

Machine Performance and Human Failure: How Shall We Regulate Autonomous Systems?

Horst Eidenmüller

“But the mind that had once rebelled against the gods was about to dethrone itself by way of its own fabulous reach. In the compressed version, we would devise a machine a little cleverer than ourselves, then set that machine to invent another that lay beyond our comprehension. What need then of us?”¹

INTRODUCTION

Machines powered by artificial intelligence (AI) are on the rise.² In many use cases, their performance today already exceeds human capabilities. However, machines do not operate flawlessly—defects and accidents occur. In 2018, an article on “Robotic Rules of the Road” in *The Economist* discussed the legal regime for self-driving cars.³ In particular, it raised the question of when such cars would be allowed to participate in regular traffic. The author suggested that “... [autonomous vehicles] will always be held to higher safety standards than human drivers.” We learn that scholars and practitioners ponder whether they should be 10%, 90% or even 99.9% safer before being allowed to cruise the roads.⁴

But why, one may ask, should this be so? Why should self-driving cars not be allowed on our roads once they are as safe as human drivers or just *marginally* safer? Would this not be an improvement compared to the status quo? Indeed, taking the “safety logic” seriously appears to suggest an even more radical question: *When do we prohibit humans from driving cars?* Elon Musk pondered this question as early as 2015. He was reported to think “... that once

¹ IAN MCEWAN, *MACHINES LIKE ME* 80 (2019).

² The literature on the subject is of course vast. See, for example, ERIK BRYNJOLFSSON & ANDREW MCAFEE, *THE SECOND MACHINE AGE—WORK, PROGRESS, AND PROSPERITY IN A TIME OF BRILLIANT TECHNOLOGIES* (2014); MARTIN FORD, *THE RISE OF THE ROBOTS—TECHNOLOGY AND THE THREAT OF MASS UNEMPLOYMENT* (2015); THOMAS RID, *RISE OF THE MACHINES—THE LOST HISTORY OF CYBERNETICS* (2016).

³ *THE ECONOMIST*, *Robotic Rules of the Road*, 12 May 2018, at 70 (71).

⁴ *Id.*

self-driving cars become widely used, traditional human-driven vehicles may need to be banned. ‘It’s too dangerous. You can’t have a person driving a two-tonne death machine,’ said Musk during an appearance at Nvidia’s annual developers conference, where he discussed Tesla’s ambitions for autonomous-cars.”⁵ Applying the same logic, one may wonder whether superior machine performance should not have a feedback effect on the level of care required from humans during the intermediate stage when humans are still allowed to drive. This too has already been suggested in the literature: “... [O]nce it becomes practical to automate, and that doing so is safer, a computer should become the ‘reasonable person’ or standard of care.”⁶

In this essay, I explore fundamental regulatory issues related to “autonomous machines”.⁷ When speaking of “regulation” I have in mind regulation by states—i.e. not self-regulation by private actors—with machines and human behavior as objects of regulation as opposed to technology as a possible regulatory tool. By “autonomous machines” I mean machines which automate or mechanize tasks previously performed by humans.⁸ I adopt an *analytical* perspective that highlights the importance of what I call the “deep normative structure” of a particular society for crucial policy choices with respect to autonomous machines.⁹ In other words: I do not attempt to develop and defend new normative principles for the regulation of autonomous machines. Rather, I seek to proceed on the basis of and relative to a “deep normative structure” that I take as a given.

⁵ THE GUARDIAN, *Elon Musk: self-driving cars could lead to ban on human drivers*, 18 March 2015, <https://www.theguardian.com/technology/2015/mar/18/elon-musk-self-driving-cars-ban-human-drivers> (last visited on June 3, 2019).

⁶ RYAN ABBOTT, *The Reasonable Computer: Disrupting the Paradigm of Tort Liability*, 86 GEO. WASH. L. REV. 1, 36 (2018).

⁷ The essay builds on a previous blog post: Horst Eidenmüller, *Machine Performance and Human Failure*, Oxford Business Law Blog, Nov. 2, 2018, <https://www.law.ox.ac.uk/business-law-blog/blog/2018/11/law-and-autonomous-systems-series-machine-performance-and-human> (last visited on June 3, 2019).

⁸ See NIKITA AGGARWAL, HORST EIDENMÜLLER, LUCA ENRIQUES, JENNIFER PAYNE & KRISTIN VAN ZWIETEN (EDS.), *AUTONOMOUS SYSTEMS AND THE LAW 1* (2019).

⁹ See HORST EIDENMÜLLER, *The Rise of Robots and the Law of Humans*, 25 ZEITSCHRIFT FÜR EUROPÄISCHES PRIVATRECHT (ZEuP) 766, 774–777 (2017).

I make two principal claims. First, a naive form of utilitarianism may indeed take us down the policy road envisaged by Elon Musk and also reflected in the statement above on the influence of machine performance on the standard of care required from humans. It is a naive form of utilitarianism because it assumes more than the utilitarian calculus can yield. It also ignores the categorical difference between machines and humans characteristic of Western legal systems. Second, taking the “deep normative structure” of Western legal systems seriously leads to policy conclusions regarding the regulation of autonomous machines that emphasize this categorical difference. Most importantly, only humans enjoy (fundamental) *human* rights, and humans are treated according to their human—not super-human—faculties. Indeed, maintaining current levels of human activities and intercourse may even require us to relax these standards in order to prevent a crowding out of the former by machine action.

I begin in Section I with a brief overview of what AI is, and what it can and cannot do, including in the realm of ethical decision-making (“machine ethics”). Against this background, the challenges of regulating autonomous machines and of potential feedback effects on the standards governing human conduct become much clearer. In Section II, I present a utilitarian dystopia which, I submit, might be the end-result of the naive utilitarianism inspiring much of the pro-machine policy rhetoric cited above. What makes this dystopia not an unrealistic prospect is the fact that utilitarian thinking in the form of simple welfare economics has a significant influence on legal policy-making, especially in the Anglo-American world. Relying on welfare economic concepts alone as guideposts for such policy-making might put us on a slippery slope which, at some point, could lead us to ignore the categorical difference between humans and machines. In Section III, I contrast this utilitarian dystopia with a richer and, I submit, more accurate account of the “deep normative structure” of Western societies. This account emphasizes the categorical difference between machines and humans, provides a convincing justification for human rights, acknowledges the failures of humans and protects them in situations of vulnerability, and it does not endeavor to make humans machine-like. I

illustrate my argument with legal applications taken from various issues in the field of contract and tort law.

I. AI AND AUTONOMOUS MACHINES

AI aims at building artificial systems that function as well as, or better than, a human—in a domain requiring intelligence. The classic assessment of whether a system functions as well as a human is the so-called “Turing test”. In this test, a human is asked to engage in a conversation with messages sent through a mechanism that does not reveal whether the party on the other side is human or not.¹⁰ If a human participant cannot distinguish the communications of an artificial system from a human, then the test is passed by that system. To pass a Turing test without any constraints around the type of conversation that could be had, the machine would need to exhibit “Artificial General Intelligence” (AGI); that is, as good as human in every dimension of intelligence. Modern AI systems do not come anywhere near AGI. Rather, the AI deployed today only has (super)human-level capability in respect of narrowly-defined functions, such as image recognition, driving vehicles in straightforward surroundings, or the classification of documents.

Over the years, various methodological/technical approaches have been pursued within AI research.¹¹ The most recent development in AI has related primarily to *machine learning* (ML). This is an approach to computing in which the solution to an optimization problem is not coded in advance, but is derived inductively by reference to data.¹² In a sense, ML turns the logic of expert systems upside down: Instead of deriving answers from rules and data, rules are developed from data and answers. The technique relies on applying computing power to very large amounts of data, the availability of which has blossomed in recent years.

¹⁰ ALAN M. TURING, *Computing Machinery and Intelligence*, 49 MIND 433, 434 (1950).

¹¹ See STUART RUSSELL & PETER NORVIG, *ARTIFICIAL INTELLIGENCE: A MODERN APPROACH* 16–28 (3rd ed. 2010).

¹² For a comprehensive treatment see, for example, ETHEM ALPAYDIN, *MACHINE LEARNING* (2016).

Progress since 2012 has largely been in a particular type of ML known as deep learning, which involves running multiple layers of representation of the data in series.¹³ A typical deep learning setup consists of an input and an output layer, with multiple hidden layers in between that lie at different levels of abstraction and are linked to each other (“artificial neural networks”). The various (hidden) layers represent different stages of abstraction of a thought process. For example, if the question is to identify at the output layer whether an image is a door, the first level would be pixels, the second would be edges, then corners, and so on up to the output layer: door (or not). The learning process of the algorithm takes place via so-called back-propagation: In the course of training the algorithm, new information is fed back from the output layer over the various hidden levels and recalibrates the settings or weights of the individual neurons with the aim of improving the accuracy of results.

The greatest practical successes with ML to date have been in the use of supervised learning techniques.¹⁴ This refers to a process that begins with a dataset that is classified or labelled by humans according to the dimension of interest (“training data”). The system analyses this dataset and determines the best way to predict the relevant outcome variable (classified by the experts) by reference to the other available features of the data. The trained model—that is, the algorithm with the set of parameters that optimized performance on the training dataset—is then put to work on a new test dataset, to see how effective it is at predicting outside the original training sample.¹⁵

¹³ See, for example, FRANÇOIS CHOLLET, DEEP LEARNING WITH PYTHON 8–11 (2018).

¹⁴ Approximately 95% of all ML applications today are based on this method: MARTIN FORD, ARCHITECTS OF INTELLIGENCE 186 (2018).

¹⁵ Two other approaches to ML that are subject of much current research activity have yet to see the same level of practical application. These are unsupervised learning and reinforcement learning, *see* ALPAYDIN, *supra* note 12, at 111–23, 125–39. Unsupervised learning relies on the model itself to identify patterns in the data. In contrast to supervised learning, this does not necessitate a labelled training data set. However, the setup of the model’s parameters becomes correspondingly more important in order to ensure that the resulting patterns are open to meaningful interpretation. In a reinforcement learning setup, the algorithm learns by trial and error. It receives a “reward” by finding the correct answer to a specified problem. A practical example is the training process relating to board games such as Go. AlphaGo by DeepMind has achieved superhuman performance within a very short time by playing very many games against itself. While reinforcement learning has much promise as a technique, its applications are currently limited to contexts in which the entire payoff structure can be specified with clarity. In most real-world settings, as opposed to game-playing, the complexity of the payoff structure

AI has also brought significant advances to the ability of machines to engage in ethical reasoning. Machines can be trained to act as if they had undertaken complex ethical considerations. In this sense, intelligent machines can be said to exhibit “moral agency”, and “machine ethics” has emerged as a new discipline at the intersection of philosophy, computing and robotics.¹⁶ The methods and technologies to teach machines to act ethically vary.¹⁷ ML based on stories of real-world human behavior in situations involving (complex) ethical decision-making is one such method.¹⁸ By contrast to such a bottom-up approach, top-down approaches seek to formalize and automate ethical rules based on theories such as utilitarianism or Kantian principles.¹⁹ Interestingly, this allows us to develop systems according to which machines apply ethical rules in a much more systematic and stable manner than humans. McEwan makes the point well when writing that “[h]umans were ethically flawed – inconsistent, emotionally labile, prone to biases, to errors in cognition, many of which were self-serving.”²⁰

AGI is not on the horizon anytime soon. The near future will be characterized by (a) “autonomous AI” (autonomous systems/machines) for certain limited functions (see above) and (b) by applications that support and improve human decision-making (“assisted AI”) or enable humans to do new things (“augmented AI”).²¹ As a consequence, we will see a lot of collaborations between (smart) humans and machines—in all kinds of professional and private domains (“semi-autonomous systems/machines”). However, for the purposes of the following discussion, I will focus on fully autonomous systems/machines. This has the advantage of more clearly revealing the fundamental policy choices

rapidly exceeds the capability of hardware for the application of reinforcement learning techniques.

¹⁶ See, for example, CATRIN MISSELHORN, *GRUNDFRAGEN DER MASCHINENETHIK* (3rd ed. 2018).

¹⁷ See, for example, SIMON PARKIN, *Teaching Robots Right from Wrong*, *The Economist* 1843, June/July 2017, <https://www.1843magazine.com/features/teaching-robots-right-from-wrong> (last visited on June 4, 2019).

¹⁸ *Id.*

¹⁹ See MISSELHORN, *supra* note 16, at 97—114.

²⁰ MCEWAN, *supra* note 1, at 86.

²¹ See THOMAS H. DAVENPORT, *THE AI ADVANTAGE: HOW TO PUT THE ARTIFICIAL INTELLIGENCE REVOLUTION TO WORK* 133—137, 190—192 (2018).

societies face with respect to regulating autonomous machines and the impact these choices have on humans and human behavior.

II. A WELFARIST DYSTOPIA

How should we regulate AI systems, especially autonomous systems? When should we, for example, allow self-driving cars to take part in regular traffic? What should be the liability regime if a self-driving car causes an accident? And should that liability regime influence the regime applicable to human drivers? Should fully autonomous cars have legal personality so that the car itself is liable?

These are important policy questions. It is clear that these questions cannot be answered without some normative conception of regulation/lawmaking. In the following Section, I hope to demonstrate that much of the discourse on these issues is heavily influenced by a welfarist conception of “the good” which has its roots in utilitarianism. It has a significant traction particularly in Anglo-American legal systems, especially in so far as commercial activities in a broad sense are concerned. At the same time, it isolates and radicalizes one element of the “deep normative structure” of Western societies, and it may lead to policy conclusions which are squarely at odds with other elements of that structure.

A. Utilitarianism and Welfarism

As an ethical theory, utilitarianism still appears to hold some appeal, also in the context of AI regulation. In a well-known textbook on machine ethics, for example, utilitarianism is discussed as the first top-down approach for teaching autonomous machines to behave morally.²² This may be because of the individualistic starting-point of utilitarianism as a behavioral theory—every individual attempts to maximize pleasure over pain—and the fact that programming machines to maximize utilities (for humans) appears to be an innocuous enough societal goal.

²² MISSELHORN, *supra* note 16, at 97—101.

At the same time, the deficiencies of utilitarianism as a *regulatory* theory are plainly obvious:²³ Utilitarianism does not tell us how to measure pleasures and pains—the utilitarian effects of a contemplated regulation—, it is not able to provide a solid foundation for fundamental human rights which we believe humans should enjoy,²⁴ and it does not care about how utilities are distributed, violating widely held views on the minimum content of a (distributively) just society; finally, utilitarianism simply does not answer the question why it is a desirable thing to maximize utility in a given society—as a normative proposition this does of course not follow from the fact that individuals engage in such a maximization exercise (assuming that they do).

For these reasons, utilitarianism never has been a serious candidate to guide legal policy-making. However, a “modern” version of utilitarianism, namely the economic analysis of law, has become just that, and for a simple reason: By substituting the utilitarian calculus with a cost/benefit assessment of the real-life consequences of legal rules and regulations, it appears to solve at least the measurement problem that plagues utilitarianism. Some economists pay lip service to a broader welfare conception that goes beyond monetizable effects,²⁵ but when it comes to practical legal-economic analyses, immaterial benefits, “psychic costs” or “moralisms” are conveniently ignored.²⁶ Fundamental human rights appear not to be a central concern at least if the scope of such analyses is restricted to commercial/corporate activities, distributive goals are claimed to be better (more efficiently) pursued by the tax and transfer system,²⁷ and the question “Why efficiency?”²⁸ is answered with a simple “Because it is

²³ See, for example, HORST EIDENMÜLLER, *EFFIZIENZ ALS RECHTSPRINZIP: MÖGLICHKEITEN UND GRENZEN DER ÖKONOMISCHEN ANALYSE DES RECHTS* 187—234 (4th ed. 2015).

²⁴ Rule utilitarianism does slightly better than act utilitarianism in this respect. However, even under a system of rule utilitarianism, the status of human rights is precarious in the sense that such rights are always *contingent* on whether they increase societal utility (or not).

²⁵ STEVEN SHAVELL, *FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW* 2 (2004).

²⁶ FRANK I. MICHELMAN, *Norms and Normativity in the Economic Theory of Law*, 62 *MINN. L. REV.* 1015, 1036 (1978); DUNCAN KENNEDY, *Cost-Benefit Analysis of Entitlement Problems: A Critique*, 33 *STANF. L. REV.* 387, 398 (1981).

²⁷ STEVEN SHAVELL, *A Note on Efficiency vs. Distributional Equity in Legal Rulemaking: Should Distributional Equity Matter Given Optimal Income Taxation?*, 71 *AM. ECON. REV. PAPERS AND PROCEEDINGS* 414 (1981).

²⁸ RONALD DWORKIN, *Why Efficiency?*, 8 *HOFSTRA L. REV.* 563 (1980).

better for everybody to live in a richer society—at least in the long run—as benefits and costs will be randomly (evenly) distributed.”²⁹

B. Welfarism and AI

Many scholars and/or practitioners who engage in the discussion on policy-making regarding autonomous systems/machines do so with little appreciation for the need to specify and defend a regulatory goal and to demonstrate why and how that goal could be achieved by some proposed regulation. Some, even though highly critical of the prospect of “superintelligence” and the “existential risks” it poses, stop short of considering legal intervention at all and confine themselves to promoting “best practices” among AI researchers—a puzzling or even contradictory position.³⁰ Others proceed on the basis of vague statements such as the following: “We’ve developed laws to incentivize and facilitate *cooperation*, so if AI can *improve* our legal and governance systems, then it can enable us to cooperate more successfully than ever before, bringing out the very *best in us*” (emphasis added).³¹ Even those who ambitiously set out to provide “... a roadmap for a new set of regulations, asking not just what the rules should be but—more importantly—who should shape them and how can they be upheld”³² end up with a hotchpot of pragmatic, utilitarian, economic, humanistic and ontological arguments that are presented in a random fashion to justify, it seems, whatever result with respect to a particular AI policy problem the author thinks is appropriate.

Hence, it is not surprising that lawmakers around the world struggle to identify some clear guideposts for AI regulation. The European Parliament is a good example for this. In a resolution adopted on February 16, 2017, “with recommendations to the Commission on Civil Law Rules on Robotics”,³³ the

²⁹ See, for example, A MITCHELL POLINSKY, *Probabilistic Compensation Criteria*, 86 Q. J. OF ECON. 407 (1972).

³⁰ NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES 218–219 (2014).

³¹ MAX TEGMARK, LIFE 3.0: BEING HUMAN IN THE AGE OF ARTIFICIAL INTELLIGENCE 105 (2017).

³² JACOB TURNER, ROBOT RULES: REGULATING ARTIFICIAL INTELLIGENCE 3 (2019).

³³ EUROPEAN PARLIAMENT, Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)),

Parliament sets out “General Principles” (at T-Y) for AI regulation from which, I submit, *any* particular outcome could be deduced. It appears that the middle ground between computer scientists and philosophers on the one hand and lawmakers on the other is not occupied. Coherent and convincing principles for AI-related lawmaking are needed, but it seems they are not yet available.

In a sense this is to be expected: AI is a (relatively) new field, certainly for lawyers. Prudent legal scholars and lawmakers would first try to understand the key issues, problems and risks before suggesting normative concepts or principles on how to deal with these. At the same time, jurisdictions worldwide have to take certain crucial policy decisions—such as, for example, whether to admit self-driving cars to practice—*now*; and for these decisions a principled approach is necessary or at least highly desirable. This is even more important as any adopted regulatory approach is likely to create path-dependencies, i.e. determine—to a certain extent—the steps that will follow in the future.

Those who do attempt to come up with coherent principles (or even a theory) for AI regulation often resort to welfare economics—either implicitly (and rather vaguely) or more explicitly. In a paper calling for a “global solution” on the issue of AI regulation, for example, the authors pose the following question: “Do these applications really make human society more *efficient, better, or safer?*”³⁴ (emphasis added). “Efficient” has a precise meaning (“better” does not), and “safer” refers to risks, i.e. costs that influence the efficiency calculus. Another scholar argues that the liability regime he proposes for autonomous machines “... would benefit the general *welfare* ...”³⁵ (emphasis added). Even more clearly, another author writes this about the appropriate “framework of liability for autonomous systems”: lawmakers should devise a framework “... with a view to maximize the net surplus for society by minimizing the costs

http://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html#title1 (last visited on June 5, 2019).

³⁴ OLIVIA J. ERDELYI & JUDY GOLDSMITH, *Regulating Artificial Intelligence: Proposal for a Global Solution*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3263992 (last visited on June 5, 2019), at p. 2.

³⁵ ABBOTT, *supra* note 6, at p. 7.

associated with personal injury and property damage.”³⁶ Finally, in a paper on “Regulating Artificial Intelligence Systems” we are reminded of the efficiency-enhancing path of the common law system and the need “... to find mechanisms for internalizing the costs associated with AI.”³⁷ Many more statements such as these could be cited.

However, one should certainly not exaggerate. AI regulation is by no means just about welfare economics or efficiency, and there is an extensive literature by scholars concerned with other fundamental issues of AI regulation such as, for example, privacy concerns and data protection or questions of “algorithmic discrimination”.³⁸ At the same time, I anticipate that the rhetoric of welfare maximization is going to become more and more important, shaping the debate about AI regulation and achieving a crucial role for AI-related policy-making. A combination of factors and considerations leads me to this assessment.

First, AI is about developing artificial systems that are at least as intelligent/rational/productive as humans and, in many instances, even more so. Hence, AI appears to increase welfare (and wealth) of societies because we humans are, with the help of AI, able to accomplish tasks more rationally, systematically and efficiently. There is a natural nexus between the nature of AI and the rhetoric of welfare economics.

Second, the debate about AI regulation is shaped to a significant degree by those who have an informational advantage with respect to the (ab)uses to which certain AI applications may be put and their real-world effects: the entrepreneurs who develop these applications in the first place (viz. Elon Musk’s statement

³⁶ GERHARD WAGNER, *Robot Liability*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3198764 (last visited on June 9, 2019).

³⁷ MATTHEW U. SCHERER, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J. OF LAW & TECH. 353, 391 (2016).

³⁸ For two seminal pieces of scholarship on these issues see, for example, FRANK PASQUALE, *THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION* (2015); CATHY O’NEIL, *WEAPONS OF MATH DESTRUCTIONS: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY* (2016).

cited in the Introduction).³⁹ Clearly these entrepreneurs are more than happy to use the rhetoric of welfare economics to help sell their products: Self-driving cars are safer, allow us to use driving time more productively, free up parking space, mobilize the elderly etc.;⁴⁰ algorithmic credit scoring provides access to credit for many who were denied credit in the past;⁴¹ and smart medical applications diagnose illnesses faster and more precisely, treat patients better and, as a consequence, allow us to lead a longer and more fulfilling life⁴²—again, this list of examples could easily be extended. Even where and to the extent negative effects of autonomous systems are discussed, the debate is framed in terms of risks, i.e. expected costs, and how these could be reduced. So AI regulation is just about benefits and costs, it seems—familiar welfare economic terrain.

Third, one of the most heavily researched areas within the economic analysis of law is risk and the regulation of risk in contract and tort.⁴³ This research has led to a highly sophisticated conceptual apparatus that appears well-suited to answer key questions of AI regulation such as: When should we allow the use of autonomous systems? Who should be liable if something goes wrong, and under what standard? Should the liability regime applicable to

³⁹ See also ALAN J. DIGNAM, *Artificial Intelligence: The Very Human Dangers of Dysfunctional Design and Autocratic Corporate Governance*,

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3382342 (last visited on June 10, 2019).

⁴⁰ See the following statement on the Tesla website: “Self-driving vehicles will play a crucial role in improving transportation *safety* and accelerating the world’s transition to a *sustainable future*. Full autonomy will enable a Tesla to be *substantially safer* than a human driver, *lower the financial cost of transportation* for those who own a car and provide *low-cost on-demand mobility* for those who do not” (emphasis added), https://www.tesla.com/de_AT/blog/all-tesla-cars-being-produced-now-have-full-self-driving-hardware (last visited on June 6, 2019).

⁴¹ See the following statement on the ZestFinance website: “ZAML® is the only way to get your ML credit models from the lab to production. *Reduce losses. Increase approvals. Fast.*” (emphasis added), <https://www.zestfinance.com/> (last visited on June 6, 2019).

⁴² See the following statement on the DeepMind Health website: “Nurses and doctors in the NHS and elsewhere across the globe simply don’t have the tools to instantly analyse each test result, determine the right treatment, and make sure that every single patient who needs complex or urgent care is escalated to the right specialist immediately. Many people think that new technology could help clinicians with *more accurate analyses*, and ultimately get *faster treatment* to the patients who need it most. We’re committed to working with those on the frontline of healthcare to build technological solutions to these problems. We work with hospitals on mobile tools and AI research to help get patients from test to treatment *as quickly and accurately as possible.*” (emphasis added), <https://deepmind.com/applied/deepmind-health/>, (last visited on June 6, 2019).

⁴³ See, for example, SHAVELL, *supra* note 25, at 175—385.

autonomous machines influence the liability standard applicable to humans and, if so, how?

What is more, in most Western legal systems technological advances have an influence on the standard of care required from humans, for example with respect to medical diagnosis and treatment. In the UK, the so-called “*Bolam* test” is one of the foundational stones of the modern law of professional negligence. The test dates back to the *Bolam* case (1957)⁴⁴ in which the court held that “... a medical man [cannot] obstinately and pig-headedly carry on with some old technique if it has been proved to be contrary to what is really substantially the whole of informed medical opinion.”⁴⁵ Hence, more efficient techniques must be used if there is a consensus on their superior properties by professionals working in the field.⁴⁶ This is “fertile ground” for the advent of smart AI-powered systems. Also, the classic Hand formula⁴⁷ for determining negligent behavior has made it—often in a revised form that considers *marginal* costs/benefits instead of absolutes—into court decisions⁴⁸ and leading textbooks and commentaries on tort in many common and civil law jurisdictions.⁴⁹ Hence, the welfare economics-inspired doctrinal apparatus to deal with autonomous systems is in place.

C. Regulatory Consequences

On this basis, a series of mostly risk-related policy-recommendations for the regulation of autonomous systems can be deduced which will, I think, shape

⁴⁴ *Bolam v. Friern Hospital Management Committee* [1957] 1 WLR 582.

⁴⁵ *Id.*

⁴⁶ See also JASON MILLAR & IAN KERR, *Delegation, relinquishment, and responsibility: The prospect of expert robots*, in: RYAN CALO, A. MICHAEL FROOMKIN & IAN KERR (eds.), *ROBOT LAW* (2016) at 102, 117: “Once there are expert robots, it will be easier to argue in some instances that they *ought* to be used to their full potential, because the evidence will suggest that in those instances they will, on average, deliver better results than human experts.”

⁴⁷ *United States et al. v. Carroll Towing Co., Inc., et al.*, 159 F.2d. 169 (2d. Cir. 1947).

⁴⁸ See, for example, *United States Fidelity & Guaranty, Co. v. Jadranska Slobodna Plovidba*, 683 F.2d. 1022 (7th Cir. 1982) (Posner, J.); German Federal Supreme Court, Case VI ZR 137/82, Judgment of 29 November 1983.

⁴⁹ See, for example, CLERK & LINDSELL ON TORTS 8-185 (22ND ED. 2018); GERHARD WAGNER, IN: FRANZ JÜRGEN SÄCKER ET AL. (EDS.), *MÜNCHENER KOMMENTAR ZUM BÜRGERLICHEN GESETZBUCH BAND 5, SCHULDRECHT BESONDERER TEIL III Vor § 823 paras. 47—57* (6th ed. 2013).

the debate about AI regulation in the years to come. Some of these recommendations are probably quite uncontroversial, others less so. Some may potentially bring about a fundamental change of our legal and political order. At the same time, it will be difficult to argue why this should not happen if one buys into the logic of welfare economics—it is a slippery slope that may well lead to unforeseen and “upsetting” consequences.

I will illustrate the analysis in the following mostly with reference to self-driving cars, for a couple of reasons: Self-driving cars will become a reality affecting our lives fairly soon—they are not science fiction; transport by autonomous vehicles has a huge economic and political importance; and the regulatory challenges posed by self-driving cars are representative of the kinds of regulatory challenges by innovative and complex autonomous systems.

Taking the issue of admitting self-driving cars to regular traffic first, the key welfare economic variable appears to be *safety*. People are concerned about potential damage caused by accidents, especially injury to humans, and so the scholars and practitioners cited in The Economist’s article in the Introduction engage in discussions and speculations on the “safety margins” we should require from autonomous cars. In the context of automobile traffic, safety can be measured by (expected) accident costs. If self-driving cars lower these costs, they are safer than human-driven cars. One reason why safety plays such a prominent role in the (economic) analysis of whether and when to allow self-driving cars on our roads is an “availability bias”: Data on accidents involving cars is readily available,⁵⁰ so this is a hard factor in any argument about traffic regulation.

Against this background, the question on the admission of autonomous cars to practice appears to have a straightforward and simple answer (see already the Introduction on this point): Autonomous cars should be allowed to operate if the expected costs from accidents are as low as or marginally lower than those from cars driven by humans. Based on a welfare economic analysis, there seems to be no reason to require autonomous cars to be (much) safer than

⁵⁰ See, for example, the various reports by the World Health Organization (WHO) on road traffic deaths and road safety, “Global Health Observatory (GHO) Data”, https://www.who.int/gho/road_safety/mortality/en/ (last visited on June 7, 2019).

humans at all—at least if we restrict the analysis to the issue of safety. At the same time, *if* autonomous cars are marginally safer, such an analysis appears to lead to the conclusion that *only* autonomous cars should be allowed to operate and that humans should be prevented from driving. Total accident costs would, it can be assumed, be lower if this policy were implemented.

This surely would be a radical measure and departure from the status quo. At the same time, market forces might drive societies to the very same end result even if humans are still able to drive alongside autonomous cars. This requires us to consider the liability regime for autonomous cars: Who should be liable if an autonomous car causes an accident, and should the liability regime be fault-based or strict? Economic analysis suggests that the appropriate liability standard with respect to autonomous cars should be strict, for two reasons: it is exceedingly difficult to precisely define the efficient level of care in this context,⁵¹ and only a strict liability regime regulates the “activity level” of the car which influences the likelihood of accidents.⁵²

The primary liability addressee of this regime should be the car manufacturer. The manufacturer is best positioned to control the risks and balance the benefits and costs of the technologies that are “driving” autonomous cars.⁵³ This is clearly so if the manufacturer develops the relevant AI applications. But even if the AI (device) producer is different from the car

⁵¹ See THE ECONOMIST, *supra* note 3, at 70: “It will take years rather than months for the industry to cohere around a standard.” On uncertainty in the finding of negligence see SHAVELL, *supra* note 25, at 224—228.

⁵² EIDENMÜLLER, *supra* note 9, at 771—773. A strict liability regime is also not necessarily bad for incentivizing innovation. Indeed, studies have shown that more liability can *increase* investments in innovations that make products safer, see ALBERTO GALASSO & HONG LUO, *Punishing Robots: Issues in the Economics of Tort Liability and Innovation in Artificial Intelligence*, in: AJAY AGRAWAL, JOSHUA GANS & AVI GOLDFARB (eds.), THE ECONOMICS OF ARTIFICIAL INTELLIGENCE: AN AGENDA (2019) at 493.

⁵³ One problematic aspect of a strict tort liability regime that holds the car manufacturer liable for accidents caused is the activity level of *owners* of autonomous cars. Just think of a taxi company on the one hand and a private car owner on the other hand—the activity levels and associated accident risks differ significantly. One can think of various potential “solutions” to this problem. One would be co-liability of owners depending on their activity profile. Another could be tying the sale of the car to liability insurance with the premium determined by (i) the manufacturer, (ii) the type of car, and (iii) the owner/user profile. Such personalized insurance is available already today, *see*, for example, <http://www.blackboxinsurance.com/> (last visited on June 9, 2019).

manufacturer, the car manufacturer controls the overall system, including all component parts. Hence, the car manufacturer probably is the “cheapest cost avoider”.⁵⁴ As far as tort liability vis-à-vis third parties is concerned, it should therefore be the only liability addressee.⁵⁵ Interestingly, this “solution” to the liability problem seems to be the one toward which the market and private contracting practice are moving. Late in 2015, for example, Volvo announced that it would take responsibility for the actions of its self-driving cars.⁵⁶

Car manufacturers would insure against this additional liability risk—maybe because of a new legal requirement to do so—,⁵⁷ and this would drive up the price for self-driving cars. Hence, those who purchase and own self-driving cars will internalize the expected accident costs of their operation. This mirrors the current situation with human-driven cars. Hence, the owners and operators of cars—either autonomous or human-driven—always bear the expected accident costs of the respective car type. And *if* autonomous cars are safer than human-driven cars, they will be cheaper to operate, which is an incentive to substitute the latter with the former.⁵⁸ It therefore does not matter that much whether the law does or does not allow humans to drive cars once we admit autonomous cars to practice. Because autonomous cars will be cheaper to operate, market

⁵⁴ If parties were able to bargain for the applicable liability rule at zero (transaction) costs, they would contract for liability of the party that is best positioned to avert the expected accident costs at the lowest costs (Coase Theorem). See RONALD COASE, *The Problem of Social Cost*, 3 J. OF LAW & ECON. 1 (1960).

⁵⁵ Of course, the car manufacturer could seek indemnity from the device producer based on their contractual relationship if a defective AI device ultimately caused an accident. Holding the car manufacturer strictly liable involves the risk of the manufacturer falling insolvent and therefore not being able to pay up. To cover this risk, manufacturers should be required by the law to purchase product liability insurance.

⁵⁶ See KIRSTEN KOROSK, *Volvo CEO: We will accept all liability when our cars are in autonomous mode*, Oct. 7, 2015, <http://fortune.com/2015/10/07/volvo-liability-self-driving-cars/> (last visited on June 7, 2019).

⁵⁷ If there is no legal requirement to buy insurance from a third party, (large) car manufacturers might also “self-insure” as they also benefit from the law of large numbers. In terms of additional costs—and a higher price for purchasers—this does not make much difference.

⁵⁸ Of course *expected liability costs* are only a fraction of *total operating costs*, and regular cars differ significantly from self-driving cars in many cost-relevant aspects. Self-driving cars will likely cost more than regular cars when first available for consumers because *scale* will be a key cost factor. Once these cars are produced and sold at significant scale, the situation might look quite different. In any event, *at the margin*, expected liability costs will certainly be a relevant factor affecting price.

participants will have an economic incentive to use them—and to quit driving themselves.

This kind of market dynamic will play out in many settings in which AI-powered machines will come to be used. For another example, think of medical treatment and liability for malpractice. In most common and civil law jurisdictions, such liability is not governed by strict liability but by a negligence standard. If human doctors, autonomous machines and human doctors assisted by smart machines are held to the exact same standard of care, superior machine performance will translate into a (expected liability) cost advantage and associated price advantage of machines or machine-assisted human doctors because these will be able to escape liability with greater certainty.⁵⁹ The effect will be even more pronounced if, following the *Bolam*-logic, human doctors are supposed to meet a higher standard of care which might be considered justified because of the superior performance of autonomous machines.

The consequences could be dramatic: a rapid crowding out of human activity in our daily lives. Superior machine performance translates into a cost advantage for autonomous machines, and this translates into a competitive advantage for AI-powered goods and services. From a welfare economic perspective, this development is as inevitable as it is welcome: If the goal is to maximize net economic welfare in a society, then surely the law should set incentives to substitute less safe conduct (and “technologies”) with safer conduct (and “technologies”). To put it differently: If machines operate more efficiently than humans, the liability system should also contribute to a process of substituting humans with machines.

Taking this analysis a step further might also yield an answer to the key policy question of granting (fully) autonomous machines legal personality and the power to hold property, conclude contracts etc. This has been suggested by several scholars,⁶⁰ and it has also already become a point of consideration for

⁵⁹ On the effect of liability costs on total costs, see *supra* note 58.

⁶⁰ See, for example, JERRY KAPLAN, HUMANS NEED NOT APPLY: A GUIDE TO WEALTH AND WORK IN THE AGE OF ARTIFICIAL INTELLIGENCE 79—92 (2015); JENS KERSTEN, *Menschen und Maschinen*, 70 JURISTENZEITUNG 1, 7 (2015).

lawmakers. In the European Parliament’s resolution of February 16, 2017, “with recommendations to the Commission on Civil Law Rules on Robotics”,⁶¹ for example, the Parliament calls on the European Commission to consider “... creating a specific legal status for robots in the long run, so that at least the most sophisticated autonomous robots could be established as having the status of electronic persons responsible for making good any damage they may cause, and possibly applying electronic personality to cases where robots make autonomous decisions or otherwise interact with third parties independently” (Principle 59 f).

A welfare economic analysis of the problem would look to the consequences of such a policy and assess its effects, i.e. it would be based on a functional account of the contemplated rule change. If autonomous machines act like humans—passing the Turing test—, and if they perform better than humans in their respective domain, according them legal personality and giving them rights (to hold property, conclude contracts etc.) not only appears to be a justifiable policy move because it would reflect the fact that autonomous machines are functionally equivalent or even superior to humans; it would also allow these machines to participate in commercial transactions, creating more and more wealth for themselves and, if appropriately taxed, also for humans—or so it could seem.⁶²

III. HUMANISM AND MACHINE REGULATION

To some, the welfarist regulatory approach to autonomous machines discussed in the previous Section might be an attractive prospect. To others it might reflect a dystopian future. But what exactly are the flaws and shortcomings of the argument presented in Section II? I submit that the welfarist regulatory approach to AI is seriously defective for three reasons: it assumes much more than the economic calculus can deliver; it turns a blind eye

⁶¹ See *supra* note 33.

⁶² For this to happen we would also need to solve what is called the “goal alignment problem” in AI research: the challenge is to determine the objective function of an application and to ensure its correspondence with desirable human values. See, for example, NICK BOSTROM, *supra* note 30, at 132—134. Further, the AI application would have to have the property of “corrigibility”: it must be able to be switched off, and it must be able to be corrected if it is doing something humans don’t like. See, for example, STUART RUSSELL, in: MARTIN FORD, *supra* note 14, at 67.

to widely held considerations of distributive justice; and it is not able to offer a convincing justification for fundamental *human* rights. This does not mean that welfare-economic analyses of AI-related regulatory problems are useless. It is the radicalization and de-contextualization of the welfarist viewpoint which creates serious problems. Central to the “deep normative structure” of Western societies is a *humanistic* approach that stresses the categorical difference between humans and machines. It protects humans where they suffer from human failures, and it seeks to avoid a crowding out of humans by machines.

A. A Critique of AI-related Welfarism

1. The Limits of the Welfarist Calculus

Analyzing the costs and benefits of regulatory measures targeting AI applications in monetary terms suggests that there exists a precise calculus that can guide policy-making on a welfare-economic basis. However, this is an illusion. The best that such an analysis can do is to highlight how certain measurable variables probably impact on the efficiency effects of a proposed regulatory measure, and often even this (modest) goal will not be reached.

Take self-driving cars again as an example. First, even if one restricts the analysis to safety and (expected) accident costs, a meaningful analysis would have to factor in the likelihood and potential damage triggered by a global cyberattack on networked cars, for example. It is obvious that, similar to the global financial crisis, such an incident could inflict extraordinary monetary and personal damage on thousands or even millions of people, far exceeding any safety benefit that self-driving cars might otherwise have. But how should we assess the likelihood and damage size of such an incident? We have no reliable data, and because we do not have data, the issue is ignored.

Second, the same will happen to other cost/benefit-relevant factors which we find difficult or even impossible to quantify. Just think of the increase in “quality” time that some will be enjoying when being driven in an autonomous

car or, on the negative side, the loss in “sheer driving pleasure”⁶³ others will bemoan if they can no longer drive themselves. Even if they can, such pleasure might be negatively affected by the number of autonomous cars on the road—how much pleasure do we get from driving if surrounded by driverless cars?⁶⁴ And what kind of price tag do we put on the “quality time” if being driven or on the loss of “driving pleasure” when not driving? Asking humans how much they would request for giving up an entitlement to “quality time” or “driving pleasure” might yield dramatically different results to asking them what they are prepared to pay to have such an entitlement.⁶⁵ There is no way to assess the (monetary) value of these effects in an “objective” manner. Hence, it is absolutely unclear how the (rising) presence of autonomous cars will affect the overall net welfare in a society—the calculus is indeterminate with respect to the regulatory question of the (safety) requirements such cars have to fulfil before we allow them to operate. It is also indeterminate with respect to the question of when we disallow humans from driving.

A similar indeterminacy will limit the usefulness for welfare economic considerations in all contexts of AI regulation in which many factors must go into the cost/benefit analysis, and we have hard (monetary) data on only a few. I fear that this will be the rule, not the exception.

2. Concerns about Distributive Justice

Welfare-economic analyses are blind to considerations of distributive justice. All that is important is aggregate welfare or maximizing the “size of the pie”. This is justified by the claim that, in the long run, everybody can expect to

⁶³ Delivering “Sheer Driving Pleasure” happens to be the key marketing slogan of one of the world’s premium automobile manufacturers, see <https://www.facebook.com/BMW/photos/sheer-driving-pleasure-/10154742975092269/> (last visited on June 8, 2019).

⁶⁴ See BMW Welcomes—Artificial Intelligence, <https://www.youtube.com/watch?v=ELEaIdy7boU>, at 1:58:40 (last visited on June 8, 2019) (Horst Eidenmüller).

⁶⁵ See KENNEDY, *supra* note 26, *passim*; MARK KELMAN, *Consumption Theory, Production Theory, and Ideology in the Coase Theorem*, 52 SOUTH. CAL. L. REV. 669 (1979); EIDENMÜLLER, *supra* note 23, at 118—133. The reason for this discrepancy are income effects and endowment effects.

benefit from a wealthier society (see Section II A *supra*). In this respect, (economic) welfarism is structurally identical to utilitarianism.

However, the argument that everybody can expect to be better off in the long run under a welfare economic policy is not convincing, for a variety of reasons.⁶⁶ First, those who have more now will likely end up having much more in the future—because of income effects in assessing cost/benefits of policy measures and because of their higher leverage to influence the political process. Second, those who argue in favor of the “compensation thesis” do not claim that everybody actually will be better off. The claim is that a random distribution of wins and losses under various policy measures will generate a positive *expected* value. However, even a positive expected value for everybody—which is unlikely, see above—is perfectly compatible with significant *actual* losses for some. Third, the long run may be much too long for many. Issues of distributive justice of policy measures must be addressed convincingly *now*.

Crucial AI-related policy issues can be used to illustrate these points. As is well-known, Big Data and AI tools enable firms to personalize offers, i.e. to calibrate them to the individual (idiosyncratic) preferences of specific consumers. This relates to features of products and services, but it also relates to price. For a long time, first-degree price discrimination, i.e. setting prices based on the preferences and reservation values of *individual* consumers, existed only in textbooks. With Big Data and smart algorithms, it has become a reality.⁶⁷

The total welfare effects of first-degree price discrimination are unclear. On the one hand, it gives certain individuals access to goods and services at (lower) prices, allowing them to buy goods and services that they were not able to buy before (at the then higher market price). On the other hand, privacy concerns may lead consumers to avoid price-discriminating firms, and rent-seeking investments by firms to capture as much value as possible and defensive

⁶⁶ See EIDENMÜLLER, *supra* note 23, at 243–251, 281–283.

⁶⁷ For an analysis, see OREN BAR-GILL, *Algorithmic Price Discrimination: When Demand is a Function of Both Preferences and (Mis)perceptions*, 86 CHICAGO L. REV. 217 (2019); GERHARD WAGNER & HORST EIDENMÜLLER, *Down by Algorithms? Siphoning Rents, Exploiting Biases, and Shaping Preferences: Regulating the Dark Side of Personalized Transactions*, 86 CHICAGO L. REV. 581, 585–592 (2019).

tactics by consumers to protect their privacy might lead to huge deadweight welfare losses.

The one thing that is clear about first degree price discrimination is that it massively redistributes transaction surplus from consumers to producers. If it works perfectly, consumers' surplus in fact shrinks to zero—consumers pay exactly their respective reservation price. At the same time, producers' surplus is maximized. Hence, it is not surprising that studies have found that consumers emphatically view individual price discrimination as unfair.⁶⁸ These fairness perceptions regarding the allocation of welfare/rents are clearly part of the “deep normative structure” of Western societies. Of course, different jurisdictions will draw the line between distributively unacceptable and acceptable consequences differently. But they will not ignore distributive questions completely—as the welfare economic perspective does.

3. Concerns about Fundamental Human Rights

Finally, in a welfare-economic policy conception for regulating autonomous systems, fundamental human rights are always *contingent* on whether or not they contribute to maximizing social welfare. They have no inherent value independent of that goal. In this respect, welfare economics suffers from the same problems as utilitarianism (see Section II A *supra*). At the same time, the contingent status of fundamental human rights is clearly at odds with widely held beliefs regarding the cornerstones of a just and free society (on normative justifications for fundamental human rights, see *infra* Section III B).

A proponent of a welfare-economic approach to regulating autonomous systems might argue that human rights concerns should not be raised in many instances of AI systems regulations, at least not if the scope of such regulations is restricted to commercial/corporate activities. This argument is clearly disingenuous. Just think of algorithmic credit scoring and the ever present

⁶⁸ See, for example, TIMOTHY J. RICHARDS, JURA LIAUKONYTE, AND NADIA A. STRELETSKAYA, *Personalized Pricing and Price Fairness*, 44 *INTERNAT. J. OF INDUSTR. ORG.* 138, 140 (2016); KELLY L. HAWS AND WILLIAM O. BEARDEN, *Dynamic Pricing and Consumer Fairness Perceptions*, 33 *J. OF CONSUMER RESEARCH* 304, 306–307 (2006).

issues of direct or indirect discrimination involved in the decision on whether to accept a credit application or not. Similar concerns will be present in almost all cases in which large amounts of (personal) data and ML applications (see *supra* Section A) are relevant for autonomous systems, and that means practically always.

Reconsidering the issue of admitting autonomous cars to our roads, for example, one might be concerned about the privacy implications of all the personal data that is needed to train the underlying ML models powering autonomous cars, plus the data that these cars will increasingly collect as part of the Internet of Things. One might also be concerned about the liberty of human drivers in the sense of autonomously choosing how to drive and interact with other drivers. Being constrained by the liberty of other humans is one thing—being constrained by machines’ actions is quite another. Clearly these concerns about liberty become much stronger still if humans are banned from driving cars.

On the other hand, context matters, and other issues of regulating autonomous cars might well be less sensitive from a human rights point of view. The privacy and liberty concerns mentioned above relate primarily to the question of whether autonomous cars should be allowed to travel on roads at all. The applicable liability regime in case defects or accidents occur does not appear to raise such fundamental concerns. Another way of putting this is to say that, for example, a Kantian probably does not hold strong views as to whether this liability regime should be strict or fault-based. If anything, he or she might be more inclined to argue for a strict liability regime as it appears to be more protective of human safety. By comparison, the marginal reduction of choice opportunities for humans—strict liability might lead to fewer autonomous cars on the road compared to a negligence regime—appears to carry less weight.

B. A Humanistic Approach

In the final Section of this essay, I would like to briefly outline a humanistic approach to the regulation of autonomous systems that better reflects the “deep normative structure” of Western societies than the welfarist conception

criticized above. I should stress (again) that my goal in this Section is analytical, not normative. I am attempting to capture what, I believe, makes up the normative fabric of our societies, and I am neither offering nor defending a new conception. Central to the following observations is the view that humans are categorically different from machines and that this difference is central to our legal systems.

1. Humans and Machines

There is no denying that AI-powered machines are becoming smarter and smarter, outperforming humans in many domains. Machines will pass the Turing test for more and more applications. At the same time, machines are able to engage in complex ethical considerations more systematically and precisely than humans (see Section I *supra*)—they exhibit “moral agency”. There is evidence that humans develop feelings towards humanoid robots and treat them like humans.⁶⁹

All this is true. At the same time, it is, I submit, missing the crucial point. We should not commit “The Android Fallacy”.⁷⁰ *Machines and humans are categorically different, and everybody knows this.* This knowledge is not so much rooted in human behavior as it is derived from the collective human knowledge of the history of humankind and of human life—how it begins and how it ends. It is also deeply ingrained in the “deep normative structure” of Western legal systems. Key aspects of this structure reflect this point.

2. Human Weaknesses and Failure

The first relates to how we treat humans and machines as legal actors. Machines may work perfectly or close to perfectly in certain settings and

⁶⁹ See, for example, KATE DARLING, *Extending legal protection to social robots: The effects of anthropomorphism, empathy, and violent behaviour towards robotic objects*, in: RYAN CALO ET AL., *supra* note 46, at 213, 226—229; RYAN CALO, *Robotics and the Lessons of Cyberlaw*, 103 CAL. L. REV. 513, 545—549 (2015) on “social valence”.

⁷⁰ NEIL M. RICHARDS & WILLIAM D. SMART, *How should the law think about robots?*, in: RYAN CALO ET AL., *supra* note 46, at 3, 18—21.

circumstances—humans do not. We are not machines and we make mistakes.⁷¹ Research in cognitive psychology has demonstrated that these mistakes are not random events. Rather, we suffer from systematic irrationalities, especially when we are “thinking fast” and follow heuristics.⁷² Humans are often weak and wish they were not. We engage in self-paternalism to strengthen those preferences we wish we had in the long run—but do not have in a moment of vulnerability or temptation.⁷³

The legal system takes humans as they are, with all their irrationalities, vulnerabilities and weaknesses. It reacts to these with an elaborate set of protective rules. Among these are disclosure duties by better informed parties, mandatory rules on substantive contract provisions—for example guarantees or liability rules—and rights to withdraw from certain contracts such as online sales or doorstep sales within a specified time-period (“withdrawal rights”)—to name just a few. The general rationale for these rules is to assist us not to engage in transactions that are potentially very harmful to us. Sophisticated AI-powered tools used by businesses have made the problem more severe: businesses systematically use these tools to exploit behavioral anomalies.⁷⁴

In tort law, humans are held to an objective standard of care.⁷⁵ But “objectivity” means no more than eliminating the personal idiosyncrasies of a particular person. The reasonable care required from a potential tortfeasor reflects what an “average” human could and would do under the circumstances. This standard punishes below-average laxity—and sets incentives to do more, better. But the law does not require us to function like machines. It is true that the law may also require us to use new technologies under certain circumstances to escape liability (see Section II B *supra*). But this does not call the categorical

⁷¹ See also NICK BOSTROM, in: MARTIN FORD, *supra* note 14, at 100: “Humans are a mess. We don’t have a particular goal from which all the other objectives we pursue are sub-goals. We have different parts of our minds that are pulling in different directions, and if you increase our hormone levels, we suddenly change those values. Humans are not stable in the same way as machines, and maybe don’t have a very clean, compact description as goal-maximizing agents.”

⁷² See, for example, DANIEL KAHNEMAN, *THINKING, FAST AND SLOW* (2011).

⁷³ See EIDENMÜLLER, *supra* note 23, at 374—385.

⁷⁴ See WAGNER & EIDENMÜLLER, *supra* note 67, at 592—597; ILLAH REZA NOURBAKHS, *ROBOT FUTURES* (2013) at 14—15 (“[p]erfectly manufactured desire”); RYAN CALO, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 1007 (2017) (“Mass Production of Bias”).

⁷⁵ See, for example, CLERK & LINDSELL, *supra* note 49, at 8-151.

difference between humans and machines into question. We do not treat humans and machines as if they were literally the same “thing”.

This categorical difference is also reflected in the standards we apply to human behavior in situations involving moral dilemmas such as accidents that will inevitably occur but will have different negative consequences for different persons depending on the action (or inaction) of a moral agent. Because machines are capable of speedier, more systematic and consistent calculations than humans, we demand more from them than from ourselves. Humans will not be held morally or legally responsible for actions or inactions in a dilemmatic situation if they have to decide in a split-second what to do and if their behavior is not considered to be completely unreasonable. We acknowledge an “all too human” reaction in a situation involving complex and difficult moral decision-making.

3. Human Rights

The second feature of the “deep normative structure” of Western societies that reflects the categorical difference between humans and machines is human rights. Currently, only humans enjoy such rights. As has been mentioned before (see Sections II A and III A 3 *supra*), human rights have a precarious status under social conceptions such as utilitarianism or welfarism which attempt to maximize a social welfare function. Indeed, one prominent philosophical justification for human rights conceives of them as “trumps” in the hands of individuals to protect themselves against majority rule.⁷⁶

Within Western legal systems, human rights have a firm and crucial status, putting them above domestic parliamentary decisions in many jurisdictions. Just think of the “Charter of Fundamental Rights of the European Union”⁷⁷ or the core international human rights treaties negotiated and

⁷⁶ See RONALD DWORKIN, *TAKING RIGHTS SERIOUSLY* (2nd ed. 1978), at 231—238, 272—278.

⁷⁷ “Charter of Fundamental Rights of the European Union”, Official Journal of the European Union, 26 October 2012, C 326/391, <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:12012P/TXT&from=DE> (last visited on June 11, 2019).

concluded under the auspices of the United Nations.⁷⁸ With respect to the regulation of autonomous systems in particular, the European “General Data Protection Regulation” (GDPR)⁷⁹ has become crucially important for data processing and privacy protection.

As a starting point under these and similar rules and regulations, fundamental human rights are enjoyed by humans, i.e. natural persons, and by humans only. The Charter of Fundamental Rights of the European Union, for example, starts out in Article 1 with stipulating that “[h]uman dignity is inviolable. It must be respected and protected.” Similarly, Article 1(1) of the GDPR stipulates that “[t]his Regulation lays down rules relating to the protection of natural persons with regard to the processing of personal data and rules relating to the free movement of personal data.”

It is true that jurisdictions world-wide also extend “human rights” protection to legal persons such as corporations. Just think of the famous (and controversial) US Supreme Court decision in *Citizens United v. FEC*,⁸⁰ granting free speech protection to corporations. More generally, Article 19(3) of the German Constitution (“Basic Law”)⁸¹ stipulates that “[t]he basic rights shall also apply to domestic legal persons to the extent that the nature of such rights permits.” However, we extend “human rights” to legal persons because our legal systems grant personhood to legal persons as vehicles for humans to exercise their fundamental human rights such as liberty and freedom of entrepreneurial activity. A corporation is a legal fiction run by humans for the benefit of humans. Even if and to the extent “self-driving corporations” become a reality,⁸² it will ultimately still be humans that devise these vehicles and benefit from them. It would be an altogether different matter to grant fully autonomous machines legal

⁷⁸ See “The Core International Human Rights Treaties”, <https://www.ohchr.org/Documents/Publications/CoreTreatiesen.pdf> (last visited on June 11, 2019).

⁷⁹ “General Data Protection Regulation”, Official Journal of the European Union, 4 May 2016, L 119/1, <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679&from=EN> (last visited on June 11, 2019).

⁸⁰ *Citizens United v. Federal Election Commission*, 558 U.S. 310 (2010).

⁸¹ See https://www.gesetze-im-internet.de/englisch_gg/englisch_gg.html#p0108 (last visited on June 11, 2019).

⁸² See JOHN ARMOUR & HORST EIDENMÜLLER, *Selbstfahrende Kapitalgesellschaften?*, 183 ZEITSCHRIFT FÜR DAS GESAMTE HANDELSRECHT UND WIRTSCHAFTSRECHT 169 (2019).

personality and “human rights” without there being any humans involved as beneficiaries or shareholders etc.

4. Distributive Justice

Finally, the “deep normative structure” of Western societies is characterized by some conception of achieving a distributively just outcome. Maximizing the size of the pie in terms of welfare is not enough. Even in capitalist societies that are closest to the welfarist manifesto, a lot of taxing and redistribution on social grounds is going on—NOZICK’S minimal state⁸³ has been implemented nowhere in the real world. As an antidote to the perceived huge problems of “global capitalism”, societies currently appear to shift towards more redistribution, not less. Prime Minister May promised a “Britain that works for everyone”⁸⁴ in a post-Brexit UK. Maybe her successor will be able to deliver on that promise.

CONCLUSION

If one wanted to sum up the features of the humanistic approach discussed above, one has a societal conception that stresses the categorical difference between humans and machines, acknowledges human weaknesses and failures and protects humans, and is characterized by fundamental human rights and by the desire to achieve some level of distributive justice. Welfaristic pursuits are constrained by these humanistic features, and the severity of these constraints differs from jurisdiction to jurisdiction.

Against this background, there is nothing wrong or problematic about, for example, requiring autonomous cars to be (much) safer than human drivers before we allow them to participate in regular traffic, and there is nothing wrong about allowing humans to drive cars even though their driving skills might fall much short of the level achievable by smart cars. There is also nothing wrong

⁸³ ROBERT NOZICK, *ANARCHY, STATE, AND UTOPIA* (1974).

⁸⁴ THERESA MAY, *Speech at Conservative Party Conference in Birmingham*, 5 October 2016, <http://www.bbc.com/news/av/uk-politics-37563510/conservative-conference-theresa-may-s-speech-in-full> (last visited on June 11, 2019).

about applying different standards of care to humans and smart machines. In fact, societies probably will (and should) consider relaxing the standards applicable to humans: applying the same standards to humans and to autonomous machines translates into a cost and price advantage of the latter and might contribute to humans being shut out of more and more domains of our daily lives—such as driving a car or just going out for a walk.

For this is the “slippery slope” of all societies which are built on foundations which reflect not only deep humanistic values but also a commitment to free markets as the main form of organizing economic activities: the jargon of welfare economics appears well-suited to analyse the chances and risks of innovative new technologies, and it is also reflected in legal doctrine on risk, responsibility and regulation. However, the welfarist narrative has an inbuilt tendency to go to extremes and shake off the humanistic constraints discussed above. What seems to be clear is that a pure welfarist conception of “the good” will tend to move a society into a direction in which autonomous systems eventually will take a prominent role—by virtue of the law.

Hence, regulating autonomous systems is a challenge that requires us to take the “deep normative structure” of our societies seriously. Our laws are an expression of the “human condition”. They reflect what we believe lies at the heart of humanity, at the heart of what it means to be human. It simply and literally would be the dehumanizing of the world if we were to treat machines like humans, even though machines may be smart—possibly even smarter than humans.⁸⁵

⁸⁵ On the threats for the human condition by Big Data, see generally CURTIS WHITE, *WE, ROBOTS: STAYING HUMAN IN THE AGE OF BIG DATA* (2015).