolving in the face of technological change. Dubbing this world 'post-digital', Ess refers to a stance that neither embraces nor rejects the digital, instead taking seriously the novel affordances it creates for our embodied, analogue existence. This raises interesting questions about the relationship between law, ethics and the political, and the extent to which virtue ethics could contribute to new types of interaction between humans and their predictive environments. Ess makes two further points, first, warning against a determinist understanding of the relationship between media and its users, and, second, providing a convincing argument of the relevance of virtue ethics in what he calls a post-digital era. The first point addresses readers versed in old school media theory, that could misread Hildebrandt's concept of affordances and her stress on the impact of smart technologies on both human agency and the Rule of Law. Though *The End(s)* does not endorse a determinist position, being rooted in a postphenomenological philosophy of technology, those not familiar with the empirical turn in philosophy of technology may misread its core claim: whether a technology is determinist or not is an empirical question and depends on whether it induces or forces, inhibits or precludes behaviour. The second point addresses both developers and users of digital technologies as moral agents, inviting them to develop the practical wisdom (*phronēsis*) that should inform the design of technologies that contribute to human flourishing and enable others to develop their practical wisdom. This is a compelling argument because practical wisdom is a key concept in some strands of legal theory, notably for those that believe 'that the life of the law is not logic but experience' (Holmes 1997, pragmatism), noting that judgement is not equivalent to calculation (Ricoeur 2003, hermeneutics), meaning that discretion is inherent in legal decision-making but should not be understood as arbitrary subjectivist decisionism (Dworkin 1982, hermeneutics).

PART II. THE POLITICAL AND THE LAW

The second set of issues considered in this book revolves around the political⁵ and the law. Big data and machine learning, together with other sociotechnical practices such as social networking and personalization, have produced a series of effects, such as fake news, filter bubbles, chaotic pluralism and a dramatic reduction in private space. Some believe these effects have actually invigorated our democracy, whereas others find them somewhat disconcerting to say the least. These are not arcane issues: we saw the debate played out as national psychodrama in 2016, when Donald Trump used Twitter and the expertise of an advisor from the Breitbart news site to rout mainstream (and not-so-mainstream) Republicans to gain the nomination, and then to outflank the Democratic candidate who showcased her great experience and traditional political virtues. Away from the campaign trail, the political is also implicated as decision-making is being distributed away from traditional representative institutions – but to whom? Though some power is devolving to individuals whose opinion can more easily be consulted, there is a great deal of power in the machines that set the parameters for interaction and discussion.

Philosopher Gerard de Vries examines the threats posed by digital technologies for both democracy and the Rule of Law, raising the question of whether the weaponizing of these technologies merely changes the way democracy operates or invokes an existential crisis for both. De Vries faces this question by asking how digital technologies generate new vulnerabilities for democracy, tested against two conceptions of democracy. One, the minimalist, defines democracy merely in terms of voting, but offers little protection against manipulation of the constituency either before or after the vote takes place. De Vries argues, instead, for a relational understanding of democracy, focused on the public interest, as argued by Dewey and Latour, with historical roots in Montesquieu's *On the Spirit of the Laws*. This entails keen attention to the web of relationships between citizens, businesses, government agencies and other institutions, highlighting the need for hard work to keep the network of interacting actors stable and in place. In the final part of his chapter, De Vries demonstrates how this relational understanding of democracy better explains the threats of digital technologies to both its *nature* (institution of government by way of popular vote) and its *principle* (care of the public interest). The automation of fake news that is enabled by digital technologies is hidden from public scrutiny, which endangers the assumptions of free and informed general elections, thus threatening the *nature* of democracy. Big data and AI, in the meantime, threaten the *principle* of democracy, as they do not merely

manipulate existing political preferences but frame and subliminally constitute publics, replacing the difficult process of negotiating the public interest with algorithmic shortcuts.

Political philosopher David Stevens takes a daring stand by developing a sophisticated argument in favour of governments' efforts to induce 'attitudes and beliefs in citizens regarding a sense of social justice and cooperation' via artificial agents such as the personal assistant (Toma) that features in the narrative opening of *The End(s)*. Whereas many readers frame Toma's interferences as a corruption of human autonomy and dignity, Stevens comes out with a justification for a specific type of subliminal influencing that boils down to a specific type of nudging, described as 'structuring the background against which individuals can be led to make better (more optimal) choices, given their desires'. He notes that Toma 'frames' choices, and 'primes' its human master, and builds a sustained argument under what conditions this could actually be justified. The main condition is that such surreptitious nudging by the digital limbic system concerns a type of behaviour that does not depend on reasonable agreement, e.g. behaviour that should be enforced irrespective of whether people find such enforcement acceptable. This concerns, notably, deliberate harming of others, so – according to Stevens – the state should be applauded for influencing people's moral beliefs or sense of justice in the direction of rejecting deliberate harming of others. Basically, Stevens follows Rawls' distinction between an individual's choosing their own version of what they consider the good life, and an individual's sense of social or public justice. He argues that whereas the state is not allowed to manipulate a person's preferences for the good life, it is allowed to interfere – even subliminally – in the formation of a person's sense of justice, as long as absence of such a sense of justice would be unacceptable. Though Stevens admits that it would be great if children were primed for justice as a matter of course through education and upbringing, to the extent that this fails, the state would be justified in employing what Hildebrandt calls the digital unconscious to nudge citizens into developing a sense of social and public justice.

Next, **computer scientist Kieron O'Hara** and **political scientist, historian and biographer Mark Garnett** take another daring stand by defending a conservative philosophy that prizes familiarity and problematizes change, instead of clinging to the liberal paradigms that seem to underline human autonomy. To begin with, O'Hara and Garnett reject right-wing appropriations of conservative thought, instead building on Burke, Tocqueville and Oakeshott. For them, conservatism is a sceptical philosophy that problematizes the rationalist underpinnings of social engineering (think smart cities and IoT, and the pre-emptive nudgings of Julie Cohen's

limbic system). Simultaneously, they emphasize the need to preserve and sustain tested institutions that protect against the volatility of innovative disruption and against untenable assumptions about a free-floating human autonomy combined with naïve notions of human perfection. Interestingly, O'Hara and Garnett are suspicious of the idea that human behaviour is predictable as long as we have access to big data, and wary of attempts to use emotion-recognition as a way to subvert first order preferences. They value organic, mostly given relationships rather than egalitarian, 'designed' and seamless interaction, based on an inquiry into the 'mystical state' heralded by Burke, the disintermediated community (multitude) detected by Tocqueville, and the agonistic pluralism put forward by Oakeshott. The latter speaks of the *individual manqué*, who submits to the lure of communal pressure instead of following their own line of action. This accords with the idea of the state as a civil association (*societas*) of independent individuals, rather than the state as a monolithic entity (*universitas*) that imposes itself on the *individuals manqués* that are the target of its public management. It is interesting to note that the conservative Oakeshott seems to beat liberals at their own game, foregrounding the individual as an independent agent that should not be treated as a manipulable pawn, whether by their government or powerful social networks (or does this mean that Oakeshott was really a liberal in disguise?). The authors suggest that this particular strand of sceptical conservatism may have more clout in protecting the values of contextual integrity and individual agency than the utilitarian vision of an atomistic liberalism.

Legal scholar Christoph Graber looks at the composition of rights in the onlife world. Graber rejects the view of lawyers who tend to take the technology as a given black box, and the perspectives of scholars of science, technology and society (STS) who tend to view the law as a closed book. Instead, building on Luhmann's and Teubner's systems theory and Gibson's concept of affordance, Graber investigates how law and technology development may interact to further the public interest. He employs the concept of affordance to refer to the enabling and constraining implications of a technology, staying close to Gibson's original notion. He then adds Ihde's notion of the multistability of a technology, arguing that technologies have multistable affordances, a notion developed by Pfaffenberger to indicate that a particular technology will have multiple affordances that depend on how the technology has been interpreted. This leads Graber to an interesting employment of Pfaffenberger's 'technological drama theory', which enables him to confront the design constituency that develops a technology and thus creates its multiple affordances with the impact constituency that constructs various interpretations of the

technology, depending on what it affords. This way Graber can sort out the actors in the drama (those create the technology, those who interpret and those who redesign it), the processes (of creation, interpretation, redesign and reconstitution) and the politics involved (dominant mythos on the intended use, a counter mythos based on actual, alternative uses, thus creating a counterartefact that results from the negotiations between the design constituency and the impact constituency). What makes this chapter so very interesting is the application of this frame of reference to the AI narrative, the role of big tech platforms and the need for counter-narratives and activism to reconstitute AI in ways that afford a proper engagement with the public interest.

Legal scholar Paul Ohm introduces two metrics to ensure that machine learning systems are sufficiently civilized to enter the domain of human intercourse. The first metric requires that for an ML system to replace a human decision-maker it should demonstrate far better performance than the human, instead of mere equivalence. The second metric requires that the training data of an ML system should be less than complete, to prevent availability of behavioural data at a scale that enables overdetermination. These metrics both aim to throttle ML systems, in order to tame their performance in line with human expectations and human empowerment, to reintroduce seamfulness. This chapter is refreshing if only because it moves from diagnosis to solutions, while rejecting the mantra that speed, completeness and automation of data-driven decision-making will by itself solve the problems they create. The idea of throttle mechanisms goes beyond the usual calls for transparency, accountability and interpretability, providing for direct hands-on criteria to intervene in the process of developing and adopting ML systems. Such intervention includes prevention, in the sense of shrinking back from investing in systems that do not clearly demonstrate added value for those subject to their decision-making. It also includes intervention that slows down the blind progress of machineries that may hold promise while simultaneously involving huge societal costs. Building on his previous work on ‘Desirable Inefficiency’, Ohm and Frankle (2018) explains how these metrics would serve as tools to slow down the inexorable march of disruptive innovation, giving us time to turn away from untested implications while taking the time to develop ML systems that actually do good work without externalizing potential costs to others.

Legal philosopher Niels van Dijk has been triggered by the European Parliament (EP)’s recommendation to consider attributing electronic personhood to specific types of AI systems. In the last chapter of this volume, he discusses and juxtaposes a range of different ‘persons’, thus hoping to clarify to what extent e-persons could make sense. He distinguishes (1)

juristic persons (legal subjects other than ‘natural persons’) as fictions with effect, (2) public persons (notably Hobbes’ state), as unifying a multitude, (3) average persons (based on the mean of an aggregate population), as statistical realities, (4) profile persons (inferred from a training set), as machine-generated group portraits, and (5) digital persons (targeting of a particular person based on inferences gained from aggregate data, or software agents), as dividual data portraits and smart agents. Having presented the context, history and background of different types of persons, Van Dijk pins down overlaps, similarities and differences. He notably discusses the composition, what is represented and what is afforded to whom. His final point is that introducing a new (juristic) person does not necessarily imply a new ontological entity with a claim to semi- or pseudo-human personhood; rather it is a pragmatic invention meant to produce specific legal effects in order to solve practical problems. Besides, the public person that initiated the discussion (the EP), should be reminded that the introduction of novel legal subjects will have myriad implications within the legal system – that should be considered before making such a move.

AN INVITATION TO OPEN THE BLACK BOX

A set of scholars such as these will naturally provoke debate, disagreement and contestation. This is, of course, the point of this book. Yet, while there are distinct viewpoints on offer, there is also surprising (and at the same time cheering) convergence on a number of key points. Most notably, the authors are as one in resisting technophilic myths of omniscience and omnipotence. They also all focus on existing social practices, institutions or ideals (different ones in each case), and consider data-driven agency not in the abstract but in the context of the particular. Seamlessness, the great ideal of technologists, outlaws interruption and contingency; but society as we know it *is* discontinuous, *is* haphazard. If it is to remain plural, then surely it must remain untidy and scruffy. And so many of our authors, like Cohen, schraefel and O’Hara and Garnett, think about reintroducing or protecting the seams, even unsightly ones; others, like Stevens and Ohm, want to make sure the power of the technology enhances something inherently desirable at a higher level than its merely local effect. Put another way, the public interest, which features in many of the chapters, is not a parameter to be optimised within systems, but a commons whose manipulation must be open to debate and control.

In his *Dialogues Concerning Natural Religion*, David Hume (2018 (1779)) expressed the problem of evil – if God is omniscient, omnipotent

and infinitely good, then how is there evil in the world? This argument seems to reappear in Silicon Valley solutionism: if we are omniscient (we have all the data), and omnipotent (we have totally effective computing power and control of an increasingly virtualized world), how can bad things happen? Answer: they mustn't, the technology should ensure happiness. But happiness can only be ensured if the technology itself is allowed to define and determine it.

The US *Declaration of Independence* promised life, liberty and the pursuit of happiness. Data-driven agency goes further – according to the solutionists it promises happiness itself. But achieving that may have severe effects on liberty, as Aldous Huxley argued almost a century ago. The French *Declaration of the Rights of Man and of the Citizen* and its subsequent revolutionary uptake promised freedom, equality and community, and data-driven agency may be as severe on these ideals, as e.g. Oscar Gandy (2010) and Frank Pasquale (2015) have argued more recently. Of the two editors, O'Hara is more enthusiastic about the American vision, Hildebrandt about the French (and we will explore our differences in more detail in the next chapter), but we are equally concerned about the potential for social harm. All our authors are clear that data-driven agency may present opportunities to enhance, through its affordances, human flourishing. Most of us, however, warn against expecting too much and elaborate in salient detail how data-driven agency presents a framing problem that surreptitiously reconfigures our choice architectures. We hope that the reader will come to appreciate the many loopholes and pitfalls that face whoever promises that we can eat our cake and have it all: life, liberty, equality, community and happiness.

NOTES

1. With credits to Aniek van den Teuling for helping to prepare this and the final chapter; her 'reading' of the chapters of this volume further clarified what makes this a salient volume.
2. Brownsword (2017), Calo (2017), Cohen (2017), Kerr (2017), Raab (2017), Hildebrandt (2017).
3. Further developed in *The End(s)*, notably in Part 1.
4. This is a modulation of the Thomas Theorem: 'If men define a situation as real, it is real in its consequences', see *The End(s)* at 26.
5. On the difference between 'the political' and 'politics' see Mouffe (2005). 'The political' refers to the dimension of ineradicable antagonism that is inherent in human society, whereas 'politics' refers to concrete attempts to institutionalize ways to cope with such antagonism.

REFERENCES

- Brooks, R. 2017. 'Machine Learning Explained', MIT RETHINK, *Robots, AI, and Other Stuff* (blog), 28 August, <<http://rodnebrooks.com/forai-machine-learning-explained/>>
- Brownsword, R. 2017. 'Disruptive Agents and Our Onlife World: Should We Be Concerned?' *Critical Analysis of Law* 4 (1), 61–71
- Calo, R. 2017. 'Technology, Law, and Affordance: A Review of Smart Technologies and the End(s) of Law'. *Critical Analysis of Law* 4 (1), 72–7
- Chalmers, M. and A. Galani. 2004. 'Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems'. In *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*, 243–52. DIS'04. New York, NY: ACM, <<https://doi.org/10.1145/1013115.1013149>>
- Cohen, J. E. 2012. *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice*. New Haven, CT: Yale University Press
- Cohen, J. E. 2017. 'Affording Fundamental Rights: A Provocation Inspired by Mireille Hildebrandt'. *Critical Analysis of Law* 4 (1), 78–90
- Cohen, J. E. 2019. *Between Truth and Power. The Legal Constructions of Informational Capitalism*. Oxford: Oxford University Press
- Diver, Laurence. 2018. 'Law as a User: Design, Affordance, and the Technological Mediation of Norms'. *SCRIPTed* 15 (1) 4–41, <<https://doi.org/10.2966/scrip.150118.4>>
- Dworkin, Ronald. 1982. 'Law as Interpretation'. *Texas Law Review* 60 (2), 527–50
- Floidi, L. 2014. *The Onlife Manifesto. Being Human in a Hyperconnected Era*. Dordrecht: Springer
- Gandy, Oscar H. 2010. 'Engaging Rational Discrimination: Exploring Reasons for Placing Regulatory Constraints on Decision Support Systems'. *Ethics and Information Technology* 12 (1), 29–42, <<https://doi.org/10.1007/s10676-009-9198-6>>
- Hildebrandt, M. 2017. 'Law as an Affordance: The Devil is in the Vanishing Point(s)'. *Critical Analysis of Law* 4 (1), 116–28
- Holmes, Oliver Wendell. 1997. 'The Path of the Law'. *Harvard Law Review* 110, 991–1009
- Hume, David. 2018 (1779). *Dialogues Concerning Natural Religion*. Greenwood, WI: Suzeteo
- Jewell, M. 2016. 'Book Review Smart Technologies and the End(s) of Law'. *SCRIPTed* 13, 215–18
- Keenan, Bernard. 2016. 'Mireille Hildebrandt, Smart Technologies and the End(s) of Law, Cheltenham: Edward Elgar'. *The Modern Law Review* 79 (4), 733–8, <<https://doi.org/10.1111/1468-2230.12209>>
- Kerr, Ian. 2017. 'The Devil is in the Defaults'. *Critical Analysis of Law* 4 (1), 91–103
- Le Vine, S. 2017. 'Artificial Intelligence Pioneer Says we Need to Start Over', *Axios*, 15 September, <<https://www.axios.com/ai-pioneer-advocates-starting-over-2485537027.html>>
- McGee, K. 2016. 'On Legal Replicants'. *Jurimetrics* 56, 305–18
- Mouffe, C. 2005. *On the Political*. Abingdon: Routledge
- Ohm, Paul and Jonathan Frankle. 2018. 'Desirable Inefficiency'. *Florida Law Review* 70, 777–838

- Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press
- Raab, Charles D. 2017. 'Hildebrandt's Onlife World: Public Goods, Design and Politics'. *Critical Analysis of Law* 4 (1), 104–15
- Ricoeur, Paul. 2003. *The Just*. Translated by David Pellauer. Chicago, IL: University of Chicago Press
- Van der Sloot, B. 2015. 'Book Review Smart Technologies and the End(s) of Law' *European Data Protection Law Review* 1, 157–9