3-4-2022

Moving Toward Personalized Law

Cary Coglianese
University of Pennsylvania Carey Law School

MOVING TOWARD PERSONALIZED LAW

Cary Coglianese*

Part 121 of Title 14 of the Code of Federal Regulations provides rules for operating commercial air transportation services. Section 121.383(e) of this part states, without equivocation, that "[n]o pilot may serve as a pilot in operations under this part if that person has reached his or her 65th birthday."[1]

The purpose behind this age restriction should be intuitive. Among other things, the risks of sudden debilitating health incidents, like heart attacks, increase with age, which thereby increases the safety risks to passengers on commercial aircraft. Yet, nothing magically transformative occurs to pilots on the very day of their sixty-fifth birthday that makes the risks they pose to the flying public different than they were on the day before. Moreover, older pilots' experience might well enable them to handle aircraft more safely than younger pilots. It was, after all, the flying experience that came with his advanced age that allowed pilot Chesley "Sully" Sullenberger in 2009 to land his U.S. Airways jetliner safely on the Hudson River after bird strikes had knocked out power to both of the plane's engines.[2]

But rules are, after all, rules. They draw lines. Just as with the age limitations on airline pilots, rules make generalizations that aim to serve lawmakers' goals in the aggregate, that are ineffectual, perverse, or counterproductive.[3] That is the concern that motivates Professors Omri Ben-Shahar and Ariel Porat's book, *Personalized Law: Different Rules for Different People*. Their book starts with the important recognition of the limitations of general rules and then paints a hopeful vision for a future that overcomes these limitations through algorithmic systems that lead to individualized determinations based on the goals of the law.

---

* Cary Coglianese is the Edward B. Shils Professor of Law and Director of the Penn Program on Regulation at the University of Pennsylvania. This essay is forthcoming in the *University of Chicago Law Review Online*.

[1] 14 C.F.R. § 121.383(e).

[2] Sullenberger was, at the time, about a week shy of his fifty-eighth birthday—which was itself near what had then been the longstanding maximum age limit for commercial pilots of sixty. Back then, federal regulations specified a limitation of age sixty on commercial pilots (although a statute adopted two years earlier had raised the limit to age sixty-five). After Sullenberger's successful landing, the U.S. Department of Transportation amended its regulations to reflect the extended age limit. The U.S. Governmental Accountability Office has failed to find evidence to indicate that safety risks to the flying public have increased with the extension of the retirement age. U.S. GOVERNMENT ACCOUNTABILITY OFFICE, AVIATION SAFETY: INFORMATION ON THE SAFETY EFFECTS OF MODIFYING THE AGE STANDARD FOR COMMERCIAL PILOTS (2009), *online at* https://www.gao.gov/products/gao-10-107r (visited Dec. 5, 2021).

[3] Frederick Schauer & Richard Zeckhauser, *Regulation by Generalization*, 1 REG. & GOV. 68 (2007).

Ben-Shahar and Porat's vision of a world with fewer general rules and more individualized determinations is increasingly possible to view as more than just science fiction. Advances in predictive analytic tools—that is, machine learning or artificial intelligence—are starting to allow more precise, individualized determinations in a variety of other settings, such as medicine and marketing. Indeed, these same analytic tools and digital technologies underlie the auto-pilot systems that control much of the operation of commercial jetliners today.

With *Personalized Law*, Ben-Shahar and Porat offer important commentary on a future in which the law could rely more on predictive analytics to make personalized judgments about legal obligations. Acknowledging that under current law "commercial airline pilots must undergo frequent health screenings starting at age forty and retire at sixty-five," Ben-Shahar and Porat note that "[i]n a personalized law regime, rather than establishing a one-for-all age threshold, these controls would be personalized, relying also on additional factors."[4] Their vision responds to a series of longstanding complaints about law and regulation: rules are blunt instruments; they can be over- and under-inclusive; and the world is diverse, making one-size rules not always fit all.[5]

These criticisms have generated reactions from lawmakers and regulators who have responded to the pull of the particular, even absent the use of advanced technologies. To take account of relevant particularities, rules are sometimes made complex so that they can fit better the complexities found in the world. With this kind of bespoke regulation, or obligation-gerrymandering, rules weave to and fro, including some entities or circumstances under their mandates, while carving others out.[6] Rules may also be designed to provide greater flexibility to accommodate differences by imposing obligations to achieve or avoid outcomes, instead of taking specific actions, thereby allowing for all individuals to select those actions that work best for them to meet the performance designated in the rule.[7] In addition, ex

---

[4] OMRI BEN-SHAHAR & ARIEL PORAT, PERSONALIZED LAW: DIFFERENT RULES FOR DIFFERENT PEOPLE at 110-111 (Oxford University Press 2021).

[5] EUGENE BARDACH & ROBERT A. KAGAN, GOING BY THE BOOK: THE PROBLEM OF REGULATORY UNREASONABLENESS (Routledge 1982); PHILIP HOWARD, DEATH OF COMMON SENSE: HOW LAW IS SUFFOCATING AMERICA (Random House 1994). For valuable conceptual analysis, see Colin S. Diver, *The Optimal Precision of Administrative Rules*, 93 YALE L. J. 65 (1983); FREDERICK SCHAUER, PROFILES, PROBABILITIES, AND STEREOTYPES (Harvard University Press 2003).

[6] Cary Coglianese, Gabriel Scheffler, & Daniel E. Walters, *Unrules*, 73 STAN L REV 885, 892 (2021).

[7] *See, e.g.,* Cary Coglianese & Lori S. Bennear, *Flexible Approaches to Environmental Regulation*, in MICHAEL KRAFT AND SHELDON KAMIENIECKI, EDS., THE OXFORD HANDBOOK OF U.S. ENVIRONMENTAL POLICY (Oxford University Press 2012).

post obligation-alleviation mechanisms—waivers, exemptions, and variances—allow for the tailoring of off-the-shelf obligations.

It is not hard to find examples of this particularity impulse in current law. Federal regulations, for example, call upon public school systems to establish an "individualized education program" for students with learning disabilities, thus creating a set of customized rules for the delivery of these students' educational services.[8] In the environmental regulatory context, officials customize the emissions limits contained in pollution permits that are individually negotiated between state officials and industrial facilities. Federal law specifically provides for variation in the water pollution limits contained in these permits when "factors relating to the discharger's facilities, equipment, [and] processes . . . are fundamentally different from the factors considered by EPA in development of the national limits."[9] Finally, in what might be considered the regulatory equivalent of couture, regulators in some countries apply what is known as safety case regulation to the approval of large industrial facilities, requiring owners to submit individualized plans for operating these facilities safely based on their distinct characteristics and circumstances.[10] Once approved by the regulator, these individualized plans become binding law imposed on the individual facility.[11]

Examples like these notwithstanding, much law remains in a form akin to Part 121 of Title 14 of the Code of Federal Regulations: namely, rules that apply across the board, such as age limitations. With technological advances, it is no longer fanciful to consider moving toward a system in which individual determinations replace general rules. Surely, that would lead to more optimal and just outcomes. For example, in the case of "an individual applying for a governmental license—say, a prospective pilot asking the Federal Aviation Administration (FAA) to grant her flight certification," instead of relying on blunt rules about age, flight training, and the like, "the FAA might instead rely on a machine-learning algorithm to improve the process of determining when a pilot's license should be granted."[12] The algorithm could comb "through all of the available data about the applicant—say, school records, medical records, social media postings, and fine-grained data from the flight recorders from

---

[8] 20 U.S.C. § 1401(14).

[9] 5 U.S.C. § 125.30.

[10] *See* Prerna Jain, Anne M. Reese, Dushyant Chaudhari, Ray A. Mentzer & M. Sam Mannan, *Regulatory Approaches - Safety Case vs US Approach: Is There a Best Solution Today?,* 46 J. LOSS PREVENTION IN THE PROCESS INDUS. 154, 155–57 (2017).

[11] In the United Kingdom, for example, offshore oil and gas facilities subject to safety case regulation "must ensure that the procedures and arrangements described in the current safety case which may affect the health and safety of persons or the environment are followed." UK HEALTH AND SAFETY EXECUTIVE, THE OFFSHORE INSTALLATIONS (OFFSHORE SAFETY DIRECTIVE) (SAFETY CASE ETC.) REGULATIONS 2015: GUIDANCE ON REGULATIONS 72 (2015), https://www.hse.gov.uk/pUbns/priced/l154.pdf.

[12] Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1, 10 (2019).

previous training flights flown by the applicant."[13] Then, the government would grant the "license when the machine-learning algorithm forecasts the applicant's risk to be below a specified threshold."[14]

Signs of movement in this direction can already be found, as cities, states, and federal agencies start experimenting with the use of machine-learning tools. These tools are already being used to help support important governing tasks that previously had been driven either by general rules or human judgment.[15] In the automated traffic light systems in Los Angeles and Pittsburgh, machine learning is being used to establish, quite literally, the rules of the road.[16]

But achieving a general and widespread system of legal decisions made through machine learning remains at least decades off. Legal change of any kind is slow and difficult to bring about, with organized interests that benefit from the status quo working to resist reform. Among the various technological, organizational, and political hurdles that stand in the way of a personalized system of law, I discuss here three potential challenges, each flagged by Ben-Shahar and Porat to varying degrees in their book. First, to work well, personalized law seems to necessitate a complete specification of the goals currently behind legal rules. Second, personalized law will necessitate a sufficient level of social consensus around the precise specification of law's goals and how different values should be traded off against each other. Finally, to achieve the full promise of individualized justice that delivers optimal outcomes for society, personalized law will need to remain current with changing values and new feeds of data. As shorthand, I refer to these as challenges of completeness, consensus, and currency.

After briefly elaborating each of these challenges, I offer two solutions: custom and competence. By custom, I mean that—as a sociological matter—the public could generally accept the use of personalized law, even absent full resolution of the challenges of completeness, consensus, and currency. By competence, I mean that responsible efforts to achieve pragmatic improvements in the law should help bring about public acceptance of personalized law. Both

---

[13] *Id*.

[14] *Id*.

[15] Cary Coglianese & Lavi Ben Dor, *AI in Adjudication and Administration*, 86 BROOKLYN L. REV. 791 (2021), DAVID ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY, & MARIANO-FLORENTINO CUELLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES (2020), *archived at* https://perma.cc/7PCS-PZRZ.

[16] Ian Lovett, *To Fight Gridlock, Los Angeles Synchronizes Every Red Light*, N.Y. TIMES (Apr. 1, 2013), online at https://www.nytimes.com/2013/04/02/us/to-fight-gridlock-los-angeles-synchronizes-every-red-light.html (visited Dec. 5, 2021); G. WANHOO LEE, IBM CTR. FOR THE BUS. OF GOV'T, CREATING PUBLIC VALUE USING THE AI-DRIVEN INTERNET OF THINGS (2021), online at https://www.businessofgovernment.org/sites/default/files/Creating%20Public%20Value%20using%20the%20AI-Driven%20Internet%20of%20Things.pdf.

custom and competence might well "solve" the completeness, consensus, and currency challenges even if we concede that these challenges can never be fully overcome.

In other words, it is possible to acknowledge that personalized law will never be perfect. Yet acknowledgement of its imperfection need not undermine its appeal. After all, today's legal system is not perfect either. It may turn out that personalized law's imperfections will prove acceptable to the public in a society increasingly dependent on digital tools, especially if personalized law performs better than the status quo that relies primarily on general rules and human judgment.[17]

* * *

Completeness. The completeness challenge for personalized law arises because it is goal-driven law.[18] It establishes particularized legal obligations for individuals based on what an algorithm determines would best achieve the relevant goals given each individual's characteristics and other factors. In principle, this is an eminently sensible approach to law, as by definition it eliminates the over- and under-inclusiveness of general rules. The rule's goals become the law of the algorithm itself, so compliance with the law necessarily advances the relevant societal goals. The challenge, though, lies with identifying a complete set of goals that should be operationalized in the algorithm underlying personalized law.

Let us return to Part 121 of Title 14. The age limitation for commercial pilots is not the goal itself, but merely a proxy for a set of goals related to the safety of air transportation. As Ben-Shahar and Porat observe, these limitations do not arise because society cares about age per se. Rather, these limitations "are concerned primarily with unreasonable risks."[19] Instead of relying on age, the government could shift to a goal-based approach that would instead subject each pilot to an individualized screening and identify the health risks for each, allowing individuals whose risk is below a specified level to qualify to fly commercial aircraft regardless of their age. Although age is often thought of as a proxy for the risks of sudden heart attacks, it is not a precise predictor. As a result, it would be better to rely on a more data-rich analytic tool—an algorithm—to make a more accurate prediction of heart failure. For that reason, peer-reviewed science recommends that pilots should be screened using "risk estimating calculators"[20] or "individualized risk scoring using an

---

[17] Cary Coglianese & Katelyn Hefter, *From Negative to Positive AI Rights*, WILLIAM & MARY BILL OF RTS. J. (forthcoming).

[18] Ben-Shahar and Porat correctly observe that "[p]ersonalized law requires lawmakers to be explicit about the goals of the law." Ben-Shahar & Porat, *supra* note 4, at 138.

[19] Ben-Shahar & Porat, *supra* note 4, at 106.

[20] Ries Simons & Rene Maire, *Extending the Age Limit of Commercial Pilots?*, 41 EUROPEAN HEART JOURNAL 2239, 2242 (2021).

appropriately calibrated prediction algorithm"[21]—with "risk" in the extant studies equating to the risk of a heart attack.

If lawmakers were to accept these recommendations and substitute an individualized forecast of cardiovascular disease for a general age limitation, the resulting test for pilot licenses would provide a better prediction of heart attacks—but it would not completely capture the goals underlying the current age limitation. Although instances of pilots becoming suddenly incapacitated are extremely rare, no more than half of such incapacitations are due to cardiac incidents. This means that other types of health crises, some also related to age, present sudden incapacitation risks that would need to be included in any individualized risk calculation. Moreover, as pilots age, they can suffer other physical limitations that can affect their ability to fly safely, such as the degradation of sensory perceptions, reflexes, and reaction times, or certain cognitive impairments. None of these would be captured in a risk calculator focused on cardiovascular health. A complete algorithm would need to take in the panoply of health risks that might affect safe piloting. It would also, of course, need to take into account other qualities—such as experience and skill—that could improve safety with age.

It may seem obvious that a complete specification of the goals behind an age requirement would need to include more than just heart attack risk. But in practice, the complete set of goals is not always so obvious for all problems. If we look to the history of regulation, it is not hard to find examples of the incomplete specification of goals. Goal-based regulation, after all, is not unique to personalized law—it even dates to the Code of Hammurabi.[22] But not infrequently, lawmakers specify a relatively simple, single-dimensional goal, based on a salient problem at hand, only later to discover that they overlooked other relevant goals. A few examples:

- Since at least the passage of the Poison Prevention Packaging Act in 1970, federal regulators sought to reduce the number of childhood poisonings from the ingestion of medicines and household cleaners.[23] They set goal-based standards that required companies to create product packaging that would be unable to be opened by most children. After child-resistant packages complying with these standards entered the marketplace, it became evident that many adults also could not easily open the packages. Once

[21] Hans Bauer, Dennis Nowak & Britta Herbig, *Aging and Cardiometabolic Risk in European HEMS Pilots: An Assessment of Occupational Old-Age Limits as a Regulatory Risk Management Strategy*, 38 RISK ANALYSIS 1332, 1344 (2018) (footnote omitted).

[22] Cary Coglianese, Jennifer Nash, & Todd Olmstead, *Performance-Based Regulation: Prospects and Limitations in Health, Safety, and Environmental Regulation*, 55 ADMIN. L. REV. 705, 706 (2003). *See also* Cary Coglianese & Jennifer Nash, *The Law of the Test: Performance-Based Regulation and Diesel Emissions Control*, 34 YALE J. REG. 33 (2017).

[23] Poison Prevention Packaging Act of 1970, 15 U.S.C. ch. 39A § 1471 et seq.

adults managed to open these packages, they tended to leave them open or transferred them to non-resistant containers—creating an increased risk of childhood poisoning. Regulators eventually modified the regulations to incorporate a twin goal: keep children from opening packages, while allowing adults to open them.[24]

- Automobile safety regulations governing air bags are based on the physical force imposed on crash test dummies. Initially, the dummies specified in these tests were sized to the average adult male. Only later did it become tragically evident that air bags designed to meet these tests deployed too forcefully against individuals smaller than the average male—killing occupants who would otherwise not have been harmed in low-speed crashes. Eventually, regulators specified a more multi-faceted regulatory test that necessitated so-called smart air bags that could deploy differently for different occupants.[25]

- A shift to a performance-based building code in New Zealand gave builders flexibility to make highly personalized choices about construction materials as long as the code's required structural strength goals were met. Only after a vast swath of new housing throughout the country began to experience serious mold problems did regulators eventually realize that the building code should have included mildew resistance goals as well as ones for structural strength.[26]

Experiences such as these provide a note of caution for anyone designing personalized law. Personalized law's completeness challenge calls for breaking free of tunnel-vision tendencies and ensuring a fully specified set of goals can be incorporated into law's algorithms.[27]

Consensus. Identifying a complete set of goals animating a law is only the first challenge. If multiple values are to be incorporated into an algorithm that generates personalized obligations, then weights and tradeoffs between these values will need to be articulated with mathematical precision. As Ben-Shahar and Porat point out, a system of "algorithmic directives would require fixing the objectives of the law in advance with precision, so as to translate it into code."[28] Yet, as Cass Sunstein observed years ago, social consensus over

---

[24] Cary Coglianese, *The Limits of Performance-Based Regulation*, 50 U. MICH. J. L. REF. 525 (2017).
[25] *Id*.
[26] *Id*.
[27] Justice Stephen Breyer has provided an instructive account of regulators succumbing to tunnel vision. STEPHEN BREYER, BREAKING THE VICIOUS CIRCLE: TOWARD EFFECTIVE RISK REGULATION (Harvard University Press 1995).
[28] Ben-Shahar & Porat, *supra* note 4, at 233.

values can be hard to come by.[29] Due to difficulties in achieving agreement, the law today is riddled with what Sunstein has called "incompletely theorized agreements." Such agreements manifest in the spongy language that populates much of law today, through words such as reasonable, undue, feasible, and adequate.[30]

Ben-Shahar and Porat point out that, under personalized law, legislators could no longer get by with such "loose" formulations[31] and instead would need to agree to "an exact weighing of their relative importance."[32] There can be, as Ben-Shahar and Porat put it, "[n]o fudging."[33]

This requirement is not difficult to meet when values are simple, clear, and unidimensional—or when the consequences of the use of personalized automation is de minimus. No societal consensus need be sought around the precise values to be used, for example, in the machine-learning algorithms that automatically read and sort postal mail.

The law can similarly rely on machine learning whenever disparate effects or values can be easily converted into a common metric and a well-accepted function for ranking against this metric exists. For example, the goal behind licensing commercial airline pilots might in principle be specified in terms of a common, single metric of accident risk. And a simple function of "minimize accident risk" could presumably be applied. In such a case, as long as the necessary data exist, an algorithm might account for how a pilot's age can both increase risks due to health limitations and decrease risks due to greater experience.[34] These risk estimates could be combined, and the algorithm could take age and other individualized characteristics and data into account to determine who becomes authorized to fly commercial aircraft.

But the relevant calculus will often not be so clear. If an algorithm were used to establish individualized emissions limits for the polluting sources affecting air quality in the Grand Canyon region, the objective function would presumably need to consider, and make tradeoffs between, a variety of relevant values: human health, visibility in the national park, costs to firms and customers, impacts on the reliability of the electricity grid, and so forth. Unfortunately, little consensus exists over the precise weighting and relationships between these values. The same can be said of many legal and policy issues in a pluralistic society.[35] Some people even see the competing

---

[29] Cass R. Sunstein, *Incompletely Theorized Agreements*, 108 HARV. L. REV. 1733, 1735 (1995).
[30] *Id*.
[31] Ben-Shahar & Porat, *supra* note 4, at 235.
[32] *Id*. at 36.
[33] *Id*. at 233.
[34] *See generally* JOHN GRAHAM & JONATHAN WIENER, RISK VS. RISK: TRADEOFFS IN PROTECTING HEALTH AND THE ENVIRONMENT (1995).
[35] Moreover, the aggregation of preferences about these values could prove unstable. Kenneth J. Arrow, *A Difficulty in the Concept of Social Welfare*, 58 J. POL. ECON. 328 (1950).

values implicated by major legal questions as completely incommensurable and thus conceptually incapable of being combined or balanced.

Throughout the law, the lack of clear social consensus will, as Ben-Shahar and Porat observe, "make some lawmaking more difficult [in an era of personalized law] and even impossible."[36] Ultimately, the absence of clear consensus over values makes it evident that "[t]he necessity of well-specified goals is a major hurdle for personalized law."[37]

Currency. A system of personalized law will also need a steady supply of data and a process for adapting to changes in individual and environmental conditions. I consider this a "currency" challenge in two senses of the word: data will constitute the medium or currency on which personalized law depends, and these data need to be kept up to date.

The first currency challenge is foundational. Data are what drive personalizing algorithms, and when data are not available, algorithmically based personalized law is not possible. Personalized law will hold its greatest promise with legal matters that involve repeated incidents or scenarios than with truly sui generis situations. It is no surprise that Ben-Shahar and Porat use driving and traffic laws as examples throughout their book.[38] Driving is a highly repeated activity for which the outcomes of principal concern—accidents—are both frequent and well-documented. The large amount of data on automobile driving and accidents is precisely why it is one of the most empirically studied areas of regulation.[39]

Yet the legal system also operates in domains or over issues for which data are not available. Consider the age of commercial airline pilots. It is completely conceivable to shift to a system that relies on algorithmic forecasts of risk for commercial pilots in an age range of, say, 30 to 65 years. But it would be much harder, if not impossible, to conceive of a system that could produce sufficiently reliable forecasts for commercial pilots over the age of 65. Why? Current rules do not permit pilots to fly commercial aircraft if they are over the age of 65, which means no data exist on commercial pilots over this age. Even if one were to look for data on noncommercial pilots, there will be much fewer data available on pilots the older they get, simply because there are fewer older pilots overall. Without sufficient data, there can be no personalized law.

The second sense of currency refers to the need for the data and algorithms underlying personalized law to remain current or up to date. Much of what makes machine learning marvelous stems from its dynamic nature.

---

[36] Ben-Shahar & Porat, *supra* note 4, at 235.

[37] *Id*. at 37.

[38] *Id*. at 1, 4, 6, 8-9, 19, 32, 62-67, 82, 106-07, 109-10, 125-26, 128, 147, 175-76, 183, 188, 193-94, 201, 203, 207, 209, 215, 234-35, 237.

[39] John Mendeloff, Evaluation in Two Safety Regulatory Agencies (2004), *archived at* https://perma.cc/DQ5U-9RBQ.

These algorithms are often updated on a regular basis (sometimes constantly), and their accuracy improves as they are fed more data over time. Their forecasts grow brittle and stale if the data upon which these algorithms train grow out of date and conditions change. Similarly, if the objective functions in these algorithms no longer reflect current needs or social values, then this too will lead a system of personalized law to generate outmoded results.

A personalized system of licensing commercial airline pilots, for example, might need new data or analysis when the technology in the airplane itself changes.[40] Perhaps new automated flying systems can compensate for some of the perceptual limitations that come with age, which might make data from a period before the adoption of these systems less applicable. On the other hand, the automated systems in the cockpits of commercial aircraft may come with more buttons, dials, and indicators that demand greater or different cognitive skills which are possessed more frequently by younger pilots. For these reasons, an algorithmic system to license pilots would likely need to re-train its algorithm regularly, and the outcomes generated by the system would vary as conditions in the world (and in the data) change. As change is endemic in the world, a personalized law system would need data to be refreshed and algorithms revisited regularly, as Ben-Shahar and Porat acknowledge, both to avoid "stale objectives and accommodate new circumstances and priorities."[41]

Consider one recent "new circumstance" relevant to a personalized system of driving law. Between 2019 and 2020, miles traveled by car in the United States dropped by about 13%, while automobile fatalities increased by about 7%—the greatest annual number of fatalities in over a decade.[42] The precise cause of this seeming anomaly can still only be speculated, but it might include: competency deficits from less frequent driving during the COVID-19 pandemic; drivers driving at higher speeds due to less congestion during the pandemic; or an increase in road rage due to heightened stress levels during the pandemic. The upshot for a personalized system of traffic laws is that, if commands were based on algorithms trained on data from a normal year, the system could generate individualized speed limits that result in too much risk under the changed circumstances.

* * *

It should be clear that despite personalized law's alluring promise to overcome the inefficiencies, frustrations, and injustices of general rules, the

[40] One of the main issues with the recent fatal crashes of two Boeing 737-Max airplanes stemmed from how their automated systems pushed the nose of the airplane down. Some pilots knew how to respond, but the pilots of the ill-fated planes did not.

[41] Ben-Shahar & Porat, *supra* note 4, at 36.

[42] NHTSA Traffic Safety Facts: Early Estimates, https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813115.

three challenges I have just outlined—*completeness*, *consensus*, and *currency*—will stand in the way of personalized law fulfilling its aspirations completely.

Ben-Shahar and Porat are thus right to suggest that a wholesale transformation to personalized law will not occur any time soon. Rather, if personalized law is to take root, it will be best to "start with small steps," as they recommend.[43] Such small steps may already be underway, reflected in the various ways that governments are already digitizing their operations and on occasion relying on algorithms to support various adjudicatory and administrative functions.[44] Some countries, such as Denmark, Korea, and Estonia, have placed a high priority on digitizing and automating many governmental functions—establishing some of the necessary building blocks for making a shift to personalized law.[45]

A shift to personalized law need not—and almost surely should not—be made all at once. Ben-Shahar and Porat advise that change should proceed incrementally, with personalized law adopted "in areas 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."[46] This is indeed how a transition to personalized law is likely to occur: incrementally, and perhaps also only partially. Making any kind of non-incremental change is, after all, exceedingly hard.[47] More importantly, because the challenges I have outlined above may well never be fully met, incremental change may serve to acclimate the public to an imperfect (but hopefully better) personalized system of law.

Nevertheless, despite the inherent limitations, a large-scale shift to personalized law may well eventually still occur in the decades to come. If it does, it will be in part because of intelligent efforts to design algorithmic systems that are defined in terms of well-considered sets of goals, with objective functions that incorporate weights and tradeoffs that seek to approximate prevailing moral sentiments in society, and with as much data as can be assembled and updated feasibly. It will also be in part because of scholarship like that of Ben-Shahar and Porat's, along with the work of many others (perhaps especially those more critical of the use of artificial intelligence in the law). Critical scholarship may make the designers of these systems better informed and more thoughtful. But even so, personalized law will never be perfect. Its success in winning and deserving public support will derive, I suspect, from two principal sources: custom and competence.

Custom. As the public grows increasingly accustomed to personalized, algorithmic decisions by private actors in other spheres of their life, they will

---

[43] Ben-Shahar & Porat, *supra* note 4, at 240.
[44] *See* note 15 and accompanying text.
[45] U.N., UNITED NATIONS E-GOVERNMENT SURVEY 2020, at 12 (2020).
[46] Ben-Shahar & Porat, *supra* note 4, at 240.
[47] Charles E. Lindblom, *The Science of "Muddling Through,"* 19 PUB. ADMIN. REV. 79 (1959).

likely become more accustomed to personalized law—indeed, they may well come to demand it. Already, we have seen this kind of acclimation effect with the so-called first stage of e-government: the building of websites and online service tools for most government agencies. After all, as it has become possible to order groceries and buy cars online, not to mention conduct any number of other social and economic activities, it has become exceedingly hard for government offices not to allow members of the public to pay their taxes or renew their driver's licenses online too. Much the same could be expected in years to come of more advanced digital systems. For example, customers at eBay who resolve a problem through the company's automated dispute resolution tool are reportedly more likely to return to eBay to do more business than those customers who have no disputes at all.[48] Perhaps a shift to fully automated court proceedings will not be far off.

Similarly, although personalized law may well suffer from some degree of incompleteness, or even seeming arbitrariness, when it comes to how algorithms are designed, the public already tolerates a good bit of incompleteness and incoherence in the current impersonal, rule-based legal system. The popularity of Philip Howard's critique of modern rule-based governance, The Death of Common Sense (which stayed on the New York Times bestseller list for months when it first came out in the 1990s), confirms that the public is not blind to the perversities of the current system.[49] The bar that personalized law would need to meet is far from a high one.

The courts already tolerate a fair degree of looseness in the joints. The rational basis test under the Constitution hardly demands any strong degree of rationality. Statutes only need to make some surface-level sense. Moreover, although the courts review agency-created rules under a standard that purports to ferret out arbitrary and capricious decisions, the reality is that judicial review under the arbitrary and capricious standard is often quite deferential—and arguably somewhat arbitrary itself. The U.S. Environmental Protection Agency, for example, can pass up on tightening standards that would save human lives in favor of tightening standards that avoid minor illnesses.[50] It is far from clear that the courts would insist on total value completeness and precision for personalized law algorithms, as long as the government had some sensible-sounding rationales for their design.[51]

---

[48] BENJAMIN H. BARTON & STEPHANOS BIBAS, REBOOTING JUSTICE: MORE TECHNOLOGY, FEWER LAWYERS, AND THE FUTURE OF LAW 113 (2017).

[49] HOWARD, *supra* note 5.

[50] In the 1990s, for example, judges upheld federal environmental regulators' decision to decline to tighten particulate matter standards in a way that would have prevented additional fatalities, while they simultaneously tightened ozone standards even though doing so yielded no mortality-related benefits. Cary Coglianese & Gary Marchant, *Shifting Sands: The Limits of Science in Setting Risk Standards*, 152 U. PA. L. REV. 1255, 1321-23 (2004).

[51] Indeed, by the logic of the Supreme Court's decision in Baltimore Gas & Electric Co. v. NRDC, 462 U.S. 87 (1983), where the Court said that judges should be most deferential to

Competence. The key is likely to be that any system of personalized law will need to work well. If it delivers on its promise of achieving outcomes that better advance the goals of the law than the current rule-based system, that is likely to be enough. When credit card companies use machine learning to identify fraudulent transactions, for example, they can do so with remarkable speed and accuracy. Credit card customers surely have few qualms about these personalized tools because they work so well. Much the same could be expected if personalized law delivers outcomes widely perceived as accurate, fair, and sensible.

Given the imperfections of the status quo, personalized law hardly needs to be perfect; it just needs to be better.[52] Can it do better? Ben-Shahar and Porat put forward a highly plausible case that it can. Not only can personalized law reduce the over- and under-inclusiveness of the current rule-based system, but it holds out the prospect of treating people more humanely by recognizing their individuality and better meeting their needs. If personalized law does deliver on the promises that Ben-Shahar and Porat articulate, it is surely an approach to embrace. As David Lehr and I have noted still earlier, "many aspects of public administration could undoubtedly benefit from the application of machine-learning algorithms."[53]

This is not to say that success is guaranteed. Although personalized law depends on mathematical functions and digital technologies that hold theoretical advantages, a system of personalized law will ultimately be constructed by humans. And just as humans can make mistakes in establishing rules, they can—and surely will—make mistakes in developing the information systems upon which personalized law can be delivered. Examples of past failures with governmental information systems are not hard to come by. The most prominent was the fiasco experienced in rolling out HealthCare.gov following the adoption of the Affordable Care Act. Similar recent failures in facial recognition software, automated forensic tools, and welfare fraud systems can also be cited.[54]

Yet with technology development of all kinds, failures are to be expected early in the innovation process. Although DNA analysis has today become an evidentiary gold standard, the early adoption of DNA analysis in criminal law cases, it bears remembering, was met with suspicion and disapproval by some courts.[55] We should not necessarily conclude that personalized law cannot succeed merely because governments have sometimes

---

agencies when they are working at the frontiers of science, governmental use of machine-learning algorithms might merit more deferential review by the courts. *See* Coglianese & Lehr, *supra* note 12.

[52] Cary Coglianese & Alicia Lai, *Algorithm vs. Algorithm*, 72 DUKE L. J. 1281 (2022).

[53] Cary Coglianese and David Lehr, *Regulating by Robot: Administrative Decision-Making in the Machine-Learning Era*, 105 GEORGETOWN L. J. 1147, 1152-53 (2017).

[54] Coglianese & Lai, *supra* note 52.

[55] Coglianese & Hefter, *supra* note 17.

failed in their deployment of information technology. The failures are much more noticeable than the successes—and there are indeed successes to trumpet. Consider the following:

- Greek border officials used a machine-learning system to screen travelers for COVID-19 infections. By using a digital algorithm, officials identified about two to four times as many asymptomatic travelers compared with traditional screening protocols.[56]

- Only a small fraction of industrial facilities subject to federal water pollution regulations in the United States can be inspected by regulators in any given year. With the aid of a machine-learning algorithm, environmental regulators can find more than seven times the number of violators while conducting the same number of inspections.[57]

- Defendants who are released from jail pose a risk of committing crimes while they are out on bail. If the decisions of a human judge were substituted with a machine-learning algorithm, such post-bail crimes could be reduced by 25 percent while maintaining the same rate of decisions to grant bail. Such an improvement is possible while at the same time reducing racial disparities in jailing rates.[58]

When a system of personalized law comes to deliver results like these more consistently throughout the legal system, and when that demonstrated competence is combined with greater familiarity with and public acclimation to algorithmic personalization in other spheres of life, public acceptance of even imperfect personalized law seems quite probable, if not even reasonably assured.

* * *

Ben-Shahar and Porat's book, *Personalized Law*, articulates a vision for a more individualized, effective, and just legal system that is based on data-driven, fine-tuned legal obligations. It also articulates the challenges that would confront a legal system transformed to deliver obligations personally rather than in a wholesale fashion through general rules. Whether society achieves Ben-Shahar and Porat's hopeful vision will depend on overcoming personalized law's challenges—but not necessarily on overcoming them completely. A system of personalized law need not be perfect to be normatively appealing. All that personalized law must be is better than the imperfect system in place today.

---

[56] Hamsa Bastani, et al., *Efficient and Targeted COVID-19 Border Testing Via Reinforcement Learning*, 599 NATURE 108 (2021).
[57] Miyuki Hino, Elinor Benami, and Nina Brooks, *Enhancing Environmental Monitoring Through Machine Learning,* 1 NATURE SUSTAINABILITY 583, 583-584 (2018).
[58] Jon Kleinberg, et al., *Human Decisions and Machine Predictions*, 133 Q. J. ECON. 237 (2017).