

2022

Two Cheers for Cyborgs

Lauren Henry Scholz
Florida State University

Follow this and additional works at: <https://ir.law.fsu.edu/articles>

Recommended Citation

Lauren Henry Scholz, *Two Cheers for Cyborgs* *University of Chicago Law Review Online* 1 (2022),
Available at: <https://ir.law.fsu.edu/articles/762>

This Article is brought to you for free and open access by Scholarship Repository. It has been accepted for inclusion in Scholarly Publications by an authorized administrator of Scholarship Repository. For more information, please contact efarrell@law.fsu.edu.

TWO CHEERS FOR CYBORGS

*Lauren Henry Scholz*¹

* * *

In *Personalized Law: Different Rules for Different People*, Professors Omri Ben-Shahar and Ariel Porat defend the desirability and justice of personalized law. Personalized law is characterized not only by the individualization of the legal commands, but also by who is doing the individualization: machines. Ben-Shahar and Porat's first, simplest definition of personalized law is as "precision law characterized by two primary features: individualization and machine-sorted information." In the ideal form of personalized law that Ben-Shahar and Porat describe, "machines will implement the will and plans of humans." These machines would be sophisticated enough not just to determine personalized outputs based on human-determined inputs but also to determine what the relevant inputs would be. Ben-Shahar and Porat acknowledge, however, developing such a personalized law is costly on the front end. Relative to traditional uniform law, "personalized law has a new, critical, component of information costs: the cost to build and, importantly, to monitor the code that promulgates the personalized commands."

Ben-Shahar and Porat refer to the book as "a work of science fiction" on its very first page, but we are not as far from being able to implement personalized law as it may seem. The contemporary practice of integrating machine learning with human contributions offers a shortcut to implementing personalized law. Collaboration between machines and humans in implementing personalized rules could bring most of the benefits of personalized law sooner rather than later. On the way to personalized law is, rather, the law of cyborgs.

Cyborgs are often portrayed in science fiction as embodied persons with fully integrated machine and biological elements, like Darth Vader. In this paper, I use a broader definition than that. By "cyborg," I mean a system that is based on fully integrated and equal cooperation between machine and human elements. This Essay compares cyborg personalized law to the fully-automated personalized law the book typically assumes. Cyborg personalized law has several advantages over personalized law, chief of which is the fact that it is technically feasible now, or will be in the near future. Does the immense potential of personalized law point toward the desirability, or even the moral imperative, of turning to cyborg personalized law in the medium term? This Essay presents cyborg personalized law not as an alternative

¹ Lauren Henry Scholz is the McConnaughay and Rissman Professor of Law at Florida State University College of Law.

to personalized law but rather a weaker form of the same concept. At minimum, cyborg personalized law could be a waypoint to fully automated personalized law—to prove the concept to skeptics and help develop the relevant technology. Or perhaps cyborg personalized law is a second-best approach if the limitations in underlying technology, political will, or information-sharing make a fully automated personalized law system unfeasible.

This Essay will proceed as follows. First, I will define what I mean by “cyborg personalized law” and will contextualize the term within the book’s framework. Then, I will provide some examples supporting the position that a cyborg system can provide most of the advantages of fully-automatic systems at much lower buy-in costs and with more limited technology requirements. Finally, I will explore the question of whether cyborg personalized law is sufficiently less precise than personalized law to limit its desirability on precision grounds. My provisional answer is probably not. I will conclude the Essay with some thoughts about how the goals of law can depend on the tools available for implementing the law.

I. Defining Cyborg Personalized Law

Personalized law is familiar in many ways. The fact that it has precedent is part of its intuitive appeal and adds to its legitimacy as a concept. Ben-Shahar and Porat acknowledge that the idea of different rules for different people exists in current law. They call earlier forms of individualized law “contextualized law.”

Ben-Shahar and Porat argue that what makes personalized law novel is how it is implemented. The main virtue of personalized law in contrast to contextualized law is that personalized law is precise in its implementation where contextualized law is “crude.” For Ben-Shahar and Porat, pre-algorithmic contextualized law is crude in at least two important ways. First, contextualized law bases different treatment on just one or a few factors, as opposed to a larger number of factors that may be relevant. Secondly, when contextualized law does incorporate a larger number of factors, it does so based on the imprecise “cognitive process” of a human judge or other trier of fact that does not “reflect[] any actual measurement.”

They also suggest that unlike most contextualized law, personalized law incorporates factors “internal” to the person, not just external factors. They write:

In old contextualized law, the precision factors characterize the external environment in which the commands operate . . . [t]hey rarely rely on internal personal differences. The law normally considers how the reasonable person—one acting responsibly in

the external circumstances—should be treated, but ignores interpersonal differences in characteristics, preferences, or experience.

I take this to be another way of framing the greater crudity of contextualized law relative to personalized law. Machine learning and big data allow us to make better approximations of internal characteristics than we were able to previously. But ultimately, even the machines will be extrapolating what they call “internal” characteristics from external stimuli. For example, the book discusses the inputs a machine may choose to use in personalizing the age of capacity to purchase liquor. One of them is “impulsivity.” A person’s impulsivity cannot be objectively measured in the way one might measure a person’s height or weight. A machine that uses impulsivity as an input is taking measurements of a variety of factors external to that person and then synthesizing them to form a conclusion regarding her impulsivity. The ability to assess the subjective characteristics is dependent on the precise assessment and analysis of the objective world, and the reason why Ben-Shahar and Porat trust that assessment is due to the superior processing ability of machines.

Ben-Shahar and Porat contend that in law, contextualization is the exception to the general rule that laws tend to be uniform in their application. The book suggests that one important reason for this exceptionalism is contextualism’s crude methods for individualizing rules. On their account, contextualized law is right in principle but flawed due to its implementation. Personalized law solves the implementation problem through incorporating machines.

Ben-Shahar and Porat acknowledge that “[p]ersonalized rules are a type of contextualized rules” but highlight that personalized rules’ manner of implementation makes a serious difference in terms of their normative justification and desirability versus earlier contextualized law. Their model of personalized law works as follows:

General rules of law would be enacted by stating the goals they seek to promote and the constraints they have to meet. Specific personalized commands that best meet these goals and constraints would then be derived for each person. The tailoring is performed with the aid of Big Data and algorithmic processing, to identify the relevant features by which people vary and which justify different treatment.

What I call “cyborg personalized law” falls within this general model of personalized law. Cyborg personalized law is distinctive because instead of machines acting alone to implement policy goals, in cyborg personalized law, humans and machines collaborate to achieve

implementation. I use the word “cyborg” to refer to a fully and seamlessly integrated human-machine system, which is completely reliant on both human and machine elements. There is no need for a cyborg system, as I use it, to be physically embodied as one visually coextensive entity.

Definitions of cyborgs that require no visible physical embodiment are commonplace. For example, Benjamin Wittes and Jane Chong suggested that in many senses the average human in the developed world is sufficiently reliant on smartphones and other technology that she could be considered a cyborg, or is at least on her way to becoming one. Their analysis took as its starting point a comment by Chief Justice Roberts in *Riley v. California* (2014) which used a cyborg metaphor to describe the relationship Americans have with handheld technology: “[M]odern cell phones . . . are now such a pervasive and insistent part of daily life that the proverbial visitor from Mars might conclude they were an important feature of human anatomy.”

What distinguishes a cyborg system from a situation where humans use technology to achieve their goals is (1) the essential role of the machine and (2) the lack of hierarchy between human and machine. The machine element is not secondary to the human element of the system; they are both of equal status. Like an embodied cyborg, a cyborg system cannot function without the machine element, and both human and machine are part of the identity of the system. Ultimately, that is why I have chosen the word “cyborg” to have a visual representation of the fundamentally integrated nature of human and machine that this system represents. And to illustrate that the way from a rule implemented by man to a rule implemented by machine may not need to be direct. Laws implemented by cyborgs could be a waypoint to governing by machine. Or perhaps cyborgs are stronger than either a human or a machine alone.

II. The Benefits of Cyborg Personalized Law

The rise of machine learning and big data has led to a proliferation of cyborg systems in the private sector. Breathless articles about the potential replacement of the workforce in the early 2000s and 2010s have given way to analysis about how big data is used to control and regulate workers.² These practices are turning human workers into elements of a cyborg system. Examples of these practices abound in industries from manufacturing to service industry to information technology industry to law practice. For example, Karen Levy’s

² IFEOMA ANJUNWA, *THE QUANTIFIED WORKER: LAW AND TECHNOLOGY IN THE MODERN WORKPLACE* (Cambridge University Press, forthcoming 2022).

extensive work on the trucking industry has shown that given the level of monitoring and machine learning intervention that occurs, truck driving more closely resembles a cyborg system than simply a human driving a vehicle based on her own skills and intuitions.³

A study by McKinsey Quarterly showed that technological feasibility is not the only factor businesses consider when they decide to delegate a task to a machine. Five factors they identify as relevant to whether a process is delegated to a machine include: technical feasibility; costs to automate; the relative scarcity, skill, and cost of workers who might otherwise do the activity; benefit (e.g., superior performance) of automation beyond labor-cost substitution; and regulatory and social-acceptance considerations.

Thought leaders in business now see machine-human collaboration as the future rather than complete delegation to machines. The results of a study reported in Harvard Business Review are instructive:

[M]any companies have used AI to automate processes, but those that deploy it mainly to displace employees will see only short-term productivity gains. In our research involving 1,500 companies, we found that firms achieve the most significant performance improvements when humans and machines work together. Through such collaborative intelligence, humans and AI actively enhance each other's complementary strengths: the leadership, teamwork, creativity, and social skills of the former, and the speed, scalability, and quantitative capabilities of the latter. What comes naturally to people (making a joke, for example) can be tricky for machines, and what's straightforward for machines (analyzing gigabytes of data) remains virtually impossible for humans. Business requires both kinds of capabilities.

It may surprise some readers to learn that machines are distinctly worse at some processes than humans, in potentially fundamental ways. Much of the argument for fully automated personalized law comes from a latent greater trust that some have in machines than humans. Yet that trust is not always well placed. As Rebecca Crootoof has observed, "human beings and machine systems process information and reach conclusions in fundamentally different ways, with AI being particularly ill-suited for the rule application and value balancing often required of human judges." What's more, it is very expensive to create systems for the tasks that machines tend to be bad at, and there still tend to be

³ KAREN LEVY, DATA DRIVEN: TRUCKERS AND THE NEW WORKPLACE SURVEILLANCE (Princeton University Press, forthcoming 2022).

errors. At least while the technology is developing, it is dramatically less expensive and more effective to simply use humans to perform such tasks.

The promise of cyborg personalized law can be summed up thusly: cyborg personalized law can get us to functional personalized law sooner than fully automated personalized law; it can help us get around some of the technical limitations, expense of development, and even the hesitations from a cultural-mores perspective that plague fully automated personalized law; and cyborg systems may produce better results than either humans alone or machines alone.

III. The Precision of Cyborg Personalized Law

While there are many benefits to cyborg personalized law, there are also limitations of cyborg personalized law relative to fully automated personalized law. There is the added difficulty of the need for human components of the system to engage and communicate with the machine components. Though compatibility costs are hardly unique to machine-human interactions, they can also occur with machine-machine interactions, as anyone who has dealt with interoperability issues can report. The more fundamental issue is that any system is only as strong as its weakest components. There is a sense in which human decision-making is less precise than machine decision-making, as discussed *supra* in Part II. So, depending on what components are delegated to the human component of the system, there could be a loss of precision in a cyborg system. A threshold observation is, of course, there is a possibility that some cyborg systems end up creating superior fit between goals for law and output legal rules. So greater precision for fully-automated personalized law over cyborg personalized law is not a foregone conclusion—at least if we care about what it is the personalized system is being precise about. But let's assume there is some slight loss to precision from the inclusion of human elements. Since cyborg personalized law is distinctly cheaper and easier to implement than fully-automated personalized law, the precision disparity would need to be quite large to rule out cyborg personalized law out of hand as an unacceptable form of personalized law.

When constructing a cyborg personalized law, we would need be mindful that there are actually substantial gains to precision relative to traditional customized law. After all, greater precision is the primary distinct feature of personalized law as a category. As Ben-Shahar and Porat write, “[m]ore precise commands, which differentiate behavior along more relevant factors and context, generally result in better advancement of the goals of the law.” This provokes the larger question about what we want the law to do. At the point when we know, or suspect, that personalizing law will do a better job at vindicating the

rights we value in a liberal society, does the law not have an obligation to pursue even reasonable second-best approaches, such as cyborg personalized law?

The book makes some gestures at how to get from the world we live in to a world with personalized law. Ben-Shahar and Porat suggest some ways to transition some personalized law into our legal system. One is “gradual personalization,” or sorting into fewer and more human-processable categories. Another is allowing personalized law to “germinate” in substantive areas particularly suited to it, such as: “where the benefits of differentiation are large, where the goals of the law are widely accepted, where the data about the relevant interpersonal differences are reliable, and where the distributive effects of non-uniformity are desirable.”

I would submit cyborg personalized law as another avenue alongside these of introducing personalized law. My previous discussion in Part II has shown that the private sector method of incorporating machines into work product is allocating parts of processes to machines and teaching humans how to work with the machines. This reflects historical practice. For example, the accounting industry initially resisted electronic spreadsheets, thinking they would make accountants obsolete, but instead the technology simply improved what could be achieved by accountants while altering the relevant skill set to do the job.

Crootof has expressed a further danger of cyborg personalized law: it risks all the downsides of fully-automated personalized law, with additional risks. If humans trust machines too much, cyborg systems lose the full advantage of the human contribution over fully automated systems. As Crootof writes:

If the human being trusts the system too much, to the extent that they endorse an algorithm’s conclusion in the face of contradictory evidence or an obviously unfair result, the human in the loop isn’t performing their needed role. Such overtrust—which is related to “automation bias”—risks both overreliance and skill fade. If we wish to elicit the benefits of human reasoning, teaming systems must be designed so that the human in the loop understands the AI program’s capabilities and limitations, has reason to exercise valued human skills, and is actively engaged in the decisionmaking process. Absent this, the person may become little more than a figurehead (or scapegoat, should something go wrong).

When humans and machines work together, humans have the tendency to blame the human involved rather than the machine, even if it can be shown the error was on the part of the machine.

So, while cyborg personalized law gets us to many of the benefits of personalized law faster, it may not ward away all of personalized law's risks. In fact, it may amplify certain risks, in particular the worrying tendency to over-rely on algorithmic results and lose the ability to fix issues that arise with the machine's result. An early Star Trek: The Next Generation episode entitled "When the Bough Breaks" examined these very issues. In it, an alien civilization eliminates labor so the entire society can focus on arts and other forms of leisure, as the busywork of society is handled by a machine, called the Custodian. The Custodian has managed the basic functions of society for thousands of years, and nobody living knows how it works. This alien world is essentially a utopia, but it comes on the radar of our heroes because they have turned to kidnapping children. Due to a computer error that the alien civilization didn't understand how to correct, radiation has led to the infertility of everyone on the planet, and to prevent themselves from dying out, the aliens had to betray the values the society held dear, turning to kidnapping children from other societies.

This hypothetical illustrates two dangers of relying on machines, either in full, or in part, with unquestioning trust and reliance on them. First, if the skills to understand the machine dissipate, there is no way to correct the machine when it is either malfunctioning or yielding results that its makers would disagree with. Second, the machine's functioning only reflects the policy judgement of society at the time when the program's priorities were set. If mores change or facts on the ground change in ways the system was not trained to anticipate, the machine will be unable to correctly make rules and solve problems. The second problem builds on the first. If humans do not understand how to update and engage with machine-learning output, they will not be able to understand or correct issues arising from moral evolution or dramatic factual change.

Raising the more realistic possibility of cyborg personalized law pushes us to question what exactly seems so utopian about fully-automated personalized law. It is not simply the greater precision provided by purely automated system that appeals to our intuitions of justice, it is the ability to factor out a human being. As Ben-Shahar and Porat observe, it may be easier to free an algorithm of bias, once the issue is observed, than it is to handle a human who is making biased choices.

Skepticism of human involvement is what underlies greater skepticism of cyborg personalized law than fully-automated

personalized law. This provokes the question of whether cyborg systems are escapable at all. There is a sense in which even the most machine-forward variants of personalized law from the book are fundamentally integrated with humans. As Ben-Shahar and Porat assure readers, the ultimate goals of the law that machines will implement even in fully-automated personalized law will come from humans:

Personalized law would require a degree of clarity and forethought in setting the objective of any law. No longer could lawmakers fudge this determination, deferring the explicit reconciliation of the law's competing goals to judges, or inviting enforcers to tease out the goals by subsequent refined inquiry. Any ambivalence, crudeness, or uncertainty over the law's goals and the costs associated with deviations from the goals would disrupt the personalization algorithm, or leave too much power in the hands of those writing the code. Moreover, lawmakers would not be able to merely state several cumulative goals of a statute; an exact weighing of their relative importance would instead be necessary.

So even though machines are doing the process of implementing values on their own, they ultimately are not deciding what values they are implementing. At that point, we can see even what I have been describing as fully-automated personalized law is inescapably both of human and of machine. Humans are always in the loop in an underlying way. Can we really get humans out of the loop at all, or are we just fooling ourselves, or worse, obscuring and legitimating human choices under the cloak of automation?

This Essay has been an extension of the proposal outlined in the book. The book presented a self-consciously "ideal version of personalized law." It posits a world where we have the technology, the political will, and the cooperation of all actors who have relevant information. Yet we do not have to live in such an ideal world to have some law that achieves some of the main advantages of personalized law.

Cyborg personalized law has many of the advantages of personalized law while being both more feasible and less worrying for those who have concerns about the complete removal of a human role in in rulemaking. It provides more precision than traditional contextualized law, while reducing the buy-in expenses relative to a fully-automated system by having people contribute skills that are too expensive or not currently possible for machines to do. Or, contribute at moral choke points that a given society may not yet be comfortable fully delegating to a machine.

Personalized law pushes us to confront the role that technical limitations place on the law we have. Advocacy for many doctrinal choices is grounded in the assumption that the capacity limitations of human judges will lead to delay and uncertainty. For just two examples, consider textualism in contract interpretation and the preference for rules over standards in legal analysis generally. Personalized law could explode the assumptions about capacity limitations of legal decision-makers. With the assistance of machine learning and big data, taking many factors into account would not need to lead to increased costs or uncertainty. At that point, the legal profession would be free to weigh the substantive merits of taking more factors into account.

Personalized law promises an approach to lawmaking and legal analysis that is focused on getting the tough moral questions right—that is, the part of the work machines cannot do for us. When algorithms can make even complex multi-factor decisions consistent and relatively inexpensive, there is no hiding from the hard choices that humans must make as to what factors the machine should consider in implementing the law.

* * *

Lauren Henry Scholz is the McConnaughay and Rissman Professor of Law at Florida State University College of Law.