Law Society of Ontario | Barreau de l'Ontario

TAB 2A

AI for Litigators

Self-Driving Laws

The Death of Rules and Standards

Will Robot Judges Change
Litigation and Settlement Outcomes? A First Look
At the Algorithmic Replication of Prior Cases

Professor Tony Casey
University of Chicago Law School

Anthony Niblett, Faculty of Law
University of Toronto

April 16, 2021

SELF-DRIVING LAWS

Anthony J. Casey and Anthony Niblett[1]

Abstract

Machines refine and improve products. Artificially intelligent machines will soon have the same effect on the law. Future developments in artificial intelligence and machine learning will dramatically reduce the costs currently associated with rules and standards. Extending this insight, we predict a world of precisely tailored laws ("micro-directives") that specify exactly what is permissible in every unique situation.

These micro-directives will be largely automated. If the state of the world changes, or if the objective of the law is changed, the law will instantly update. The law will become "self-driving."

The evolutionary path toward self-driving laws will be piecemeal and incremental. At first, machine-driven algorithms will merely be used to guide humans; but, over time, law will increasingly reflect principles and prescriptions developed by machines.

We explore three extensions. First, we examine the possibility that the technology is not merely used to provide information about the law, but is used as means of command by the state. Second, we ask how these technological changes will affect contracting behaviour. Third, we examine the effect of micro-directives on social norms.

---

1 Anthony J. Casey is an Assistant Professor of Law and Mark Claster Mamolen Teaching Scholar at The University of Chicago Law School; Anthony Niblett is an Assistant Professor and Canada Research Chair in Law, Economic, & Innovation at the University of Toronto, Faculty of Law. We wish to thank Matt Levine at Bloomberg View whose discussions of our previous work inspired the title.

INTRODUCTION

Machines refine and improve products. Artificially intelligent machines will soon have the same effect on the law. In this paper, we ask how future developments in artificial intelligence, machine learning, and big data will affect the production of law and the structure of law. We predict that these new advancements will fundamentally change the way we, as a society, choose to govern behaviour.

We have elsewhere argued that these technological advancements will lead to the death of rules and standards.[2] To be more precise, we argued that these advancements will lead to the death of the costs of rules and to the death of the costs of standards. Rules are clear, but they are static and rigid. Rules are likely both over- and under-inclusive. They can be improved by taking into account particular facts of a situation. Standards are flexible and allow the lawmaker to take into account the circumstances of a specific case. But, standards are judged after the citizen has acted, giving rise to legal uncertainty at the time of action. Further, subjective biases of human judges may generate legal inconsistency. Standards can be improved by alleviating the uncertainty and informing citizens how to comply with objectively stated laws.

We envision a world where lawmakers use machines to refine the law, improving on both rules and standards. Ultimately, law will exist in a catalogue of precisely tailored directives, specifying exactly what is permissible in every unique situation. In this world, when a citizen faces a legal decision, she is informed of exactly how to comply with every *relevant* law before she acts. The citizen does not have to weigh the reasonableness of her actions; nor, does she have to search for the content of a law. She follows a simple directive that is optimized for her situation. We call these refined laws, "micro-directives."

These micro-directives will be largely automated. If the state of the world changes, or if the objective of the law is changed, the vast array of micro-directives will instantly update. These laws will be better calibrated, more precise and more consistent. The law will become, for all intents and purposes, "self-driving."

---

[2] Anthony J. Casey & Anthony Niblett, 'The Death of Rules and Standards' (2017) 92 Ind. L. Rev. (forthcoming).

In Part 1, we outline how micro-directives will be used to govern behaviour in the future. In Part 2, we explore the likely evolution towards a world of micro-directives. We examine how reductions in the cost of information will first change human behaviour and then how the law, over time, will become machine-produced directives. In Part 3, we explore three extensions. We explore the possibility of automatic penalties, ask how technology will change contracting behaviour, and examine the effect of micro-directives on social norms. A final part concludes.

1. AUTOMATED MICRO-DIRECTIVES

1.1 The reduced cost of information

At the heart of our thesis is *information.* The cost of information drives the lawmaker's choice between using a rule and using a standard. Where citizen behaviour is frequent and predictable, rules (such as speed limits) are preferable because lawmakers have the ex ante necessary information to regulate behaviour. Where citizen behaviour is infrequent and heterogeneous, standards (such as reasonable care) are preferable because lawmakers can ex post take into account additional information to determine whether the behaviour complied with or violated the law.

Technological advances will result in a dramatic reduction in the cost of acquiring and using information. Such technologies will allow lawmakers to better predict outcomes and human behaviour. As the differences in information costs fall away, the distinction between rules and standards will erode. The lawmaker's decision between rules and standards will become unnecessary. A new form of law, the micro-directive, will emerge. The micro-directive provides ex ante behavioral prescriptions finely tailored to every possible scenario.

Micro-directives update automatically. If relevant circumstances change, the micro-directive changes. No longer will we need rigid rules; the law adapts to the new environment. But citizens will no longer have to operate in a world of legal uncertainty, waiting for a judge to determine whether the behaviour was reasonable. Citizens will be informed immediately of what is permissible and what is not.

The technological changes will allow the law to be more precise, better calibrated, more flexible, more consistent, and less biased. The machine-driven algorithms will allow the law to become self-driving. In the same way that self-driving cars anticipate the changes in surrounding circumstances and provide the optimal response, we hypothesize that laws will take specific circumstances into account and provide a tailored statement of what is permissible.

1.2 *Using predictive technology to make law*

Consider how improvements in prediction – *predictive technology* – will foster the rise of micro-directives. Innovations in big data and artificial intelligence will make it increasingly easy to predict outcomes. The costs of collecting, storing, processing, and analyzing data will fall. New machine learning techniques outperform traditional regression approaches to prediction.[3] Algorithms based on these approaches, using big data, will form the backbone of precise and finely calibrated laws.

Citizens, armed with increased power of forecasting, will use predictive technology to assess whether their behaviour complies with the law. Corporate directors can use predictive technology to assess whether or not their actions will violate their fiduciary duties to shareholders. Uber drivers can use predictive technology to determine whether or not they are independent contractors or employees for tax purposes. Corporate entities can assess whether a proposed merger will violate antitrust laws. Here, compliance information is precise and tailored to each citizen's particular circumstance.

To some, this is not "the law," but rather a description of the law at one point in time. Critics may be concerned that if these algorithms are seen as the law, then the algorithms will not change with different states of the world, nor take into account special or unforeseeable circumstances. But any machine-produced law can be re-calibrated to take into account new circumstances.

A further concern is that an algorithm will simply entrench biases in the law. But predictive technology will not just be used to inform citizens of the existing state of the law. The technology will also be used to change the contours of the law, improving precision and consistency.

---

[3] See, e.g., Jon Kleinberg, Jens Ludwig, Sendil Mullainathan, & Ziad Obermeyer, 'Prediction Policy Problems' (2015) 105 Am. Econ. Rev. 491.

Take, for example, how predictive technology will be used to decide whether or not to grant bail to a defendant accused of a crime. Currently, a human judge must weigh up many factors, including the seriousness of the alleged crime, whether the defendant has jumped bail before, and the defendant's social and family ties. Based on the information about this particular defendant, the judge must assess whether the defendant will skip bail. The decisions of human judges have been shown to be inconsistent across different judges and infused with racial bias.[4]

Society can improve upon this situation by using analytics of big data and machine learning technology. We have millions of observations about how criminal defendants actually behave once they are granted bail. Why would we ignore this information? Predictive algorithms give a much more precise and accurate answer as to whether the defendant will skip bail. Not only are these algorithms more accurate than human judges, they are also more objective, more consistent, and less prone to bias.

Predictive technologies will fill gaps in the law. Micro-directives will be available for every hypothetical situation, eradicating the grey area of law. Justice Cardozo, in a famous contracts case, contended that the dividing line between an important and a trivial omission resulting in a breach of a condition "cannot be settled by a formula" and "precise boundaries are impossible."[5] But, in the near future, predictive technologies will be used to discern these boundaries. Deep learning technology will find hidden connections in the law, elucidating principles that do – and, more importantly, should – underpin the law.

Importantly, the new machine-learning techniques will update and adapt to new situations. These models absorb new information and factor in new circumstances. In order to better calibrate predictions, evolutionary algorithms in machine learning operate using principles similar to those used in randomized trials in medicine. Micro-directives, based on these predictive algorithms, will update automatically as the state of the world changes. Laws will update automatically. We move toward a world of self-driving laws.

---

[4] *See* Shaila Dewan, *Judges Replacing Conjecture with Formula for Bail*, N.Y. Times (June 26, 2015), Available at: http://www.nytimes.com/2015/06/27/us/turning-the-granting-of-bail-into-a-science.html.

[5] Jacob & Youngs, Inc. v. Kent, 129 N.E. 889, 892 (1921, N.Y.).

Human policy makers will still play a crucial role. Just as self-driving cars will determine the safest and fastest route to a destination selected by humans, self-driving laws will determine the optimal way to achieve a policy objective chosen by humans. Even though the micro-directives are automated and update in real time, human lawmakers will be required to set the broad objectives of the law. These broad objectives may look like a standard, but the predictive technology will take the objective and engineer a vast catalogue of context-specific directives for every situation.[6]

1.3 Using communication technology to better inform citizens

Simply having a better calibrated and automatically updating law is not enough, however, for micro-directives to flourish. These new laws must also be *accessible* to citizens. Imagine a "rulebook" for doctors that contained micro-directives covering every possible scenario. This rulebook would be enormously detailed, but unwieldy. The cost of complying with such detailed rules would be exorbitant.

This is where *communication technology* comes in. The cost of communicating specific information that updates in real time continues to fall dramatically. This technology will be able to identify which specific micro-directive applies to a particular situation and inform the regulated actor how to comply with the law.

Regulators will be able to provide instantaneous information about the legality of proposed actions. For example, let's say that an individual wishes to know whether she is an employee or an independent contractor. Under the current system, the individual may ask the regulator for an advance tax ruling, by providing all information to the regulator.[7] But this process can take weeks or even months. In the near future, predictive and communication technologies will enable these "rulings" to be provided within seconds.[8]

---

[6] For more on the role that human policymakers will play in this new system, *see* Casey & Niblett, 'The Death of Rules and Standards', supra note 2, Part I.D.

[7] *See, e.g.,* Canada Revenue Agency website: http://www.cra-arc.gc.ca/tx/hm/xplnd/rlng-eng.html.

[8] This mechanism of immediate and definitive responses to what are now considered grey areas of law, especially in the field of tax, is a commonly feature of what Benjamin Alarie calls the "legal singularity". *See* Benjamin Alarie, 'The Path of Law: Toward Legal Singularity' (2016) UTLJ.

These advancements in prediction and communication will be reinforced by other technological advancements in fact gathering and verification. As machines get better at gathering and verifying facts, more and more data will be generated and analyzed. The predictive power will be further enhanced. These fact-gathering technologies will also improve the precision of the communicated micro-directive. The micro-directives will be better tailored as the lawmaking machines absorb more information about particular scenarios.

2. THE EVOLUTION TOWARDS SELF-DRIVING LAWS

2.1 Incremental change

The death of rules and standards will be piecemeal and incremental. An analogy can be drawn to the evolution of self-driving vehicles. Vehicles will not suddenly shift one day from completely human operated to completely self-driving. The evolution is progressing, incrementally. Many aspects of self-driving vehicles are already standard features in new models. These features include self-parking, lane keeping, automatic braking, adaptive cruise control, and accident avoidance.

The technology, at first, simply provided drivers with information. In the 1990s, for example, technology provided drivers with warnings that they were too close to other parked cars. As this technology became standard, newer models provided self-parking technology. Similarly, technology has been introduced warning a driver that she is not keeping to her lane. Soon, the driver will be presented with the option of using technology to automatically stay in the lane. Over time, with increasing acceptance, vehicles will become entirely self-driving.

We predict that the evolution of the law toward micro-directives will follow a similar pattern. At first, technology will be used to provide general information to citizens. Then, with increasing acceptance from citizens and lawmakers, the predictions will become the law. We provide three examples of how we expect the evolution to play out.

First, consider the example of judges granting bail. A computer-driven algorithm to predict the likelihood of a defendant skipping bail is already being used in some jurisdictions in the United States. But this algorithm has not completely replaced human judges, yet. The transformation will take time. The algorithm is currently used to provide human judges with a better forecast of the risk of flight. Soon, we imagine, the algorithm will provide recommendations as to how the judge should decide. These recommendations could be followed or ignored by the human judge. But as more information is generated, and the evolutionary algorithm updates and becomes a better forecaster, we imagine that judges will increasingly rely on advice of the algorithm. Over time, with increased acceptance, the algorithm will *become* the law. The algorithm will effectively replace the judge.

Second, consider how the law of medical malpractice will begin to mirror predictive machine-driven algorithms. Initially, these predictive algorithms will simply provide information, perhaps outlining the likelihood of adverse outcomes if a particular action is taken. But, over time, the machines will provide recommendations on how to best proceed or warnings on how not to proceed. As these recommendations and warnings become increasingly accurate, and doctors increasingly rely on predictive algorithms to guide their practice, the algorithms will become enshrined in the law. In the same way that it would be negligent for a doctor to ignore an x-ray today, it will become negligent to ignore the advice of the machine. Over time, the algorithm will *become* the law of medical malpractice.

Third, consider how regulators may use the technology to provide the law directly to regulated actors. A tax regulator could, for example, use machine-learning programs to automatically process questions of tax residency. Predictive programs would analyze how judges have resolved these questions in the past, and would allow the regulator to process questions asked of them by taxpayers. As the regulator becomes more confident in the automated responses, the technology will be made available to taxpayers directly. Taxpayers receive instantaneous legal advice about their affairs. Again, over time, the algorithm *becomes* the law.

These three examples illustrate the incremental nature of the evolution of the law away from rules and standards and towards automated micro-directives. The speed of the change will depend on the type of law. The evolution will likely be fastest where the costs both of legal uncertainty and of poorly calibrated laws are high. Such costs are likely greatest in commercial fields such as tax, corporate law, securities, and antitrust. The pressure to automate laws will be strongest in these spheres. The push toward automation will also be greater where data is already abundant (e.g., granting bail) and where the law is more inherently stable. The principles underpinning the law of whether a worker is an independent contractor or an employee for tax purposes have remained relatively stable in Canada for years; but, determining whether particular laws violate the protection of freedom of expression section 2 of the *Charter* may, initially, prove more difficult for a machine-driven predictive algorithm.

2.2 Human skepticism

The incremental and piecemeal nature of the evolution toward micro-directives is not simply a matter of feasibility. While some predictive algorithms may take some time to update and improve, there are other barriers.

Humans are skeptical creatures. In the same way that driverless cars are frightening to some, the idea of automated machine-produced law is also terrifying. How can we trust the machines to get the "right" answer? How can we trust an algorithm to deliver a law that is just?

Throughout history, humans have held a deep distrust of automated technology. When automated elevators were first introduced, they too were scary. Elevators had, for years, been "driven" by human operators to guide them to the right level. When elevators with automatic stopping were invented in 1900, some people refused to ride them. Automatic elevators were truly terrifying. How can you trust an automated machine to lift you hundreds of feet above the ground in a tiny metal box? Automatic elevators did not become standard until after the Second World War because of this skepticism. Today, few in the developed world today are frightened of automatic elevators. It took time, but we overcame our skepticism of the technology. As laws become increasingly automated, we believe the skepticism to machine-produced law will also fade away.

But can a machine actually do the tasks currently performed by legislators, regulators, judges, and lawyers? Almost everyone thinks his or her profession is special. Humans instinctively believe that their judgment and reasoning is special and that technology cannot replicate or replace their particular skill. Doctors, teachers, and baseball scouts all believe that they uniquely possess special skills that cannot be automated.[9]

---

[9] *See generally* Michael A. Bishop & J.D. Trout, *Epistemology and the Psychology of Human Judgment* (Oxford: Oxford University Press, 2005) at 24-53 (humans instinctively deny or ignore the success of such technology because of deep-seated cognitive biases, such as overconfidence in our own abilities and judgments). On medicine, *see e.g.*, Samuel W. Bloom, 'Structure and Ideology in Medical Education: An Analysis of Resistance to Change' (1988) 29 J. Health & Soc. Beh. 294. On education, *see* Francoise Blin & Morag Munro, 'Why Hasn't Technology Disrupted Academics' Teaching Practices? Understanding Resistance to

Lawyers are no different.[10] The belief that the legal profession is special and that lawyers and judges are immune from displacement by technological advances hinges on a bias that leads one to believe that only a human can deliver such wise judgments and decisions.

But human decision makers are flawed and biased. The biases and inconsistencies found in individual judgments can largely be washed away using advanced data analytics. The judgment of one human judge is outweighed by the wisdom of a decision generated by predictive technology that takes into account millions of judgments and decisions.[11] Even if a machine-produced law is not perfectly unbiased, as long as it is *less* biased than a law produced by individual humans, the result will be net beneficial. Plus, the decisions generated by the machine will be far more consistent than human judgments. Finally, even if a machine-produced law were to entrench biases, reprogramming a machine to correct for bias will be far easier than reprogramming and de-biasing many human judges.

2.3 *Other roadblocks and concerns*

The rise of micro-directives will bring enormous institutional upheaval and autonomy concerns that may present additional roadblocks in the evolution of the law. The death of rules and standards produces a shift in the balance of our political institutions, greatly diminishing the power of the judiciary. As the number of cases and controversies litigated falls and the interpretation of policy becomes unnecessary, the opportunity afforded to judges to use cases to make policy statements and impact opinion will diminish. On the other hand, the opportunities for judges to inject bias and error will also diminish.

Change Through the Lens of Activity Theory' (2008) 50 Computers & Ed. 475. On baseball scouts, *see* Michael Lewis, *Moneyball: The Art of Winning an Unfair Game* (New York, N.Y.: Norton, 2003).

[10] *See* Jeffrey M. Lipshaw, 'The Venn Diagram of Business Lawyering Judgments: Toward A Theory of Practical Metadisciplinarity' (2011) 41 Seton Hall L. Rev. 1; Cass R. Sunstein, 'Of Artificial Intelligence and Legal Reasoning' (2001) 8 U. Chi. Law Sch. Roundtable 29 (suggesting that computer programs do not reason analogically the way humans do.)

[11] *See* James Surowiecki, *The Wisdom of Crowds: Why the Many are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations* (New York, N.Y.: Anchor House, 2005).

The normative concern here raises a separate question from whether machine-aided algorithms can implement policy objectives. The question is whether there is an independent branch of government with the power to question the policy decisions of the ex ante lawmakers. When the lawmakers decide on legislative objectives and parameters for the machine algorithms, do we want a separate branch of government to review those decisions? If we do, the reduced role of the judiciary is troubling.

There are also broader consequences for individual citizens. Privacy would no doubt be affected, as machines need to gather data about human behaviour in order to make decisions. The capability of machines to invade privacy will increase. These concerns are exacerbated when a government uses the information it gathers in conjunction with technology to predict future actions by an individual.

Individuals may choose to ignore micro-directives, in the same way that many individuals today choose not to have cell phones and other communication devices. While the micro-directive merely provides information about how to comply with the law rather than a command,[12] there are ethical questions of holding individuals liable for laws that may change rapidly and individuals are not informed of these changes.

Automated laws also affect human autonomy. Human autonomy may be increasingly constrained as more and more ethical decisions are shifted from the purview of flawed humans to consistent machines. Moral atrophy may ensue. Individual citizens who simply follow rules and directives may become robotic, mere automatons who fail to appreciate the moral choices that should underlie their actions.

The trend toward micro-directives will be real as the cost of prediction and communication falls. The consequences relating to morality, privacy, and autonomy should be addressed before micro-directives arrive.

3. BROADER IMPLICATIONS

In this section, we explore three extensions to our thesis. First, we explore the possibility that micro-directives are used not merely to provide information about the law, but used to enforce the law. Second, as the cost of information falls, the cost of contracting will also fall. We suggest that

---

[12] Below in Part 3.1, we discuss how micro-directives might be used as commands.

citizens will increasingly use micro-directives when ordering their private affairs. Third, we explore how the falling cost of information may lead to formal micro-directives replacing informal social norms.

3.1 Automatic penalties for violation

We have set out a vision of a world where citizens are informed about the contours of law pertaining to their situation. The micro-directive merely provides a highly tailored rule, not a specific command. Upon receiving the micro-directive, the individuals may still elect to violate the law.

For example, upon receiving a micro-directive from a tax regulator that you are an employee, you may still elect to file your taxes as an independent contractor in order to claim more deductions. There will be some probability that you will not be audited and your violation will remain unpunished. Similarly, a doctor may receive a micro-directive that says surgery is not required, but may disagree with the law. She may perform the surgery. If the patient is not harmed, the doctor will suffer no consequences for ignoring the directive. In the language of Thomas Hobbes, the micro-directive is merely "counsel," rather than "command."[13]

But let's suppose that the lawmaking authorities can impose an automatic fine or punishment for violating the micro-directive. In this world of command, fact gathering and verification technologies may permit immediate notification of a violation of a micro-directive. Here, citizens' actions could attract scrutiny and punishment irrespective of the ultimate consequences.

Penalties could become immediately payable for individuals not following the micro-directive. A doctor who wishes to perform surgery in spite of a micro-directive forbidding surgery would immediately pay an automatic fine for disobeying the directive. A jaywalker may have a fine immediately deducted from her bank account. These fines operate as a price for violating the law. There are benefits of such policies. Through this mechanism, the machine would learn about "efficient violations" of the law. The

---

[13] Thomas Hobbes, Richard Tuck ed., *Leviathan* (Cambridge: Cambridge University Press, 1996), 176 (*"Command is where a man saith, Doe this, or Doe not this, without expecting other reason than the Will of him that sayes it. From this it followes manifestly that he that Commandeth pretendeth thereby his own Benefit… Counsell, is where a man saith, Doe, or Doe not this, and deduceth his reasons from the benefit that arriveth by it to him to whom he saith it."*)

evolutionary algorithm harnesses increasing amounts of information from citizens.

Such commands come at a cost. As Frederick Hayek noted in *The Road to Serfdom*, "commanding people which road to take" is different to providing signposts; it is coercion.[14] Automatic penalties for violation would pose additional ethical questions that will need to be addressed before the arrival of micro-directives. Would the stigma of illegal behaviour disappear if rich citizens were able to simply pay a fine at the time of acting?[15] Should these automatic penalties be different for the rich and the poor?

A far more dystopian vision is one where lawmakers turn micro-directives into *physical restraints* on behaviour. Rather than commanding which action should be taken, the individual is restrained from undertaking actions that do not comply with the law. Instead of simply telling the doctor that surgery is not the wisest course of action and that performing surgery will constitute negligence, imagine now that the medical technology required to perform the surgery is automatically switched off, denying the doctor the possibility of performing the surgery. From an ethical and policy perspective, the move from micro-directives to automatic restraint and strict coercion is enormous. While there may be increased compliance and greater certainty, the costs to individual autonomy would be great. Further, a complete ban on violations would be deeply inefficient, as it would dull the ability of a machine-driven algorithm to learn about how well calibrated the law is.

3.2 Micro-directives in contracts

The improvements in predictive technology will not just change the way that law is produced by legislators, regulators, and the judiciary. As the cost of information falls, and the accuracy of forecasts improves, the way contracts are produced will also change. Currently, contracts are designed to trade off certainty and flexibility. But in a world with greater certainty about the future, the problems of incomplete contracting will begin to fade away.

When contracting parties have poor information about future contingencies, parties commonly use vague standards to guide future behaviour. Parties use terms such as "best efforts" or "reasonable efforts." But, as information about future states of the world improves, the obligations of each party can

---

[14] Frederick A. Hayek, *The Road to Serfdom* (London: Routledge Press, 1944) at 74.
[15] Uri Gneezy and Aldo Rustichini, 'A Fine Is Just a Price,' (2000) 29 J. Leg. Stud. 1

be directed with greater precision. Rules and standards will give way to extremely precise courses of action. In each state of the world, the parties will be informed how best to act in order to preserve the intent of the contract. Contracts will fully specify how each party should behave in any state of the world.

Micro-directives in contracts do not need to be understood and agreed to at the time of contracting, however. Suppose that the parties simply agree to "maximize joint surplus" and agree upon a general principle for the splitting the dividends. With this guiding principle, a machine-driven algorithm will be able to automatically update and inform parties of their obligations as the state of the world changes. The contract that governs the behaviour of the parties will essentially be self-driving.

The evolution towards self-driving contracts will, of course, not happen overnight. At first, contracting parties may use information from predictive technologies to provide better advice on how surplus can be maximized. As confidence in the results increases and the benefits of using machine-learning predictions are realized, contracting parties will increasingly rely on algorithms to provide the guidance on how to behave. Over time, the algorithms will *become* the contracts.

This vision of contracting is, on one view, a radical departure from the world of contract law as we know it. Contracting parties no longer need to assent to the particulars of a contract. As long as parties agree to the broad vision, an algorithm will fully describe the obligations of the parties.

But, on another view, this is a continuation of the evolution of contracting that we have witnessed over the past few centuries. As the length of contracts continues to grow, covering more and more contingencies, the likelihood that all parties have read and understand all the terms of a contract becomes slimmer. Courts, though, have held that long, unread contracts are still enforceable, provided the terms are reasonable. Our vision of micro-directives in contracting simply extends this principle one step further.

3.3 Laws and norms

While we have argued that the technologies will enable greater specification and precision of the law, we do not necessarily predict "more" law. Rather, within the spheres of action where we, as a society, have chosen to govern human behaviour through law, the law will be more efficient and better calibrated. We have focused on scenarios where micro-directives replace grey areas of law. They are simply replacing vague laws with more certain and better-tailored laws. On this view of the future, the law will not encroach on or infiltrate all aspects of human behaviour. Indeed, on one argument, there may be less law, as fewer cases will be litigated.

But given that the cost of producing law is falling, one might expect to see more formal laws produced. When technologies can prescribe behaviour at low cost, the benefits of using informal mechanisms, such as social norms, may fade in comparison. Under this view, the appropriate boundaries between formal laws and social norms will change. Indeed, one might argue that these technologies lead not only to the death of rules and standards, but also lead to the death of norms. There will be no "norms" of driving when all vehicles are self-driving. All current norms will either vanish or be entrenched in the vehicles' algorithms.

CONCLUSION

The exponential growth of technology in the coming years will greatly reduce the cost of information. This cost reduction will have a deep and profound impact upon the way that laws are made and communicated to citizens. In this paper, we have suggested that, as predictive technologies continue to evolve and improve, the law will increasingly reflect principles and prescriptions developed by machines. Further, technological advancements will mean that these laws, micro-directives, will update automatically.

There will, of course, be skepticism and fear. People will be skeptical that machines could ever replicate human judgment. And people will, initially, be frightened of following a law that has been developed by a machine. But, in the same way that vehicles will soon be self-driving, we predict that laws, too, will be self-driving.

The Death of Rules and Standards

ANTHONY J. CASEY & ANTHONY NIBLETT*

Scholars have examined the lawmakers' choice between rules and standards for decades. This Article, however, explores the possibility of a new form of law that renders that choice unnecessary. Advances in technology (such as big data and artificial intelligence) will give rise to this new form—the microdirective—which will provide the benefits of both rules and standards without the costs of either.

Lawmakers will be able to use predictive and communication technologies to enact complex legislative goals that are translated by machines into a vast catalog of simple commands for all possible scenarios. When an individual citizen faces a legal choice, the machine will select from the catalog and communicate to that individual the precise context-specific command (the microdirective) necessary for compliance. In this way, law will be able to adapt to a wide array of situations and direct precise citizen behavior without further legislative or judicial action. A microdirective, like a rule, provides a clear instruction to a citizen on how to comply with the law. But, like a standard, a microdirective is tailored to and adapts to each and every context.

While predictive technologies such as big data have already introduced a trend toward personalized default rules, in this Article we suggest that this is only a small part of a larger trend toward context-specific laws that can adapt to any situation. As that trend continues, the fundamental cost trade-off between rules and standards will disappear, changing the way society structures and thinks about law.

---

* We thank Benjamin Alarie, Omri Ben-Shahar, Spencer Funk, Stavros Gadinis, Saul Levmore, Ariel Porat, and participants at the Junior Business Law Conference, the STILE conference in Rome, the Annual Meeting of the Canadian Law and Economics Association, the Law and Big Data International Symposium in Paris, the Law, Economics and Business Workshop at Berkeley Law, the Law & Economics workshop at the University of Michigan Law School, the Society of Institutional and Organizational Economics meetings in Paris, France, the Granular Norms conference at Villa Vigoni, Italy, and workshops at the University of Chicago Law School, the University of Connecticut School of Law, the Faculty of Law, University of Toronto, and Washington University St. Louis for helpful conversations and comments. We also thank Michael Alcan, Elizabeth Kiernan, Erica Yang, and Stephanie Xiao, for excellent research assistance. The Richard Weil Faculty Research Fund, the Paul H. Leffmann Fund, and the SSHRC provided generous support. In the interests of disclosure, one of the authors (Niblett) is a co-founder of a start-up bringing machine learning to the law.

INTRODUCTION

Imagine a world where lawmakers enact a catalog of precisely tailored laws, specifying the exact behavior that is permitted in every situation. The lawmakers have enough information to anticipate virtually all contingencies, such that laws are perfectly calibrated to their purpose—they are neither over- nor underinclusive. Now imagine that when a citizen in this world faces a legal decision, she is clearly informed of exactly how to comply with every relevant law before she acts. This citizen does not have to weigh the reasonableness of her actions, nor does she have to search for the content of a law. She just obeys a simple directive. The laws at work in this world are not traditional rules and standards. Instead, they take a new form that captures the benefits of both rules and standards without incurring the costs. This new form—we call it the microdirective—is the future of law.

When lawmakers enact laws today, they must choose between using rules and using standards to achieve a desired goal.[1] This choice requires a trade-off between certainty and calibration. Rules provide certainty through clear ex ante statements of the content of the law.[2] But rules are costly to design because lawmakers must, at the

---

1. The trade-off occurs on a particular level. Any given law may use rules for some components and standards for others. Louis Kaplow, *Rules Versus Standards: An Economic Analysis*, 42 DUKE L.J. 557, 561 n.6 (1992); *see* John O. McGinnis & Steven Wasick, *Law's Algorithm*, 66 FLA. L. REV. 991, 1027 (2014) ("[I]n the real world, rules and standards rarely exist as perfect Platonic forms."). For demonstrative purposes, we follow convention in discussing the rules-standards decision as a binary choice.

2. The literature on this distinction is vast. *See* WARD FARNSWORTH, THE LEGAL ANALYST: A TOOLKIT FOR THINKING ABOUT THE LAW 163–71 (2007); JOSEPH RAZ, PRACTICAL REASON AND NORMS (1990); Cass R. Sunstein, *Problems with Rules*, 83 CAL. L. REV. 953,

Electronic copy available at: https://ssrn.com/abstract=2693826

outset, identify and analyze all the various scenarios to which rules might apply. Rules can also be imprecise and error prone. Because they are defined ahead of time, they can be poorly calibrated[3] to the events as they actually occur.[4]

Standards, on the other hand, are adjudicated after the fact. As a result, lawmakers avoid high up-front design costs. Moreover, when applied after the fact, standards can be precisely tailored or calibrated to a specific context as it actually arose.[5] But they also generate ex ante uncertainty because regulated actors do not know up front whether their behavior will be deemed by the adjudicator to comply with the standard.[6]

We suggest that technological advances in predictive and communication technologies will render this trade-off between rules and standards unnecessary. A new form of law, the microdirective, will emerge to provide all of the benefits of both rules and standards without the costs of either. These microdirectives will provide ex ante behavioral prescriptions finely tailored to every possible scenario.

The first technology to consider is *predictive technology.* Innovations in big data and artificial intelligence will make it increasingly easy to predict the outcomes that certain behavior will produce. Lawmakers will ultimately have the ability to cheaply gather information and use predictive algorithms and big data to update the law

---

961–62 (1995); *see also* Isaac Ehrlich & Richard A. Posner, *An Economic Analysis of Legal Rulemaking,* 3 J. LEGAL STUD. 257 (1974); Kathleen M. Sullivan, *The Supreme Court, 1991 Term—Foreward: The Justices of Rules and Standards*, 106 HARV. L. REV. 22 (1992). *See generally* Kaplow, *supra* note 1 (explaining the distinction from an economic perspective); Frederick Schauer, *The Tyranny of Choice and the Rulification of Standards,* 14 J. CONTEMP. LEGAL ISSUES 803, 803 n.1 (2005) (same).

3. We use the term calibration to denote the fit of a law to its legislative purpose. For example, a fifty-five miles-per-hour speed limit may be poorly calibrated because it is too low when the weather is perfect and the roads are clear and too high when the weather is bad and the roads are crowded. Another term could be "inclusiveness." The fifty-five miles-per-hour speed limit is both under- and overinclusive because it prohibits some desirable behavior (driving sixty miles per hour on a sunny day) and allows some undesirable behavior (driving fifty miles per hour on a rainy day). *See* Colin S. Diver, *The Optimal Precision of Administrative Rules*, 93 YALE L. J. 65, 73–74 (1983) (exploring the costs of rulemaking and defining under- and overinclusiveness); Kaplow, *supra* note 1, at 565; McGinnis & Wasick, *supra* note 1 at 1030–31; Frederick Schauer, *The Convergence of Rules and Standards*, 2003 N.Z. L. REV. 303, 305–09; Schauer, *supra* note 2, at 803–04; Sunstein, *supra* note 2, at 992. What we call calibration is similar to what Diver calls congruence. Diver, *supra*, at 67.

4. More formally, in law-and-economics terms, a rule introduces high ex ante decision and error costs because it is costly to predict and set rules for every possible scenario. *See infra* Part I.A.

5. This precision is less costly for standards because the adjudicator only has to figure out the context-specific applications for cases that actually arise, whereas an ex ante rule has to address all possible applications. *See* Kaplow, *supra* note 1, at 562–63; McGinnis & Wasick, *supra* note 1, at 1031; Sunstein, *supra* note 2, at 1003–04.

6. Kaplow, *supra* note 1, at 569, 575 n.42, 587–88; Sunstein, *supra* note 2, at 974–77; *see also* Richard Craswell & John E. Calfee, *Deterrence and Uncertain Legal Standards*, 2 J.L. ECON. & ORG. 279 (1986) (modeling the costs of uncertain standards); Duncan Kennedy, *Form and Substance in Private Law Adjudication*, 89 HARV. L. REV. 1689–701 (1976). Our comparisons here assume unbiased lawmakers and judges. We discuss bias in Part II.

instantly based on all relevant factors.[7] In effect, this lowers the cost of designing precise, finely calibrated laws.

The second technology to consider is *communication technology*. Ubiquitous and instantaneous communication capabilities will reduce the uncertainty of law. From the vast catalog of rules generated by predictive technology, communication technology will be able to identify the rules applicable to an actual situation and inform the regulated actor exactly how to comply with the law.[8] It will be able to translate all the information into a single behavioral directive that individuals can easily follow.

To see how the mechanism might work, consider the regulation of traffic speed. In a world of rules and standards, a legislature hoping to optimize safety and travel time could enact a rule (a sixty miles-per-hour speed limit) or a standard ("drive reasonably"). With microdirectives, however, the law looks quite different. The legislature merely states its goal. Machines then design the law as a vast catalog of context-specific rules to optimize that goal. From this catalog, a specific microdirective is selected and communicated to a particular driver (perhaps on a dashboard display) as a precise speed for the specific conditions she faces. For example, a microdirective might provide a speed limit of 51.2 miles per hour for a particular driver with twelve years of experience on a rainy Tuesday at 3:27 p.m. The legislation remains constant, but the microdirective updates as quickly as conditions change.

In this Article, we explore whether this example could become the model for law more broadly. Our long-run prediction is that microdirectives will become the dominant form of law, culminating in the death of rules and standards. But even if that full evolution does not happen, microdirectives are certain to become a viable alternative for many laws. This short-run phenomenon is of great importance, as even a limited spread of microdirectives has the potential to change the way laws are structured and thought about generally.

This advent of microdirectives may take various paths. In the simplest story, the legislature uses the new technology and communicates the command to the citizen. We use this example to illustrate the concept. More realistically, however, the technology will often be implemented at the administrative level by regulators and enforcement agencies. Lawmakers may still enact standards, but administrative agents will convert them to microdirectives. A third possibility is that private citizens will generate the microdirectives. Citizens using private predictive technology may

---

7. The optimal rate of change in law (from the perspective of social welfare or some other legislative goal) may be slower. But that optimal rate could be factored in by the technology. *See infra* Part III.B. Courts currently update standards over the span of many years. *See* Oliver Wendell Holmes, Jr., Justice of the Supreme Judicial Court of Massachusetts, The Path of the Law, Address at the Dedication of the New Hall of the Boston University School of Law (Jan. 8, 1897), *in* 10 HARV. L. REV. 457 (1897).

8. We define communication function to include two steps. First, there is the communication of the context of an actual scenario to the machine. Second, there is the communication of the legal directive from the machine to the individual. *See infra* Part I.B. The first step might alternatively be called fact gathering. We lump them together because, in practice, the technology facilitating the information flow in each direction is likely to be the same or closely related.

inform themselves of the most reasonable action in any particular situation. As that private technology gets better, two things will happen. First, failure to use the technology will become a per se violation of a legal standard. And, second, the technology will be able to predict judicial outcomes. Both effects will result in citizens using private technology to derive a simple microdirective for how to comply with the law.

For all of these paths, the result is that laws that look like standards to the legislatures will appear as simple and easy-to-follow directives to the regulated individual. This form of law is neither a standard nor a rule. It provides the certainty of a rule and the calibration of a standard, with none of the decision costs associated with either. Moreover, the law, in application, morphs from a standard (for the legislature) to a set of complex rules (within the machine process) to a simple command (for the citizen).

We describe the rise of microdirectives as the death of rules and standards. One might alternatively frame the coming change simply as the death of standards. After all, microdirectives are ex ante rules that govern behavior. The driver in our example is told exactly how to behave ex ante. In that framing, technology has reduced the cost of precise ex ante rule making. Rules will no longer be over- and underinclusive. As a result, the rationale for using standards goes away. That is consistent with the conventional law-and-economics definition of a rule as having ex ante content (relative to the regulated actor). But the lawmakers are not enacting rules. The lawmakers need not spend the time to prescribe precise rules. They can enact broad standards and let the machines do the rest. Indeed, from the perspective of the lawmakers, it is the death of rules. The framing is less important than the recognition that microdirectives will change the foundational nature of law.[9]

Our analysis is positive rather than normative. One might think of perfect calibration of laws to legislative goals as problematic in a system with multiple branches and checks and balances. Indeed, our analysis implies a reduced role for judges and perhaps the need for institutional reforms to preserve important aspects of our current system. Others may view microdirectives as a threat to privacy and autonomy. The easier it is for the government to learn information about the behavior of an individual and use technology to predict outcomes, the more the government can micromanage to achieve desired social results. Finally, some may have concerns about ethics and moral health in a world where many important decisions are automated.[10] We do not take a side on these normative questions. We do, however, try to flag the areas where the thorniest normative questions will arise.

The primary contribution of this Article is to explore the most far-reaching effects

---

9. This question of framing suggests an interesting semantic deficit in the way legal academics talk about rules and standards. Readers of our earlier drafts have been equally split on what it means to call something a rule. Some infer that the label "rule" denotes an ex ante statement of content from the lawmaker. Others infer that it denotes an ex ante instruction for the regulated individual. That disconnect does not matter much with traditional lawmaking. But as microdirectives proliferate, the tension will come to the forefront. As a result, not only do actual rules and standards die, but so too does the meaningful use of those words to label the laws that exist.

10. *See, e.g.*, Seana Valentine Shiffrin, *Inducing Moral Deliberation: On the Occasional Virtues of Fog*, 123 HARV. L. REV. 1214, 1222, 1244 (2010) (standards provide for ethical decision making important to moral health).

of technology on the general structure of law. This contribution builds on and connects with two strands in the law-and-technology literature. The first strand looks at the effects that predictive technology has on the legal services industry.[11] The second strand looks at the nature of personalized default rules.[12]

We suggest, however, that these strands understate the momentous effect that the coming technological revolution will have on law.[13] By connecting the growing literature on technology and the law to the literature on rules and standards, we show that the same technology that will bring us automated compliance lawyers and personalized default rules will also bring us the microdirective.[14] And that change in the form of law will have broader consequences than retail personalization of law. Indeed, microdirectives have the potential to bring wholesale institutional changes to our entire system of laws and the way we choose to regulate behavior.

---

11. Some scholars predict the effects of technology on legal services. RICHARD SUSSKIND, THE END OF LAWYERS? RETHINKING THE NATURE OF LEGAL SERVICES (2010); RICHARD SUSSKIND, TOMORROW'S LAWYERS: AN INTRODUCTION TO YOUR FUTURE (2013) [hereinafter SUSSKIND, TOMORROW'S LAWYERS]; Daniel Martin Katz, *Quantitative Legal Prediction—or —How I Learned To Stop Worrying and Start Preparing for the Data-Driven Future of the Legal Services Industry*, 62 EMORY L.J. 909, 914–15 (2013). Others explore the current trends in legal markets and provide guidance for how law schools should respond to these trends. William D. Henderson, *A Blueprint for Change*, 40 PEPP. L. REV. 461 (2013); Bruce H. Kobayashi & Larry E. Ribstein, *Law's Information Revolution*, 53 ARIZ. L. REV. 1169 (2011); *see also* William D. Henderson, *From Big Law to Lean Law*, 38 INT'L REV. L. & ECON. 5 (2013) (exploring the changing trends in markets for legal services); Larry E. Ribstein, *The Death of Big Law*, 2010 WISC. L. REV. 749 (predicting the demise of big law firms); Brian Sheppard, *Incomplete Innovation and the Premature Disruption of Legal Services*, 2015 MICH. ST. L. REV. 1797 (predicting the consequences of technological innovation on the legal services industry).

12. *See, e.g.*, Ariel Porat & Lior Jacob Strahilevitz, *Personalizing Default Rules and Disclosure with Big Data*, 112 MICH. L. REV. 1417 (2014). Porat & Strahilevitz provide a theory of personalized default rules in a world of big data. We jump off from that point to explore the wholesale effects of technological advances on law more generally. *See* Omri Ben-Shahar & Ariel Porat, *Personalizing Negligence Law*, 91 N.Y.U. L. REV. 627 (2016); George S. Geis, *An Experiment in the Optimal Precision of Contract Default Rules*, 80 TUL. L. REV. 1109 (2006); Cass R. Sunstein, *Deciding by Default*, 162 U. PA. L. REV. 1 (2013).

13. The closest work to ours is that of John McGinnis and Steven Wasick. McGinnis & Wasick, *supra* note 1. Though they reach strikingly different conclusions, McGinnis and Wasick begin in the same place as we do, asking how technological advances that reduce information costs will affect the balance of rules and standards. Focusing primarily on legal search technology and the ability to predict judicial outcomes, they predict a world where standards and dynamic rules are favored over simple rules. *Id*. at 1049–50. Building on this analysis, we add in the effects of communication technology and machine learning to show that standards and rules (simple and dynamic) will no longer be viable forms of law.

14. Porat & Strahilevitz note that the dichotomy of personal and impersonal rules is not the same as the dichotomy of rules and standards. Porat & Strahilevitz, *supra* note 12, at 1457–58. Personalized defaults can be rules or standards. And impersonal defaults also come in both forms. *Id*. Beyond that observation, Porat and Strahilevitz focus their attention on the personal-impersonal dichotomy. Our analysis suggests, however, that all laws—both personal and impersonal—will ultimately gravitate toward microdirectives that transcend the distinction between rules and standards.

The Article proceeds as follows. Part I sets out our general theory of microdirectives and provides demonstrative examples. Part II explores the feasibility of the technologies behind microdirectives. Part III discusses implications and broader consequences of the rise of microdirectives and the death of rules and standards. A final section concludes.

I. THE EMERGENCE OF MICRODIRECTIVES AND THE DECLINE OF RULES AND STANDARDS

In this Part, we spell out how technology will affect the administration of law and the structure of legal content. We outline two types of technology that will lead to a dramatic reduction in the cost of calibrating and communicating ex ante legal directives, thereby eliminating the need to choose between rules and standards. The analysis is presented in three sections. First, we briefly review the distinction between rules and standards and outline the cost choices presented by the dichotomy. Second, we set out our core theory that technology will fundamentally change those cost choices. We provide two examples to demonstrate how predictive and communication technologies will pave the way for microdirectives that capture the benefits of both rules and standards. Third, we discuss how the emergence of microdirectives can take place through different branches of lawmaking or can be driven by private actors with access to predictive technology.

A. Background: Rules and Standards

Rules are precise and ex ante in nature. Rules indicate to an individual whether certain behavior will violate or comply with the law. When a rule is enacted, effort must be undertaken by lawmakers to give full and precise content to the law before the individuals act. Standards, on the other hand, are imprecise when they are enacted.[15] The exact content of the law comes after an individual acts, as judges and other adjudicators determine whether the individual's specific behavior in a particular context violates the standard.

Generally, lawmakers incur both error costs and decision costs when enacting a law. Error costs arise when a law is over- or underinclusive; the law allows behavior that should be prohibited, or prohibits behavior that should be allowed.[16] Errors can be reduced as lawmakers exert greater effort to get the law right. But this requires information and deliberation. Reducing error costs imposes decision costs on the lawmakers. Additionally, regulated individuals face a cost in figuring out whether their behavior complies with the law. When the application of the law to a particular situation cannot be easily predicted, the individual incurs cost of legal uncertainty.

Error, decision, and uncertainty costs arise in different ways for rules and standards. The classic models in the rules-versus-standards literature conclude that, for

15. Standards are found wherever vague and ambiguous terms such as "reasonable," "material," or "excessive" are used in the law. *See* Schauer, *supra* note 3, at 308–09; Schauer, *supra* note 2, at 804–05.

16. McGinnis & Wasick, *supra* note 1, at 1031.

several reasons, standards tend to perform better when the behavior of the regulated actors is infrequent and heterogeneous.[17]

First, when behavior of regulated actors is infrequent, standards generate lower decision costs because the content of the law only needs to be decided in the infrequent event that the relevant context actually arises. Rules, on the other hand, require ex ante decisions about all future possible scenarios. Where behavior is infrequent and heterogeneous, lawmakers must make many more decisions if they want to write rules that are as precise in application as a standard that is adjudicated ex post would be. Rules do, however, impose lower decision costs when behavior is frequent and homogeneous. Economies of scale kick in and a law need only be enacted once rather than litigated over and over again.[18]

Second, error costs for standards are lower when behavior is infrequent and heterogeneous because the adjudicator determining the content of the law ex post has more information than the ex ante lawmaker. The adjudicator has additional context not available to the ex ante lawmaker and has the benefit of hindsight in identifying which factors are relevant.

On the other hand, adjudicator competency and bias complicate this simple model of error costs.[19] Ex post adjudication may suffer from hindsight bias[20] and from biases based on the personal characteristics of particular individuals.[21] Such biases can manifest themselves in arbitrariness, political favoritism, covert influence, inconsistency, and discretionary justice[22] even when judges believe they are being

---

17. *See, e.g.*, Louis Kaplow & Steven M. Shavell, *Economic Analysis of Law*, *in* 3 HANDBOOK OF PUBLIC ECONOMICS 1665, 1744–45 (Alan J. Auerbach & Martin Feldstein eds., 2002).

18. Strict application of the doctrine of precedent also introduces economies of scale for standards, but it does so in a way that turns the standard into a rule. *See* Holmes, *supra* note 7; Anthony Niblett, *Case-by-Case Adjudication and the Path of the Law*, 42 J. LEGAL STUD. 303, 310 (2013).

19. *See, e.g.*, RAZ, *supra* note 2; FREDERICK SCHAUER, PLAYING BY THE RULES: A PHILOSOPHICAL EXAMINATION OF RULE-BASED DECISION-MAKING IN LAW AND IN LIFE (1991); *see also* Kaplow, *supra* note 1, at 609 (discussing institutional competence generally).

20. *See generally* DANIEL KAHNEMAN, THINKING, FAST AND SLOW (2013); Christine Jolls, Cass R. Sunstein & Richard Thaler, *A Behavioral Approach to Law and Economics*, 50 STAN. L. REV. 1471, 1523–27 (1998); Jeffrey J. Rachlinski, *A Positive Psychological Theory of Judging in Hindsight*, 65 U. CHI. L. REV. 571 (1998).

21. *See, e.g.*, Jeffrey J. Rachlinski, Sheri Lynn Johnson, Andrew J. Wistrich & Chris Guthrie, *Does Unconscious Racial Bias Affect Trial Judges?*, 84 NOTRE DAME L. REV. 1195 (2009) (finding evidence of judicial bias based on race); *see also* Christine Jolls & Cass R. Sunstein, *The Law of Implicit Bias*, 94 CAL. L. REV. 969 (2006) (discussing implicit biases that individuals hold against disadvantaged groups).

22. *See, e.g.*, Thomas J. Miles & Cass R. Sunstein, *The New Legal Realism*, 75 U. CHI. L. REV. 831 (2008) (finding that political preference, race, gender, and other demographic characteristics sometimes have effects on judicial judgments); Anthony Niblett, *Tracking Inconsistent Judicial Behavior*, 34 INT'L REV. L. & ECON. 9 (2013) (finding that judges in California decide unconscionability cases inconsistently with precedent). *See generally* Jeffrey A. Segal, *Judicial Behavior*, *in* THE OXFORD HANDBOOK OF POLITICAL SCIENCE 275 (Robert E. Goodin ed., 2011). It has been argued that these flaws of judges may be partially responsible for the increased flight to agency regulation over the past twenty to thirty years, in spite of the

unbiased.[23]

Ex ante lawmakers and regulators may, of course, also be biased.[24] But the biases exhibited in ex post adjudication are particularly costly. Hindsight bias is more pervasive and difficult to minimize for ex post adjudication. Additional biases based on personal characteristics of an individual are also more likely for ex post adjudication and may be particularly pernicious and harmful to social objectives.[25] The presence of biased adjudicators, thus, alters the error-cost trade-offs between rules and standards and weakens any claims that standards have lower error costs.

A third cost comparison is also relevant when assessing the relative merits of rules and standards: the uncertainty cost imposed on the regulated actor in understanding whether her behavior complies with the law. Uncertainty about the content of a law is greater with standards than with simple rules. When regulated by a simple rule, an individual will more likely know whether her behavior is allowed or prohibited.[26] When regulated by a standard, on the other hand, the individual does not know how any particular judge with wide discretion will apply the standard to the facts. She may not know what behavior a judge will consider reasonable.

The choice between using a rule or a standard to achieve a particular policy objective is therefore a question of weighing and trading off these costs. We predict that advances in technology will fundamentally change that trade-off.

---

many well-recognized and well-documented flaws of regulators and economic costs of regulation. ANDREI SHLEIFER, THE FAILURE OF JUDGES AND THE RISE OF REGULATORS (2012); *see also* Joshua Schwartzstein & Andrei Shleifer, *An Activity-Generating Theory of Regulation*, 56 J.L. & ECON 1 (2013) (modeling the choice between ex ante regulation and ex post judging where courts commit errors).

23. *See generally* KAHNEMAN, *supra* note 20; Jolls & Sunstein, *supra* note 21, at 970–71; Rachlinski et al., *supra* note 21, at 1201–04 (exploring the effects of unconscious or implicit biases).

24. *See, e.g.*, Stephen M. Bainbridge, *Mandatory Disclosure: A Behavioral Analysis*, 68 U. CIN. L. REV. 1023, 1056–58 (2000) (noting the lack of attention to the behavioral biases of regulators); Stephen J. Choi & A.C. Pritchard, *Behavioral Economics and the SEC*, 56 STAN. L. REV. 1, 20–36 (2003) (cataloguing the biases affecting SEC regulators).

25. For our purposes, the important observation will be that machine-created rules are less likely to be biased than humans in making rules or applying standards. Our analysis suggests that given a legislative goal, machines will more faithfully implement that objective. *See infra* Part II. It is possible, still, that judges are debiasing bad legislative policy (though some empirical evidence suggests otherwise). In that case, judges have the power to override and influence policy in a way that may be socially beneficial. That power will be lost as standards die. We address these issues in Part III.

26. This assumes that judges (and juries) follow rules. They may, however, import exceptions that turn rules into standards—or ignore the rules altogether. *See* Schauer, *supra* note 3, at 312–14. For the most part we bracket the possibility of such rule nullification. But it is worth noting that the developments we explore make nullification less likely as well. *See infra* Part III.A (discussing the diminished capacity for judges to influence and change legal substance and policy). This is yet another way that law will become more rule based from the perspective of the regulated individual.

B. Technology Will Facilitate the Emergence of Microdirectives as a New Form of Law

Two types of technology will lead to the death of rules and standards and the rise of microdirectives: *predictive technology* and *communication technology.* The first will facilitate lawmakers' efforts to craft precise ex ante context-specific rules that provide the nuance and specificity traditionally associated with standards. The second will allow for the translation of those nuanced and specific laws into simple directives that are communicated to the regulated actors in a timely manner.

Predictive Technology. Predictive technology, driven by ever increasing computational capacity, will allow lawmakers to sculpt more perfect ex ante laws.[27] Computation power is growing at exponential rates. The consistent trend of the last fifty years suggests that that power will, by the end of this century, be more than one trillion times greater than what it is today.[28] With even a fraction of that processing power, tomorrow's computers will be able to gather and analyze more facts than any human lawmaker or judge. Lawmakers will be able to direct a machine to analyze a massive amount of data instantly to predict which rules can precisely achieve a policy objective.

Relying on the machines to observe and analyze more relevant facts, lawmakers will make better predictions about the impact of a law and will face reduced error costs. Lawmakers will no longer have to think up rules to enact laws. Judges will no longer have to examine citizens' decisions on a case-by-case basis in order to apply laws. And the laws will be highly calibrated to policy objectives with no chance of judges introducing bias or incompetence. Of course, the calibration need not be perfect, it only needs to be better than the calibration associated with the alternatives of legislated rules and adjudicated standards.

As a practical matter, the result will be a new hybrid form of law that is both rule and standard. The lawmaker can set a broad objective, which might look like a standard. But the predictive technology will take the standard and engineer a vast catalog of context-specific rules for every scenario. But that is only the first half of the story.[29]

---

27. In a different context, Professor Michael Abramowicz identified the power of predictive decision making to "take[] advantage of the best of both the world of standards and the world of rules." Michael Abramowicz, *Predictive Decisionmaking*, 92 VA. L. REV. 69, 74 (2006). Our analysis is consistent with and builds on Abramowicz's important insight. When the power of prediction that he identified is harnessed and amplified by technological advances and coupled with new communication technologies, the law-making process fundamentally changes.

28. *See infra* Part II.A.2. This estimate is based on a trend known as Moore's Law. *See generally* Mark Lundstrom, *Moore's Law Forever?*, 299 SCIENCE 210, 210 (2003) (explaining Moore's Law and its implications for electronic systems); McGinnis & Wasick, *supra* note 1, at 1041 (describing Moore's Law); Gordon E. Moore, *Cramming More Components onto Integrated Circuits*, ELECTRONICS, April 19, 1965, at 114 (setting out the premise of Moore's law).

29. The discussion of predictive technology here and throughout this Article assumes a consequentialist approach to law. For a consequentialist, the content of the law is driven by a prediction of the outcome of behavior. For nonconsequentialist theories, the use of the technology is slightly different. But the trend toward microdirectives will likely be the same. For example, imagine that a lawmaker wants to prohibit certain behavior she deems immoral

Communication Technology. In the second half, the communication technology will simplify that context-specific catalog of rules into clear microdirectives for the regulated individuals. Without that simplification, the catalog of rules would be too complex and pose significant compliance challenges. It would be impossible for people to learn, remember, and process all of the requirements contained in the catalog. But advances in communication technology will produce microdirectives that reduce or eliminate those compliance costs and prevent uncertainty costs that might otherwise arise.

The mechanism for translation is straightforward. Communication technology will gather and transmit information about the scenario in which the individual finds herself,[30] identify the applicable rule from the vast catalog, and then translate that into a simple directive that is communicated back to the individual when she needs it. In this way, microdirectives will turn hundreds or thousands of context-specific, machine-generated rules into simple directives that are easy to understand and follow. The law controlling a particular scenario may take into account hundreds or thousands of factors,[31] but the individual will receive a simple command like a red or green light. When the output from the predictive technology is translated into a microdirective, citizens will be able to act *as if* they are taking into account more relevant factors than are humanly possible.[32]

* * *

To summarize, these technologies will combine to do the following. First, they

---

regardless of the consequences of that behavior. She does not want to list out all permutations of immorality, so a rule will not work. Instead, she can start with a standard—immoral activity is prohibited—and then identify samples of immoral behavior to feed into a machine. The machine can then use analytic and pattern recognition technology to determine whether other new scenarios would be deemed immoral by the lawmaker. We discuss below a similar process of pattern recognition for the question of pornography that a lawmaker knows is pornography when she sees it. *See infra* Part II.A.2.

30. We include this fact-gathering function in our analysis of communication technology because the key innovation is the communication of the factual scenario from the specific context to the analytic process. The technology facilitating this communication is likely to be the same or related to the technology facilitating the communication (in the other direction) of the final microdirective from the process to the individual.

31. To the extent that certain factors like race and gender are considered out-of-bounds, the machines can be programmed to ignore those factors. Indeed, it is easier for a machine to affirmatively ignore a prohibited factor than for a human.

32. These microdirectives share some important characteristics with McGinnis and Wasick's "dynamic rules." McGinnis & Wasick, *supra* note 1, at 1039–45. Both can be very precisely calibrated to specific conditions. But McGinnis and Wasick envision that at least the algorithm is "fixed by a rule" that must be changed if the "world may change in a way that makes another weighting of factors achieve the legislature's original objectives." *Id.* at 1047–48. We suggest instead that one of the core functions of a microdirective is the ability to learn from data and automatically update the weighting of factors the way a judge would update her application of a standard. In this way, microdirectives are not rules. They update automatically and continuously to account for such changes in the weighting of factors. But unlike standards they can be communicated ex ante with certainty.

will take a standard-like policy objective, analyze its application in all possible contexts, and create a vast catalog of legal rules—each of which is tailored to best achieve the objective in a specific scenario. Second, when a regulated actor is in any actual scenario, the technologies will search the vast catalog and identify the specific rules that are applicable. Third, they will translate those rules into a simple microdirective on how the regulated actor can comply with the law. Fourth, they will communicate that microdirective to the regulated actor in a timely and efficient manner.

C. Examples

To demonstrate the point, we present two stylized examples.

1. Example 1: Predictive Technology in Medical Diagnosis

In this subsection, we provide an example that demonstrates how improved *predictive technology*—technology that allows lawmakers to better predict the outcomes of actions—will foster microdirectives.

Suppose you are a legislator. You are charged with determining when doctors should be liable for performing a risky surgery on a patient. How can you best regulate doctors' behavior? How can you best draft a statute that will help doctors understand when their behavior complies with or violates the law? How many of the specific details should you include in the statute? How many of these details can be postponed until we have more information about how doctors behave in each case?

One option is to provide doctors with a clear and simple bright-line rule that dictates the circumstances under which surgery should or should not be conducted. This simple rule provides great certainty to the doctors and is easily enforced; either a doctor complied with the rule or she didn't. A simple, precise ex ante rule would be your preferred method if similar patients frequently present with the same symptoms.[33] Under these circumstances, a rule would be preferred because the content of the law can be established just once, and there are enormous benefits from economies of scale.

But a doctor's decision to operate on a patient frequently turns on many different factors. A "one size fits all" rule here would likely not be optimal. Any simple bright-line rule you enact will likely be overinclusive and underinclusive compared to an optimal decision rule. There will be some patients who receive surgery who do not need it (type I errors); there will also be other patients who do not receive surgery who do need it (type II errors).

To overcome these errors, you may try to write a more complex rule. To formulate this rule, you may try to think up many different scenarios, where you imagine different types of patients presenting with various symptoms. A complex rule is preferred if the cost of thinking and writing the rules is very low *and* the cost of doctors understanding and being able to comply with such a complex rule is also low. But it is often very costly for legislators to think up and write down all contingencies.

33. *See* Kaplow, *supra* note 1, at 573–77 (discussing the importance of frequency in assessing the desirability of rules).

Further, the more complex the rule you write, the more difficult it becomes for a doctor to follow.[34]

Rather than implement a rule, another option you have is to enact a *standard* and evaluate the conduct of a doctor after the decision to operate (or not operate) has been made. That is, the decision to hold a doctor liable would be made once all the circumstances of the particular case are known.[35] For example, the legal standard might stipulate that all doctors must take "reasonable care" in determining whether to operate on patients. This provides doctors with greater flexibility to decide whether or not the patient needs surgery.[36] But it also provides an ex post adjudicator with the flexibility and discretion to determine what is meant by "reasonable."

If a patient suffers harm as a result of a doctor's decision, then a judge can look at all the facts as they actually occurred and make an informed decision as to whether the doctor took reasonable care. A standard would be better than a rule if patients and symptoms are heterogeneous and the likelihood of two patients with the same background and symptoms is very low.

There are, of course, costs associated with implementing and enforcing a standard. First, the cost of deciding each case is not zero. There are decision costs of learning the best course of action the doctor should have taken in the circumstances. Second, a judge may apply the standard incorrectly, either due to error or to bias. Third—and importantly—unlike a clear rule, a vague standard creates a great deal of uncertainty for the doctor. A doctor may not know how a judge will decide any given case; further, different judges may decide inconsistently. If doctors are risk averse, a vague law can chill socially desirable behavior,[37] and the uncertainty may generate considerable expense in the form of compliance costs.

But in our hypothetical situation, let's suppose that a standard is optimal. Let's assume that the question of surgery rarely arises and that patients are highly diverse, both in terms of health backgrounds and in terms of the symptoms they present. Formulating detailed rules that cover all those situations and being able to communicate these complex rules to doctors would be difficult, and a simple rule would create high error costs. Case-by-case adjudication is not costless but it is preferred in our example because the infrequent cost of determining the content of the law ex post is

34. Louis Kaplow, *A Model of the Optimal Complexity of Legal Rules*, 11 J.L. ECON. & ORG. 150, 151 (1995) (modeling the trade-off between complexity and regulated actors' ability to comply); *see also* Diver, *supra* note 3, at 73–74 (noting the trade-off between precision and ease of applying and following a law).

35. As Henry Hart and Albert Sacks note: "The wise draftsman . . . asks himself, how many of the details of this settlement ought to be postponed to another day, when the decisions can be more wisely and efficiently and perhaps more readily made?" HENRY M. HART, JR. & ALBERT M. SACKS, THE LEGAL PROCESS: BASIC PROBLEMS IN THE MAKING AND APPLICATION OF LAW 157 (1958).

36. *See* John Braithwaite & Valerie Braithwaite, *The Politics of Legalism: Rules Versus Standards in Nursing-Home Regulation*, 4 SOC. & LEGAL STUD. 307 (1995). After comparing nursing homes in rule-based United States and standard-based Australia, Braithwaite & Braithwaite conclude that the flexibility of standards in Australia allows health care professionals to respond to their patients' needs better than professionals merely applying strict rules. *Id.*

37. *See* Craswell & Calfee, *supra* note 6, at 298–99 (concluding that uncertainty can reduce socially desirable behavior).

lower than the costs of trying to specify the law up front in all potential situations, many of which will never arise.

Now let's examine how technology will eliminate this trade-off between rules and standards. Suppose that you learn of the existence of a diagnostic machine that is designed to predict when surgery is required. The machine takes into account relevant facts about the patient[38]—her history, the symptoms, and other relevant information—to provide a best guess as to whether the patient requires surgery.

You, the legislator, have access to this machine. How does this predictive machine affect your decision to enact a rule or a standard? The answer turns on two factors. First, how good is the machine at accurately predicting outcomes? If the predictive technology is very powerful and the machine is able to provide precise and accurate information, then this points in favor of using the machine to create a rule, rather than relying on a judge to adjudicate a standard. Second, can this information be easily communicated to a doctor? That is, can lawmakers provide the doctor with timely notice of what behavior will comply with or violate the law?

Consider two scenarios:

> *Scenario 1: A terrible predictor*
>
> In scenario 1, the machine is very poor at predicting when a patient requires surgery. The machine essentially randomizes patients for surgery. The machine generates both type I and type II errors. One might think of the technology as a simple coin toss: heads for surgery, tails for no surgery.
>
> *Scenario 2: A perfect predictor*
>
> In scenario 2, the machine can predict with 100% accuracy whether a patient requires surgery or not. The machine instantly examines the patient's history and symptoms, analyzes millions of prior cases, and reads all articles in medical journals. It then makes a perfect prediction. It is better than any human at determining whether it is optimal to have surgery. There are no type I errors: patients who do not need surgery are not designated for surgery. There are no type II errors: patients who need surgery are designated for surgery.

Under scenario 1, the technology should have no effect on your decision as a regulator to implement a rule or a standard. You should implement a standard and determine liability on a case-by-case basis, learning more about doctors' behavior over time.

Under scenario 2, however, the optimal form of the law will be different. The machine's predictions provide the exact content of the law. The machine provides microdirectives for each and every scenario. The over- and underinclusivity associated with simple rules have disappeared. There are no errors (type I or type II) in this scenario. And the costs incurred in thinking up and formulating such a complex rule have already been incurred in the development of this machine.[39] The justification

---

38. In practice, the machine would actually take into account relevant information about the doctor as well, such as his track record with surgeries of the relevant type.

39. In reality many of the costs for developing the machine may have been incurred by

for relying on ex post adjudication of standards—reducing the error costs of rules—is gone. Further, we have an added benefit of eliminating uncertainty for the doctors. If they follow the directive of the machine, they know they will not be held liable.

The emergence of microdirectives and the death of rules and standards as we know them do not rely on *perfect* predictive technology. Rather, as the predictive technology gets better and better, we move away from the world of scenario 1 and towards the world of scenario 2. There will come a point where the technology is *good enough* that the costs of using a microdirective are sufficiently low so that there is no longer any need to use traditional rules or standards.

A caveat is necessary. This tipping point can only be realized if the rules generated by the machine can be easily communicated to doctors. That is, the legislator has to be able to provide the doctor with a quick and simple answer to the question of whether the patient requires surgery.

Doctors would find it difficult to follow complex, computer-derived rules. Regulated actors have neither the desire nor the time to thumb through thousands of pages of legislation and understand complex algorithms. Rather, lawmakers need some form of technology to allow a doctor to easily input all the relevant facts about a patient and receive an instant output that dictates whether or not the patient requires surgery. One might imagine a web-based program or mobile app, where the doctor can quickly and easily enter all relevant facts, submit the information, and instantly receive a binding ex ante opinion. Such technology is emerging and will be able to transform the complex rules generated by machine prediction into a simple directive that the doctor can follow.[40] The costs to the doctor in understanding the complex rule will be dramatically reduced as this technology improves. Even though the rule will be highly complex and based on a sophisticated algorithm, from the perspective

---

industry for the nonlegal benefits that the machine brings. In that sense, the marginal costs of using it for law are negligible. Moreover, even if the machine had to be developed specifically for law, that is a fixed cost that can be averaged across all applications when calculating the per rule cost.

40. *See, e.g.*, Robert McMillan & Elizabeth Dwoskin, *IBM Crafts a Role for Artificial Intelligence in Medicine*, WALL ST. J. (Aug. 11, 2015, 12:04 AM), http://www.wsj.com/articles/ibm-crafts-a-role-for-artificial-intelligence-in-medicine-1439265840 [https://perma.cc/DXL7-TUWG] (describing IBM's planned move into artificially intelligent diagnostics for cancer and other diseases); Joseph Walker, *Can a Smartphone Tell if You're Depressed?*, WALL ST. J. (Jan. 5, 2015, 7:03 PM), http://www.wsj.com/articles/can-a-smartphone-tell-if-youre-depressed-1420499238 [https://perma.cc/6AWE-C4AU] (describing tests of a new generation of "health-surveillance technologies" that can gather information to diagnose illness and assess physical and mental well-being); Ron Winslow, *Patients Seeking Alternatives to Statins May Undergo Rigorous Vetting,* WALL ST. J. (July 27, 2015, 4:40 PM), http://www.wsj.com/articles/patients-seeking-alternatives-to-statins-may-undergo-rigorous-vetting-1438029636 [https://perma.cc/79SD-9EDM] (describing a software application that guides doctors through the decision to put patients on nonstatin cholesterol treatment). *See generally* William M. Grove & Paul E. Meehl, *Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy*, 2 PSYCH., PUB. POL'Y & LAW 293 (1996) (noting how simple formal algorithms frequently outperform human predictions).

of the doctor, the rule will be simple: operate or do not operate.[41] We explore communication technology further in our second example.

2. Example 2: Communication Technology in Traffic Laws

In this subsection, we highlight the way improved *communication technology* will facilitate microdirectives. Machines can almost instantaneously gather information, process it, and produce a useable output that directs how individuals should behave.

Traffic lights provide an example of this type of technology. They communicate the content of a law to drivers at little cost and with great effect. This notice technology—combined with technology for predicting traffic patterns and driver behavior—creates an environment where lawmakers are able to replace vague standards and simplistic rules with crisp and increasingly complex microdirectives.

Electric traffic lights communicate to drivers precisely when they are required to stop and when they may proceed. Traffic lights appear to generate very simple rules: if the light is red, you must stop; if the light is green, you may go. But these rules are simple only from the perspective of the driver. From the perspective of the lawmakers, the underlying rules are complex. The simplest underlying rule may dictate that cars must stop during regular, alternating time intervals. In more complex examples, the time intervals can vary by intersection, direction of traffic, or time of day.

If promulgated without traffic lights, these rules would be far too complex. Drivers would have to consult tables that matched intersections, times, and directions with prescribed intervals of stopping. They would also have to consult precise clocks to determine when the intervals start and end.

The traffic light translates complexities into a simple command. From the driver's point of view, the lights provide a directive that is easily understood. And the lawmaker's cost of giving notice is low.[42] Electric traffic lights take advantage of significant economies of scale that enable lawmakers to make complex rules, translate them into simple directives, and deliver notice of the required behavior to many drivers.

Moreover, while the command of the traffic light remains simple, the substance of the underlying rules is becoming more complex. Predictive analysis facilitates this process. Stopping at a red light when an intersection is deserted is wasteful and

---

41. It may seem odd at first that lawmakers are in the business of diagnostic technology. But this is no different from what judges do in medical litigation. Judges hear expert testimony and decide ex post whether certain behavior was reasonable. In our example, lawmakers just use expert technology to do that ex ante. It is true that the role of the doctor has changed—diagnostic judgment is less important—but that is the inevitable result of advances in diagnostic technology. Our point is simply that in the hands of lawmakers the technology also changes the role of law. When the technology is only available to the private actors—the doctors in this example—then the evolution of rules into standards takes a slightly different path. We discuss this *infra* Part I.D.

42. These stop-go rules would be far more costly if humans operated traffic lights. Indeed, the first gas-powered traffic light used in 1868 in London, United Kingdom, was operated by humans. *The Man Who Gave Us Traffic Lights*, BBC: NOTTINGHAM (last updated July 22, 2009, 11:57 AM), http://www.bbc.co.uk/nottingham/content/articles/2009/07/16/john_peake_knight_traffic_lights_feature.shtml [https://perma.cc/4M28-DWR7].

costly.[43] That rule is overinclusive. It would be better if the directive to the driver could change depending on the circumstances (as it would with a standard). To address this, traffic lights in some jurisdictions already contain sensors that detect and predictively analyze traffic flow and adjust the timing of red and green lights accordingly.[44] Some traffic lights contain detectors allowing emergency service vehicles to "preempt" the signal and expedite their journey.[45] In the near future, these systems will take into account more variables, such as the number of cars, speed of travel, or type of intersection. They might even take into account personal characteristics of a vehicle's driver or passengers.[46] In the not-so-distant future, a traffic-light system may know that a passenger in a regular vehicle requires medical attention and give the rushing driver a series of green lights all the way to the hospital.

The progress of traffic lights shows how lawmakers can define optimal policy outcomes (for example, travel times and accident rates) and machines can generate a catalog of rules and exceptions to achieve those outcomes. And yet—even while the lawmakers enact a standard and the machines generate an increasingly complex catalog of rules underpinning the operation of traffic lights—from the perspective of the driver, the law will remain constant and straightforward: a simple stop-go directive.

This phenomenon is not limited to traffic law. The forces at work here are ubiquitous. The invention and mass adoption of Internet technology has facilitated instantaneous and cheap communication between individuals across all domains.[47] It also, importantly, allows for immediate communication between lawmakers and individuals.

D. The Different Channels Leading to the Death of Rules and Standards

We have, until now, spoken generally of lawmaking by a legislature. That is by no means the only avenue. Microdirectives can emerge through two other channels: (1) nonlegislative (regulatory or judicial) law making; and (2) private use of technology by regulated actors.[48] We discuss each in turn.

---

43. There are other potential costs such as the increase in the number of rear end traffic accidents caused by cars braking as lights turn yellow. We argue that these costs will also die out as the rule becomes more context specific.

44. *See, e.g.*, Ian Lovett, *To Fight Gridlock, Los Angeles Synchronizes Every Red Light*, N.Y. TIMES (Apr. 1, 2013), http://www.nytimes.com/2013/04/02/us/to-fight-gridlock-los-angeles-synchronizes-every-red-light.html [https://perma.cc/SS66-29A7] (describing Los Angeles's $400 million system of synchronized traffic sensors aimed at controlling traffic flow and reducing gridlock); *see also* Diane Cardwell, *Copenhagen Lighting the Way to Greener, More Efficient Cities*, N.Y. TIMES (Dec. 8, 2014), http://www.nytimes.com/2014/12/09/business/energy-environment/copenhagen-lighting-the-way-to-greener-more-efficient-cities.html [https://perma.cc/3F2D-YS68] (noting Copenhagen's use of lights and sensors aimed at easing mobility and cutting use of fuel as well as achieving more ambitious goals).

45. U.S. DEP'T TRANSP., TRAFFIC SIGNAL PREEMPTION FOR EMERGENCY VEHICLES A CROSS-CUTTING STUDY 1-1 (2006), http://ntl.bts.gov/lib/jpodocs/repts_te/14097_files/14097.pdf [https://perma.cc/9SDF-BRC8] (noting the signal preemption programs in various jurisdictions).

46. *Cf.* Porat & Strahilevitz, *supra* note 12 (noting the value of personalized laws).

47. *See infra* Part II.

48. Here, we discuss different channels through which technology will affect the law. In

1. The Production of Microdirectives by Nonlegislative Lawmakers

Legislatures are not the only lawmakers with access to technology. In many cases, the lawmaking power is entrusted to a regulator or enforcement agent.[49] In other cases, judges make law. Those entities can also use technology to create and communicate microdirectives to regulated actors.

Regulatory microdirectives. It is likely to be more politically feasible for regulators to develop microdirectives than legislators. The legislative path to enacting a computer algorithm is complicated. Pork barrel and horse-trading amendments to an algorithm do not make for successful programming. On the other hand, a regulator tasked with enforcing some legislated standard might easily adopt an algorithm-driven system of microdirectives.[50]

The pressures on a budget-constrained regulatory body will push the agency toward adopting technology. Likewise, trends towards cost-benefit analysis and requirements that regulations be shown to be cost justified[51] are likely to accelerate agency adoption. Predictive technology facilitates such cost-benefit analysis, reduces uncertainty costs to the regulated actors, and cuts down on ex post adjudication costs.

Congress could enact a standard and direct that these standards be administered by an algorithm-based system of microdirectives overseen by regulators or the regulators could themselves decide to implement the standard in that manner.[52]

Advance tax rulings provide an example of an area for regulators to use microdirectives.[53] As it currently stands, taxpayers may seek clarification of vague standards in the law by asking the tax authority to examine their tax arrangements and determine whether they comply with the code.[54] A taxpayer may ask the tax authority to give a ruling on a matter that takes into account a number of factors such as: Am I a resident of the United States for tax purposes? Or, are my workers independent contractors or are they employees?[55]

---

other work, we have discussed the incremental nature of these changes. *See* Anthony J. Casey & Anthony Niblett, *Self-Driving Laws*, 66 U. TORONTO L.J. 429 (2016).

49. From the legislature's perspective, the delegation to an agency or enforcer takes the form of a standard. *See* Schauer, *supra* note 3, at 310. The legislature sets a broad goal and gives the agency the power to fill the content of rules.

50. McGinnis & Wasick, *supra* note 1, at 1042 (discussing the use of algorithms by rule makers); *cf. id.* at 310–12.

51. *See generally* CASS R. SUNSTEIN, THE COST-BENEFIT STATE: THE FUTURE OF REGULATORY PROTECTION (2002) (exploring the rise of cost-benefit analysis in administrative agencies).

52. This is not the same as traditional convergence predictions where rules become standards or standards become rules. Schauer, *supra* note 3, at 310–12. With microdirectives, laws take a new form that has some of the benefits of rules (more certainty) and some of the benefits of standards (better calibration) but fewer of the costs associated with either.

53. *See generally* CARLO ROMANO, ADVANCE TAX RULING AND PRINCIPLES OF LAW (2002); Yehonatan Givati, *Resolving Legal Uncertainty: The Unfulfilled Promise of Advance Tax Rulings*, 29 VA. TAX. REV. 137, 144–47 (2009) (describing how advance tax rulings reduce uncertainty).

54. 26 C.F.R. § 601.201 (2017); Givati, *supra* note 53, at 149–52 (outlining the process and implications of obtaining an advance tax ruling).

55. ROMANO, *supra* note 53, at 80.

These advance tax rulings bind the tax authority to the tax arrangements set out in the ruling, but only for the one specific taxpayer.[56] Essentially the taxpayer is asking the tax authority to turn an ex post standard into a specific rule that applies solely to her circumstances. These advance rulings have a variety of benefits. Most prominently, they provide greater legal certainty to the taxpayer.[57] They eliminate the uncertainty costs of the standard.[58] But such rulings can be costly to generate.[59] The tax authority is essentially engaged in personalized rule making. It is incurring high ex ante decision costs by enacting a rule that applies to just one taxpayer.[60]

Now imagine the tax authority could create a system where a taxpayer simply turns to a machine to answer her tax questions. She could, for example, turn to an agency website or a mobile app. She could ask the machine whether her tax arrangements will expose her to liability and the machine could quickly read the entire tax code, all relevant cases, all associated regulations, and all relevant advisory opinions. The machine could immediately provide an answer to the taxpayer's question.[61]

The tax authority, thus, could use this artificially intelligent machine to provide advance tax rulings. Depending on the underlying objective of the legislature, the tax authority could use the machine to identify optimal rules that allow it to generate more revenue with greater efficiency and fewer distortions on market behavior. It could use this technology very broadly to choose very specific rules that are highly calibrated to legislative objectives without introducing compliance costs that would otherwise be associated with such complexity.

If regulators adopt these technologies, the answers provided by the tax authority would essentially become the red or green lights of tax law. Even though the underlying tax laws would be very complex, the directives provided to an individual would be simple. Any enforcement agent could adopt technology of this kind.[62] As

56. In the United States, these rulings ("private letter rulings") are "binding on the IRS if the taxpayer fully and accurately described the proposed transaction in the request and carries out the transaction as described." *Understanding IRS Guidance—A Brief Primer*, IRS (last updated July 6, 2016), https://www.irs.gov/uac/Understanding-IRS-Guidance-A-Brief-Primer [https://perma.cc/RGV3-APFP]; *see also* 26 C.F.R. § 601.201(a)(1)–(2), (l); Givati, *supra* note 53, at 149–50.

57. ROMANO, *supra* note 53, at 77–78.

58. Givati, *supra* note 53, at 147.

59. ROMANO, *supra* note 53, at 277–80.

60. *See* Givati, *supra* note 53, at 149. As a formal matter, the rulings only resolve the relationship between the tax authority and one specific taxpayer. Further, they have no formal precedential effect for future taxpayers. As a practical matter, however, the tax authority is required to treat taxpayers consistently and so a de facto precedential value arises—but this does not rise to the level of a binding rule for all future cases. *Id.* at 158–61 (discussing the nuances of the precedential value of advance rulings and surveying the legal scholarship on the matter).

61. For more on this process, see Benjamin Alarie, *The Path of the Law: Towards Legal Singularity*, 66 U. TORONTO L.J. 443 (2016); Benjamin Alarie, Anthony Niblett & Albert H. Yoon, *Using Machine Learning to Predict Outcomes in Tax Law*, 58 CANADIAN BUS. L.J. 231 (2016).

62. The Securities and Exchange Commission has a program similar to advance tax rulings where it provides "no-action" letters that state that the staff will not recommend enforcement actions against the individual or entity seeking guidance. The letter has no binding effect on other individuals or entities, and the SEC reserves the right to change its

predictive technology makes it easier to automate such regulatory advance rulings and ensure their accuracy, they will become a common mechanism for the adoption of machine-generated microdirectives.[63]

Judicial microdirectives. Aside from the legislator and the regulator, there is, of course, a third potential rule maker: the judge. But, as they currently function, judges do not quite fit into this model of law making. To be sure, judges could use artificial intelligence and big data to apply standards or complex rules.[64] But judges are not—at least in a formal sense—regularly in the business of providing ex ante notice of the outcomes of hypothetical scenarios.

For better or worse, advisory opinions are frowned upon by the American judicial system. Judges might use the predictive technology to refine the law ex post. But without notice to the regulated actors, those specific rulings impose some of the same costs as standards. For example, if judges announce that all negligence cases will be decided using a computer algorithm,[65] a regulated actor without access to the algorithm would still be faced with nothing more than a standard that imposes uncertainty

---

position. *See No Action Letters*, U.S. SEC. & EXCHANGE COMMISSION, http://www.sec.gov/answers/noaction.htm [https://perma.cc/HP8Q-WCWB] (last updated Sept. 21, 2012). *See generally* Donna M. Nagy, *Judicial Reliance on Regulatory Interpretations in SEC No-Action Letters: Current Problems and Proposed Framework*, 83 CORNELL L. REV. 921 (1998) (describing the no-action letter process).

63. There will be some areas in law where the provision of advance directives is problematic. Tax provides a salient example. For some things, the lawmaker and the individual have aligned incentives. The individual wants to comply with the lawmaker's policy objectives and certainty makes compliance more likely. But for other things, the individual wants formal compliance with law but would prefer to avoid the policy objective. In other words, the individual is looking for a loophole. If the law provides a clear rule and the regulated individual would prefer to circumvent that rule, then certainty provides a road map for avoidance. *See* David A. Weisbach, *Formalism in the Tax Law*, 66 U. CHI. L. REV. 860, 882–84 (1999) (describing the use of "anti-abuse" standards to deal with rule avoidance).

In these spaces, it is difficult to predict how microdirectives will fair. On the one hand, the use of microdirective technology to craft a precise law may shrink the space for avoidance. After all, a perfectly calibrated law will provide no space for avoidance. On the other hand, if the law has any imperfection and the regulated individual has superior private technology, she may use the technology to find the imperfections and craft her behavior (such as creating elaborate tax avoidance mechanisms) to avoid application of the microdirective. These arms race scenarios—where avoidance creates private benefits and private technology is in competition with the lawmakers' technology—suggest areas where standards, such as the anti-abuse standard, will survive to supplement microdirectives. Still, the problem could be solved without standards. Revelation of the microdirective could simply be delayed until immediately after the regulated individual took action. This prevents evasion but also commits the government to the rule ahead of time to avoid bias. On the general idea of delaying the revelation of rules to prevent evasion, see Saul Levmore, *Double Blind Lawmaking and Other Comments on Formalism in the Tax Law*, 66 U. CHI. L. REV. 915 (1999).

64. *Cf.* Porat & Strahilevitz, *supra* note 12, at 1436 ("Under certain circumstances, we want the courts (and advocates in the courtroom) to embrace the science of Big Data as a means of deciding what terms ought to be imported into an ambiguous contract or will.").

65. The hypothetical scenario is not as fanciful as it may sound. The algorithm here is just a more precise amalgamation of the expert opinions that courts routinely rely on in deciding cases.

about how the judges will apply that standard. It would make little difference to the individual that the actual judge happens to be a computer.[66]

Things change if the regulated actors have access to the algorithm that judges will use. In that world, the regulated actors can predict the outcome with precision. If judges commit to using a certain technology that is available to the public, that would be equivalent to providing advance rulings.[67] This would essentially shift the judge's role to that of ex ante regulators. While not implausible, we think the avenues of legislative and regulatory rulemaking will be more pervasive.

There is another way that judges could be involved in the promulgation of microdirectives. Just as legislatures could set a broad policy objective and delegate the rule making to an agency, so too could the courts. In deciding cases, courts can announce a standard that blesses any rule that results from a process aimed at the correct policy objective and that takes into account the relevant factors. The agency could then create an algorithm that does exactly that. This "second-order regulation" by the court would send a message to the agencies on how to design the algorithm to ensure compliance.[68] Here again it would be the agencies and enforcers who have the ultimate responsibility for implementing the machine algorithm to promulgate microdirectives.

2. An Alternative path: Private Use of Technology by Regulated Actors

Predictive technology will be available to private actors. And, in some cases, private actors may have more advanced proprietary technology than legislatures, regulators, or courts.[69] Private use of predictive technology will lead to the emergence of microdirectives. There are two main ways this can occur.

The first path is through the interplay of reasonableness, industry standards, and technology. In our medical example above, imagine that the machine that predicts medical outcomes is available not to lawmakers but only directly to doctors. As the

---

66. It is worth noting that ex post error and inconsistency costs are likely to be lower if the judge is using the algorithm. See the example of bail, *infra* Part II.A.2.

67. In a slightly different but related context, one commentator has suggested that judges could bind themselves to textualist interpretations of statutes by using computers to derive the meaning of text. Betsy Cooper, *Judges in* Jeopardy!*: Could IBM's Watson Beat Courts at Their Own Game?*, 121 YALE L.J. FORUM 87 (2011), http://yalelawjournal.org/forum/judges-in-jeopardy-could-ibms-watson-beat-courts-at-their-own-game [https://perma.cc/8DDC-ZC9L].

68. *See* John Rappaport, *Second-Order Regulation of Law Enforcement*, 103 CALIF. L. REV. 205, 214 (2015) (defining a second-order judicial decision as one that "states its obligations in terms of ultimate goals that must be achieved. The [agent] is then free to achieve those goals in any appropriate way" (quoting STEPHEN BREYER, REGULATION AND ITS REFORM 105 (1982))). This can be done for all standards including those that the court applies pursuant to the Constitution. *See infra* Part III. Rappaport's example of the court's second-order regulation of Fourth Amendment searches, Rappaport, *supra*, at 220–22, is an area where the death of rules and standards will be swift. Machine algorithms will be able to easily determine probable cause, exigent circumstances, bias of officers, and the like better than humans.

69. For example, private traders of securities have technology that permits higher frequency trading than regulators can observe and monitor in real time. *See* Eric Budish, Peter Cramton & John Shim, *The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response*, 130 Q.J. ECON. 1547 (2015).

technology becomes more accurate we can expect more and more doctors to use it. At some point, it is likely that courts will begin to deem it *per se* unreasonable to *not* use such advanced technology. Imagine an orthopedic practice today that did not use an x-ray machine[70] or a colorectal specialist who refused to perform colonoscopies in diagnosing colon cancer.[71] As technology becomes more accurate and widespread, the likelihood that courts will base a reasonableness standard on the use of that technology increases. The proliferation of these technologies across industries will cause behavior that complies with standards to function exactly as if it were complying with a microdirective promulgated by the predictive technology.

The second path is a softer version of our main thesis. This path does not require lawmakers to use technology. Individuals can use predictive technology to provide predictions on how judges will apply a standard.[72] In this way, technology improves on the role of lawyers as compliance advisors.[73] When lawyers provide compliance advice, they are, in part, predicting how ex post adjudicators will apply a standard.[74] As computers can gather and analyze more and more prior cases, they will outperform lawyers at this task. On first blush, this advance would appear to reduce the compliance cost of standards. But it does so in a way that effectively turns the standard into a microdirective, as it reduces the costs of legal uncertainty because it tells the individual exactly how to behave. As Oliver Wendell Holmes noted, a prediction of a judicial outcome *is* the law.[75]

Advances in big data and artificial intelligence will spawn intelligent machines that can predict legal outcomes with great accuracy.[76] In our traffic example, imagine

---

70. Doctors have frequently been held negligent for failing to order x-rays. *See, e.g.*, Rudick v. Prineville Mem'l Hosp., 319 F.2d 764 (9th Cir. 1963) (x-rays would have revealed fractured vertebrae); Webb v. Lungstrum, 575 P.2d 22 (Kan. 1978) (x-rays would have revealed small metal fragment in wound); Betenbaugh v. Princeton Hosp., 235 A.2d 889 (N.J. 1967) (negligence found where doctor failed to order x-ray of the injured part of the spine).

71. Doctors who failed to order a colonoscopy have been held negligent. *See, e.g.*, Morse v. Davis, 965 N.E.2d 148 (Ind. Ct. App. 2012).

72. McGinnis & Wasick, *supra* note 1, at 1033–39 (discussing the use of technology to predict judicial outcomes).

73. *See generally* Michele DeStefano Beardslee, *Taking the Business Out of Work Product*, 79 FORDHAM L. REV. 1869, 1874–81 (2011) (arguing that it is difficult to distinguish between business and law in corporate practice); Christine Parker, *Lawyer Deregulation via Business Deregulation: Compliance Professionalism and Legal Professionalism,* 6 INT'L J. LEGAL PROF. 175 (1999) (exploring the role of lawyers as compliance officers); Gregory J. Millman & Samuel Rubenfeld, *Compliance Officer: Dream Career?*, WALL ST. J. (Jan. 15, 2014, 8:13 PM), http://www.wsj.com/articles/SB10001424052702303330204579250722114538750 [https://perma.cc/E5AW-2Z54] (examining the rise of compliance officers).

74. On the law-and-economics of ex ante legal advice, see Louis Kaplow & Steven Shavell, *Private Versus Socially Optimal Provision of Ex Ante Legal Advice*, 8 J.L., ECON. & ORG. 306 (1992). *See also* Lynn M. LoPucki & Walter O. Weyrauch, *A Theory of Legal Strategy*, 49 DUKE L.J. 1405 (2000) (discussing the law-and-economics literature of legal culture and legal strategy).

75. Holmes, *supra* note 7, at 461 ("The prophecies of what the courts will do in fact . . . are what I mean by the law."); *see also* Abramowicz, *supra* note 27.

76. *See* Katz, *supra* note 11 (exploring the power of big data to predict legal outcomes); *see also* Porat & Strahilevitz, *supra* note 12, at 1436 (same).

that traffic is regulated only with yield signs that impose a reasonableness standard.[77] But in this world, consumer technology has advanced to a stage where it can predict when a court will deem yielding to be required under the standard. This private technology provides a mechanism for informing the driver when she must stop under the law. The technology gathers the relevant facts, applies the standard to those facts as a judge would, and provides predictive analysis.

Even though we have standards and private technology, the resulting behavior looks *as if* we had public traffic lights with underlying complex rules. And compliance is as simple for the driver as it would be with a microdirective. The driver simply gets a message saying stop. She does not have to even take mental note of the underlying facts. As technology makes ex post adjudication more predictable, citizens treat a prediction as a rule. They receive directives ex ante and have little uncertainty about how the law requires them to behave.

This may lead lawmakers to simply enact those predictions as law. It is possible, though, that lawmakers may deem fully predictable ex post adjudication to be the satisfactory equivalent of a microdirectives and not take the final step to formalize the microdirectives into law. But from the individual's perspective the transformation will be already complete. Drivers will know to stop when the technology in their car gives a signal—the equivalent of a red light.[78]

II. FEASIBILITY

In this Part, we examine the feasibility of using technology to generate microdirectives. This Part is divided in two main sections. First, we examine the feasibility of predictive technology. We look at examples where big data and artificial intelligence have been used to generate better predictions and insights than humans ever could provide, and we look to where the technology is headed. We look at how such predictive technology has dramatically diminished the need for human discretion.

Second, we examine the feasibility of communication technology. For the most part, this technology is already here and steadily improving. Mobile devices are becoming our first port of call for information. Individuals can easily and quickly communicate with other individuals, and—more importantly for our argument—lawmakers can easily and quickly communicate with regulated actors.

A. The Feasibility of Predictive Technology

There are two key takeaways from this section: (1) machines are, in many areas,

---

77. *See, e.g.*, 625 ILL. COMP. STAT. 5/11-904(c) (West 2016) (requiring that a driver at a yield sign slow to "a speed reasonable for the existing conditions" and stop "if required for safety"); MASS. GEN. LAWS 89 § 9 (LexisNexis 2012) (same); Pierce v. Coltraro, 252 So.2d 550, 552–53 (La. Ct. App. 1971) (noting the standard that applies at a yield sign).

78. It is possible that judges, knowing about the predictive technology, will (consciously or unconsciously) respond by changing their behavior. If that were true, and assuming that advanced algorithms could not account for the changes when making predictions, that would suggest that technology for predicting judicial outcomes would lag behind other predictive technology in effectiveness. This alternative path toward the death of standards would, therefore, be less likely than the paths through legislative and regulatory rule making.

already better at predicting outcomes and behavior than any human; and (2) this technology is improving so rapidly that the superiority of machines in predicting outcomes will continue to grow at an exponential rate.

Machines can process billions of data points instantly to determine an optimal course of action. Even the most competent, objective humans cannot compete with algorithms generated by big data and artificial intelligence. We are producing and analyzing ever-increasing stores of data that will provide the backbone of predictive technology. It may be difficult to envision these longer-term trends, but as Bill Gates has noted: "We always overestimate the change that will occur in the next two years and underestimate the change that will occur in the next ten."[79] One can only imagine the extent to which we underestimate the change that will occur in the next twenty years, or by the end of this century.

In this section, we first explore how technological developments will improve the prediction of human behavior by better understanding and analyzing millions of hypothetical situations. We foreshadow the future growth of cognitive computing, artificial intelligence, and evolutionary algorithms to show how these powerful new technologies will facilitate the emergence of microdirectives. We then look at how human discretion is being replaced by computer-based rules in all professions and argue that law is no different.

1. The Power of Predictive Technology

Big data and artificial intelligence have reached a stage where likely outcomes can already be predicted in many aspects of human life.[80] By the end of the last century, computing machines were able to defeat the best grandmasters in chess.[81] A decade later, an artificially intelligent machine destroyed the grandmasters of the television trivia show *Jeopardy!*.[82] Indeed, machines outperform humans in many areas of life.[83] They can predict consumers' taste[84] and advise clients on financial

---

79. BILL GATES, THE ROAD AHEAD 316 (1996).

80. *See generally* Katz, *supra* note 11.

81. In 1997, IBM's Deep Blue defeated Garry Kasparov 3½ games to 2½. For commentary on and descriptions of the match, see BRUCE PANDOLFINI, KASPAROV AND DEEP BLUE: THE HISTORIC CHESS MATCH BETWEEN MAN AND MACHINE (1997).

82. John Markoff, *Computer Wins on* 'Jeopardy!'*: Trivial, It's Not*, N.Y. TIMES (Feb. 16, 2011), http://www.nytimes.com/2011/02/17/science/17jeopardy-watson.html [https://perma.cc/CM9S-7CQ4].

83. *See* MARTIN FORD, THE RISE OF THE ROBOTS: TECHNOLOGY AND THE THREAT OF A JOBLESS FUTURE (2015).

84. *See generally* THOMAS W. MILLER, MODELING TECHNIQUES IN PREDICTIVE ANALYTICS: BUSINESS PROBLEMS AND SOLUTIONS (2014). As an example, artificially intelligent machines can predict wine prices better than wine connoisseurs. *See Quants and Quaffs*, THE ECONOMIST (Aug. 8, 2015), http://www.economist.com/news/science-and-technology/21660405-artificial-intelligence-may-beat-connoisseurship-quants-and-quaffs [https://perma.cc/7PCL-X6U9]. Companies such as Amazon, Netflix, and Match.com have all used machine learning algorithms to better understand consumers' tastes. Pedro Domingos, *Why Businesses Embrace Machine Learning [Excerpt]*, SCI. AM. (Oct. 29, 2015), https://www.scientificamerican.com/article/why-businesses-embrace-machine-learning-excerpt/ [https://perma.cc/U5Y8-FNQG].

opportunities.[85] In the field of medicine, computers can analyze images and predict the likelihood of cancer.[86]

But today's use of big data and algorithms to predict outcomes is just the beginning. The capacity of computers to process information and collect and store data continues to explode.[87] The Director of Engineering at Google, Ray Kurzweil, recently noted: "There's a very smooth exponential increase in the price-performance of computing going back to the 1890 census."[88] As economist Professor William Nordhaus notes, the increase in computer power over the course of the twentieth century was "phenomenal," [89] improving manual computing power by a factor of between 1.7 trillion and 76 trillion times with an explosive trend beginning only after the Second World War.[90]

The growth in computational power has closely tracked "Moore's Law" over the past fifty years.[91] Moore's Law is the observation that the number of transistors in a dense integrated circuit doubles approximately every two years.[92] This observation has proved remarkably accurate and is now used as a guide to understanding where computing will be in the future.[93] If the trend continues, then within twenty years, computing power will be 1000 times what it is today.[94]

That trend will allow computing technology to expand its influence. In the same way that city planners have already developed computers that track aggregate traffic flows,[95] governments will likely be able to collect and use data on how humans

---

85. Brad Power, *Artificial Intelligence Is Almost Ready for Business*, HARV. BUS. REV. (Mar. 19, 2015), https://hbr.org/2015/03/artificial-intelligence-is-almost-ready-for-business [https://perma.cc/ZA2K-57R3].

86. IBM's Watson, the same artificially intelligent process that defeated the grandmasters of *Jeopardy!*, has been used in the medical context. Carl Zimmer, *Enlisting a Computer To Battle Cancers, One by One*, N.Y. TIMES (Mar. 27, 2014), http://www.nytimes.com/2014/03/27/science/enlisting-a-computer-to-battle-cancers-one-by-one.html [https://perma.cc/89VR-LSYN].

87. *See* Martin Hilbert & Priscila Lopez, *The World's Technological Capacity To Store, Communicate, and Compute Information*, 332 SCIENCE 60 (2011) (estimating the growth of computing power and capacity); *see also Data, Data Everywhere,* ECONOMIST (Feb. 25, 2010), http://www.economist.com/node/15557443 [https://perma.cc/ZJZ2-8MHN].

88. *Ray Kurzweil on the Price-Performance of Computing*, WALL ST. J. ONLINE (Aug. 20, 2013, 1:29 AM), http://www.wsj.com/video/ray-kurzweil-on-the-price-performance-of-computing/C1F2B611-4B92-469C-AA33-3129587EC113.html [https://perma.cc/X8DZ-A9FL].

89. William D. Nordhaus, *Two Centuries of Productivity Growth in Computing*, 67 J. ECON. HIST. 128, 128 (2007).

90. *Id*. at 142–47.

91. *See* Lundstrom, *supra* note 28; Moore, *supra* note 28.

92. Lundstrom, *supra* note 28, at 210.

93. Indeed, some suggest that Moore's Law is akin to a self-fulfilling prophecy. *E.g.*, Harro van Lente & Arie Rip, *Expectations in Technological Developments: An Example of Prospective Structures To Be Filled in by Agency*, *in* GETTING NEW TECHNOLOGIES TOGETHER: STUDIES IN MAKING SOCIOTECHNICAL ORDER 206 (Cornelis Disco & Barend van der Meulen eds., 1998).

94. If this exponential trajectory continues to hold, by the end of this century, computing power will be over one trillion times what it is now.

95. *See, e.g.*, Todd Litman, *Generated Traffic: Implications for Transport Planning*, ITE J., Apr. 2001, at 38; *see also* Thomas Liebig, Nico Piatkowski, Christian Bockermann &

behave in almost all aspects of life. But the growth of data collection and analytics will not be uniform in all areas of law. The evolution will be fastest where regulated actors' behavior is more frequent and more homogenous. In these situations, lawmakers will have more data on how individuals behave. Where behavior is less frequent and more heterogeneous, the predictability of behavior will initially be weaker.

In the long run, however, artificially intelligent machines will not be bound by the limits currently facing big data.[96] Artificially intelligent machines are not simply programmed with a given structure to anticipate every possible contingency and every possible answer. Rather, artificially intelligent machines are trained to predict, infer, and intuit behavior and adapt to new and unique situations.[97]

Artificially intelligent machines find "hidden" or "deep" connections in unstructured data to provide stronger predictions.[98] In some sense, these machines are capable of "learning."[99] They update to take into account whether their best guesses are correct or not. In doing so, they amalgamate the wisdom of crowds.[100] Artificially intelligent machines marshal this wisdom better than traditional statistical techniques because the machines craft their own learning rules, rather than relying on a potentially biased structure imposed by humans.[101]

2. Predictive Technology Will Displace Human Discretion

In the near future, more perfect algorithms will begin to displace lawmaker

---

Katharina Morik, *Dynamic Route Planning with Real-Time Traffic Predictions*, 64 INFO. SYS. 258 (2017).

96. *See, e.g.*, Daniela Rus, *The Robots Are Coming: How Technological Breakthroughs Will Transform Everyday Life*, FOREIGN AFF., July/Aug. 2015, at 2. *But see* Martin Wolf, *Same as It Ever Was: Why Techno-Optimists Are Wrong*, FOREIGN AFF., July/Aug. 2015, at 15 (providing a skeptical view).

97. *See generally* STUART RUSSELL & PETER NORVIG, ARTIFICIAL INTELLIGENCE: A MODERN APPROACH (3d ed. 2009).

98. *See, e.g.*, Geoffrey E. Hinton, Simon Osindero & Yee-Whye Teh, *A Fast Learning Algorithm for Deep Belief Nets*, 18 NEURAL COMPUTATION 1527 (2006).

99. Machine-learning algorithms learn by recognizing features, concepts, principles, and ideas that humans instinctively recognize but find difficult to program or code. Rather than having to structure a program in order to code rules, the rules are crafted and understood by the artificially intelligent machine.

100. *See* JAMES SUROWIECKI, THE WISDOM OF CROWDS: WHY THE MANY ARE SMARTER THAN THE FEW AND HOW COLLECTIVE WISDOM SHAPES BUSINESS, ECONOMIES, SOCIETIES AND NATIONS (2004).

101. Bias may affect algorithms that are based on traditional statistical techniques if some errors are not observable. If the bias is not corrected, errors can be replicated and reinforced by using the algorithm. But recently, a branch of artificial intelligence called *evolutionary computation* has been developed to deal with such problems. Evolutionary algorithms, based on techniques used by evolutionary biologists, use elements of trial and error to search for globally optimal solutions, rather than simply optimizing with the existing space. Candidate solutions are tested using an iterative process. *See* David B. Fogel, *Introduction to Evolutionary Computation*, *in* EVOLUTIONARY COMPUTATION 1: BASIC ALGORITHMS AND OPERATORS (Thomas Bäck, David B. Fogel & Zbigniew Michalewicz eds., 2000).

discretion. While this displacement of human discretion may appear novel in the legal sphere, it is simply a manifestation of the *Moneyball* phenomenon highlighted by Michael Lewis.[102]

In the book *Moneyball*, Lewis explores the use of data in major league baseball to elucidate the idea that statistics and data, used correctly, are superior to human judgment. Scouts and coaches in baseball previously relied on the "look" of the player to predict whether a player would make it in the big leagues.[103] But they were wrong. Their hunches were really just manifestations of years of inherited biases, prejudice, and outdated modes of thinking. Taking advantage of this, the Oakland A's used statistical analysis to consistently outperform rivals who had greater financial resources.[104]

The lesson here is that humans and their hunches are unreliable.[105] Examples can be found everywhere. From bankers assessing loan applicants[106] and employers hiring prospective employees[107] to commercial pilots flying planes,[108] humans increasingly place their trust in machines and discover that outcomes predicted by big data are systematically better than human intuition.

The phenomenon is starting to permeate the field of law. Consider how judges set bail. The decision to set bail has historically been based on a standard. The judge weighed a number of factors, such as the seriousness of the alleged crime, the likelihood of guilt, whether the defendant had jumped bail before, the defendant's social ties and employment situation, the defendant's mental condition, and so on.[109] The

---

102. MICHAEL LEWIS, MONEYBALL: THE ART OF WINNING AN UNFAIR GAME (2003).

103. *Id.* at 32.

104. With limited financial resources, the A's were able to make the playoffs year after year by performing detailed statistical analyses of players to build a cost-effective, winning team that outperformed other teams with far higher payrolls. LEWIS, *supra* note 102.

105. *See* KAHNEMAN, *supra* note 20, at 4.

106. Bank managers who must decide whether or not to give a customer a loan have seen their discretion dissolve. Banks have turned to predetermined rules about who can borrow and how much they can borrow. The human bank manager is left with little discretion. The algorithm outperforms any individual bank manager in determining the viability of a customer.

107. Employers that use statistical analyses when hiring workers make better hiring decisions than humans that make hiring decisions based on a one-hour interview. *See, e.g.*, Chen-Fu Chien & Li-Fei Chen, *Data Mining to Improve Personnel Selection and Enhance Human Capital: A Case Study in High-Technology Industry*, 34 EXPERT SYS. WITH APPLICATIONS 280 (2008); Nathan R. Kuncel, David M. Klieger & Deniz S. Ones, *In Hiring, Algorithms Beat Instinct*, HARV. BUS. REV., May 2014, at 32.

108. Commercial airline pilots rely heavily on autopilot technology and are instructed not to take control of the airplane under certain circumstances. For example, the 2009 crash of Air France 447 into the Atlantic Ocean would have likely been prevented if the copilot did nothing and did not touch the controls when the plane encountered turbulence. BUREAU D'ENQUETES ET D'ANALYSES POUR LA SECURITE DE L'AVIATION CIVLE, FINAL REPORT ON THE ACCIDENT ON 1ST JUNE 2009 TO THE AIRBUS A330-203 REGISTERED F-GZCP OPERATED BY AIR FRANCE FLIGHT AF447 RIO DE JANEIRO–PARIS (2012), http://www.bea.aero/docspa/2009/f-cp090601.en/pdf/f-cp090601.en.pdf [https://perma.cc/WYZ3-PA6H].

109. See, for example, the standard in Massachusetts where bail is determined by examining the alleged crime, the likely penalty, the likely flight risk, history of defaults, family in the area, employment status, and previous criminal records, among other criteria. MASS

list of potentially relevant factors is almost inexhaustible.[110]

But now some jurisdictions are turning to predictive technology to reduce uncertainty and inconsistency in judges' decisions, as well as to reduce the time taken to set bail.[111] Algorithms have been developed that seek to predict when particular defendants will likely skip bail.[112] The predictive power of this algorithm, which takes into account data on the defendant's characteristics, far exceeds that of any individual judge.[113]

This output from the data is more systematic and reliable than an individual judge's hunch. The algorithm reduces error costs (it is better at assessing the likelihood of a defendant jumping bail) and decision costs (judges can simply apply the algorithm). Judges without the algorithm have less information and cannot process the information they do have as efficiently.

Moreover, judges introduce bias into the system by considering irrelevant factors. A well-meaning judge may not even know when she is considering irrelevant factors. A machine does not suffer from this problem. Relatedly, machines can be instructed to ignore factors that we do not want the law to consider. Thus a machine can be told to ignore race, gender, religion and the like even if they are relevant to an outcome objective. It is much harder for a judge to affirmatively ignore subconscious impacts of such factors.[114]

These observations run counter to the idea that there is something "special" and "human" about the law and legal reasoning.[115] Almost every profession thinks their profession is special.[116] In the same way that most drivers believe that they are above average,[117] humans reflexively believe that their judgment and reasoning is special

---

GEN. LAWS ANN. ch. 276, § 57 (West 2015).

110. *Id.*

111. Shaila Dewan, *Judges Replacing Conjecture with Formula for Bail*, N.Y. TIMES (June 26, 2015), http://www.nytimes.com/2015/06/27/us/turning-the-granting-of-bail-into-a-science.html [https://perma.cc/B267-HDQR].

112. *Id.* Algorithms of this type are already in use in twenty-one jurisdictions across the United States, in places such as Arizona, Illinois, New Jersey, and Pennsylvania.

113. Recent work has illustrated the value of machine learning in reducing errors in granting bail. Crime can be reduced by up to 24.8% with no change in jailing rates, or jail populations can be reduced by 42.0% with no increase in crime rates. Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig & Sendhil Mullainathan, *Human Decisions and Machine Predictions*, (Nat'l Bureau Econ. Res., Working Paper No. 23180, 2017), http://www.nber.org/papers/w23180.pdf [https://perma.cc/QRG3-ZY6W].

114. For machine algorithms, the instruction to ignore prohibited factors is not perfect. Other allowed variables may perfectly correlate with and therefore inadvertently proxy for out-of-bounds factors. These are sometimes called "clones." To the extent a lawmaker wants to exclude a factor from calculation, a programmer has to account for the correlation of clone variables.

115. *See* Joseph Raz, *Reasoning with Rules*, 54 CURRENT LEGAL PROBS. 1 (2001) (asking what is special about legal reasoning).

116. *See generally* MICHAEL A. BISHOP & J. D. TROUT, EPISTEMOLOGY AND THE PSYCHOLOGY OF HUMAN JUDGMENT 24–53 (2005) (humans instinctively deny or ignore the success of such technology because of deep-seated cognitive biases, such as overconfidence in our own abilities and judgments).

117. Iain A. McCormick, Frank H. Walkey & Dianne E. Green, *Comparative Perceptions*

and that technology cannot replicate or replace their particular skill. In *Moneyball*, baseball scouts thought that their ability to pick a future major league star would outperform any statistical analysis.[118] Doctors similarly think that doctors possess special skills.[119] The same is true of teachers.[120] Lawyers too.[121] Perhaps more surprisingly, legal scholars accustomed to looking for biases share this skewed we-are-special belief.

Legal philosophers often contend that law is *necessarily* vague and indeterminate.[122] Others argue that legal reasoning is different from other types of reasoning.[123] Professor Cass Sunstein, for example, suggests that legal reasoning requires an understanding of the principles that underpin reasoning by analogy and has been skeptical that artificial intelligence will be able to replicate this understanding.[124]

Professor Dan Kahan's 2006 address to the graduating class of Yale Law School provides a nice example of the argument that there is something "special" about law.[125] He contends that in order for lawyers to truly understand and evaluate legal reasoning they need years of learning from "grandmasters"[126]—such as professors and senior lawyers—who inculcate students with the power of legal intuition and judgment.[127]

---

of Driver Ability—A Confirmation and Expansion, 18 ACCIDENT ANALYSIS & PREVENTION 205, 206 (1986) (about eighty percent of drivers believe that they are better than the median driver).

118. *See* LEWIS, *supra* note 102, at 29–42.

119. *See, e.g.*, DONALD E. POLKINGHORNE, PRACTICE AND THE HUMAN SCIENCES: THE CASE FOR A JUDGMENT-BASED PRACTICE OF CARE (2004); Samuel W. Bloom, *Structure and Ideology in Medical Education: An Analysis of Resistance to Change*, 29 J. HEALTH & SOC. BEHAV. 294 (1988).

120. *See* Françoise Blin & Morag Munro, *Why Hasn't Technology Disrupted Academics' Teaching Practices? Understanding Resistance to Change Through the Lens of Activity Theory*, 50 COMPUTERS & EDUC. 475 (2008).

121. *See* Jeffrey M. Lipshaw, *The Venn Diagram of Business Lawyering Judgments: Toward a Theory of Practical Metadisciplinarity*, 41 SETON HALL L. REV. 1 (2011).

122. *See, e.g.*, TIMOTHY A. O. ENDICOTT, VAGUENESS IN LAW 1 (2000) ("Although not all laws are vague, legal systems necessarily include vague laws.").

123. Lipshaw, *supra* note 121 (arguing that algorithmic judgment cannot replicate legal reasoning).

124. Cass R. Sunstein, *Of Artificial Intelligence and Legal Reasoning*, 8 U. CHI. LAW SCH. ROUNDTABLE 29, 32–34 (2001) (suggesting that computer programs do not reason analogically the way humans do); *see also* Brian Sheppard, *Incomplete Innovation and the Premature Disruption of Legal Services*, 2015 MICH. ST. L. REV. 1797, 1870 (suggesting that machines are limited in their abilities to understand complex law and cannot perform well and "would not play nice" where standards are currently the dominant form of law).

125. Dan M. Kahan, Deputy Dean and Elizabeth K. Dollard Professor of Law, Yale Law School Commencement Remarks (May 22, 2006), http://digitalcommons.law.yale.edu/cgi/viewcontent.cgi?article=1007&context=ylsca [https://perma.cc/C9ZA-QB4C].

126. *Id.*; *see also* Dan M. Kahan, David Hoffman, Danieli Evans, Neal Devins, Eugene Lucci & Katherine Cheng, *"Ideology" or "Situation Sense"? An Experimental Investigation of Motivated Reasoning and Professional Judgment*, 164 U. PA. L. REV. 349, 355 (2016) (testing whether judges have a unique "situation sense" expertise based on training and experience).

127. Kahan, *supra* note 125.

Professor Kahan compares the profession of lawyers to the profession of chick sexers who determine the gender of one-day old chicks. To the untrained eye, there is nothing discernibly different about newborn male and female chickens. And yet some people with training from a "chick-sexing grandmaster" can examine a chick and tell whether it is male or female with ninety-nine percent accuracy. Amazingly, no one (not even the chick sexers themselves) can say exactly what these experts are looking for. They simply know the difference when they see it. Professor Kahan claims that this "special power to intuitively perceive the gender of a newborn chick" is analogous to how lawyers determine the difference between "good and bad decisions."[128] Professor Kahan argues that lawyers learn how to reason in a special way, and that is what makes the craft of good lawyering so "distinctive" from other professions.[129]

Within four years of Professor Kahan's address, the world of chick sexing had changed dramatically. Predictive technology had been developed that could accurately determine the gender of a chick before birth.[130] Just as the machines defeated the grandmasters of chess and *Jeopardy!*, this new predictive technology bested the grandmasters of chick sexing.

Professor Kahan's address was not about the effect of technology on law. But the fate of chick sexers illustrates a major point: there is nothing so special about individual human intuition—at least in practice—that makes it immune to displacement by technology. Whether pure (unbiased and unrushed) human intuition is special and beyond replication is an unresolved philosophical question. But in application, human intuition is imperfect and biased. And machine technology can, it turns out, do as well as humans even when the individuals themselves cannot adequately describe their intuitive process.

The shortsighted belief that the legal profession is special and that lawyers and judges are immune from displacement by technological advances hinges on a bias that leads one to believe that only a human can deliver such wise judgments and decisions. Yes, lawyers require judgment. Yes, judges require judgment. But, the judgment of one human is outweighed by a decision generated by technology that takes into account millions of judgments and decisions.[131]

To see where this is all going for law, consider how artificially intelligent machines may turn one of the most classic statements of a standard in U.S. legal doctrine into a microdirective. In *Jacobellis v. Ohio*, Justice Potter Stewart found it very difficult to precisely pin down what distinguished pornography from nonpornography in determining the threshold test of obscenity.[132] Instead, he simply wrote: "I know it when I see it."[133]

---

128. *Id.*

129. *Id.*

130. *Hey Little Hen*, THE ECONOMIST: ONLINE EXTRA (Feb. 9, 2010), http://www.economist.com/node/15491505 [https://perma.cc/8BLH-5P3W]. This new predictive technology relies on the detection of estrogen in the yolk of an egg. *Id.*

131. *See* SUROWIECKI, *supra* note 100. One might also note that judgment is about making the right decision in the absence of full information about outcomes. Once the outcomes are known, that sort of judgment is in a sense unnecessary.

132. Jacobellis v. Ohio, 378 U.S. 184, 197 (1964) (Stewart, J., concurring).

133. *Id.*

Justice Stewart's view suggests that distinguishing between pornography and nonpornography is something that humans can do, but it is difficult to write an ex ante rule that clearly defines the line. Justice Stewart preferred to leave the determination as a standard, to be resolved later.

Artificially intelligent technology can already recognize and analyze images.[134] It is not just a judge who can "see it." In the near future, artificially intelligent machines will be able to develop highly complex rules that generate immediate and simple predictions of the legality of particular materials ("this image is/is not pornographic"). Just imagine that Justice Stewart identified fifty pornographic images for the computer. At that point, the artificial intelligence programs can find deep connections to identify the pattern that is driving the distinction, but that Justice Stewart could not articulate. Indeed, such pattern recognition is one of the areas where this technology is already way ahead of humans. And it is why the technology is thought to be so valuable as a diagnostic tool.

If such technology were implemented, the law on the books might still look like a standard; but an individual could refer to the machine output to get advance micro-directives and behave as though she were governed by a rule.

B. The Feasibility of Communication Technology

Can lawmakers adequately give timely notice of the law to regulated actors? Can they provide these individuals with instant notice of how best to comply with the law? In the same way that traffic lights let a driver know that she should stop, communication technology can give rise to a world where all laws are reduced to stop-go directives that are instantly communicated to regulated actors. In this subsection, we discuss the types of technology and infrastructure that will facilitate this.

The costs of communication have been almost obliterated by the Internet. The so-called Internet of Things[135] is an interconnected network of physical objects and devices that are embedded with electronics and sensors to allow products to be controlled and used remotely by the user or manufacturer. Recent estimates suggest that between 50 billion to 100 billion objects and devices will be embedded with such technology by the year 2020.[136] Mobile applications are becoming the first port of

---

134. Four Microsoft researchers have developed a visual recognition program that has an error rate of 4.94%, less than the 5.1% error rate of a human. Kaiming He, Xiangyu Zhang, Shaoqing Ren & Jian Sun, *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*, 2015 IEEE INT'L CONF. ON COMPUTER VISION 1026, http://www.cv-foundation.org/openaccess/content_iccv_2015/papers/He_Delving_Deep_into_ICCV_2015_paper.pdf [https://perma.cc/PC3A-87SQ]. According to one report, Japanese cameras are even being used to identify whether subway passengers are intoxicated. Amber Bouman, *Clever Cameras Detect Drunken Railway Passengers in Japan*, ENGADGET (Aug. 13, 2015), http://www.engadget.com/2015/08/13/clever-cameras-detect-drunken-railway-passengers-in-japan/ [https://perma.cc/RW3E-DVKQ].

135. Kevin Ashton, *That 'Internet of Things' Thing*, RFID J. (Jun. 22, 2009), http://www.rfidjournal.com/articles/view?4986 [https://perma.cc/46LY-7SCS].

136. Maria Farrell, *The Internet of Things—Who Wins, Who Loses?*, GUARDIAN (Aug. 14, 2015, 10:48 AM), http://www.theguardian.com/technology/2015/aug/14/internet-of-things-winners-and-losers-privacy-autonomy-capitalism [https://perma.cc/8T63-WR3F] (estimating 100

call for gathering and processing information.[137]

The Internet of Things and mobile applications are not, however, simply ways to improve the consumer experience. Lawmakers can use this technology. The Internet will facilitate *immediate* communication between lawmakers, regulators, individuals, and corporations. Take, for example, the field of environmental regulation. Regulators could more easily monitor emissions of factories through the Internet of Things. Regulators could instantly determine when factories are exceeding their limits and quickly inform firms operating those factories of the violation.

The example of advance tax rulings above suggests that the IRS will be able to provide immediate compliance information to individuals and corporations using similar technology to the Internet of Things. Regulated actors could enter information into a web-based or mobile application and receive a ruling on a device (like a phone) or some wearable technology (like a watch) from the regulator in a timely manner.

Such infrastructure already exists. For example, cardiologists today can simply refer to an app, enter in relevant information, and be given the optimal response for a patient.[138] As technology improves on the fact-gathering front, individuals may not even be required to enter much data into the programs; rather, devices will simply recognize the contours of the factual situation and give notice of whether the individual is complying with the law.

Indeed, lawmakers could even provide notice of microdirectives when the need for it is immediate. Individuals could wear items, such as contact lenses, that instantly analyze a situation and give an immediate directive as to the legality of a potential action. The military is already experimenting with this type of technology for identifying combat targets.[139] But more mundane uses are also probable: Can you turn left at an intersection? Can you cross the street? Can you attempt to board that subway car? And so on.

The Internet of Things drastically reduces the cost of a lawmaker communicating

---

billion embedded objects); Philip N. Howard, *Sketching out the Internet of Things Trendline,* BROOKINGS: TECHTANK. (June 9, 2015), https://www.brookings.edu/blog/techtank/2015/06/09/sketching-out-the-internet-of-things-trendline/ [https://perma.cc/JZS6-ZVY8] (estimating 50 billion embedded objects).

137. For example, there are 2.2 million different applications at the Apple App store and 2.8 million Google Play applications. *See, e.g.*, *Number of Apps Available in Leading App Stores as of March 2017*, STATISTA, http://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores/ [https://perma.cc/QCL8-NUPW].

138. Winslow, *supra* note 40 (noting the availability of such an app launched by the American College of Cardiology in June 2015).

139. In an attempt to avoid (or, at least, minimize) friendly fire and fratricide, military scientists have developed combat identification technology (known as "Identification Friend or Foe" or "IFF") that can more easily and more quickly identify whether combatants are friendly or enemies. For example, Lockheed Martin recently announced certification to produce an IFF system for aircrews for the United States Department of Defense. *MEADS System Gains Full Certification for Identifying Friend or Foe Aircraft*, LOCKHEED MARTIN (May 21, 2014), http://www.lockheedmartin.com/us/news/press-releases/2014/may/mfc-052114-mead-system-gains-full-certification-identifying-friendfoe-aircraft.html [https://perma.cc/BU3H-6V3N].

with regulated actors. But the Internet of Things does not just facilitate communication. There is a feedback effect. The devices that form part of the Internet of Things also *collect* data on how individuals and corporations behave. Lawmakers can generate even better predictions of human behavior by harnessing such data.[140] In doing so, the Internet of Things will further reinforce the feasibility of the predictive technology.

III. IMPLICATIONS AND CONSEQUENCES

In this Part, we explore some implications and consequences of our predictions. If microdirectives become the dominant form of all law, these implications will be profound. But even a short-run trend in isolated fields of law will have major impacts on the way we think about law. We will focus on the long-run version of our prediction to demonstrate the scope of possible changes.

We suggest that the microdirectives will emerge as a new form of law that reduces the uncertainty costs of standards and the decision and error costs of rules. There are, however, other costs that may arise from a world of microdirectives. While the law will generate less uncertainty and fewer errors, it may be deficient in other ways. Here, we identify four areas where the consequences and potential costs of the emergence of microdirectives may be substantial. First, it will change the broad institutional balance of power in our political and legal system. Second, it may change the development and substantive content of legislative policy. Third, it will transform the practice and training of law. Fourth, it will have moral and ethical consequences for individual citizens, altering their day-to-day decision-making process and changing their relationship with lawmakers and government. In the remainder of this Part, we explore these implications and consequences in general terms. We conclude by noting how the existence of these costs may or may not affect our prediction.

A. The Death of Judging? Institutional Changes to the Legal System

The death of rules and standards will produce a shift in the balance of our political institutions. The proliferation of clear microdirectives largely obviates the need for ex post adjudication. This reduced role diminishes the ability of judges to influence the law and increases the power of ex ante lawmakers.[141] The change in the structure of the law does leave some room for some ex post adjudication of evidentiary questions, but even that will be reduced as the technology for observing facts ex ante improves.

---

140. *Cf.* Porat & Strahilevitz, *supra* note 12 (discussing fact-gathering technology for personalized default rules).

141. *See* Ehrlich & Posner, *supra* note 2, at 261 ("The legislature's choice whether to enact a standard or a set of precise rules is implicitly also a choice between legislative and judicial rulemaking."); *see also* Schauer, *supra* note 3, at 310 ("According to the conventional wisdom, therefore, the choice between rules and standards . . . is an important and powerful implement of institutional design, determining much of who decides what in a complex and multi-institutional society."). In this sense, the death of rules and standards precedes the death of the judicial function.

This is potentially concerning because when judges decide cases they do more than simply apply rules or standards. They also have the ability to shift and modify the law. This can happen in at least three different ways. First, judges can interpret a law (rule or standard) differently than the ex ante lawmakers intended—assuming those lawmakers even had an identifiable intent.[142] Judges can also choose to ignore rules and standards altogether in the guise of interpretation.[143]

Second, judges can influence popular and institutional views about policy objectives. Judges can impact popular opinion by highlighting a particular issue in a case, using their position to make policy statements, or by issuing incremental holdings that generate support for movements that have broader consequences.[144] Additionally, given the U.S. federal system, decisions of courts in one jurisdiction might have larger social consequences that impact nonjudicial change to policy objectives in other jurisdictions.[145] Many think that this role of the courts in challenging stale and

---

142. On the ability of judges to add their own interpretation, see—among many others —STANLEY FISH, IS THERE A TEXT IN THIS CLASS? (1980); Robert W. Bennett, *Objectivity in Constitutional Law*, 132 U. PA. L. REV. 445 (1984); Anthony D'Amato, *Can Legislatures Constrain Judicial Interpretation of Statutes?*, 75 VA. L. REV. 561 (1989); Michael S. Moore, *The Semantics of Judging*, 54 S. CAL. L. REV. 151, 251–52 (1981); Frederick Schauer, *An Essay on Constitutional Language*, 29 U.C.L.A. L. REV. 797 (1982). *But see* William N. Eskridge, Jr., *Overriding Supreme Court Statutory Interpretation Decisions*, 101 YALE L.J. 331 (1991). The idea that legislative intent exists is not obvious. *See, e.g.*, Kenneth A. Shepsle, *Congress Is a "They," Not an "It": Legislative Intent as Oxymoron*, 12 INT'L REV. L. & ECON. 239 (1992).

143. *See, e.g.*, Anthony Niblett & Albert H. Yoon, *Friendly Precedent*, 57 WM. & MARY L. REV. 1789, 1795 (2016) (finding that Court of Appeals judges lean towards citing precedents that align with the political composition of the panel); Anthony Niblett & Albert H. Yoon, *Judicial Disharmony: A Study of Dissent*, 42 INT'L REV. OF L. & ECON. 60, 64–67 (2015) (showing that different judges writing different opinions in the same case cite different precedents and lean toward precedents that align with each judge's political preference); William Hubbard & M. Todd Henderson, *Do Judges Follow the Law? An Empirical Test of Congressional Control Over Judicial Behavior* (Coase-Sandor Inst. L. & Econ., Working Paper No. 671, 2014). *But see* Anthony Niblett, *Do Judges Cherry Pick Precedents to Justify Extra-Legal Decisions?: A Statistical Examination*, 70 MD. L. REV. 234 (2010).

144. The legitimation hypothesis suggests that public opinion will begin to converge toward the opinion of the court after a court has handed down a decision. *See, e.g.*, Brandon L. Bartels & Diana C. Mutz, *Explaining Processes of Institutional Opinion Leadership*, 71 J. POL. 249 (2009); Valerie Hoekstra, *The Supreme Court and Opinion Change: An Experimental Study of the Court's Ability To Change Opinion*, 23 AM. POL. Q. 109 (1995). There is prominent literature discussing "backlash" to court decisions that have the opposite effect of the legitimation hypothesis. *See, e.g.*, Michael J. Klarman, *How* Brown *Changed Race Relations: The Backlash Thesis*, 81 J. AM. HIST. 81 (1994). Others suggest that the effect is more constrained. *See, e.g.*, GERALD N. ROSENBERG, THE HOLLOW HOPE: CAN COURTS BRING ABOUT SOCIAL CHANGE? (1991). Further, court decisions can polarize opinion. *See* Charles H. Franklin, & Liane C. Kosaki, *Republican Schoolmaster: The U.S. Supreme Court, Public Opinion, and Abortion*, 83 AM. POL. SCI. REV. 751 (1989); Timothy R. Johnson & Andrew D. Martin, *The Public's Conditional Response to Supreme Court Decisions*, 92 AM. POL. SCI. REV. 299 (1998).

145. For example, in 2003, the Massachusetts Supreme Judicial Court held that same-sex marriage was legal. Goodridge v. Dep't of Pub. Health, 798 N.E.2d 941 (Mass. 2003). Scholars

entrenched views has a salutary effect on our democracy.[146]

Finally, judges can outright declare policy objectives to be improper or unconstitutional.[147] This judicial review of legislative policy is considered by many to be an integral part of our system of checks and balances.[148]

These lawmaking roles of judges will change along with the fundamental nature of law. Judges will lose much of their oversight and lawmaking power. For nonconstitutional questions, the interpretive role may disappear entirely.[149] They will no longer have the power to reinterpret or ignore laws. The policy objectives of law will be set by the ex ante rule makers (legislative or regulatory). And the judiciary—at least if it maintains its current form and structure—will have little or no occasion to question or change those policy objectives. The opportunities for statutory interpretation and filling in the gaps in vague standards will dry up as citizens are simply instructed to obey simple directives.

---

have debated the effect that this decision had on public opinion and whether the subsequent change in public opinion set the law on a new path, culminating in the Supreme Court of the United States finding a constitutional right to same-sex marriage in 2015. *See* MICHAEL J. KLARMAN, FROM THE CLOSET TO THE ALTAR: COURTS, BACKLASH, AND THE STRUGGLE FOR SAME-SEX MARRIAGE 89–118 (2013); Patrick J. Egan, Nathaniel Persily & Kevin Wallsten, *Gay Rights*, *in* PUBLIC OPINION AND CONSTITUTIONAL CONTROVERSY 234, 239–245 (Nathaniel Persily, Jack Citrin & Patrick J. Egan eds., 2008); Thomas M. Keck, *Beyond Backlash: Assessing the Impact of Judicial Decisions on LGBT Rights*, 43 LAW & SOC'Y REV. 151 (2009); Jane S. Schacter, *Courts and the Politics of Backlash: Marriage Equality Litigation, Then and Now*, 82 S. CAL. L. REV. 1153 (2009).

146. *See generally* ALEXANDER M. BICKEL, THE LEAST DANGEROUS BRANCH (Yale Univ. Press 2d ed. 1986) (1962); Robert A. Dahl, *Decision-Making in a Democracy: The Supreme Court as a National Policy-Maker*, 6 J. PUB. L. 279 (1957).

147. Marbury v. Madison, 5 U.S. (1 Cranch) 137 (1803); *see also* Saikrishna B. Prakash & John C. Yoo, *The Origins of Judicial Review*, 70 U. CHI. L. REV. 887, 887 (2003) ("In [*Marbury v. Madison*], as it is often taught in law schools, the Supreme Court created its authority to declare federal statutes unconstitutional.").

148. *See* JOHN HART ELY, DEMOCRACY AND DISTRUST: A THEORY OF JUDICIAL REVIEW (1980); Barry Friedman, *The Importance of Being Positive: The Nature and Function of Judicial Review*, 72 U. CIN. L. REV. 1257 (2004) (discussing the role that judicial review plays in our legal system); David A. Strauss, *The Modernizing Mission of Judicial Review*, 76 U. CHI. L. REV. 859 (2009) (explaining the role of judicial review in democratic government). There are, of course, strong critics of judicial review. *See, e.g.*, Larry D. Kramer, *Putting the Politics Back into the Political Safeguards of Federalism*, 100 COLUM. L. REV. 215, 233–37 (2000); Larry D. Kramer, *The Supreme Court 2000 Term—Foreword: We the Court,* 115 HARV. L. REV. 5 (2001) (critiquing modern conceptions of judicial review and judicial supremacy); William W. Van Alstyne, *A Critical Guide to* Marbury v. Madison, 1969 DUKE L. J. 1, 38–45 (collecting sources on the question of the legitimacy of judicial review).

149. Deciding nonconstitutional questions is the bulk of what judges do. Constitutional questions, while high profile, reflect a small fraction of the judicial caseload. Even in the Supreme Court of the United States, the percentage of cases has not exceeded fifty percent in recent years. *See* RICHARD A. POSNER, HOW JUDGES THINK 271 (2008). The number is far lower in state courts. *See, e.g.*, Robert A. Kagan, *Constitutional Litigation in the United States*, *in* CONSTITUTIONAL COURTS IN COMPARISON 25, 28 (Ralf Rogowski & Thomas Gawron eds., 2002) (noting that over the period 1940–1970, only 14.6% of state court cases had constitutional issues).

The concern here is a separate question than whether machine-aided algorithms can implement policy objectives. The question is whether there is an independent branch of government with the power to question the policy decisions of the ex ante lawmakers. When the lawmakers decide on legislative objectives and parameters for the machine algorithms, do we want a separate branch of government to review those decisions? If we do, the reduced role of the judiciary is troubling.

Moreover, the number of cases litigated will plummet. The question in most cases will simply be whether or not the citizen complied with the simple directive. The case will have two questions: Was the light red? And did the citizen stop? The evidence to answer those questions will continue to be more readily accessible. As the number of cases and controversies litigated falls and the interpretation of policy becomes unnecessary, a judge's opportunities to use a case to make policy statements and impact opinion will diminish. On the other hand, the judge's opportunities to inject bias and error will also diminish.[150]

Things are a little more complicated for questions of constitutional law. In theory at least, constitutional standards[151] are no different from the standards we have discussed throughout this Article. A machine could easily be programmed to tell us whether a particular search was unreasonable,[152] whether certain speech was pornography, whether microdirectives are valid under the commerce clause or some other provision, and so on.

There are institutional structures, however, that may appear to be barriers to the promulgation of microdirectives for constitutional review. As our regime currently stands, neither Congress nor any agency can dictate that a machine algorithm will decide the constitutionality of laws. If Congress creates an algorithm that takes into account race or results in the prohibition of speech, the courts—and not a machine—would declare those algorithms unconstitutional.

As long as *Marbury v. Madison*[153] remains good law, the constitutional decision on an algorithm would have to come from the judiciary. But the courts could, of course, bless the use of particular types of algorithms going forward, deeming these to be constitutionally proper. For example, an exigent search that was conducted pursuant to a machine directive could be presumed to be reasonable if the machine used a judicially blessed algorithm. This precedential guidance to regulatory agencies[154] could essentially provide the policy objectives that must guide the microdirective technology for constitutional review. This delegation would facilitate the promulgation of microdirectives in the constitutional law space.

Some, of course, may argue that reducing judicial power over policy is a good thing. As the democratically elected branches become more powerful, for example, fears about overreaching by unelected judges will be dampened.[155] Others will

150. *See supra* Part II.

151. The U.S. Constitution generally operates through standards. Schauer, *supra* note 3, at 308.

152. *Cf.* Michael L. Rich, *Machine Learning, Automated Suspicion Algorithms, and the Fourth Amendment*, 164 U. PA. L. REV. 871 (2016).

153. 5 U.S. (1 Cranch) 137 (1803).

154. We discuss this type of second-order regulation above. *See supra* Part I.C.1.

155. *See, e.g.*, LARRY D. KRAMER, THE PEOPLE THEMSELVES: POPULAR CONSTITUTIONALISM AND JUDICIAL REVIEW (2004); *see also* sources cited *supra* note 148.

disagree.[156] For those who advocate the active role of judges, alternative mechanisms for that role might be pursued. Perhaps a judiciary that provides advisory opinions on legislative and regulatory policy decisions could preserve the judiciary's oversight and influence on society.[157]

B. The Development and Substance of Policy Objectives

In addition to policy decisions moving away from judges, the process by which legislatures and regulators make those decisions will change.

Broader objectives. As we noted above in Part I.D., regulatory agents will be the primary force behind the shift to microdirectives. Legislatures may continue to enact standards, but they will leave the machine-aided implementation to regulators. Those standards may be nothing more than a statement of the policy objective that should guide the rule-making machines. Regulators will then translate that broad objective into specific sets of rules generated by machines.[158]

Additionally, the ability to achieve broad goals through machine-derived microdirectives will potentially allow legislatures to state their objectives at increasingly higher levels of abstract social policy. Rather than concern themselves with details of implementation, the legislature will be able to concentrate on the bigger picture. For example, instead of worrying about specific speed limits, the legislature will focus on the purpose of traffic law: does society want laws that reduce accidents, minimize travel time, reduce fuel consumption, or some perfect mix?

At its extreme, learning algorithms aided by big data could be asked to prescribe a vast set of microdirectives covering multiple fields to achieve an even broader social goal, such as maximizing welfare, minimizing accidental death, minimizing wealth inequality, or (more likely) some combination that sets certain acceptable thresholds for these and other social values. This becomes possible because lawmakers no longer have to figure out and set out each precise rule and its connection to other rules. Instead, machines will work out the millions of connected microdirectives that achieve a stated slate of policy objectives. This complex catalog of microdirectives can be targeted to small instances of behavior that fit within a larger web of behavioral actions to achieve a broad goal.

Faster change. Algorithm-driven laws will automatically and rapidly adapt to the circumstances, optimizing according to the objective of the law. But changes to the law result in winners and losers.[159] Frequent changes to the law may impose additional risks on individuals and may affect the willingness of individuals to invest in

---

156. *See* sources cited *supra* note 148.

157. On the other hand, it is possible that the influence on society requires actual cases and controversies to make judicial rulings more salient.

158. *See* Benjamin Alarie, Anthony Niblett & Albert H. Yoon, *Regulation by Machine*, 18 J. MACHINE LEARNING & RES. (forthcoming 2017).

159. *See, e.g.*, MICHAEL J. TREBILCOCK, DEALING WITH LOSERS: THE POLITICAL ECONOMY OF POLICY TRANSITIONS 1–3 (2014) (discussing how changes in legal policy that are not compensating reduces the incentive to make optimal investments); Louis Kaplow, *An Economic Analysis of Legal Transitions*, 99 HARV. L. REV. 509 (1986) (providing an economic analysis of the gains and losses created by legal transitions); Louis Kaplow, *Government Relief for Risk Associated with Government Action*, 94 SCANDINAVIAN J. ECON. 525 (1992) (exploring and

projects that may be subject to legal uncertainty.[160] A smart machine will, however, be able to take into account any effects on the values of reliance investments to find a global optimum, rather than merely a local optimum.

Moreover, predictive technology can be used to advise lawmakers on other potential unintended consequences of certain policy objectives. Under today's system, laws frequently have unintended consequences. Laws that change behavior in unexpected ways that undermine the law's goal or disrupt some unrelated area of human behavior in unexpected ways are common.[161]

With the current state of technology, it often takes years before the consequences of a policy decision are fully understood. But as big data and predictive technology improve, lawmakers will be able to more accurately identify these consequences at the time when they make the rules.[162]

Unintended consequences. The potential for unintended consequences highlights an important facet of the death of rules and standards. Those who set the broad policies in this new world will need to have deep understandings of both social objectives and the way that these technologies work. Machine-generated microdirectives can only reduce unintended consequences if the machines are programmed to identify the types of consequences about which policymakers or society care. The humans who instruct machines to create microdirectives must communicate policy objectives to the machines in ways that do not distort the message, and they must "ask" the machines to provide assessments of the consequences of proposed objectives.

If lawmakers do not have a deep understanding of policy consequences and programming, the machines may distort rather than further law.[163] Such concerns animate many science-fiction movies about the fears of artificial intelligence.[164] The

---

modeling the costs imposed by government transitions).

160. The fact that such uncertainty leads to a reduced ex ante investment is a manifestation of the hold-up problem. *See generally* Paul A. Grout, *Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach*, 52 ECONOMETRICA 449 (1984); Oliver Hart & John Moore, *Foundations of Incomplete Contracts,* 66 REV. ECON. STUD. 115 (1999); Jean Tirole, *Procurement and Renegotiation*, 94 J. POL. ECON. 235 (1986).

161. *See generally* Robert K. Merton, *The Unanticipated Consequences of Purposive Social Action*, 1 AM. SOC. REV. 894 (1936). An example of a recent legislative change with unintended consequences occurred when the Ontario government increased access to the small claims court, which had a regressive effect, with richer plaintiffs displacing poor plaintiffs. *See* Anthony Niblett & Albert H. Yoon, *Unintended Consequences: The Regressive Effects of Increased Access to Courts*, 14 J. EMPIRICAL LEGAL STUD. 5 (2017).

162. The technology may, however, get ahead of itself. While predictive technology reduces the chance of unintended consequences for any given law, it also increases the rate at which laws can be promulgated. If the rate of promulgation increases fast enough, unintended consequences may increase even as laws become more accurate. We must know more about how eager lawmakers will be to promulgate microdirectives to understand how significant this risk is.

163. In the literature on artificial intelligence, this is referred to as "perverse instantiation." *See, e.g.*, NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES, 146–49 (2014).

164. Such dystopian visions of the future are found in many popular books and movies. *See, e.g.*, ISAAC ASIMOV, I, ROBOT (1950); 2001: A SPACE ODYSSEY (Metro-Goldwyn-Mayer 1968); ROBOCOP (Orion Pictures 1987); TERMINATOR 2: JUDGMENT DAY (Carolco Pictures 1991). *See generally* Illah Reza Nourbakhsh, *The Coming Robot Dystopia: All Too Inhuman,*

current debate about Google driverless cars highlights these concerns. Many have questioned how one should program a self-driving car to deal with ethical decisions about the value of life.[165] Can Google cars, they ask, deal with ethical questions that face human drivers? One commentator notes that humans believe that avoiding a collision with a dog is more important than avoiding a collision with animals that are not pets (like squirrels).[166] But machines—if programmed correctly—could replicate that value judgment and, given the advances in predictive technology, would execute the judgment with greater accuracy than a human. On the other hand, if the commentator is wrong—and squirrels are to be avoided with the same care as dogs—then the program can be changed accordingly.[167] The key, then, is in the lawmaker's ability to program that value into the machines.

Still, some appear to worry that poorly programmed cars will implement a frightening system of social values where they swerve to kill the "wrong" people.[168] Implicit in this critique, however, is the false idea that *human* drivers always swerve to kill the "right" people. It would seem that the trick in getting all of this right is not in programming the computer, but in somehow agreeing on which people are the "right" ones to kill. That is an age-old moral problem to which we still do not have an agreed-upon answer. Thus, the so-called "trolley problem"[169] is, indeed, a real one for self-driving cars. But that is a familiar critique on the limits of *human* ethics, not on the limits of self-driving cars. In other words, it is still a problem for human-driven cars too.

FOREIGN AFF., July–Aug. 2015, at 23.

165. *See, e.g.*, Chris Bryant, *Driverless Cars Must Learn To Take Ethical Route*, FIN. TIMES (Mar. 1, 2015), http://www.ft.com/intl/cms/s/0/4ab2cc1e-b752-11e4-981d-00144feab7de [https://perma.cc/D6V9-BEFS]; Patrick Lin, *The Ethics of Autonomous Cars*, ATLANTIC (Oct. 8, 2013), http://www.theatlantic.com/technology/archive/2013/10/the-ethics-of-autonomous-cars/280360/ [https://perma.cc/QQ7V-D4C3].

166. *A Point of View: The Ethics of the Driverless Car*, BBC: MAG. (Jan. 24, 2014), www.bbc.com/news/magazine-25861214 [https://perma.cc/A847-2JRJ].

167. Janet D. Stemwedel, *Building Self-Driving Cars That Drive Ethically*, FORBES (Aug. 5, 2015, 4:45 PM), http://www.forbes.com/sites/janetstemwedel/2015/08/05/building-self-driving-cars-that-drive-ethically/ [https://perma.cc/2QRH-3U5B] (noting that Google is consulting with moral philosophers).

168. *See* Tanay Jaipuria, *Self-Driving Cars and the Trolley Problem*, HUFFINGTON POST: BLOG (June 1, 2015, 12:13 PM), http://www.huffingtonpost.com/tanay-jaipuria/self-driving-cars-and-the-trolley-problem_b_7472560.html [https://perma.cc/CP4M-5GC4] (asking whether cars can make ethical decisions that must value different lives and whether they should favor the life of their owner); Tim Worstall, *When Should Your Driverless Car from Google Be Allowed To Kill You?*, FORBES (June 18, 2014, 8:27 AM), http://www.forbes.com/sites/timworstall/2014/06/18/when-should-your-driverless-car-from-google-be-allowed-to-kill-you/ [https://perma.cc/ZL3N-4A9H] (same).

169. The trolley problem can take different forms but usually presents the question of what one would do if a trolley is on track to kill a group of people and the observer can pull a lever that will divert the trolley to a different course that will kill one (different) person, thus saving the group. On the trolley problem, see generally, PHILIPPA FOOT, *The Problem of Abortion and the Doctrine of the Double Effect*, *in* VIRTUES AND VICES AND OTHER ESSAYS IN MORAL PHILOSOPHY (1978); Judith Jarvis Thomson, Comment, *The Trolley Problem*, 94 YALE L.J. 1395 (1985).

In any event, lawmakers of the future must be able to translate society's values into programmable objectives for the machines. The task of identifying those values, it seems to us, will remain a human one.[170]

C. Changes to the Practice of Law

The observations thus far lead naturally to the next related observation: the death of rules and standards will fundamentally transform the practice of law. For years, a chorus of scholars have been pointing out that technology will disrupt and transform the practice of law.[171] We join this chorus to note that as lawmakers adopt laws that are translated and communicated to citizens as simple microdirectives, the role of lawyers will change dramatically. The role of compliance and litigation lawyers will diminish, while the role of a lawyer as lobbyist or policy advisor will grow.

The compliance lawyer today serves as an intermediary who advises a client on how best to comply with complex rules or vague standards. Part of the expertise of a compliance lawyer is in predicting how an ex post adjudicator will likely apply the relevant standard to a certain set of facts.

Thus, in our tax example, a client might ask a lawyer whether or not her business arrangement complies with the standards of the tax code. In our medical example, a doctor might ask a lawyer whether her diagnostic procedures would be deemed reasonable under the controlling legal standard. The lawyer reads the relevant law and exercises her judgment—based on education, experience, and other expertise—to provide a prediction. The lawyer may go beyond a yes or no answer and suggest creative ways that a client could alter behavior to increase the likelihood that the adjudicator would find the client in compliance.

Technology will reduce the need for such compliance lawyers. The citizen will simply be told directly whether behavior complies with the law or not. There is no need to consult a lawyer to ask whether a traffic light is green or red. Similarly, litigators will no longer be in the business of arguing about the application of standards, and judges will no longer be in the business of applying them.

There will be skeptics. As discussed above in Part II, even though technology has already displaced many labor markets, there is a common sentiment that many hold that *their* profession is different and somehow immune to technological disruptions.[172] But simply noting that a compliance lawyer's role as information middleman will disappear is not to say that the entire profession of law will be automated. Rather, there will be a shift in the types of tasks that lawyers are charged with. Lawyers will be forced to adapt to the new environment.

Setting the policy directives of a machine algorithm is complicated. To tell a machine that its objective is to minimize traffic accidents, without more, would lead to

---

170. There is a possibility that machines could simply observe human behavior and from that deduce what objectives the majority of persons would do and follow that behavior. That would eliminate even the need for human policy considerations. We reject that possibility —not because the computers cannot do it, but because few would agree that entrenching observed majoritarian behavior is the appropriate objective of law.

171. *E.g.*, SUSSKIND, TOMORROW'S LAWYERS, *supra* note 11; Henderson, *supra* note 11; Katz, *supra* note 11; Ribstein, *supra* note 11.

172. *See supra* Part II.A.2.

standstill traffic—or, more absurdly, the prohibition of motor vehicles. Instructing the machine to minimize travel times could lead to an abundance of car accidents. A machine can only write rules to meet the objective as it is presented. As we have discussed, the humans who set the objective must be able to understand the consequences of different objectives and must be able to understand which objectives are desirable.

Understanding the implications of different objectives requires not only an understanding of the technology, but also a highly interdisciplinary understanding of human behavior and the goals of our regulatory state. The trend of the last fifty years toward interdisciplinary legal education,[173] with an emphasis on understanding topics such as economics, psychology, philosophy, history, and so on, is one that will serve this new role of lawyers well. We note in passing that recent countertrends toward so-called practical lawyering[174] are likely to be wasteful. The idea of training lawyers solely in practical skills provides little benefit when the skills required are likely to change rapidly. The understanding of legal policy should remain the focus of the legal endeavor because human individuals will set the high-level policy objectives for the law.[175]

D. The Broader Consequences of These Technologies on Individuals

The death of rules and standards will raise major concerns about privacy, autonomy, and the ethics of human decision making.

---

173. *See generally* ROBIN L. WEST, TEACHING LAW: JUSTICE, POLITICS, AND THE DEMANDS OF PROFESSIONALISM (2014); Harry T. Edwards, *The Growing Disjunction Between Legal Education and the Legal Profession*, 91 MICH. L. REV 34, 34–35 (1992) (documenting the rise of "law and" movements being taught at law schools); Alex M. Johnson, Jr., *Think Like a Lawyer, Work Like a Machine: The Dissonance Between Law School and Law Practice*, 64 S. CAL. L. REV. 1231 (1991); Anthony D'Amato, *The Interdisciplinary Turn in Legal Education* (Nw. Pub. L. & Legal Theory, Research Paper No. 06-32, 2006), https://ssrn.com/abstract=952483 [https://perma.cc/C54C-2XVM].

174. *See* WILLIAM M. SULLIVAN, ANNE COLBY, JUDITH WELCH WEGNER, LLOYD BOND & LEE S. SHULMAN, EDUCATING LAWYERS: PREPARATION FOR THE PROFESSION OF LAW 12 (2007); R. Michael Cassidy, *Beyond Practical Skills: Nine Steps for Improving Legal Education Now*, 53 B.C. L. REV. 1515 (2012); Joe Palazzolo, *Law-School Program Emphasizes Practical Skills*, WALL ST. J. (Jan. 4, 2015, 7:51 PM), http://www.wsj.com/articles/law-school-program-emphasizes-practical-skills-1420419113 [https://perma.cc/Y6X6-B2HU].

175. To be clear, this is not a technological limitation. Policy will be set by humans rather than machines because that is the one area where humans will resist technological advances most strongly. There are, of course, many who think all aspects of life will inevitably be controlled by artificial intelligence. *See, e.g.*, Ray Kurzweil, THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY (2005). Others are skeptical. *See, e.g.*, Peter Murray, *Leading Neuroscientist Says Kurzweil Singularity Prediction a "Bunch of Hot Air"*, SINGULARITY HUB (Mar. 10, 2013), https://singularityhub.com/2013/03/10/leading-neuroscientist-says-kurzweil-singularity-prediction-a-bunch-of-hot-air/ [https://perma.cc/M68W-HX9V]. In any case, the power over ultimate policy objectives will be one of the last things that humans cede to machines.

1. Privacy

Most obviously, as with all applications of big data, the use of data gathering to predict outcomes raises privacy concerns.[176] These concerns have been addressed extensively in other contexts.[177] In our context, the potential for invasions of privacy is high. Government-controlled machines will be gathering data about individual behavior and using that information in two ways. First, they will use the information to assess an individual's behavior and provide a legal directive. Second, they will use the information as part of its aggregated data that goes into setting the microdirectives. Stoplight cameras and GPS tracking already create the ability for the government to know a citizen's comings and goings. These capabilities to invade privacy will increase. And the concerns become greater when the government uses the information it gathers in conjunction with technology to predict future actions by an individual.[178]

There is a trade-off here. The more limitations placed on the government's ability to gather information, the weaker will be its ability to create precise microdirectives.[179] Moreover, there may be privacy-based calls for the halting of microdirectives because the mere prediction based on aggregate data violates principles of privacy.

The debate and policy choices on privacy here are likely to track general debates and choices about privacy and big data. One can also expect that as individuals continue to waive privacy in private-law contexts,[180] public law will be given additional freedom to gather information that facilitates the evolution of microdirectives.

---

176. *See, e.g.*, Farrell, *supra* note 136 (investigating the effect of the Internet of Things on privacy and autonomy, suggesting that they will become the preserve of the powerful).

177. The literature on this new topic is already vast. *See* PRIVACY, BIG DATA, AND THE PUBLIC GOOD (Julia Lane, Victoria Stodden, Stefan Bender & Helen Nissenbaum eds., 2014); Lisa Austin, *Privacy and the Question of Technology*, 22 LAW & PHIL. 119 (2003); Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. REV. 1701 (2010); Paul Ohm, *Sensitive Information*, 88 S. CAL. L. REV. 1125 (2015); Paul Ohm, Response, *The Underwhelming Benefits of Big Data,* 161 U. PA. L. REV. ONLINE 339 (2013); Porat & Strahilevitz, *supra* note 12, at 1467–68; Richard A. Posner, *Privacy, Surveillance, and Law*, 75 U. CHI. L. REV. 245 (2008); Paul M. Schwartz, *Information Privacy in the Cloud*, 161 U. PA. L. REV. 1623 (2013); Daniel J. Solove, *Data Mining and the Security-Liberty Debate*, 75 U. CHI. L. REV. 343 (2008); Omer Tene & Jules Polonetsky, *Big Data for All: Privacy and User Control in the Age of Analytics*, 11 NW. J. TECH & INTELL. PROP. 239 (2013); Omer Tene & Jules Polonetsky, *Privacy in the Age of Big Data: A Time for Big Decisions*, 64 STAN. L. REV. ONLINE 63 (2012).

178. The U.S. government has started to investigate the benefits and costs of using big data. *See* EXEC. OFFICE OF THE PRESIDENT, BIG DATA: SEIZING OPPORTUNITIES, PRESERVING VALUES (2014).

179. *Cf.* Porat & Strahilevitz, *supra* note 12, at 1467–68 (noting the trade-off between privacy protections and "granular personalized default rules"). As predictive technology gets better, less and less personal data will be necessary to create precise microrules. But some information gathering will always be necessary.

180. S*ee id.* at 1468 (noting that "most consumers bring strongly pragmatic perspectives to privacy tradeoffs, and they are increasingly willing to share information about themselves when the benefits from sharing are increased and the threats from sharing are diminished").

2. Autonomy

As lawmakers promulgate more precise microdirectives to advance broad policy objectives, the scope of law can expand. Take, for example, a broad policy objective that seeks to increase productivity. In the hands of a powerful algorithm, microdirectives aimed at a broad goal like that could dictate virtually every decision in a citizen's life. Smart traffic lights could decide who goes first based on productivity levels. Smart restaurants could dictate what a citizen is allowed to eat for breakfast.[181] This presents real concerns for individual autonomy.[182]

But these concerns are not direct objections to the use of predictive technology. Rather they are objections to reckless lawmaking or to overreaching. Lawmakers have to understand what objectives to use in setting microdirectives. Improving productivity might be one policy objective, but there may be other constraining objectives that should be factored in, such as respecting certain spheres of individual decision making. If principles of human autonomy require the law to allow humans to make certain decisions even when those decisions are inconsistent with other social values, then the lawmakers must be aware of those principles and avoid encroaching on them when they set policy objectives. This reinforces the importance of lawyers and lawmakers as interdisciplinary policy experts.

The well-trained expert lawmaker might still overreach. The technologies we have described provide the tools for almost limitless lawmaking. A goal to increase productivity at all costs is difficult to enact through legislation today—the information costs are too high. But that will not be the case with microdirectives. As the information limits on lawmaking fall, it will only be political costs that restrain those in power. As in our discussion of the diminished role of judges, this once again counsels in favor of attention to institutional structures.

A final and perhaps even deeper concern is that lawmakers may turn the microdirectives into *actual* physical restraints on action. Rather than tell you that the light is red, the technology of the future may simply prevent your car from moving. A self-driving car with no driver override could be entirely in the control of the lawmaker.[183]

---

181. There is no doubt that such outcomes would be controversial, as the debate over the "broccoli" analogy in the Affordable Care Act litigation demonstrated. *See* James B. Stewart, *How Broccoli Landed on Supreme Court Menu*, N.Y. TIMES (June 23, 2012), http://www.nytimes.com/2012/06/14/business/how-broccoli-became-a-symbol-in-the-health-care-debate.html [https://perma.cc/J9LD-PJWG]. Another example can be found in New York City's "big-soda ban." *See* Michael M. Grynbaum, *New York's Ban on Big Sodas Is Rejected by Final Court*, N.Y. TIMES (June 26, 2014), http://www.nytimes.com/2014/06/27/nyregion/city-loses-final-appeal-on-limiting-sales-of-large-sodas.html [https://perma.cc/D9TE-VD4Z].

182. The concept of autonomy in law and philosophy is deeply controversial. *See, e.g.*, SARAH CONLY, AGAINST AUTONOMY: JUSTIFYING COERCIVE PATERNALISM (2013). Professor David Strauss has noted that "autonomy is a notoriously vague notion; there is a danger that any attempt to justify a principle in terms of autonomy will slip into question-begging assertions about the nature of truly free and rational human beings." David A. Strauss, *Persuasion, Autonomy, and Freedom of Expression*, 91 COLUM. L. REV. 334, 354 (1991). Still, there is no question that individual autonomy is implicated by the power of the state to create limitless microdirectives to achieve virtually any legislative objective.

183. Additional concerns would arise if individuals do not know who owns the controlling technology. *See* Dan Gillmor, *In the Future, the Robots May Control You, and Silicon Valley*

As the Internet of Things continues to expand, this could be true of most daily actions.

From a technological perspective, the move from microdirectives to automatic restraint is small. From an ethical and policy perspective, however, it is enormous. The benefits of such a move include increased compliance and increased certainty, while the costs arrive by way of a large loss of individual autonomy. One might think such a rule were appropriate if the restraint kept a gun from firing in a situation that would be murder. But things would be different if the rule related to something less malicious like parking in an illegal spot outside a hospital in an emergency or something mundane like crossing a neighborhood street. Our prediction about the death of rules and standards does not necessarily imply prior restraint, which is a topic that needs to be addressed separately.[184]

3. Ethics

There may be additional concerns that the death of rules and standards will erode moral decision making. Some argue that individuals who solely follow rules and directives will become robotic—mere automatons who fail to appreciate the moral choices that should underlie their actions.[185] This is a point made by Professor Shiffrin.[186] Forcing individuals to engage in moral deliberation may be important to the moral health of individuals or of a democratic society. If this is true, the death of rules and standards will bring with it significant costs. We are skeptical that anyone could stop this evolution, so the appropriate response is likely to seek out alternative outlets for human moral deliberation and take that into account in the process of determining the appropriate boundaries of the law.

Finally, people are generally uncomfortable with allowing machines to make important ethical decisions. As discussed above, the debate about Google driverless cars demonstrates this.[187] In that context, many have already begun to ask whether it is acceptable for a machine to make complex ethical decisions about life or death.[188] If we bracket Professor Shiffrin's concerns about moral atrophy, the source of concern appears to arise from a sense that humans have a unique ability to make ethical

---

Will Control Them, GUARDIAN (May 13, 2014, 6:45 AM), http://www.theguardian.com/commentisfree/2014/may/13/internet-of-things-software-privacy-silicon-valley [https://perma.cc/FMA3-TYL8] (noting that autonomy, security, and privacy seem to be an afterthought of the move towards the Internet of Things).

184. *See* Michael L. Rich, *Should We Make Crime Impossible?*, 36 HARV. J.L. & PUB. POL'Y 795 (2013).

185. *See, e.g.*, Evan Selinger & Brett Frischmann, *Will the Internet of Things Result in Predictable People?*, GUARDIAN (Aug. 10, 2015, 11:56 AM), http://www.theguardian.com/technology/2015/aug/10/internet-of-things-predictable-people [https://perma.cc/29C7-5CUR] (noting that people will essentially become programmable, like machines).

186. Shiffrin, *supra* note 10; *cf.* LARRY ALEXANDER & EMILY SHERWIN, THE RULE OF RULES: MORALITY, RULES, AND THE DILEMMAS OF LAW (2001) (investigating the dilemma created between individual moral judgment and rules that restate moral principles in concrete terms).

187. *See supra* Part III.B.

188. *See supra* Part III.B.

decisions.[189] It should be noted, however, that even when a machine is making an algorithmic calculation these are human decisions. Humans decide which values the machine considers. Humans tell it what its objective is.

But that does not necessarily alleviate the concern. There is perhaps an ethical value in having a human making important instant decisions rather than placing ourselves on a course of action that cannot be reviewed in the actual moment. We suspect that a large part of what is going on here is lingering skepticism about accuracy concerns, which we addressed in Part II. People often trust human hunches more than complex machine decisions even in the face of evidence that the machines are more accurate. Perhaps it is a fear of the unknown. But as we have noted above, there is little evidence that humans will be systematically better at making these decisions than machines.

Still a deeper philosophical problem remains. Something that makes us human might be lost when lawmakers use machines to make all of our collective value judgments in advance (even if those judgments are accurate).[190] This may be at the root of the fear with which people view artificial intelligence.

These are pressing ethical problems that will face lawmakers in the future. The current trend is toward microdirectives that reduce in-the-moment ethical decisions. To understand whether that is a good thing, lawmakers must engage with philosophers and ethicists on these questions as the evolution to machine-derived microdirectives progresses.

* * *

Before concluding, it is worth noting an implicit assumption in our prediction: the implications and consequences we discuss here will not themselves prevent the death of rules and standards. One might think that if the institutional upheaval and autonomy concerns are great enough, lawmakers will reject the move to microdirectives. We do not see this happening. The growth of predictive technology is robust. The lure of accuracy ("getting things right") and the regulated actors' desire for certainty are powerful forces that will dominate political and legal debates. The more nuanced considerations we discuss in this Part will, we think, be sidelined.

In that sense, our prediction is about the law's current course. Those who believe the costs of that course are unacceptable should focus on methods of alleviating these costs or finding means to intervene and change that evolutionary path.

CONCLUSION

As machines become increasingly intelligent, and continue to outperform human judgment, the influence of artificial intelligence will spread far and wide. The technologies we have discussed are already being used by doctors to detect cancers, by consumers to optimize their search for products, and by financial advisors to provide advice.

189. *See, e.g.*, Lin, *supra* note 165 (noting that humans are presumed to be able to make ethical judgments, whereas computers have an untested track record).

190. *See generally* MACHINE ETHICS (Michael Anderson & Susan Leigh Anderson eds., 2011).

The legal system will not be immune from this trend. We have suggested throughout this Article that this technological revolution will dramatically alter the foundational structure of law as we know it. Predictive technology will generate greater ex ante information that can be used by lawmakers to write highly specific, complex laws. And individuals will receive notice of these complex laws in a simple form thanks to technological advances in communication. This will be the death of rules and standards and the rise of microdirectives.

These developments will have profound implications for the role of judges, legislators, regulators, lawyers, and individuals in the legal system. But beyond that, we will have to change the way we think and talk about law. Take, for example, the classic debate between legal realists and legal formalists.[191] Without ex post adjudication, this debate changes radically. As standards disappear and judges have progressively less influence, legislative intent will be entrenched and concretized in the catalog of microdirectives.

Technological changes that vastly improve ex ante information will also breathe new life into old law-and-economics models that began with an assumption that lawmakers and citizens have full information. Friction in these models caused by imperfect and asymmetric information has provided a fertile source of material for critics, both inside and outside the field of law and economics. But these models will be given renewed importance. Similarly, the public choice literature will have an increased emphasis on how legislators choose objectives, rather than how they implement laws, while academic interest in subjects such as judicial behavior will dissipate.

All of this is to say that legal institutions of all types will change radically. We are witnessing an information revolution. And, like other technological revolutions, it will precede a legal revolution. The industrial revolution, for example, saw human labor replaced by machine labor and the cost of transportation fell markedly with inventions such as the steam engine. It greatly reduced transaction costs and had widespread impact on all spheres of law including contract law,[192] property law,[193]

---

191. *See generally* BRIAN Z. TAMANAHA, BEYOND THE FORMALIST-REALIST DIVIDE: THE ROLE OF POLITICS IN JUDGING (2010); Steven M. Quevedo, *Formalist and Instrumentalist Legal Reasoning and Legal Theory*, 73 CAL. L. REV. 119 (1985). On American legal realism, see GRANT GILMORE, THE AGES OF AMERICAN LAW (2d ed. 2012); WILFRED E. RUMBLE, JR., AMERICAN LEGAL REALISM (1968); ROBERT SAMUEL SUMMERS, INSTRUMENTALISM AND AMERICAN LEGAL THEORY (1982); Brian Leiter, *American Legal Realism*, *in* A COMPANION TO PHILOSOPHY OF LAW AND LEGAL THEORY (Dennis Patterson ed., 2010); Karl N. Llewellyn, *Some Realism About Realism—Responding to Dean Pound*, 44 HARV. L. REV. 1222 (1931). On formalism, see generally Frederick Schauer, *Formalism*, 97 YALE L.J. 509 (1988); Ernest J. Weinrib, *Legal Formalism: On the Immanent Rationality of Law,* 97 YALE L.J. 949 (1988).

192. *See, e.g.*, P. S. ATIYAH, THE RISE AND FALL OF FREEDOM OF CONTRACT (1979); GRANT GILMORE, *Origins*, *in* THE DEATH OF CONTRACT 5 (Ronald K. L. Collins ed., 2d ed. 1995).

193. For example, to see the feedback effect between intellectual property and the industrial revolution, see Joel Mokyr, *Intellectual Property Rights, the Industrial Revolution, and the Beginnings of Modern Economic Growth*, 99 AM. ECON. REV. 349 (2009).

employment law,[194] criminal law,[195] and tort law.[196]

The information revolution has already resulted in dramatic changes in the world of commerce. For example, companies such as YouTube, Uber, and Airbnb have disrupted and uprooted heavily regulated and stable industries. The coming technological revolution will lead to similar disruption of the legal services industry, but the effect on law will be much deeper and far wider. It will affect the very structure of legal commands and the way we, as a society, choose to govern the behavior of citizens.

---

194. *See, e.g.*, Clark Nardinelli, *Child Labor and the Factory Acts*, 40 J. ECON. HIST. 739 (1980).

195. *See, e.g.*, Douglas W. Allen & Yoram Barzel, *The Evolution of Criminal Law and Police During the Pre-Modern Era*, 27 J.L. ECON. & ORG. 540 (2011).

196. *See, e.g.*, Joel Franklin Brenner, *Nuisance Law and the Industrial Revolution*, 3 J. LEGAL STUD. 403 (1974).

Will Robot Judges Change Litigation and Settlement Outcomes?
A First Look at the Algorithmic Replication of Prior Cases

Anthony J. Casey
University of Chicago Law School

Anthony Niblett
University of Toronto, Faculty of Law

June 5, 2020

Introduction

The promise (or threat) of so-called Robot Judges has captured the attention of popular media and legal scholarship.[1] Recently, the Estonian government made a splash by announcing its plan to use artificial intelligence to decide small claims cases.[2] And similar programs exist or are under development in China,[3] the Netherlands,[4] and other jurisdictions.[5] Despite the hype and attention, little has been said about what Robot Judges are or how they actually work.

Consider three potential candidates for automating judicial decisions: (1) structured code that explicitly stipulates "if x, then y," identifying whether certain factual elements are present, outputting a binary judicial decision; (2) advanced artificial intelligence that

---

[1] See for example, Christopher Markou, *Are We Ready for Robot Judges?*, Discover Magazine (May 15, 2017) available at https://www.discovermagazine.com/technology/are-we-ready-for-robot-judges; Larry Mantle, *Can a Robot Make a Fair Verdict?* Airtalk Podcast available at https://www.scpr.org/programs/airtalk/2019/04/01/64335/can-a-robot-judge-make-a-fair-verdict/.

[2] Eric Niler, *Can AI Be a Fair Judge in Courts? Estonia Thinks So*, Wired (March 25, 2019) available at https://www.wired.com/story/can-ai-be-fair-judge-court-estonia-thinks-so/; Victor Tangermann, *Estonia is Building A "Robot Judge" to Help Clear Legal Backlog*, Futurism.com (March 25, 2019) https://futurism.com/the-byte/estonia-robot-judge.

[3] Monisha Pillai, *China Now AI-Powered Judges*, RADII (Aug. 16, 2019) available at https://radiichina.com/china-now-has-ai-powered-robot-judges/; Chris Young, *China has Unveiled an AI Judge that Will 'Help' With Court Proceedings*, Interesting Engineering (Aug 19, 2019) available at https://interestingengineering.com/china-has-unveiled-an-ai-judge-that-will-help-with-court-proceeding; see also Jingting Deng, *Should the Common Law System Welcome Artificial Intelligence: A Case Study of China's Same-Type Case Reference System*, 3 Geo. L. Tech. 223 (2019); Tom Fish, *AI shock: China Unveils 'Cyber Court' Complete with AI Judges and Verdicts via Chat App,* Express.com (Dec. 6, 2019) available at https://www.express.co.uk/news/science/1214019/ai-china-cyber-court-artificial-intelligence-judges-verdicts-chat-app.

[4] Henriette Nakad-Westrate, Ton Jongbloed, Jaap van den Herik, Abdel-Badeeh M. Salem, *Digitally Produced Judgements in Modern Court Proceedings*, 6 Intn'l J. of Dig. Soc. ational Journal of Digital Society, 1102 (2015).

[5] See generally Tania Sourdin, *Judge v Robot? Artificial Intelligence and Judicial Decision-Making*, 41 UNSW L. J. 1114 (2018); *Council of Europe Adopts first European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems*, Council of Europe (Sept. 13, 2019) available at https://www.coe.int/en/web/cepej/cepej-european-ethical-charter-on-the-use-of-artificial-intelligence-ai-in-judicial-systems-and-their-environment.

can determine the "right" outcome on its own; or (3) predictive litigation algorithms that use historical data about the outcomes of prior cases to determine how the new case fits within the contours of existing law. The first method is of limited use—it can only be used to decide the very simplest of cases where there exists no ambiguity about the relevant factual elements. The second method requires technology that does not exist today. The third method, however, has promise. Litigation assessment algorithms do currently exist, and their prediction outputs could—at least in theory—be converted into automated judicial decisions. Still the question remains as to how exactly one goes about converting predictive outputs into judicial decisions? This question, which has received little attention, is the focus of this paper. And, as it turns out, the answer has major effects on settlement and litigation outcomes.

This article, therefore, explores the effect of using litigation assessment algorithms in judicial decision making. These algorithms use data from past judgments to predict the likely judgment that will result in a new case.[6] They are already in use by lawyers,[7] and several scholars have pointed out their potential value as judicial aids.[8] Indeed, some jurisdictions appear to already be using algorithms to guide judges.[9] And in the extreme, the predictions could even be converted into automated judicial decisions.[10]

---

[6] See, for example, Dru Stevenson & Nicholas Wagoner, *Bargaining in the Shadow of Big Data*, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2325137 (2014); Ben Alarie, Anthony Niblett, & Albert Yoon, 58(3) Canadian Business Law Journal 231 (2016); Sourdin, *supra* note __ at 1125.

[7] Harry Surden, *Artificial Intelligence and Law: An Overview*, 35 Ga. St. U. L. Rev. 1305, 1307 (2019), Rebecca Crootoff, *"Cyborg Justice" and the Risk of Technological-Legal Lock-In*, 119 Colum. L. Rev. Forum 233 (2019). For a general discussion of law firm use see Bernard Marr, *How AI and Machine Learning Are Transforming Law Firms and the Legal Sector*, FORBES, https://www.forbes.com/sites/bernardmarr/2018/05/23/how-ai-and-machine-learning-are-transforming-law-firms-and-the-legal-sector/#6308a31c32c3.

[8] Jingting Deng, *Should the Common Law System Welcome Artificial Intelligence: A Case Study of China's Same-Type Case Reference System*, 3 Geo. L. Tech. 223 (2019); Eugene Volokh, *Chief Justice Robots*, 68 Duke L. J. 1134 (2019); Crootoff, *supra* note __. Separately, judges have also started using algorithms that provide information about other dimensions of litigation. See Megan Stevenson, *Assessing Risk Assessment in Action*, 103 Minn. L. Rev. 2018); Alex Albright, *If You Give a Judge a Risk Score: Evidence from Kentucky Bail Decisions,* (2019) available at https://thelittledataset.com/about_files/albright_judge_score.pdf; see also Rachel Traughber, *Finding a Link to the Human in Algorithms Setting Justice*, (2018) available at https://news.harvard.edu/gazette/story/2018/05/grad-discovers-algorithms-in-justice-system-dont-always-compute/.

[9] Courts in China are reported to being using smart "intelligent case-deciding programs" to provide guidance to judges to ensure that similar cases are decided in the same manner. Deng, *supra* note __.

[10] Deng, *supra* note __;Volokh, *supra* note __; Richard Re & Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 22 Stan. Tech. L. Rev. 242 (2019); Crootoff *supra* note __. Estonia has also proposed the use of automated "robot judges" in small private disputes. Though it is not obvious whether the system will be based on prior case data or simply a coded decision algorithm.

Here, we show that the use of algorithms to assist or automate judicial decision making can distort litigation and settlement outcomes.[11] The particulars of that distortion depend on the methods that judges actually use to translate predictions into judgments. Yet, while scholars have identified the advent of algorithm-assisted judging, few have explored the mechanisms by which a judge might translate the algorithmic output into an actual legal decision. We address that question. We use a simple model of litigation to explore various methods for converting algorithmic assessments into judicial decisions and then demonstrate how the choice of method affects both the legal outcome of the case and the settlement dynamics among the parties.

Importantly for our analysis, legal decisions over liability are typically binary in nature (negligent or not negligent, guilty or not guilty), but algorithmic assessments of liability are typically provided in probabilistic terms (a 70% likelihood). On the surface, this distinction may seem unimportant, and converting probabilities into binary outcomes may *appear* to be a fairly intuitive exercise. After all, in other contexts, courts simply use a balance-of-probabilities or preponderance-of-the-evidence standard. But the conversion of probabilistic outcome predictions into binary legal decisions actually involves difficult choices with major consequences. And yet researchers who study the interplay between algorithms and law – including ourselves – have hitherto overlooked these critically important effects and questions.

This article proceeds in four sections. The first section presents and defines the concept of litigation prediction tools. The second section presents a simple background model of litigation and settlement that will serve as a baseline comparison for our primary analysis. The third section presents a model of various ways judges can use predictive tools and how the choice of method affects litigation and settlement outcomes. The fourth section presents a general discussion.

1. Algorithmic assessments of liability

The possibility of courts using predictive algorithms to guide or even automate judicial decisions is real. And discussions around automated judging have garnered much attention in popular media where the label of "robot judges" has caught on. Estonia has announced plans to implement a system of automated judgments for small claims, and the governments of other countries such as the Netherlands, the United Kingdom, and China are reportedly implementing or looking into implementing similar systems. There are, however, unanswered questions about what exactly it means for a decision to be rendered or even assisted by predictive algorithms or automated data analytics. And the answers to those questions will affect outcomes and settlement dynamics in litigated cases.

In this article, we focus on the judicial use of one specific category of data-driven legal tool: algorithms that use data about past case outcomes to predict the outcome of a

[11] We will refer to "judges," "courts," and judicial decisions throughout this paper. The results of our analysis apply similarly to other arbiters resolving two-party legal disputes. Thus, administrative judges and other regulatory tribunals as well as arbitrators in binding arbitration would fall under the analysis.

new or future case. We call any type of tool that fits this description an algorithmic assessment of liability. These tools can provide an "objective" answer to the question: How does the existing law apply to the facts of the case at hand?[12]

For example, suppose a worker wishes to know whether she is an independent contractor or an employee for the purposes of employment law. In this instance, suppose that *if* the worker was found to be an employee then she would be entitled to an additional payment upon leaving the firm. The legal question is a binary classification problem. In Canada alone, there have been over a thousand legal cases that have provided judicial answers to this question. Suppose now that the information in these cases can be represented in the form of structured data. That is, every single case that has previously answered this legal issue has been coded in a structured way such that the existing law is represented in a dataset. Predictive techniques from statistics or machine learning can then be used to generate the likely outcome — "*independent contractor*" or "*employee*" — for any new hypothetical set of facts. This is a simple binary classification algorithm. It generates a prediction of the most likely outcome as well as a probability associated with that outcome. That is: "In your case, the algorithm predicts you are an *employee* with probability 70%."

At their core, these algorithmic assessments of liability use the data from past decisions to make predictions about how a human judge will apply the law to a given set of facts in the next case. But things can be taken a step further. The predictions might themselves be used by judges to inform their judgments in new cases. And in the extreme the predictions could, themselves, be converted into an automated judicial decision. Others have recognized the possibility of this momentous step. But few have considered how it would work. And, as we demonstrate below, the conversion implies complicated choices with major effects on case outcomes and in turn on the ex ante behavior that the law is intended to regulate. Indeed, even if one assumes an unrealistically simple environment with a perfectly accurate algorithm, the use of algorithmic assessments to automate or merely inform judicial decisions can dramatically change litigation and settlement outcomes.[13]

---

[12] To be sure, there are other tools that might assist judges in reaching a decision. Two in particular come to mind. First, for simple cases, a structured code might simply assess whether certain defined factual elements are present and output a binary judgment. This is certainly possible for very simple cases. (For example, was the car parked in a prohibited spot? If yes, liable.) But for cases with more variables the task of creating such code becomes more difficult and the rigidity becomes an obstacle for any cases that turn on a standard rather than a rule. (For example, was the behavior reasonable?). Second, in theory an advanced machine learning algorithm could be programmed to achieve certain objective and adjudicate cases based on that objective. Anthony J. Casey & Anthony Niblett, *A Framework for the New Personalization of Law*, 86 U. Chi. L. Rev. 333 (2019); Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 Ind. L. J. 1401 (2017). Such technology does not currently exist in any form that could be practically implemented. We have focused on the predictive algorithms based on prior-case data because this technology is currently in use by lawyers and even by judges in some foreign jurisdictions.

[13] The same analysis can be applied to fact-finding. But the analysis—and likely the algorithms—are more complicated, and so we bracket that application for now. We return to the factfinding application in Section 4.

While we assume accuracy in our model, we do not claim that such accuracy exists in the real world. It is well understood that logistical challenges make accurate predictions difficult to achieve. For example, these algorithms need data from a sufficiently large numbers of cases dealing with the relevant factual situations and presenting a sufficient level of consistency. A more heterogeneous area of law will result in lower out-of-sample prediction accuracy. A low number of cases will result in uncertainty about the prediction. And the usual issues of selection bias, publication bias, and coder bias present obstacles to accurate predictions.

Those accuracy and bias questions are important. For the purposes of this paper, however, we bracket them. Our goal here is to highlight another important challenge that exists *even if* the problems of accuracy and bias were solved. Assuming that the predictions are "accurate" in their probabilistic assessment of liability, allows us to isolate the challenge of translating the algorithmic outputs into judgments, which has been ignored in the literature. Like accuracy and bias, this problem is worthy of attention. Despite objections based on accuracy and bias, litigation prediction tools are currently available and are being used by law firms, insurance companies, and litigation finance firms in many contexts. And the move to use them in some courts is already underway and will only expand as more data become available, technologies improve, and accuracy increases. It is important, therefore, to understand the issues around translating predictions into judgment before judicial use of algorithms becomes widespread. This paper is a first step toward that understanding.

2. A baseline model of settlement and litigation

To see how algorithmic assessments of liability change litigation and settlement outcomes, it is first necessary to establish the baseline—that is to say, how cases are decided and settled in the absence of algorithmic assessments. We begin therefore with a simplified baseline model derived from the existing literature where settlement occurs unless there is asymmetric information or optimism on the part of at least one of the parties to the litigation.[14]

---

[14] *See generally* Andrew F. Daughety & Jennifer F. Reinganum, *Settlement*, *in* 8 ENCYCLOPEDIA OF LAW AND ECONOMICS 386, 386–71 (Chris W. Sanchirico ed., 2nd ed. 2012). While the model we present is simplified, it is intended to capture the essential elements of uncertainty, optimism, and asymmetric information from prior literature. *See, e.g.*, Robert Mnookin & Robert Wilson, *A Model of Efficient Discovery*, 25 Games & Econ. Behav. 219, 220 (1998); Steven Shavell, *Suit, Settlement, and Trial: A Theoretical Analysis under Alternative Methods for the Allocation of Legal Costs*, 11 J. Legal Stud. 55, 63 (1982); Richard A. Posner, *An Economic Approach to Legal Procedure and Judicial Administration*, 2 J. Legal Stud. 399, 422 (1973); J. J. Prescott et al, *Trial and Settlement: A Study of High-Low Agreements*, 57 J.L. & Econ. 699 (2014); Abraham L. Wickelgren, *Law and Economics of Settlement*, *in* RESEARCH HANDBOOK ON THE ECONOMICS OF TORTS 331–32 (Jennifer Arlen ed, 2013).

For simplicity, we focus on questions of law.[15] We assume that the parties have agreed on (or at least stipulated to) a set of facts and are asking a judge to apply uncertain law to those facts. In that way, the question can be viewed as akin to a conventional motion for summary judgment or an appeal of a legal issue.[16] Because the facts are stipulated, the only information the parties do not know is how the court will view the law.

Every model of settlement and litigation in the law and economics literature begins with the assumption of a probability of plaintiff success.[17] We include in our model the idea of an "accurate" probability of plaintiff success, p^I. This represents the most accurate representation of the objective probability of plaintiff success given the stipulated set of facts. Later p^I will become relevant when we assume that the predictive algorithm can reveal it to the parties.

The baseline model in this section begins with a world where no party or legal decision maker has access to algorithmic prediction tools. A legal battle in our model is fought between two risk neutral parties, *P* and *D*. *P* is the plaintiff. She claims damages of *L* from D, the defendant. The magnitude of *L* is not in dispute. The dispute lies only in the question of legal liability: given the stipulated facts, is *D* liable for the losses of *P*?

P, the plaintiff, believes she has probability p^P of winning the case. *D*, the defendant, believes that *P* has a probability of p^D of winning. These are subjective probabilities. We also assume that the parties' subjective probabilities are correlated with the accurate and objective probability, p^I. This assumption captures the real-world fact that the parties each possess valuable but imperfect information about the case. With that information they are unable to known p^I exactly, but their estimates of p^P and p^D will reflect some valuable information tying them loosely to p^I . We also assume that both parties are optimistic relative to the objective probability of plaintiff success. To keep things simple, we therefore assume the following relationships:

$$p^P = p^I + \delta$$
$$p^D = p^I - \epsilon$$

This assumption has little importance in the baseline model, but it will become important later on.[18]

We define Δp as the area of disagreement, the difference between the two subjective probabilities of the parties, $p^P - p^D$ (which, by definition, equals $\delta + \epsilon$.)

[15] We discuss applications to fact-finding below in Section 4.
[16] We also abstract away from any investments in the quality of the argument or investments to shine particular light on the facts of the case to make them more favorable.
[17] *See generally* Daughety & Reinganum, *supra note* __.
[18] The function here is an obvious oversimplification. But it captures the relevant elements of the problem that will be necessary later to demonstrate the effects of judicial use of predictive algorithms.

We presume the costs of reaching a settlement are zero, but the costs of a legal determination are positive. The costs for each party are c^P and c^D. For simplicity, these costs are not endogenous. The costs are fixed and are only incurred if settlement fails and the parties go to court. We also assume the "American rule" for awarding cost, where each side bears their own costs.[19] Further, we suppose that each party has information about the other party's beliefs and costs.

The plaintiff and the defendant play a simple game.[20] The timing of this game is as follows:

- Stage 1 – *P* decides whether or not to bring suit.
- Stage 2 – *D* makes a take-it-or-leave-it offer of settlement, S^D.
- Stage 3 – *P* decides to accept or reject *D's* offer. If she accepts, the game is over. If she rejects, both parties pay their respective costs, c^P and c^D, and liability is determined by a neutral third-party judge, *J*. The judge is not a player in the game, merely determining liability after applying the law to the facts. If the judge finds in favor of *P*, she awards compensatory damages of *L,* which must be paid by *D*. In line with our assumption that p^I is an accurate assessment of the likelihood of liability, *D* is found liable with probability p^I.

The equilibrium of this game is simple and intuitive. Working backwards, in stage 3, *P* will accept *D's* offer if and only if the settlement offer, *S*, is greater than the net return *P* would (subjectively) expect to receive from going to court. That is, *P* will settle if and only if:

$$S^P \geq p^P \bullet L - c^P$$

The defendant, in stage 2, will offer a settlement amount up to her net expected losses from litigating the case. The defendant will never offer more than these losses that she (subjectively) would expect to pay at court:

$$S^D \leq p^D \bullet L + c^D$$

We frame the equilibrium in terms of the difference in subjective probabilities. If the two parties have relatively similar assessments of the merits of the plaintiff's case, then it will be in the interest of both parties to settle. If, however, the parties substantially disagree – and p^P is much greater than p^D – then settlement will fail and the parties will proceed to court.

Settlement only occurs if this condition is satisfied:

---

[19] At this stage of analysis, little turns on this assumption.

[20] This is a common type of game theoretic model in the law and economics literature. See e.g., Lucian A. Bebchuk, *Litigation and Settlement Under Imperfect Information*, 15 RAND J. Econ. 404 (1984); I. P. L. P'ng, *Behavior in Suit, Settlement, and Trial*, 14 Bell J. Econ. 539, 541 (1983); Jennifer F. Reinganum & Louis L. Wilde, *Settlement, Litigation, and the Allocation of Litigation Costs*, 17 RAND J. Econ. 557 (1986).

$$\Delta p \leq \frac{c^P + c^D}{L} \qquad (1)$$

For the sake of simplicity and consistency throughout the rest of the paper, we call the right-hand side of this equation the cost-damage ratio. **Figure 1** illustrates the case where the area of disagreement is too large for settlement to be possible in equilibrium.

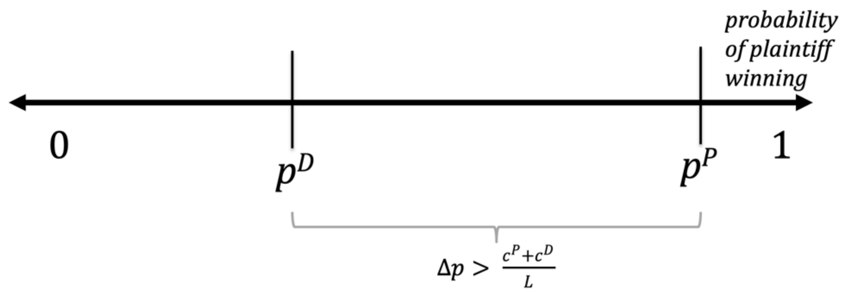

Figure 1: The subjective beliefs of the plaintiff and defendant, showing an area of disagreement greater than the cost-damage ratio. Under these conditions, the two parties are unable to reach a settlement.

But what about when the area of disagreement is not greater than the cost-damage ratio? That is, what actually happens when parties settle? We keep the dynamics of settlement in this model very straight forward. If the settlement condition in equation (1) is satisfied, the defendant in stage 2 makes a take-it-or-leave offer, S^*, where $S^* = p^P \bullet L - c^P$. In equilibrium, this offer is accepted by the plaintiff in stage 3. The dispute is settled; it does not proceed to a judge.

Given our assumption about the correlation between the subjective probabilities p^P and p^D and the objective probability p^I, the magnitude of the settlement offer is also correlated to the accurate measure of the objective probability of the plaintiff's success:

$$S^* = (p^I + \delta) \bullet L - c^P \qquad (2)$$

This relationship between S^* and p^I is continuous. The greater the likelihood of the plaintiff winning, the larger the settlement amount. This is illustrated in **Figure 2**.

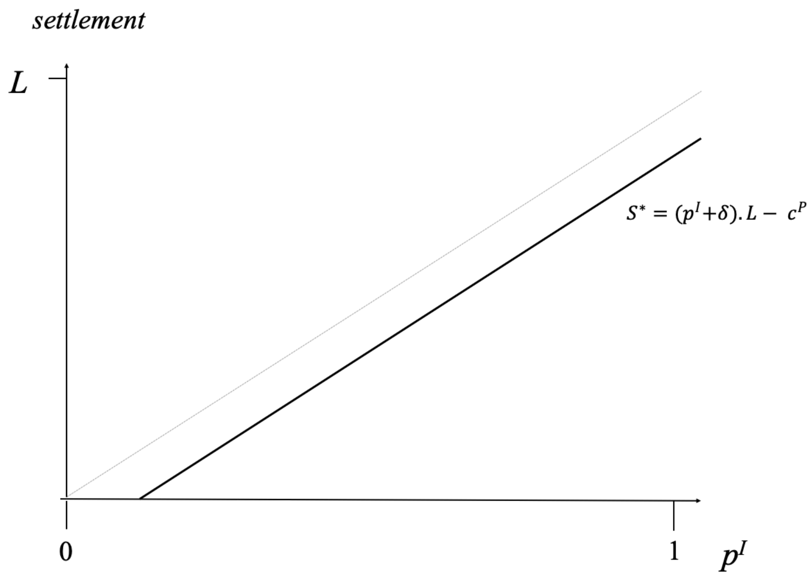

Figure 2: Settlement offers are correlated to p^I. The relationship is linear. These offers are accepted if the condition in equation (1) is satisfied.

We will refer to this as "the baseline model with subjective probabilities." Nothing in this section is new. These results can be found, in one form or another, in the existing law-and-economics literature on litigation and settlement.[21] We wish, however, to emphasize two key takeaways. First, settlement fails in this model when the plaintiff's subjective view of winning, p^P, is sufficiently greater than the defendant's belief that the plaintiff will win, p^D. If both parties are overly optimistic, then no settlement offer is accepted in equilibrium. Second, when settlement does occur, the magnitude of the settlement offer will be a function of the objective probability of liability, p^I. That is, $S^* = f(p^I)$, where $f'(p^I) > 0$. The greater the objective probability liability, the larger the settlement offer.

From here it is straightforward to show what results when the parties themselves have access to predictive algorithms that reveal p^I. If the parties understand that the algorithm is revealing the accurate objective probability, then they both believe that objective probability and every case settles with the settlement payment tied to p^I:

$$S_i^* = p^I \bullet L - c^P$$

We will refer to this as the "baseline model with objective probabilities." This is illustrated in **Figure 3**.

[21] *See* Bebchuk, *supra* note __, at 407; Alan E. Friedman, *An Analysis of Settlement*, 22 Stanford L. Rev. 67, 92 (1969); P'ng, *supra* note 4, at 544; George L. Priest & Benjamin Klein, *The Selection of Disputes for Litigation*, 13 J. Legal Stud. 1, 13 (1984); Shavell, *supra* note __, at 63; William M. Landes, *An Economic Analysis of the Courts*, 14 J.L. & Econ. 61, 66 (1971); John P. Gould, *The Economics of Legal Conflicts*, 2 J. Legal Stud. 279, 284–86 (1973); Posner, *supra* note __, at 419; Joel Waldfogel, *Reconciling Asymmetric Information and Divergent Expectations Theories of Litigation*, 41 J. Legal Stud. 451, 455 (1998).

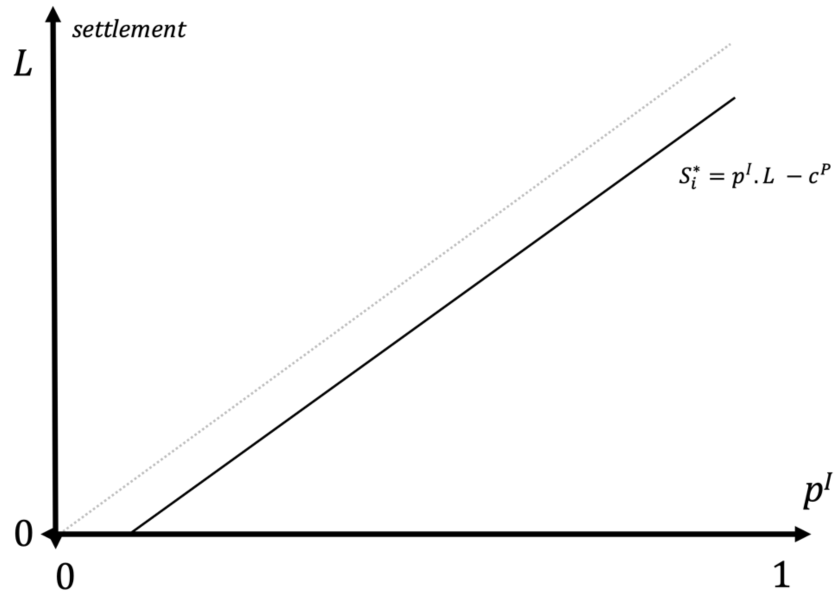

Figure 3: The Baseline Model with Objective probabilities

More realistically, the parties will have some doubts about the algorithm and thus assume that the output is a valuable but imperfect estimate of p^I. In that case, the parties will update their subjective priors, p^P and p^D, based on the weight they place on the algorithmic output. In most cases, this will increase the likelihood of settlement. Though in some "atypical" cases it may reduce settlement. These results are not essential to our analysis, but we have provided them in the appendix.

Less straightforward, is what happens when a judge has access to predictive algorithms that reveal p^I. In the next section, we introduce that scenario and show how it affects litigation and settlement outcomes.

3. Judicial use of algorithmic assessments of liability

We now introduce the idea of algorithmic assessments of liability into the model. We suppose that the algorithmic assessment of liability is accurate and provides the judge with an independent and objective probability of a plaintiff succeeding, p^I. But the probabilistic assessment is not a legal decision. It becomes an input into a legal decision. When a judge observes p^I, there are many different methods for translating that probability into a final judgment. The choice of method has major implications. It can change the litigated outcomes and alter fundamental characteristics of an adjudication system. Moreover, it can alter settlement dynamics.

Before we discuss the different options available to judges to convert these algorithmic predictions into liability rules, we wish to point out some of the tradeoffs of using such tools to make decisions. First, if judges are using these types of algorithms as the basis for a legal decision, the judges must be satisfied that the previous decisions are accurate representations of what they perceive to be "right." That is, they must be comfortable with the content of existing law. If, for some reasons, judges believe that the existing law needs to be changed, then reliance on these types of algorithms would not be sensible.

Second, we have made the rather strong assumption that the algorithms are "accurate." One of the more prominent reasons for inaccuracy is the fact that datasets that describe and explain the law are often based on existing case law, which themselves are subject to selection bias. Not all disputes turn into legal disputes; not all legal disputes are litigated; not all litigated cases are determined by judges; and not all judges' decisions are published.

With those caveats, we present six basic options for converting probabilistic assessments into legal outcomes. There may be other options available, but this analysis provides a starting point that looks at the most likely options to be proposed. For each option we show how the use of algorithmic assessments of liability changes both litigation outcomes and settlement dynamics. We assume for each option that the litigants also have access to the prediction tools.[22]

3.1 Option 1: Binary outcome determined by *ex ante* balance of probabilities

The first instinctive reaction many people have when talking about automated judgments is that judgment against the defendant will be based on whether the data tells us the defendant is *more likely than not* to be liable.[23] That is, if the prediction tool says it is 51% likely that the plaintiff has established liability, then the plaintiff wins, and full damages are awarded. If the prediction tool says it is only 49% likely, then the defendant wins, and no damages are awarded. Thus:

- If $p^I \leq 50\%$, the defendant wins, and the plaintiff is awarded zero damages;
- If $p^I > 50\%$, the plaintiff wins and is awarded L.

These outcomes are represented visually in **Figure 4**.

[22] We do not address the scenario where judges use the tools, but parties do not. The model there would be similar to the baseline model with subjective probabilities. But the parties would no longer be predicting simply what the judge will do, but rather what the judge utilizing the tool in a certain way will do. The structure of settlement will stay the same, but the parties will be settling different cases at different rates because the prediction of what a judge in a particular case will do will differ when the judge has the tools from when she does not.

[23] This statement is based on anecdotal conversations. We are currently devising a survey to more rigorously test people's intuition.

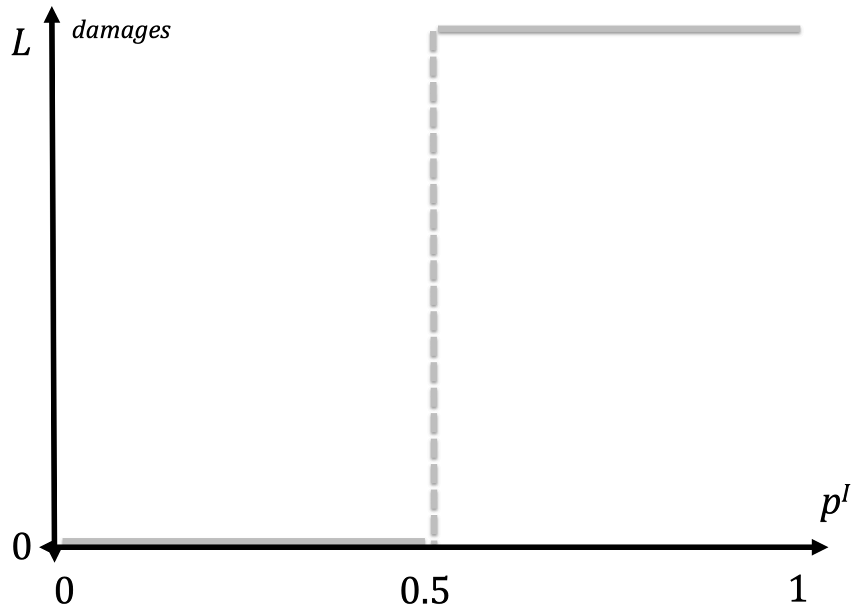

Figure 4: Where p^I is less than or equal to 50%, the defendant wins. No damages are awarded. Where p^I is greater than 50%, the plaintiff wins and is awarded full damages, L.

The intuition behind this mechanism might be its similarity to concepts around burdens of proof in the fact-finding context. But those concepts do not fit here. The more-likely-than-not standard that lawyers know so well is about factual determinations. But we are asking for a prediction about legal determinations. Judges do not announce the law in terms of a likelihood standard.[24]

3.1.1 Litigation outcomes

How does this use of algorithmic assessments compare to litigation outcomes without the algorithm? When a judge decides a case without the assistance of an algorithm, plaintiffs expected return is $p^I \bullet L - c^P$. (Recall that by assumption p^I is both the output from the algorithm and the accurate objective probability).

[24] Moreover, even if we were modeling fact finding, the prediction tool would not be telling you the likelihood that a particular fact is true. In a generic civil case, a prediction that a court is 51% likely to find that a certain factual contention is true is akin to a prediction that "it is more likely than not that a court will think that a fact is more likely than not to be true." Translating that to liability for all cases where the tool produced a 51% probability would be a drastic change in the burden of proof. That added layer of probability makes interpreting things more complicated. For example, what if you have one case with a 40% probability that a court will find a fact has an 80% probability of being true and another case with an 80% probability that a court will find a fact has 40% probability of being true? It is hard to say which is a normatively stronger case. But from a litigating plaintiff's point of view, the first is much stronger than the second because it has at least a 40% chance of winning whereas the second has at least an 80% chance of losing. We discuss the fact-finding application briefly in Part 4.

The outcomes when judges use the binary model of algorithmic assessment are different. The outcomes follow a knife edge around 50%. This changes the outcomes of cases. Take, for example, a case where the plaintiff is assessed to be 70% likely to be found liable ($p^I = 70\%$). Remember that without the algorithm, when cases go to court they are decided by human judges and 30% of them result in a defendant victory. But, under the binary model of algorithmic assessment, 0% of cases result in defendant victory. Thus 30% of cases are decided differently under the binary model. As a further example, if p^I were equal to 10%, then the binary model of algorithmic assessment gives a different outcome in 10% of cases.

The magnitude of differently decided cases is greatest where the algorithm predicts that the case is "close"—when p^I is close to 50%. The exact proportion of all cases that are "decided differently" depends on the underlying distribution of cases under p^I. But, for sake of explanatory ease, suppose that the distribution of cases that go to court follow a uniform distribution from $p^I \in [0,1]$. Under uniform distribution, the final outcome on liability would be different in one-quarter of all cases.[25]

3.1.2 Settlement outcomes

Mapping probabilities into binary outcomes in this way radically changes the nature of settlement. Settlement outcomes track the mapping of outcomes from the predicted probabilities. Using backward induction: consider how the plaintiff behaves in stage 3. If $p^I \leq 50\%$, then the plaintiff knows that she will pay c^P to go to court and has no chance of winning. Thus, in stage 3 she is willing to settle for any amount greater than zero. If $p^I >$ 50%, the plaintiff knows that she will win and recover $L - c^P$ if she rejects settlement. Knowing this, in stage 2, the defendant will offer zero if $p^I < 50\%$ and offer $L - c^P$ if $p^I >$ 50%. The plaintiff will not bring a claim in stage 1 if $p^I < 50\%$. The equilibrium can be stated simply:

(1) No cases are brought by the plaintiff when $p^I \leq 50\%$
(2) All cases where $p^I > 50\%$, the plaintiff settles the case for $S = L - c^P$

[25] The proof of this is relatively straightforward. For any given p^I, where p^I is less than 50%, the fraction of cases where the results are different = p^I. For any given p^I, where p^I is greater than 50%, the fraction of cases where the results are different = $1 - p^I$. The total fraction is found by taking the integral of this inverted-v curve. This is equal to ¼.

Of course, one might consider alternative distributions of the cases. For example, the distribution might instead take a barbell form with cases clustering around 1% and 99%. This would suggest that most cases are easy cases. If that were the case, the number of cases decided differently and the magnitude of the difference in most cases would go down.

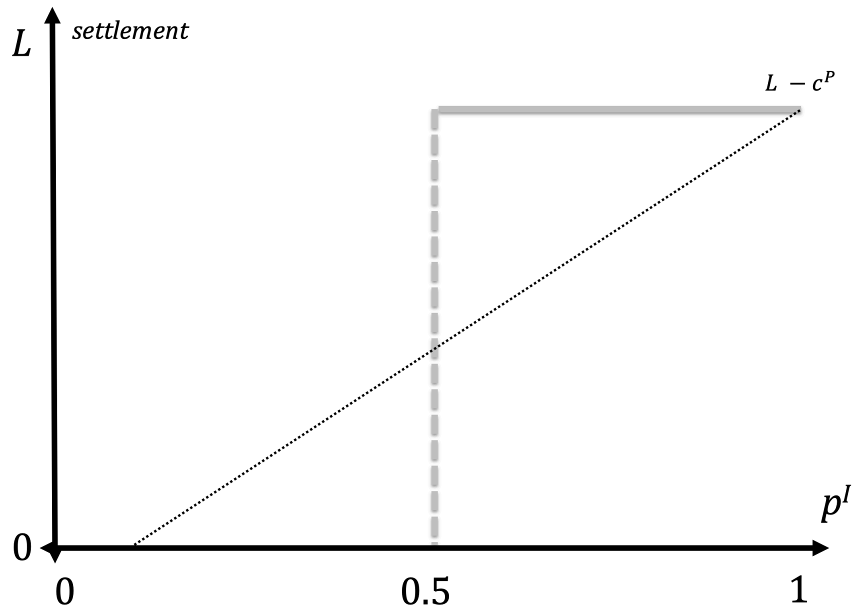

Figure 5: Where p^I is less than or equal to 50%, the case is not brought. Where p^I is greater than 50%, the case is brought and the plaintiff settles for the full amount minus costs.

Accepted settlement offers, in equilibrium, are shown graphically in **Figure 5**. The thick gray line represents the amount the plaintiff recovers in a world where judges use algorithmic assessment to determine liability. Where $p^I \leq 50\%$, the plaintiff does not bring a case. At $p^I = 50\%$, there is a knife-edge jump. For all probabilities above 50%, the plaintiff recovers (nearly) the full amount. Compare the thick gray line to the dotted black line. The dotted black line might represent the settlement outcomes in the baseline model with subjective probabilities (see **Figure 2**) or the baseline model with objective probabilities (see **Figure 3**). There are four important differences between the outcome here and the outcomes in the baseline models:

(1) In both baseline models, the dotted black line "tracks" the independent assessment of the probability of the plaintiff being liable. The gray line also tracks the probability but in a much coarser fashion. Before the binary model of algorithmic assessment was used, the settlement behavior of litigants did not distinguish greatly between a 49% plaintiff and a 51% plaintiff. But in a world where this knife edge is law, the difference between these two cases is stark. At 51% the plaintiff recovers as if she were at 100% in the baseline model with objective probabilities. At 49% she recovers nothing. Further, in a world where judges do not use the tool, a defendant who is 100% likely of being liable behaves very differently in settlement to a defendant who is 51% likely of being liable. Here they have the same incentives to settle.

(2) In the baseline model with subjective probabilities (the world without algorithms) the dotted black line is *conditional* on settlement being possible. This is dependent upon the area of disagreement between the two parties being sufficiently small. The gray line here comes with no such conditions. Settlement

always occurs when a case is brought because there is no ex ante uncertainty about the outcome of liability.

(3) In the baseline model with objective probabilities, parties settle every case only because they all trust the algorithm to be accurate. That trust is not required here. The parties' belief about the algorithm's accuracy is unimportant. The important things are they know that the court follows the algorithmic output and they know what that output is.

(4) Plaintiffs do not bring disputes when the probability is less than 50%. Cases either settle at nearly fully liability or they are not brought at all.

3.1.3 Effect on ex ante behavior

To the extent liability is a deterrent with regard to ex ante behavior, the switch from the baseline model to the binary model of algorithmic assessment are large and problematic. Essentially the switch creates an inefficient liability cross-subsidy from defendants just over the 50% mark to those just below it. This will likely lead to a discontinuity in ex ante behavior where potential defendants tend to cluster their behavior around 49% and completely avoid behavior that is around 51%. Someone who is doing something at the 51% mark has strong incentives to alter their behavior toward the 49% mark.[26] On the other hand, someone who is doing something at the 49% mark has no expected liability, and therefore no incentive to expend any effort to move from 49% toward 1%. Similarly, someone who is at 100% percent liability has no incentive to expend any effort to move toward 51%. Thus, the binary model creates a large shift in deterrence and cannot be viewed as merely automation of the existing system. This is a rather dramatic shift in the substantive effect of the law.

3.2 Option 2: Continuous outcomes – expected damages

An alternative proposal is to award expected damages. If the plaintiff rejects the defendant's settlement offer, the judge uses the independent assessment to award damages of $p^I \bullet L$ (shown graphically in **Figure 6**).

[26] Or let themselves drift toward the 100% mark.

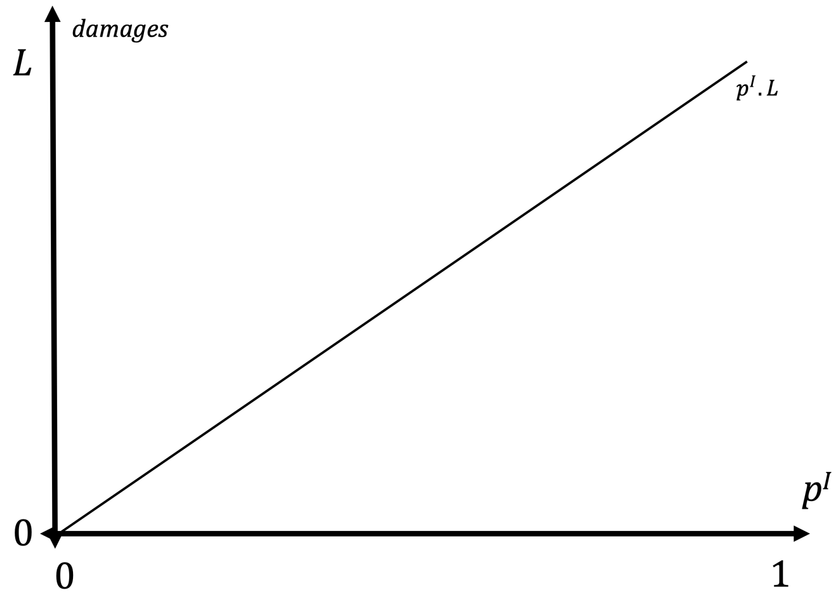

Figure 6: Damages awarded are expected damages, based on the independent assessment of the plaintiff's case.

3.2.1 Litigation outcomes

This is an intuitive application of the data, but this mapping of probabilities into outcomes would represent a radical change in the way we think about liability and compensation in the Anglo-American legal system. Importantly, it holds defendants liable (partially) for losses even when there is a low likelihood of being found liable. If the algorithm stipulates that the plaintiff has a 5% chance of winning on the merits, should we hold the defendant liable for 5% of the losses? On the one hand, it may seem at odds with our intuitions to hold defendants liable when there is such a small chance of the plaintiff's case succeeding (even though they are only liable for a small percentage of the loss). On the other hand, as we shall see, this option simply entrenches in law what already happens with settlement in a world where parties (but not the judges) use the algorithm and perfectly Bayesian update.

3.2.2 Settlement outcomes

Here, the settlement equilibrium is straightforward. The defendant offers $p^I \bullet L - c^P$ in stage 2, which is always accepted by the plaintiff in stage 3. (The plaintiff only brings cases where $p^I \bullet L - c^P > 0$.) This settlement offer by the defendant is the same as the offer in the baseline model with objective probabilities shown graphically in **Figure 3** above. In equilibrium, all cases settle. The settlement offer, in equilibrium, reflects the probability of the plaintiff winning. The plaintiff who has an 80% probability of winning recovers 80% of the loss (minus costs).

3.2.3 Effect on ex ante behavior

Beyond that, this method does not change ex ante incentives or deterrence, and so the expected damages model can be viewed as merely adding information to and automating the results from the existing system. In other words, it produces the same effect on ex ante behavior as the baseline model with objective probabilities.

3.3 Option 3: Hybrid approach

A slightly modified approach might be to combine the first and second options. Here, the plaintiff would be liable for *expected* damages if $p^I > 50\%$ and no damages if $p^I \leq 50\%$, the defendant wins.

3.3.1 Litigation outcomes

This approach is perhaps more in line with Anglo-American traditions. It is graphically represented in **Figure 7**.

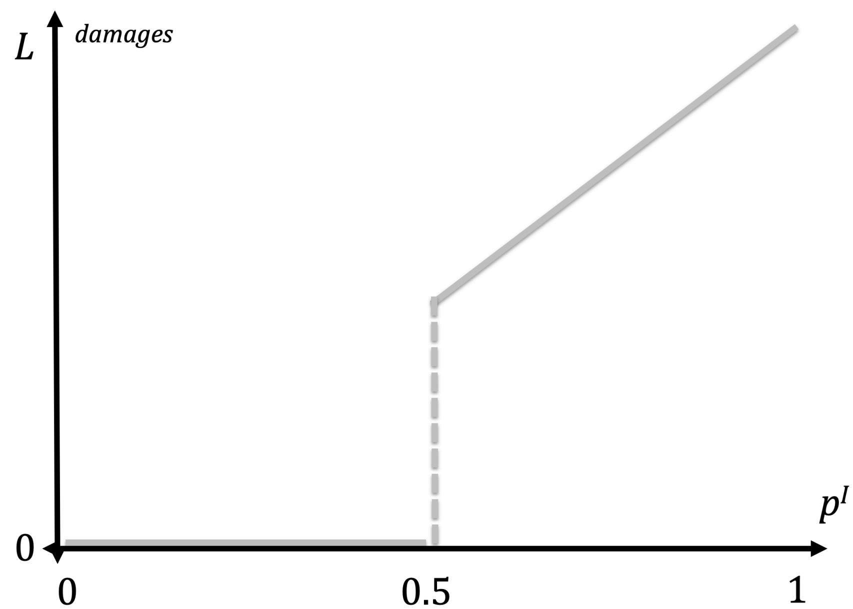

Figure 7: No liability is found when p^I is less than or equal to 50%. Where p^I is greater than 50%, expected damages are awarded.

3.3.3 Settlement outcomes

The settlement outcomes are straightforward. They track the judgment outcomes less costs.

3.3.3 Effect on ex ante behavior

While this may seem like a pleasing compromise, this conversion of predictions into outcomes would produce underdeterrence in ex ante behavior. By cutting off the possibility of liability for all cases where $p^I \leq 50\%$, this approach reduces the *ex ante* expectation of compensation to be paid by the defendant. As a result, shifting to this model is a major change in the substantive effect of the law.

3.4 Option 4: Probabilistic liability with full damages

As a fourth option, suppose that courts try to mimic the world without algorithms through probabilistic liability. Thus, when p^l is 70%, the court would impose full liability on the defendant 70% of the time. The idea is like flipping a weighted coin to determine liability where the weight of the coin is determined by p^l. In the 70% example, the coin would be weighted to come up heads (liability) 70% of the time.

3.4.1 Litigation outcomes

This approach would restore the expected liability that exists in the baseline model with objective probabilities. Graphically, the outcomes of judgments here, in expectation, looks the same as in **Figure 6**.

Even though expected liability is the same as it is in the baseline model with objective probabilities, this approach would likely meet with heavy opposition. One might object that the rule of law is violated by allowing such random and arbitrary considerations to determine liability.[27] Assessing whether the approach is an improvement turns on views about the existing variation in judicial rulings. In practice, this option converts *unexplained* variation in legal decisions into *purely arbitrary* variation.[28] If one thinks that the unexplained variation is a result of nefarious bias, this is an improvement. If one thinks that the unexplained variation is a result of arbitrary judicial whims, nothing has changed—you have replaced one arbitrary method with another. But if one thinks that behind judicial variation there are valuable case-specific human judicial intuitions that cannot be explained by data, then this would produce worse results by ignoring those reasons. In reality, the answer probably involves a combination. Any particular decision likely results from a mix of measurable information, biases, arbitrary distinctions, and unmeasurable intuitions.

3.4.2 Settlement outcomes

Importantly, however, none of the critiques about variation from the prior subsection matter if cases settle. And in this model all cases do settle. The parties have access to the algorithmic assessment, and no disagreement about outcomes. As a result, cases will settle at the expected liability point (minus costs), thus creating the same settlement rates as option 2 and as in the baseline model with objective probabilities.

---

[27] For an in-depth discussion of the potential role of randomness in law more generally see, Neil Duxbury, RANDOM JUSTICE: ON LOTTERIES AND LEGAL DECISION-MAKING, (Oxford University Press 1999).

[28] Some "variation" in outcomes might be explained by judge-specific effects. We do not discuss those effects here. But we return to them – and to the question of whether entrenchment of judge-specific effects in such an algorithm for this purpose is socially desirable – in Part 4.

It is worth emphasizing that when settlement is ubiquitous, the source of variation that we discussed in the previous subsection is unimportant. Rational actors settling a case only care about the expected rate of liability.[29] As long as they are powerless to change that rate, they do not care about its causes.

3.4.3 Effect on ex ante behavior

This result does not change ex ante incentives and deterrence. The probabilistic liability model produces the same effects on behavior as the expected damages model and the baseline model with objective probabilities.

3.5 Option 5: Triage "easy" cases

Another intuitive option for mapping probabilistic outcomes into legal outcomes is to use prediction tools to triage "easy" cases. That is, the judge can use the independent assessment, and rely on the algorithm to determine the outcome only when the case is "easy." Easy cases are those where the probability of one side winning is close to 1. If the independent probability assessment is close to 0 or 1, then the outcome can be determined *ex ante*. In this way, courts can reduce their caseload by triaging easy cases from their list, focusing on those cases where the law is less clear.

Take, for example, the situation where the judiciary uses a 5% threshold: cases where one side has less than a 5% chance of winning according to the independent assessment are automatically determined by the independent assessment. If $p^I < 5\%$, then the plaintiff recovers zero; if $p^I > 95\%$, then the plaintiff recovers the full amount. For all other cases, where $p^I \in [5\%, 95\%]$, the case proceeds to (human) judicial determination if settlement fails. For the subset of easy, triaged cases, the litigation outcomes mirror those in option 1. For those cases where the algorithm is not used, the judge has leeway to decide the outcome. The larger the subset – that is, the higher the threshold for what is considered an "easy" case – the more this option begins to reflect the litigation outcomes of option 1. An example, with a 5% threshold, is shown graphically in **Figure 8**.

[29] This is the likely state of real-world litigation today. Most cases do settle. And they settle based on predictions about the outcome, which take into account judicial variation regardless of its source.

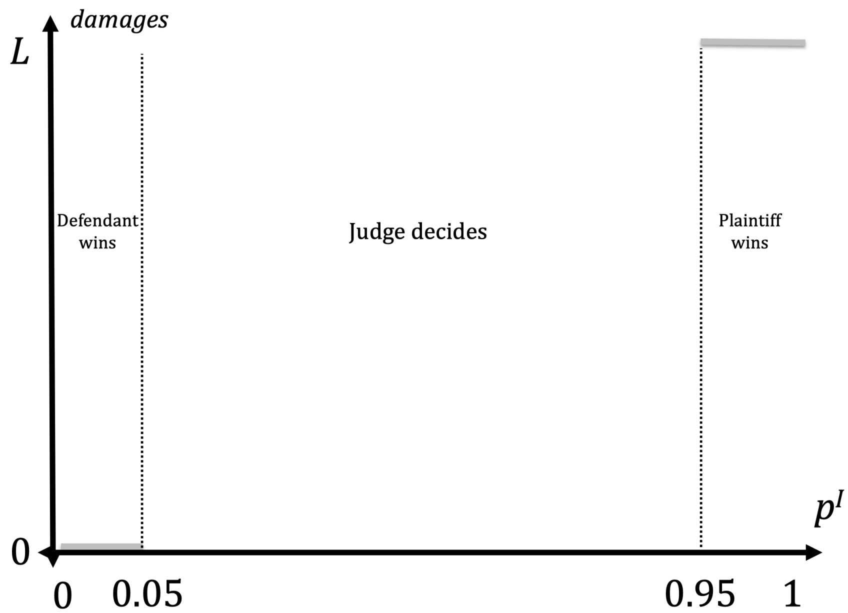

Figure 8: Easy cases where one party has less than 5% of winning are triaged. Liability in the intermediate cases is determined by the judge.

3.5.1 Litigation outcomes

The number of cases that are "decided differently" is clearly reduced compared to option 1.[30] Because cases at the two extremes are highly likely to be decided the same way, the distortion in outcomes by using the algorithm is minimized. Of course, the higher the threshold, the greater the proportion of cases that are decided differently.

Many perceived benefits of triaging cases in this way may not be readily observable given the way we have set up our simple model. Frivolous suits (where the probability of a plaintiff winning is zero) are possible in the real world, even when the plaintiff knows the probability is zero.[31] Frivolous suits don't occur in equilibrium in our model, but it is not difficult to generate models to account for them. Algorithmic triage in these types of cases would provide deterrence against frivolous cases.

3.5.2 Settlement outcomes

The effect on settlement of this type of triage depends on whether the independent assessment falls within the triaged zone or not. Let's say the threshold is 5%. If $p^I < 0.05$, the case is never brought. If $p^I > 0.95$, then the case settles for close to the full amount

[30] The magnitude of the reduction turns on the distribution of cases. Triage is particularly effective if most cases are easy and the distribution has a barbell form.
[31] For discussion on how settlement bargaining affects incentives to bring frivolous suits, *see* Avery Katz, *The Effect of Frivolous Lawsuits on the Settlement of Litigation*, 10 Int'l. Rev. L. & Econ, 3–27 (1990).

claimed. For intermediate cases, settlement results would follow our analysis for in the baseline model with subjective probabilities.

3.5.3 Effect on ex ante behavior

Similarly, the effect on ex ante behavior depends on whether a case is in the triaged zone or not. The effects of ex ante behavior will be changed in the tails. But this change might be small. Moving a 98% liability expectation to 100% and a 2% liability expectation to zero is a small—perhaps trivial—change in ex ante expectations. For the intermediate cases, the effects again remain the same as in the baseline model with subjective probabilities.

3.6 Option 6: Judges use the prediction for guidance only

The previous five options have all involved human decision makers deferring to the algorithmic assessment to some degree. For options 1 to 4, the algorithm is the ultimate arbiter of all cases, and for option 5, the algorithm is the arbiter of easy cases. But complete deference to the algorithm is unlikely to be the first step in the process. Rather, judges will likely have much discretion—at least initially—to accept or reject the algorithmic assessment.

In option 6, we suppose that the judge has access to the algorithmic assessment of liability, but there is no decision rule. There is no formalistic mapping of probabilities onto outcomes. Instead, the judge has the option of referring to the probability assessment, but only uses the algorithm's assessment for guidance.

3.6.1 Litigation outcomes

How will this change litigation outcomes in the event that cases go to court? The degree to which this affects outcomes will depend on the propensities of the individual judge using the tool. Suppose the judge ignores the suggestions or recommendations of the algorithm? In that case, very little will change. The judge relies on her own assessment of the case, as indeed she would in the absence of any such prediction. But, to the extent that the judge does begin to lean on the assessments of liability, then the question of how litigation outcomes are affected will depend on which of the five previous options best describes how the judge is using the tool.

3.6.2 Settlement outcomes

If parties know that the judge has access to the algorithmic assessment, but they don't know exactly her propensity to follow the guidance of the algorithm, how will this affect settlement? Upon revelation of the algorithm's prediction, the plaintiff and defendant must re-assess their priors knowing that the judge too has access to the prediction.

The tool provides an independent assessment of the likelihood of a plaintiff succeeding, p^I. Before stage 1, both parties update their prior subjective probabilities to reflect how they believe that the judge will update:

$$p_i^P = (1-\theta)p^P + \theta p^I$$
$$p_i^D = (1-\tau)p^D + \tau p^I$$

The degree to which parties update, θ or τ, depends on the degree to which they each believe that judge will rely on the algorithm. If parties believe that the judge will be highly influenced by the independent assessment, then θ and τ are close to 1. If parties are skeptical about the judge using the algorithm, θ and τ are close to 0.

When judges use the algorithmic assessments of liability, τ and θ represent the parties' beliefs about whether and how judges will update. In short, the updating by the plaintiff and defendant will turn on their beliefs about (1) how much the judge will update; and (2) how the judge will incorporate the probabilities into her decision. If parties believe that the judge will fully update her priors and use the binary model (in 3.1), then there is almost a self-fulfilling prophecy – the settlement offers in equilibrium mirror the knife-edge outcomes in option 1.

Depending on how parties perceive the judge's decision rule, there may be counter-intuitive effects on settlement. For example, let's say that the plaintiff initially believes that she has a 90% chance of winning. But the independent algorithm suggests that chance is only 60%. If the plaintiff believes that the judge will be faithful to the prediction *and* believes that the judge will employ the binary model, then the plaintiff will update from her priors of 90% to a posterior of 100%. If only the parties used the prediction tool, the equilibrium settlement offered by the defendant and accepted by the plaintiff would likely *decrease* upon revelation of a lower objective probability.[32] But when the judge has access to the tool, the equilibrium settlement offer may *increase* upon revelation of a lower objective probability.

4. Discussion and extensions

Settlement models in law and economics are based on the plaintiff's probability of victory.[33] So too is the output of a litigation prediction tool. But what, exactly, does that probability mean? We might think of the probability of p^I = 80% in a tort case as meaning the following: If that case were litigated 100 times, the court would on average find in favor of the plaintiff in 80 of those cases. That statement does not map on to ideas like burden of proof for fact finders. It is not the same as saying that there is an 80% chance that defendant committed the tort. To see why, imagine a case where everyone in the world agrees that factually there is a 60% chance that defendant committed the tort. There will be liability in that case 100% of the time. In another case imagine that half of the judges in the world

[32] See appendix.
[33] *See* Daughety & Reinganum, *supra* note __.

think the defendant is 99% likely to have committed the tort and the other half think the defendant is 49% likely. There will be liability in 50% of the cases.

Because these concepts are distinct, when you translate probability of plaintiff's success into a judicial rule you cannot mechanically import a 51% threshold for liability. As we have shown, the precise mapping or translation from p^I to liability and damages rules can have important effects on litigation and settlement outcomes. Additionally, we have shown that even the mere use of prediction tools by litigants and judges can have large effects on settlement outcomes and ex ante behavior.

We now consider some extensions and discuss what happens when you add potential complexities to the model.

4.1 Factual determinations

We have till now avoided the application of our model to fact finding. But surely some algorithms might predict how a judge or jury would decide certain facts in light of the available evidence and other characteristics of the case. The prediction here is more complicated and requires a more advanced algorithm with more data because it must take into account jury composition as well as the many small variations possible in the presentation of evidence.[34] But there may be some cases where a judge is the factfinder and the available evidence is similar across a large number of cases. For example, some small claims cases may have these characteristics. In any event, if algorithms could predict the outcome of fact finding, the analysis would be similar to what we have presented.

To see why, assume a case where the law is settled and the parties have stipulated all but one fact. In that case, the decision on that one fact will determine the outcome of the case. If the factfinder rules one way, the plaintiff wins. If the factfinder rules the other way the defendant wins. The parties will base their settlement strategy on their prediction of how the factfinder will rule on that one fact. That now presents essentially the same scenario from above, and the same analysis would apply.[35]

4.2 Variation and uncertainty

When p^I is less than 100% there is uncertainty. What does that uncertainty represent? Judicial bias or inconsistency? Missing variables in the analysis? Random human variation? Or meaningful human insight that data cannot pick up? The answer to those questions may change litigation outcomes. But it may not matter in a world where all cases settle. A settling party might not care about theoretical explanations for variation. If it cannot be measured it is simply litigation risk and will be priced into settlement.

4.3 Judge specific variation and other "excludable" factors

[34] Parties might also be able to expend resources to develop new evidence that will change the probabilities. But the same can be said of expenditures to develop new legal arguments.
[35] Note that the burden of proof does not affect the analysis here. It may change the probability of the judge ruling one way or the other. But beyond that it has no effect.

Courts may reach different outcomes across similar cases for various reasons as discussed above. This may be a function of inconsistency or it may simply be based on factors that don't show up in the data. One factor that likely will show up is the judge's identity. We have treated courts as a monolithic institution so far. But each judge is an individual. One judge might treat an identical case differently than another judge would. To the extent the data can identify the outcome probabilities that are attributable to the identity of the judge, this poses interesting questions.

First, it certainly complicates Option 6 to ask how a judge would react to and implement a p^I if she knows that the inputs into p^I include her own idiosyncratic preferences.

Second, we might want to design the algorithm—to the extent possible—to exclude or control for judge-specific effects. To the extent an outcome is influenced by the mere identify of a particular judge, rule-of-law considerations might counsel in favor of de-biasing results, correcting for that influence.[36] It is worth noting that doing so would affect some of the results in our models. Litigants using algorithms to predict success today, of course, do not want to exclude judge-specific effects. To the contrary, they are quite interested in the preferences of the judge to whom the case has been assigned.[37] Thus, current real-world settlement estimates include and may be primarily driven by judge-specific effects. Backing those effects out of the judgment will produce results that are different from the baseline models. Again, Option 6 becomes very complicated if the parties in settlement negotiations are predicting judge-specific reactions to a model that may or may not have corrected for judge-specific effects.

Similarly, the algorithms may reveal other undesirable factors that are driving outcomes in past cases. For example, prior outcomes may have been driven by improper factors such as race, gender, or the identify of one party's lawyer, or by trivial factors such as the date of filing.[38] To the extent these factors can be identified and corrected for, the use of algorithmic assessment tools can add greater value.

4.4 Accuracy: Selection bias

Accuracy is the elephant in the room that we have avoided until now. What if the litigation prediction tool is not an accurate reflection of the way a court would have

---

[36] For a related ideas, see Daniel Chen, *Machine Learning and the Rule of Law*, (2019) available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3302507&download=yes.

[37] The question of predicting idiosyncratic judge effects is controversial in some jurisdictions. In France, a recent law prohibited any parties from using predictive litigation algorithms that take into account the individual identities of judges. Jason Tashea, *France Bans Publishing of Judicial Analytics and Prompts Criminal Penalty*, ABA Journal, June 7, 2019 at https://www.abajournal.com/news/article/france-bans-and-creates-criminal-penalty-for-judicial-analytics.

[38] See, e.g., Anthony Niblett, *Algorithms as Legal Decisions: Gender Bias and Employment Law in the 21st Century,* Univ. of N. Bruns. L. J. (2020, forthcoming).

resolved the dispute? That is, what if the past decisions are not predictive of the future judgments? The information for a litigation prediction tool is based only on cases that *did* go to judgment. This leads to issues of selection bias. Perhaps only cases with close and perhaps confounding factual situations end up in court; straightforward cases may not. Thus, the predicted probability may generate accurate outcomes conditional upon cases going to court, but those probabilities may not reflect the full picture. This affects settlement outcomes. Lawyers and judges who use these tools therefore still need some understanding of the body of case law upon which these predictions are made. If the case at hand is of a different kind to all the prior cases, or if there are new facts that haven't been addressed before by the judiciary, this will weaken the predictive value of the algorithm. This is, of course, true of any predictions of outcomes based on precedent, irrespective of whether data analytics are used.

4.5 Accuracy: The choice of statistical model

We have assumed one objective p^I that is produced by the algorithmic assessment tool. In reality, there is no one p^I. Moreover, different the choice of statistical model in the algorithm can change the output. This in turn will change the outcome of cases under the various models discussed. The effect is likely to be the greatest in Option 2. The dollar value of a suit there turns directly on the independently assessed probability at all levels of p^I, not just around the neighborhood of $p^i = 50\%$. To see why this may matter particularly in Option 2, consider the choice between a *probit* or logistic regression model. The imputed probabilities that are returned by using *probit* or logistic regression models often return similar probabilities, especially in the neighborhoods of 0, 0.5, and 1, but they are not the same. For some fact patterns, logistic regression classification models will present a higher predicted probability than a *probit* model. Which model should be used here? More advanced machine learning classification models will return even more different probabilities. And these probabilities will vary with regularization and tuning. The choice of one model or the other to generate probabilities would favor one side. Simply put, with every modeling choice, there will be winners and losers.

4.6 Legal rigidity and stale precedent

If litigation prediction tools lead more cases to settle or if judges use them to decide cases, the result will be a reduction in the production of judicial precedent. This may be costly. Case law may benefit from being dynamic and frequently updated.[39] And litigation prediction tools may impede those updates leading the law to become stale. This could have two effects. The case law may no longer be a good fit for the world. A common example is that traffic rules developed in a horse and buggy age being a bad fit for a world with automobiles. This problem would grow when judges are themselves using the litigation prediction tools. Alternatively, judges may ignore older stale precedent (and the litigation prediction tools based on that precedent). If that happens, the litigation prediction tools—which use precedent as their input—will become less accurate and parties will have

[39] See, for example, Shavell supra note __. on the social benefits of trials and precedent setting.

renewed reasons to litigate. This could lead to a cycling effect. The tools may be very accurate and highly utilized for a time, depreciate as precedent becomes stale, go out of use, and then gain usefulness as precedent is updated.

4.7 Fee shifting: Another potential use

Rather than using the prediction tool to make determinations on outcomes, judges may use the tools to help them make determinations about awarding costs. The judge may use the *ex ante* independent assessment of the plaintiff's likelihood of victory to award costs to the ultimate victor. If the algorithm suggests that a losing plaintiff had a very low likelihood of victory all along the judge may elect to award costs to the defendant.

Conclusion

Algorithms don't make decisions. Rather, humans can make decisions that take into account algorithmic assessments. The way in which judges use algorithmic assessments of liability is not a simple yes/no decision. There are different choices that judges must make in order to convert algorithmic predictions into the legal decisions. We have endeavored to explore some of the choices that are possible.

The simple model presented here reveals that the use of algorithmic assessments of liability and the advent of automated judging will have complicated and dynamic effects on settlement practices and litigation outcomes. In turn those effects will alter the deterrence and incentive effects that laws have on ex ante behavior. Further models, allowing for endogenous costs, sequential bargaining, asymmetric access to technology, and hidden information, will no doubt complicate these effects further. We view this model as an important starting point for exploring and understanding this new technology.

Appendix: Settlement where Litigants—but not judges—use Algorithm Liability Assessment

Parties use litigation prediction tools[40]

Suppose the parties have access to litigation prediction tools (but judges and other decision makers do not). Let's suppose that use of such tools by litigants is costless. The tool provides an independent assessment of the likelihood of a plaintiff succeeding, p^I. Before stage 1, both parties update their prior subjective probabilities in light of this new independent assessment:

$$p_i^P = (1-\theta)p^P + \theta p^I$$
$$p_i^D = (1-\tau)p^D + \tau p^I$$

The degree to which parties update, θ or τ, depends on the degree to which they each believe that the independent assessment accurately reflects the law. If parties are highly influenced by the independent assessment, then θ and τ are close to 1. If parties are skeptical of the algorithm, and trust their own intuition, θ and τ are close to 0.

Typical scenario – likelihood of settlement increases

The effect of this updating turns on the relative position of p^I compared to the two subjective priors of the parties. Initially, we assume that the independent probability assessment of the plaintiff's chance is at least as great as the defendant's subjective view and at most as high as the plaintiff's subjective view:

$$p^I \in [p^D, p^P]$$

We call this the "typical" scenario, as at least one of the parties is likely overly optimistic in their subjective assessments of the law. Under these assumptions, the likelihood of settlement increases. Recall that settlement is more likely to fail when the area of disagreement, Δp, is large. Both posterior probabilities of the plaintiff and defendant (weakly) converge towards p^I. The defendant's posterior probability is greater than her prior; the plaintiff's posterior probability is lower than her prior. The new area of disagreement after the probabilities have been updated, Δp_i, is no greater than before.

It is trivial to show that, when at least one of θ or τ are greater than zero, there are disputes that will settle when parties use litigation prediction tools that would not settle when parties do not have such access. The area of disagreement is shrinking while the cost-damage ratio ($\frac{c^P+c^D}{L}$) remains constant. This brings more cases within the conditions of equation (1). We show this graphically in **Figure 9**. Conversely, there are no disputes that

[40] The analysis in this appendix is generally consistent with Stevenson & Wagoner, *supra* note _ and other settlement models. It is presented here for comparison to our new analysis of judicial use of algorithms in the main text.

would have settled in a world without litigation predictions tools that do not settle once litigation tools are introduced.

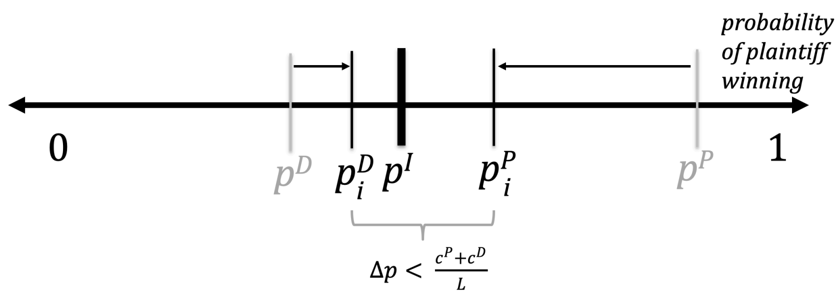

Figure 9: The independent assessment of the plaintiff's probability of victory is p^I. The defendant updates her prior up from p^D to $p^D{}_i$. The plaintiff revises her subject probability down from p^P to $p^P{}_i$. The area of disagreement decreases. Here, if the new area of disagreement is smaller than the cost-damage ratio, the two parties settle in equilibrium.

In equilibrium, if settlement is possible, then the defendant makes a take-it-or-leave-it offer in stage 2 of S_i^*:

$$S_i^* = ((1-\theta)p^P + \theta p^I).L - c^P \qquad (3)$$

The plaintiff accepts this offer. The equilibrium settlement amount increases in p^I, with a slope of θ. If the plaintiff disregards the new information and does not update ($\theta = 0$), then the equilibrium settlement is constant at $S_i^* = p^P.L - c^P$. But if the plaintiff treats the independent assessment as gospel and perfectly updates ($\theta = 1$), the settlement offer is $S_i^* = p^I.L - c^P$. This latter case is shown in **Figure 10.** These settlement equilibria are conditional on settlement being possible (i.e., the condition in equation (1) is satisfied).

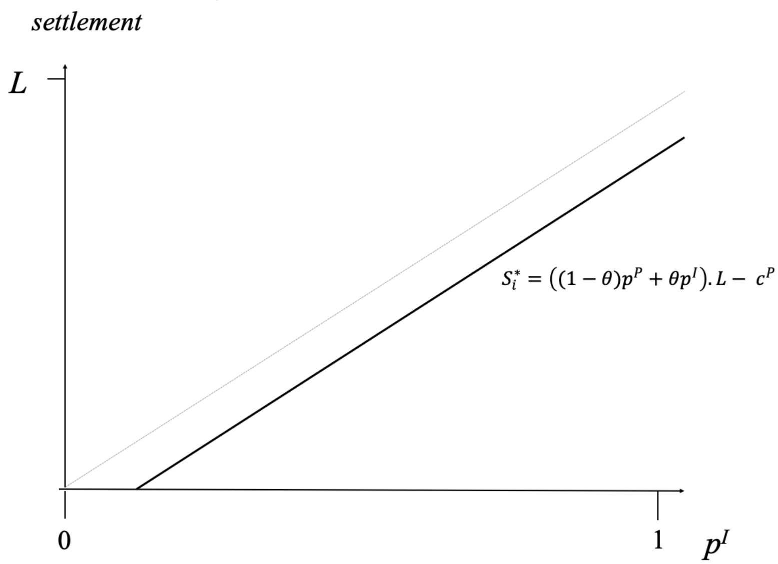

Figure 10: Settlement offers when $\theta = 1$. The settlement offers track p^I.

Atypical scenario—likelihood of settlement decreases

There are situations where prediction tools reduce the likelihood of settlement. Here, we relax the assumption that the independent assessment falls between the two subjective priors. Take, for example, a situation where the defendant's prior reflects a relatively pessimistic view (high p^D) and the independent assessment is even lower than the defendant's view of the plaintiff's case:

$$p^I < p^D$$

There are now situations where the parties would have settled in a world without litigation prediction tools, but they no longer do. To see this, assume that the defendant's pessimistic prior is the same as the plaintiff's optimistic prior, such that:

$$p^D = p^P$$

Without litigation prediction, settlement would be certain because the area of disagreement is zero and equation (1) will be satisfied in all cases. But the introduction of the independent assessment p^I has the potential to upset this equilibrium. Let's say that $p^I < p^D$. If the defendant is a strong updater (τ is high) and the plaintiff is a weak updater (θ is low), then after updating:

$$p^I < p_i^D < p_i^P$$

The area of disagreement increases. More generally, for any set of p^P and p^D , the area of disagreement will increase as long as $p^I < p^D$ and τ is greater than θ.

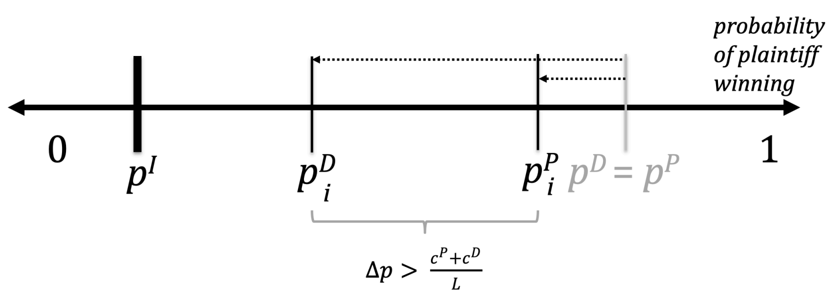

Figure 11: Initially, the two parties have the same subjective belief, $p^D = p^P$. There is no area of disagreement. With no access to litigation prediction tools, the dispute will settle. The litigation prediction tool, however, predicts that the probability of plaintiff victory is much lower, p^I. The two parties update their priors. This creates an area of disagreement. In this illustration, the new area of disagreement is sufficiently large to prevent settlement.

Similarly, settlement opportunities will be reduced when plaintiff is pessimistic relative to the independent probability assessment ($p^P < p^I$), and she strongly updates her prior and the defendant weakly updates.

The atypical scenario relies on this condition: the pessimistic party (defined as the party whose initial belief in her own success is worse than the independent assessment) updates *more* than the optimistic party (whose initial belief in her own success is better than the independent assessment). Readers may question the likelihood of this condition. But there are good reasons, from a behavioral perspective, to think that Bayesian updating here may be asymmetric. Optimism bias on the part of the parties can produce these results. Parties may update their priors more weakly in the face of bad news than in the face of good news.[41] For example, plaintiffs may be willing to gravitate toward higher independent assessments of success than lower assessments. If this is true, then *provided* the independent assessment falls outside the bound of the two priors, the assessment may reduce the opportunity for settlement. Our point here is not that settlement will always fall; it is merely that there are (atypical) situations where the opportunity for settlement falls.

The number of disputes brought

There may be cases that would be brought but for the litigation prediction tools (fewer disputes). The plaintiff files suit in this model when her expected return is greater than zero:

$$p^P.L - c^P \geq 0$$

In the typical scenario, we would expect that the litigation tool would temper the plaintiff's optimistic prior, reducing the likelihood of bringing suit. But, as above, there may be situations where the converse is true. If the plaintiff was sufficiently pessimistic in a world without litigation prediction tools, the use of such tools may encourage plaintiffs who would not have brought a claim to do so. It may even be used as tool to discover potential successful cases of which the plaintiff was not aware (imagine for example a plaintiff who did not realize a certain tort was legally actionable). Essentially, the tool informs the plaintiff of her legal rights.

To the extent that the independent assessment is correlated with the actual outcome, these effects should be welfare enhancing. Disputes where the plaintiffs have a weak claim (low p^I) are less likely to be brought, while disputes where the plaintiffs have a strong claim (high p^I) are now more likely to be brought.

[41] Many studies have shown that when having a preference for a certain future outcome, people tend to exhibit confirmation bias by updating their priors more frequently when they receive new evidence consistent with the preferred outcome. See e.g., Tali Sharot et al., *How Unrealistic Optimism is Maintained in the Face of Reality*, 14 Nature Neuroscience 1475 (2011); Matthew Rabin & Joel L. Schrag, *First Impressions Matter: A Model of Confirmatory Bias*, 114 Q. J. Econ. 37, 37–38 (1999); David Eil & Justin M. Rao, *The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself*, 3 Am. Econ. J.: Microeconomics 114 (2011).

Law Society of Ontario | Barreau de l'Ontario

TAB 2B

AI for Litigators

AI for Lawyers
How Artificial Intelligence is Adding Value, Amplifying Expertise, and Transforming Careers

Noah Waisberg, CEO and Co-Founder
Kira Systems

Dr. Alexander Hudek, CTO and Co-Founder
Kira Systems

April 16, 2021

cpd

NOAH WAISBERG AND DR. ALEXANDER HUDEK

AI FOR LAWYERS

HOW ARTIFICIAL INTELLIGENCE IS ADDING VALUE, AMPLIFYING EXPERTISE, AND TRANSFORMING CAREERS

WILEY

PRAISE FOR *AI FOR LAWYERS*

AI for Lawyers pulls together a series of easy-to-read vignettes that cut through the mystique, noise and bullshit surrounding AI for legal. It provides excellent guidance for lawyers who don't know which way to travel when they finally arrive at the intersection of legal services and technology—which is most of the profession!

—Mitchell Kowalski, author of
The Great Legal Reformation: Notes from the Field

Noah Waisberg and Dr. Alexander Hudek have taken a complex topic and made it accessible and enjoyable. Like it or not, artificial intelligence and machine learning, particularly when combined with 5G connectivity, computing on the edge of networks and eventually quantum computing, will advance by leaps and bounds to automate and change the way we practice law. It is also leveling the playing field between lawyers practicing in big firms vs. small firms. Wherever, whatever and however you are currently practicing, *AI for Lawyers* will open your eyes and make you feel excited and empowered to be part of the future.

—Louis Lehot, founder,
L2 Counsel, P.C.

Alex and Noah have written a demystifying AI book which will help lawyers take advantage of AI technology to create new customer value. They cover the key resources and processes needed to deliver value, which will help all lawyers capture this AI-driven value in their go-to-market approaches, enabling them to develop new ways to solve old problems.

—Michelle Mahoney, Executive Director,
Innovation, King & Wood Mallesons

There is little doubt that the legal industry has experienced a cataclysmic extinction moment, where yesterday's ways of working are tomorrow's fossilised memories. The changing expectations of both the consumers of legal services, and the next generations of lawyers, has seen to it that the practice of law has been changed forever by the arrival of advanced technologies.

In *AI for Lawyers*, Noah and Alex have created the definitive guide on the role of technology in the legal industry. No two authors are better qualified to commentate on how our world is changing. This is a must-read for anyone in the industry and those planning on living a life within the law.

—Justin North, Managing Director,
Morae Global Corporation

The intersection of science fiction and lawyering is both a terrible idea for a movie and a very real problem for attorneys. The terror that artificial intelligence will replace human lawyers and spew steam from the keyboard while trying to define “love” during an ill-fated document review terrifies some folks. And that’s unfortunate because when stripped of its sci-fi mystique, “artificial intelligence” here in the real world is both non-frightening and entirely essential to a thriving 21st century law practice. Waisberg and Hudek’s book provides lawyers a friendly, brass tacks introduction to this oft-misunderstood technology and provides straightforward examples of how AI can advance your practice . . . and, sometimes, how it’s already advanced your practice without you even knowing it.

—Joe Patrice, Senior Editor, Above the Law

Although many lawyers have strong views on the use of AI in the law, very few in fact have a solid grasp of the potential and limitations of this technology. Worse, some lawyers even have the temerity to use ‘AI’ as a verb, claiming—almost arbitrarily—that ‘you can AI’ this or that legal task. Into this world of bold confusion and brazen conjecture, I therefore extend a heartfelt welcome to *AI for Lawyers*. This book brings the clarity, deep technical expertise, practical experience, and commercial insight that are sorely needed in the field.

—Richard Susskind, author of *Tomorrow’s Lawyers* (2017),
***The Future of the Professions* (2015), *The End of Lawyers* (2008), and**
***Expert Systems in Law* (1987)**

Noah and Alex clearly show that the use of AI-embedded software in the legal world will soon be as ubiquitous as the use of word processing. The authors (a Who’s Who of experts in legal technology) cover an extraordinarily broad range of AI-software types and applications—from machine learning to expert systems. The book is an essential read for solo practitioners all the way up to those practicing in the lofty heights of the elite firms around the world and for the technology gurus who enable them. To succeed in law in the coming years, you will need to use AI. To be prepared to use AI, reading this book is a must.

—Harris Tilevitz, Chief Technology Officer, Skadden

I loved this book! AI is increasingly becoming a driver of success for high performing lawyers and law firms. This book is a quick, easy introduction to it. Every lawyer should read it.

—Kent Zimmermann, strategic advisor to law firms

To those driving law practice forward.

Contents

Introduction

Lorie Waisberg kept checking his watch as he waited for the typist to finish the document. She was making the standard three copies using whitener and two pieces of carbon paper. He was anxious because he knew that getting extra copies would take time. When she handed the pages off to Lorie, he took off out the doors of his dad's law firm and down three flights of stairs, across the street, and continued his pace for two blocks, dodging traffic as he made his way to Sudbury's City Hall. They had one of the only copy machines in town, and the Waisbergs could use it in emergencies. Lorie had two concerns as he ran: one was that city hall closed at 4:00 p.m. promptly. The other was that he might not be able to find the person who held the only key to the copier room.

As Lorie made his way into the building, he saw that the clock in the lobby was closing in on 4:00 p.m. He found Gary, the chief engineer, better known to many as "the guy with the copier room key." Gary was grabbing his jacket to head out for the day.

"Gary, it's just three copies, please," panted Lorie. Gary smiled. "Okay, just for you," and, with that, he unlocked the copier room.

It was 1959, and technology was a far cry from where it is today. Yet it was the year that US President Dwight Eisenhower first sent a message to Canadian Prime Minister John Diefenbaker by means of a radio signal bouncing off the moon as a forerunner of modern satellite communications. Such long-range communication would be one of many new technologies that Lorie Waisberg would see during this long legal career. After starting at what was then known as Goodman & Goodman (a small firm at the time, and today one of Canada's leading firms), he witnessed a parade of new technology, from the popular IBM Selectric typewriters to the new correctable models that made errors fixable. In the early 1970s, the Lexis service was introduced, which allowed lawyers to search case law on computers rather than laboriously poring through books. Fax machines became widely used in the early 1980s, spitting documents out at one to two pages printed per minute. This was a big improvement on waiting for couriered documents, especially when working with others far away. Shortly thereafter, word processors replaced typewriters. Then Lorie got a computer on his desk, then got the internet. "People didn't trust email at first; they wondered who else could see it," recalls Lorie. Eventually, email became a preferred means of communication. Lorie got a BlackBerry.

There were large technology changes over my dad's career, a lawyer for more than 30 years. His father, Harry, a lawyer and then a judge, started his legal career in the mid-1930s and saw new technology and other changes over his many years in law practice.

When I became a lawyer in 2006, email, the internet, and electronic legal research were standard, but we still regularly used physical books to look up information. "The printers" was an actual physical place. And, while virtual data rooms were popular, I had pleasant in-person due diligence trips to St. Louis and Pittsburgh. (“Pleasant” because the host company inevitably shut its doors at some civilized hour, as opposed to my New York Biglaw firm.) As a corporate lawyer, I had little to no specialized technology. We used email, Word (souped-up with some fancy toolbars), Excel, and PowerPoint (rarely!), the internet, virtual data rooms, and document comparison software. Someone passed around a link to an online version of the securities “Redbook,” but we mostly used the hefty physical version, and we (or our assistants) would diligently insert update pages into it as they arrived. If you asked, you could get Acrobat Professional. And, with some real effort and a partner’s permission, the firm would even give you a second computer monitor and, maybe, a laptop. VoIP phones were apparently coming soon, meaning we could take calls from home and have no one be the wiser. We could remotely access our work computer via Citrix. I really appreciated my fancy telephone headset. Things are different now.

Obviously, the legal profession has advanced quite a bit since my grandfather and father’s days as lawyers, and even since mine. Yet, challenges remain part of the job. I recall having to push hard as I started my law career, sifting through what seemed like endless pages of contracts, balancing multiple deals running simultaneously, and worrying that more work was coming when I saw my BlackBerry’s light flashing red. I recall working an all-nighter and sending a draft out just after 6 a.m.; almost immediately I received comments back from a hedge fund client who had gotten to his desk early.

Despite the ongoing changes in legal technology, widespread misconceptions remain that (i) lawyers are loath to adopt new technology, and (ii) technology has historically not been a major factor in law. Yet lawyers have regularly adopted new technology at near-ubiquitous levels, and technology has played a key role in changing how law is practiced. For years, technology has made many lawyer tasks easier to complete, raising the performance bar and allowing lawyers to focus more attention on the needs of their clients. Today, artificial intelligence is the latest step in driving the practice of law forward. AI is getting heavily used in law. It offers real advantages for lawyers who embrace it, and perils for those who don’t. I’m happy to be a part of this change, and, just for the record, my dad is happy for me.

—Noah Waisberg

The Evolution of Kira

Noah Waisberg and Dr. Alexander Hudek first got together in January 2011, introduced by a friend-of-a-friend. At the time, Noah had recently quit his job as an M&A lawyer at Weil, Gotshal & Manges, a very large New York City firm. Alex had recently gotten a Computer Science PhD from the University of Waterloo. Alex was doing post-doctoral research at the time.

For years before leaving Weil, Noah wrestled with the inefficiencies he (and friends at other firms) struggled with. Junior corporate lawyers spent vast amounts of time doing work they hated, weren't very good at, and clients hated paying for. All at—back then—over $300/hour. It seemed unsustainable. And perhaps an opportunity. Noah thought, "What are things junior corporate lawyers spend a lot of time on? Can they be done better?" He played with several ideas, but they didn't seem like they would make great businesses. Then, in conversation with his wife one crisp November day, he started to think about contracts. He realized three things:

1. People spend a ton of time reviewing contracts.
2. They make lots of mistakes in this work, even when they are top graduates, from top schools, who have been through extensive training.
3. People often review contracts for the same things over and over. In M&A, it can be change of control, assignment, exclusivity, and the like. In securities, maybe it's restricted payments baskets or asset sale covenants. In real estate, it might be base and additional rent, subletting, or maintenance responsibilities. And so on.

Since people looked through contracts for the same things over and over, Noah thought it might be possible to build software to help lawyers find and extract this information. He needed a technical partner, and teamed up with Alex to solve the problem. Based on talking with Alex and other Waterloo computer science PhD grads, they thought it would take them four months to harness available machine learning and apply it to this problem. They thought it might take them six months to raise money to pursue their idea, and decided to just plow forward; they could raise money later.

After six months, the software was not working properly—it just wasn't accurate—and there was little chance it would improve anytime soon. As Alex learned more, he realized the state of the art technology didn't work well on their problem. They faced scientific uncertainty. They might crack the problem in three months, but it could take up to 10 years. At that point, they certainly didn't think they could raise any money. Telling a venture capitalist that they thought they would lick the problem in a decade didn't seem like it would make a very compelling pitch, especially when the end product would make lawyers faster at their work.

They Just Kept Building

By 2013, two-and-a-half years later, the software was finally accurate. Early customers found they could do contract review in 20% to 90% less time, with the same, or greater, accuracy.

Sales were *sloooooow*; few people were paying to use the software. Two-and-a-half years of operations, a hard technical problem solved, but little revenue to show for it, selling to lawyers (who were reputed to be anti-tech and anti-efficiency) seemed like a hard VC pitch. So they stayed focused on improving the product and getting people to pay to use it. By 2014, there was more interest in the software, and Alex built a crude version of a long-desired feature that allowed users themselves to teach the software to find new concepts. Now, a person could teach the system without feeling the need for a technical expert at their side. This was huge. Clients could highlight and tag provisions in a document, press a button, and it would learn what to look for. This, plus a market that was getting more and more focused on efficient legal work, ignited the sales of Kira. The company grew from 4 to 8 people in 2014, up to some 35 in 2016, as the customer base also grew. In summer 2018, bootstrapped Kira Systems reached 100 team members and took its first outside funding. As we write this Introduction in summer 2020, there are 240 Kirans.

A healthy majority of the world's biggest and best law firms subscribe to Kira's AI contract analysis software, including 19 of the top 25 M&A firms, 7 of the "Vault 10" most prestigious US law firms, 11 of the UK's top 12 firms by revenue, 5 of Canada's "Seven Sisters," and leading firms in countries including Brazil, Denmark, Germany, India, Norway, and Portugal. It's not just giant firms using Kira. Law firms ranging from solos and smalls to several of the top few firms in places like Missouri or Tennessee subscribe, too. So do most Big Four firms, sometimes for their lawyers, but also for thousands of accountants or consultants to use. Plus, a growing number of corporates, which sometimes use the software to help in-house lawyers, but they often deploy it to help them understand what their contracts say to help with business problems or to augment contract management systems.

Why Are Noah and Alex Worth Reading on the Use of AI in the Legal Industry?

Why are we well qualified to be a guide through this industry? In some ways, we're not. We run a legal AI software company and so may be biased. On the plus side, we have been working on legal AI for almost a decade, meaning we're among the longest-active people in the industry. We have built among the most successful businesses in legal AI. And we bring individual advantages to the table, too. Noah has practiced law, giving him empathy for what it's like to be an attorney. Alex has deep technical knowledge. He began programming computers at age 8, and since has worked on the human genome project, gotten his PhD in computer science, and worked heavily with machine learning on text, as well as formal logics.

In the Pages Ahead

We hope you will come away from this book with two learnings:

1. AI is here in law practice, like it or not. It is already in heavy use in parts of the legal industry, and this will only grow. In time, its use will be ubiquitous.
2. AI can be great for lawyers, if they let it. It can help them do more, better work, generating happier clients; give them more interesting and fulfilling careers; and help them make more money.

This book is not intended to be an exhaustive review of everything happening in legal AI. We are not going to tell you about all areas where AI is being used in law, or which vendors are best. Honestly, it's changing quickly, and we hope this book will be helpful for years into the future. But there's a deeper reason we wrote this book. We believe that if you come away believing that AI can help your legal career, you'll be able to take the next steps to figure out how. Think of it as more like *A Year in Provence* or *Paris to the Moon* than the *Michelin Guide*. More *The Old Patagonia Express* or *In Patagonia* than the *Footprint South America Handbook*. We aren't going to tell you where to get the best socca in Nice, or where to stay in Ushuaia. But, hopefully, we will inspire you to go. Of course, this book is about legal AI, not France, and we're no Paul Theroux or Bruce Chatwin when it comes to writing. Nevertheless, we are optimistic you will find this book worth spending your valuable time with.

Among the many specific points addressed, *AI for Lawyers* will focus on:

- Why AI is now so vital in the legal workspace and how you can expand your opportunities through AI and technology.
- How to amplify legal knowledge through the use of AI.
- The various types of AI tools available including eDiscovery, legal research, contract analysis software, expert systems, and litigation analytics.
- How to incorporate AI into large, mid-sized, or small practices.

While Noah and Alex are among the most knowledgeable people in the world on contract analysis software and why lawyers should embrace AI, others know more than they do about some areas under the legal AI umbrella. So, along with the expertise of the authors, you will also find significant contributions by leading industry experts on some topics. This includes Carolyn Elefant on AI for solo and small-firm lawyers; Mary O'Carroll, Jason Barnwell, and Corinne Geller on modern legal jobs; Dera Nevin on AI in eDiscovery; Jake Heller, Laura Safdie, and Pablo Arredondo on AI in legal research; Joshua Walker and Anthony Niblett on litigation analytics; Amy Monaghan and Alicia Ryan supplementing Alex and Noah on contract analysis; and the magisterial Michael Mills on expert systems. Their background, experience, and insights add to the book's depth.

You needn't read this book chapter by chapter. Some chapters may be relevant for you in your practice, others not. Chapters 1, 2, and 5 are more general interest, primarily focused on objections to and opportunities from adopting AI. Chapter 4 focuses on how AI is creating new types of legal jobs. Chapter 6 discusses ethical issues around legal AI. Chapter 3 should be interesting for solo and small-firm lawyers, but not as useful for Biglaw or in-house readers. Part II (Chapters 7–11) focus on specific areas where AI has caught on in law practice. If you're a corporate or tax lawyer, Chapter 10 (contract analysis software) and Chapter 11 (expert systems) should be most relevant for you. If you're a litigator, Chapter 7 (eDiscovery), Chapter 8 (legal research), and Chapter 9 (litigation analytics) will be more interesting. Part III (Chapter 12) focuses on adopting AI into practice. The Conclusion is more general audience.

This book includes many quotes from people we think have something to add. Unless the source is attributed in an endnote, these quotes come from correspondence with the authors.

AI is here to stay and is changing how lawyers work. It can significantly benefit your career. If you're not already onboard, the time is now. *AI for Lawyers* can position you to get front and center in this new era of law practice. Let's go!

AI FOR LAWYERS

PART I

The Point

AI in law is here to stay. It’s time to take advantage

CHAPTER 1

How Lawyers Learned to Stop Worrying and Love AI

Simon G. is a 46-year-old corporate partner in a major New York–based law firm. He had been a partner for nearly 10 years when he took over as the relationship lead with one of the firm's top clients, a prominent Fortune 500 corporation.

This client was a major source of revenue for Simon's firm and several others. For many years, the firm was on the client's "panel" of legal service providers. To do any legal work for this company, you had to be on its panel. Each firm on the panel was designated for specific types of engagements and projects, and each would form its own deals with the client.

Everyone at the firm who worked on this client's "team" knew in-house lawyers and executives there very well. They had longstanding bonds formed over weeks-upon-weeks cooped up in conference rooms working on deals, as well as dinners, drinks, Yankee games, theater nights, parties, and more. The families of the partners and those of the corporate executives also got to know each other and would be invited to weddings and other family events. One senior partner at the firm even bought a summer house to be near a bunch of executives from this client.

Every three years, the client would go through the process of reselecting its panel of law firms to represent the firm. During each selection process over the decade in which Simon had been a corporate partner, the process had proceeded seamlessly, without even a hiccup.

Now, several of the firm's senior partners were beginning to transition into retirement. Simon was in a position to take on the leadership role of this major client relationship. This was everything he had worked toward. But, as he prepared to take over the leadership role, he quickly found himself in a major predicament.

This time, something was very different in the panel selection process. Instead of Simon's firm and other top-tier firms offering their typical 10–20% discounts, several top-notch firms, including a few that had never served on the panel before, were offering crazy discounts, some as much as 50% below their normal rates. Simon knew that these were excellent firms; he couldn't knock their quality, and he couldn't understand how they could afford to offer such low rates. Worse, he knew his firm could not afford to compete against these offers. Simon's heart sank. He realized that despite decades of great work and strong relationship development by Simon and his mentors, it was painfully clear that the firm was going to be priced out of working with this important client.

Shocked by how the panel selection was going, Simon immediately got on his computer and started doing what he should have done years prior to the panel review—discovering how law practice was changing, rather than assuming the longstanding relationship with this client would simply continue uninterrupted.

Simon spent hours over the next several days studying the competitive landscape, learning about what he and the retiring senior partners had missed. They had overlooked a very important aspect of today's legal industry: the greater drive for efficient work. Now Simon would have to figure out how to make up for falling so far behind his competitors. What he learned was that his competitors, thanks to innovations like AI, were able to do better work in less time. Through tracking and analyzing the time spent to do tasks as well as realization rates, Simon's competitors could figure out how to offer lower unit prices and still make money. Simon's firm was plenty sophisticated when it came to their legal skills, but, Simon was coming to realize, they were seriously outgunned when it came to the modern practice of law. To remain competitive, Simon and his firm would have to embrace technology in a big way to win over major clients and potentially impress their (now former) biggest client in three years at the next panel review.

Simon's problem was not uncommon, and not unique to Biglaw.

If you're a solo estate planning lawyer, how do you compete with online legal solutions like LegalZoom, who offer a will for $179?

If you're a small firm litigator, how do you compete with a bigger firm that has access to case data that's not as easy for you to obtain?

If you have a high-volume practice, how do you compete with firms that spend less time on customer intake because they use software that shortens the intake process and provides clients with self-help?

Now the question for Simon and his law firm was, could they do it? Could they get back in good favor with their most prestigious client?

AI has been a godsend for countless young law firm associates who once toiled late into the night to gather and review data, but has it played a more significant role across law practice? Let's find out. Before launching into the pros and cons of AI and the resistance and opportunities we have encountered, let's explain our definition of AI.

What Is AI?

For the purposes of this book, we consider AI to be any task a computer does that shows "human-like" intelligence or better. The precise edges of this definition are less important to us than the overall impact that AI and similar technologies have on society and the practice of law. To illustrate, let's talk about a few prominent types of AI tasks and techniques.

The field of AI encompasses many subdisciplines, including machine learning, expert systems, and other reasoning technology. At different points in history, a particular technique might be the face of AI. Although expert systems were once all the rage, today deep learning (a type of machine learning) is extremely popular.

In fact, not too long ago, arithmetic was considered an intelligent activity that only humans could perform. The term *computer* originally referred to people who did arithmetic and other math, not a machine that runs software (see Figure 1.1).

We wouldn't consider arithmetic to be artificial intelligence today, but 70 years ago, seeing a machine do this was magic. This shows how the definition of AI has a tendency to change over time. As tasks that we once considered untouchable by computers become routine, our definition of "human-like"

FIGURE 1.1 Early "computers" at work: Dryden Flight Research Center Facilities.
Source: From the Dryden Flight Research Center Photo Collection

intelligence becomes narrower. It's no longer news that computers can dominate at games of chess, and many people today take it for granted that they can speak to their phones. Self-driving cars exist and might become equally ordinary in the years to come.

AI can replicate certain aspects of human intelligence, such as pattern matching or categorization, and can often do such tasks much faster and more accurately than humans. However, AI doesn't have motivation and emotion like a human, and is generally not able to do things it wasn't designed to. The notion of a rogue AI is pervasive in popular culture and movies, but the reality is much less frightening. The AI that can learn languages is different from the AI that can hit a tennis ball, and there is no general connection between abilities. You can't assume that just because AI can win at *Jeopardy,* it will, therefore, make an amazing courtroom advocate. Those are different things. Doing one thing well doesn't mean it can do the other. Although we tend to promote the idea of AI having human intelligence by giving it human names such as Siri, Alexa, or Hal, it's still unable to emulate most of the human thought process, for better and for worse.

All that said, AI is able to do many remarkable things, such as understanding human speech, articulating responses, even writing passable text! How does it do this? It uses expert systems, machine learning, and constantly emerging innovation.

First let's talk about expert systems. These are computer systems that emulate the decision-making process of a human expert by asking a cascading series of questions. For example, an expert system might mimic what your doctor would do when they're making a diagnosis. It may ask: Do you have a fever? Do you have headaches? Do you feel dizziness? And so forth, then propose a diagnosis based on the answers you provided. The questions and decision trees in these systems must be handcrafted by human experts, generally falling into the "rule based" or "reasoning" subfield of AI. Expert systems are a good tool for a variety of tasks, but in many areas they are being replaced by machine learning.

Most of the AI you see in the news today is based on machine learning, including all the various deep-learning advances. Machine learning techniques allow computers to learn to perform tasks simply by observing data provided to them. It doesn't need experts to manually write complex rules, though it still does need to observe people to learn from them. Although the origins of machine learning are as old as those of expert systems, machine learning techniques didn't become widely effective until computers became more powerful. These systems excel at modeling unpredictable and complex tasks and can learn at a rate and scale far beyond what humans manually encoding knowledge in rules could achieve.

From driving a car, to serving as personal assistants, to face recognition, to web translation, to recommending a comedy you might like on Netflix, various types of AI are part of our world in big and small ways. In this book, the technology we discuss falls under our definition of AI. Others may have slightly different definitions of what "AI" is, but we would rather talk about its impact in law practice than debate the exact boundaries of the terms.

In the legal world, AI is being used for contract drafting, negotiation, and review; litigation document review and analysis; predicting case outcomes; suggesting courses of action; organizing legal research; time keeping; and lots more. It is opening up possibilities never before imagined and allowing lawyers to spend more time on law and less time on repetitive activities. AI is partnering with lawyers, rather than replacing them.

Appropriate Skepticism

Most people are averse to change, and lawyers are often perceived as being more change-averse than average. In fact, Dr. Larry Richard (a psychologist focused on lawyer behavior) has found that "skepticism" is consistently the highest-scoring personality trait among lawyers. According to Richard, lawyers have an average skepticism score around the ninetieth percentile, meaning they tend to be skeptical, even cynical, judgmental, argumentative, and self-protective. The general public tends to be at the fiftieth percentile on this trait, which means they'll be generally accepting of others, more trusting, and often give others the benefit of the doubt. Being skeptical is not necessarily a bad attribute for an attorney; helping clients mitigate risk is often a big part of the job. Therefore, it's especially understandable that lawyers have concerns when new technology lands on their doorstep.

"Why should we tamper with success? We've done it that way for 50 years and look where we are today." While a senior partner making that statement is not wrong, they miss that—despite many things staying the same—a lot has changed in the practice of law over the years. Change is inevitable, and today, technology is leading that change. It's no longer a matter of choice but a necessity for those who care to stay relevant.

While lawyers may be skeptical, history illustrates that when it comes to adopting, and even embracing, technology, the legal profession has often overcome initial reluctance and aggressively jumped on board.

For example, the 1970s saw the influx of computer technology. Law firms were able to use the Lexis UBIQ terminal, which later allowed lawyers to search case law online. This opened the door to numerous advances in the union between law and computer technology. Steve Carlotti, an eminent Rhode Island corporate lawyer, tells of their experience at Hinckley, Allen & Snyder LLP with early computer adoption: "We installed our first computer to handle time recording, billing, and accounting in 1976. Since then, profits per partner have risen more than 1200%, at least part of which is due to the ever-increasing use of computers and related software to deliver client services."

By the 1990s, eDiscovery had emerged with litigation support and courtroom management software. This made it possible for legal professionals to quickly process, review, and produce electronic documents for research and to use for cases. It simply would not be possible to manage the discovery process of a large litigation—like the Microsoft antitrust case—without it. More recent examples of near-ubiquitous technology adopted by law firms over the past few decades have include PCs, laptops, email, BlackBerrys, document comparison software (aka redlining / blacklining / DeltaView), and virtual data rooms.

AI is just the latest in an ongoing succession of technological advances that have gained acceptance and approval by legal practitioners. However, like technical innovations that have come before, AI needed to meet industry standards, and it's a high "bar," so to speak.

Common Lawyer Objections to AI

In the course of pitching our own legal technology, we have had a lot of conversations with lawyers about using AI in their practice (so much so that Noah eventually wrote a children's book explaining machine learning in 256 rhyming words). While many lawyers have been enthusiastic or curious, lots had questions and reservations. Over the years, we've seen the same objections recur. Some are issues specific to our software. Many are more general, and could come up almost irrespective of the legal AI software in question.

Recurring issues we've found lawyers raise regarding AI are:

- "How can I trust AI software?"
- "What if our associates use the tech to 'cheat'?"
- "How are new lawyers going to grow into great lawyers with technology doing their work?"
- "Will using AI software impact (i) my duty to keep client information confidential, or (ii) lawyer–client privilege?"
- "Do I have to invest a lot of time in training AI to get value out of it?"
- "If AI makes lawyers way more efficient, will we need fewer lawyers?"
- "How does being more efficient work out for me if I bill hourly?"
- "How do I justify the extra expense of the software?"

"How Can I Trust AI Software?"

Back in 2014, an elite law firm partner explained his trust issue to us this way, and it stuck with us:

> *A couple of people at his firm had been sued for something that went wrong on a deal more than ten years ago. They spent over a decade fighting this lawsuit. His perspective was, "I know the manual way that I do it right now. I did it that way when I was a junior. I know the people who do the work too: I helped hire them, and I've trained them. I know how they work, and that they work hard. Even if it's not perfect, I know it and I know them. I trust them, knowing that my house and my professional career are on the line. How can I trust this new way of doing things?"*

Some find they can get to trust through seeing performance data. They run a test comparing their lawyers doing work the traditional way to those using the software, and see what the results are.

A TEST IN TRUST

By Meredith Williams-Range, Chief Knowledge and Client Value Officer at Shearman & Sterling LLP

I don't trust people in general; as a lawyer, I'm trained not to. If I don't trust people, then I won't trust technology. How do you overcome that sentiment among young lawyers to get them to adopt new technology? Well, you have to take a journey with them. You have to educate them, and you have to bring them along gradually in an effort that should result in them working the way you need them to work.

My experience is that lawyers often start from skepticism with technology, AI or not. So you should recognize that going into any conversation with a lawyer, it will be psychological. It's not about the piece of technology that you're trying to get them to use, it's simply trying to overcome the psychological burden within that individual, on an innate level. What we try to do at Shearman & Sterling is build trust through sponsorship. We have three critical business units: Disputes, Finance, and Corporate. If we're going to go down the path of bringing in a piece of AI, we have to build trust within those groups.

One of the things that we have adopted at the firm is what we call a proof of value, or POV. Why not a proof of concept (POC)? Well, POCs are great, because they prove that a piece of technology actually works. But working is table stakes. To us, the real questions are does it bring value to the partner or associate who will use it? Does it bring value to the client? At Shearman & Sterling, we run extensive POV programs. We measure—side-by-side with the status quo way of doing the task—whether the technology drives value. These tests generate numbers and data, and the results drive trust.

When we evaluated Kira, the POV ran for a full year. Our M&A teams used Kira to perform due diligence with past and live deals. Capital Markets teams tested it as a better way to capture data points. We did the same thing with many of our corporate teams. Different use cases, different purposes, but running the AI hand-in-hand with the young lawyer who was actually using it. Our administrative teams tried it, too. They were looking to review our heap of outside counsel guidelines and to understand some of our own contractual obligations.

When it comes to trust, one of the biggest objections you'll hear from partners is, "Well, how accurate is it?" Our response, after our POV, became, "How accurate are the associates and technology separate, but, also, how accurate are they combined?"

These can be hard questions to answer, but when we run our side-by-side POVs, we find there's more likelihood of human error than there is of AI error. When human and machine hold hands together, we found we did even better than either alone. The combination got us close to 100% accuracy. This is what our testing proved. That helped us create trust. Though this process was more data-heavy, it is pretty similar to how partners come to trust a new-to-practice associate. They see them in action, hear reports from others, and eventually come to trust them (or not).

In law, as well as in other industries, building trust is not an easy process. With over 200 partners, getting buy-in for an AI solution can sometimes feel like trying to get a piece of legislation through the House and Senate. But this is where having a practice like a POV enables you to win over partners quickly. The POV can demonstrate exactly how it's going to alleviate some of the burdens that you have that you're not being paid for. In my experience, that's a good way to build trust.

In the earlier days of our company, nearly all of our prospective customers ran a proof of concept like this, so much so that we once had several team members with the title "Proof of Concept Manager." Today, lawyers increasingly are willing to accept that if many of their peer firms are using Kira, it probably works roughly as expected (over 60% of the Global 100 law firms subscribe to Kira).

Numbers aren't enough to make some people comfortable, though. For them, we are happy to report that you may not need to trust AI to benefit from it. For example, contract reviewers using Kira can still read through agreements page-by-page using the system's built-in document viewer, the same way as they would if doing this work the traditional way. In Kira, however, the reviewers have the advantage of being supplemented by AI. In Figure 1.2, Kira shows the original document, with highlights of information users asked it to find overlaid.

Finally, knowing AI helps build trust. Understanding the possibilities and limitations allows users to learn that AI is not magic, it's software, and software sometimes makes mistakes. Software can do a job or perform a task very

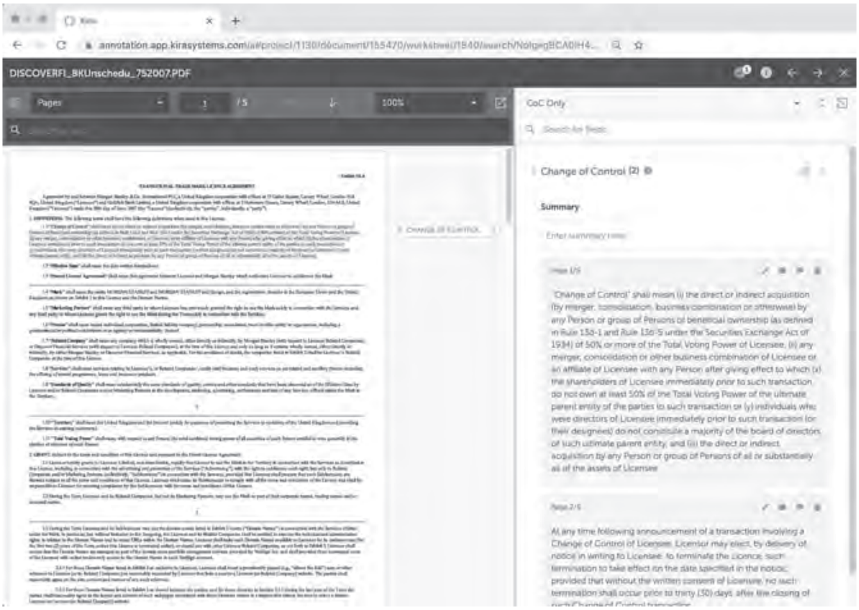

FIGURE 1.2 Kira document viewer.

efficiently as long as you trust in yourself to provide the necessary information, train it to perform specific tasks, and review the results. Trust is earned. Many skeptics rely little on AI when they first start using it, but come to trust it more as they learn how and when it delivers (and when they should count on it less).

Today, AI enhances lawyers, rather than replacing them. AI is helping lawyers do work that they never would have been able to do before. Instead of framing the decision as whether to trust AI or a human lawyer, consider whether you should trust a lawyer doing work the same old way over a technology-enhanced one. We wouldn't.

"What If Our Associates Use the Technology to 'Cheat'?"

Some lawyers worry that—instead of using AI as a supplement—their junior associates will rely heavily on AI to do the work. In other words, they will "cheat" at doing the work. The truth is, there are already non-AI ways to "cheat" at many junior lawyer tasks, and (some) associates use them.

Over our years in and around due diligence contract review, we have heard lots of suggestions on non-AI-ways to do the work faster. You could do a keyword search (ctrl-f) for relevant words such as "assignment" or a phrase like "most favored customer." The problem is that important concepts like "change of control" or "exclusivity" are frequently phrased in nonstandard ways, which makes keyword searching risky. Worse, many contracts for review come in the form of poor quality scans. Keyword searches are hard on text that appears like this, post-OCR:

> *Mengesnorter iigernent or Control*
>
> *If-any-material change occurs in the management or control of the Supplieror_the_Business,save accordance-with-the provisions of this Agreement.*

Instead, some look at the contract's table of contents for relevant sections. However, based on our years of experience in and around the details of contracts, we can assure you that details sometimes turn up in unexpected places. Section titles work as guides, except when they don't. You could also review a company's filings and financial statements to find where to review. This may work as a supplement but, if used without independent review, you are dependent on the company getting it right the first time.

Essentially, there are non-AI ways to cheat at junior lawyer work. But, they have real limitations. If your associates are going to cheat, they're going to cheat. It's about them, and how they believe law should be practiced, not the technology. Pre-AI, you needed to teach them about why they needed to review documents page-by-page, and not just rely on ctrl-f or the like. Same now, with the popularization of AI. You need to train (and convince) your team to do reviews the right way, whatever that means to you.

"How Are New Lawyers Going to Grow into Great Lawyers with Technology Doing Their Work?"

Talk to an "old timer" (by which we mean anyone from a 30-year partner to a third-year associate) and you're likely to hear about how things were different back when they were getting started, and how that molded them into the amazing lawyer they are today.

Many lawyers care deeply about how the next generation will learn the trade. It's no surprise that they worry that AI will harm lawyer training. With contract review software, for example, we have often heard:

> *I learned so much about how contracts work and where problems lurked from reading them through, over and over again. How will junior lawyers pick up this same critical skill set?*

There are three parts to the answer to this question:

1. *Change is constant*. Lorie Waisberg (Noah's father) joined a 10-lawyer firm in 1970. He learned to be a business lawyer at the elbow of a partner who had been at it for some time. Things were busy, and Lorie was given a lot of responsibility early on. They did every type of corporate law back then,

from incorporating businesses, to securities filings, to M&A. Eventually, they did corporate governance, insolvency, and antitrust. By the time the generation after Lorie joined the firm, it had grown to over 125 lawyers. They were more specialized in subareas; M&A and securities had become different disciplines. While associates still got independent work, stakes were now higher and their scope of independent operations was more constrained. They still developed into excellent lawyers. Noah now has stellar lawyer teammates who learned from someone who learned from Lorie way back when. The way lawyers learn is constantly changing. But they still often turn out all right.

2. *Consider the old way of doing due diligence contract review*. A junior lawyer reads through agreements, page by page, looking for consistent data points (e.g., change of control, assignment, restrictive covenants). Or the old way of doing discovery: junior (or temporary) lawyers scan document after document, saying which are relevant, or which are privileged. Today, thanks to AI, things are different. In contract review, AI directs lawyers to passages that might be relevant, as opposed to spending significant time finding the passages in the first place. Rather than spending lots of time trying to find on-point wording (and sometimes missing it), AI makes users consider whether "Customer will buy 100% of its requirements of paper from Dunder Mifflin" is an exclusivity obligation.
3. *AI is here to stay in law practices*. Many lawyers are using AI now. In 10, 15, or 20 years, when today's junior lawyers become senior lawyers, AI will be a standard part of practicing law. Early experience with AI on the ground level, working elbow-to-elbow (so to speak) with AI tools will equip today's juniors to more fully understand the nuances of AI; they will know what it can and cannot do, when it is more likely to make mistakes, and how to most effectively train it. Even though AI will change and improve over time, "AI-enabled-native" lawyers should have a leg up, as they will be able to understand the technology at a deeper level. Firms that dither about getting on AI now are putting their juniors at a disadvantage for the future.

Legal AI by the Numbers

AI-enabled practice is the way of the future. Juniors need to learn to work this new way. Today, AI is becoming the "market" way much legal work is done. A large number of firms and enterprises use technology assisted review (TAR) in their eDiscovery work. Some 80% of the Global 50 firms use contract analysis software (though within-firm adoption varies). Thousands of firms use AI-powered legal research software. These numbers have grown dramatically in recent years, and will continue to grow.

"Will Using AI Impact (i) My Duty to Keep Client Information Confidential, or (ii) Lawyer–Client Privilege?"

Lawyers have a duty to keep information provided by their clients confidential. Lawyer–client communications are also protected by attorney–client privilege (also known as legal professional privilege, among other names). According to *Black's Law Dictionary,* attorney–client privilege is a "client's right to refuse to disclose and to prevent any other person from disclosing confidential communications between the client and the attorney." Lawyers take this very seriously, and sometimes worry that using AI could cause problems here.

In this respect, AI is no different than many other technologies, like email. Do lawyers worry that they breach client confidentiality or risk the protection of attorney–client privilege by sending confidential information in unencrypted email? No.[1] Does the answer change whether the email is sent via a system hosted on the lawyer's premises or by using a cloud-based application like Gmail or Hotmail? No. AI is just a computer program, so it should be treated identically to email, Word, or Excel. Users put data into Excel, and, using formulas, can even have Excel transform the data. AI software is basically the same: you put data in, and it spits out judgments.[2,3]

Does using cloud technology violate attorney confidentiality obligations or impact privilege? Most, but not all, legal AI software is cloud-hosted. For example, over 85% of Kira subscribers use it in the cloud, though it is also available for on-premises deployment. In nearly every jurisdiction, lawyers are ethically allowed to use cloud software, as long as they take reasonable steps to ensure confidentiality. For example, New York State Bar Association's Committee on Professional Ethics Opinion 842 (from 2010) concludes that "a lawyer may use an online 'cloud' computer data backup system to store client files provided that the lawyer takes reasonable care to ensure that the system is secure and that client confidentiality will be maintained." It went on to list steps that may be included in "reasonable care."[4]

We find that most law firms take security (including doing diligence on vendors) very seriously. To assuage their (understandable) worries, technology vendors take steps, including becoming certified under data security frameworks like SOC2 or ISO 27001.[5] What of attorney–client privilege? Again, using New York as an example, under Section 4548 of New York's Civil Practice Law & Rules, "No communication privileged under this article shall lose its privileged character for the sole reason that it is communicated by electronic means or because persons necessary for the delivery or facilitation of such electronic communication may have access to the content of the communication."[6] In short, lawyer use of AI should not raise any special confidentiality or privilege issues. Since that 2010 ruling, cloud computing has become widely accepted. As of the 2018 ABA Legal Technology Survey Report, the majority of lawyers (55%) have now used cloud computing software tools for law-related tasks.[7]

"Do I Have to Invest a Lot of Time in Training AI to Get Value out of It?"

A common misconception about AI is that it takes a lot of effort training a system to get the most out of it, and that you may not be able to train a system without developers or data scientists involved. While this is sometimes true, it depends on which AI system you are using, and what you need the system to do. While some legal AI requires training, plenty do not. Where training is required, it may be done by using a simple user interface. In other cases, training might need to be done with the assistance of technical experts.

Many problems that lawyers need to solve with the help of AI are fairly common problems having answers that can be defined, such as trying to figure out if someone is an employee or an independent contractor; how a specific judge is likely to rule on a motion; which of a set of documents might be privileged; what a pile of contracts says about data points such as change of control, exclusivity, or confidentiality; or how to know what to do when customer data is breached. If you're in need of help on such common issues, you're not alone. Lots of lawyers—from Biglaw to small firms—need the same answers, and there are well-defined pathways to getting those answers. This has led to a lot of legal AI that comes pretrained to work for common use cases such as litigation analytics, legal research, giving HR law guidance, and contract analysis. The use of out-of-the box trained systems is less common around eDiscovery—where the determinations of what is relevant can be more case-facts specific—though there are pretrained privilege determining systems available.

How comprehensive and robust are these off-the-rack capabilities? Do real lawyers use them? We're most familiar with our own situation at Kira, so we will talk from our experiences. As of September 2020, Kira comes able to identify 1,123 provisions out of the box (e.g., assignment, auto-renewal, additional rent, incurrence of indebtedness covenant), across 40 different thematic groups (e.g., M&A, real estate, employment, banking, accounting, or noncontract use cases like UCC financing statements). Kira also comes pretrained to identify numerous document types and languages. As Figure 1.3 shows, Kira's built-in capability has expanded rapidly in recent years. We expect this to continue. Many of our law firm customers heavily use Kira out of the box. On average, over 75% of their usage is with built-in smart fields.

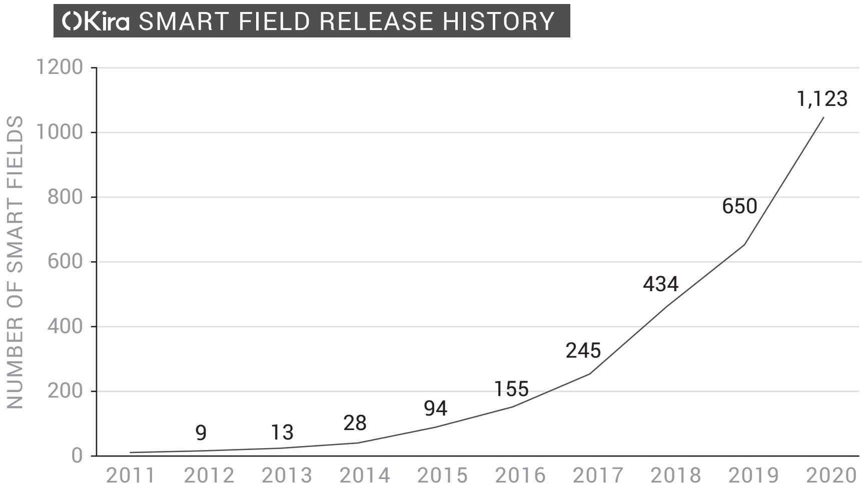

FIGURE 1.3 Kira built-in intelligence: smart-field growth.

Our own built-in smart fields are only part of the story: Kira's users have trained the system to find more than 15,000 additional data points. In the future, we expect many lawyers to choose to share these with others, further expanding Kira's pretrained functionality. So, who trains Kira? We find three main groups:

1. *Firms in some foreign locations have extensively trained Kira for their local language*. This makes sense: today, almost all of Kira's built-in knowledge is for English documents. Kira users in countries like Germany, the Netherlands, Norway, and Brazil have trained Kira to work on their home contracts. Kira has users in over 40 countries, including 10 non-English-first jurisdictions, so there are a number of firms who need to train the system to get maximum use out of it. Though Kira generally does not yet come pretrained in languages beyond English, there are local contract analysis software vendors offering AI software with pretrained clause models in languages including German, French, and Japanese.
2. *Once people get familiar with Kira, they are inspired to teach it more business, or industry-specific concepts*. Industry-specific needs can be regulatory in nature (e.g., for financial institutions that need to comply with global, national, and regional rules) or simply needs that are specific to a vertical (e.g., inventory distribution terms within apparel retailing). There are also endless business-specific needs (e.g., in manufacturing / supply chain management, to find all instances of a particular manufacturer part number or description). Companies are only beginning to

explore all the ways that custom training can benefit their businesses. We believe there is a booming "long tail" of use cases as AI tools like Kira are deployed deeper within companies.

3. *Many people train Kira to represent and capture a specific point of view.* There is tremendous value in training a model specifically on "acceptable" or "standard" language in an agreement. This can allow you to weed out language that doesn't need to be reviewed and save a ton of time. Most law departments now have "playbooks" to encapsulate their points of view on their negotiating position for every major point. Automating these playbooks to correctly route issues for review can save time.

   The other—more profound—reason why you might train Kira to capture a specific point of view is because, frankly, experts do not always agree with each other. You do not need to have attended law school to be familiar with this phenomenon. While most people can agree on the difference between a dog and cat, it takes an expert to give an opinion on whether a photo is of an Alaskan Malamute or a Siberian Husky—and with an imperfect picture, even experts will disagree. In the realm of law, where there is regularly such ambiguity, we see this all the time: one lawyer drafts a termination clause to prevent a customer from canceling their contract early, but then another lawyer sees it as full of holes with easy opportunities for early termination. How should that clause be classified? Does it permit early termination or not? Experts can and do train models to capture their unique insights and expertise. This is something we'll explore further in Chapter 5.

Not every legal AI is the same as Kira. Some will offer more out-of-the-box functionality, some less. For example, legal research and judicial prediction AIs often do not offer users an ability to train them; they just work. On the other hand, many AIs offer or require training.

Why Customize?

Customizing AI can offer advantages. First, it may help you do a task where the AI does not or (like in eDiscovery) could not come pretrained. Second, it can help lawyers amplify their expertise and differentiate against competitors. This is the focus of Chapter 5, so we won't discuss it further here.

When it comes to training AI, there are three possibilities (as shown in Figure 1.4):

1. *It comes pretrained for everything you need.* Examples: litigation analytics, legal research. Many contract analysis, expert systems, and legal prediction systems will come heavily pretrained, but may also offer training interfaces.

PRE-TRAINED	TRAINABLE VIA USER INTERFACE	TRAINING REQUIRED WORKING WITH TECHNOLOGISTS OR PROFESSIONALS
EXAMPLES: LITIGATION ANALYTICS LEGAL RESEARCH	EXAMPLES: KIRA EXPERT SYSTEMS eDISCOVERY TAR SYSTEMS	EXAMPLES: MANY CUSTOM AI PROJECTS

FIGURE 1.4 AI training capabilities.

2. *It is trainable via a user interface*. No data scientists or other technologists are required to intermediate with the trainer's work. Examples: Kira, expert systems, eDiscovery TAR systems.
3. *Training requires working with technologists or professional services*. Often, AI systems require the training be done through working with technologists. Systems built this way sometimes give impressive results on specific narrow tasks (because they have been tailored to work on these), but performance can be brittle (not able to work well beyond the exact intended use case), and further extensions will require working with technologists again. This is not a particularly scalable approach. Examples: many custom AI projects.

The Business Case for AI

Over the past pages, we have covered a number of reasonable, recurring objections to using AI in law practice. Three related objections remain:

- "If AI makes lawyers way more efficient, will we need fewer lawyers?"
- "How does being more efficient work out for me if I bill hourly?"
- "How do I justify the extra expense of the software?"

Our experience has been that these are critical. Where partners are convinced that adopting AI makes good business sense, we often see other objections melt away. Think of a manufacturer like GM or Toyota questioning whether to adopt new technology that enables them to produce an important car component like an engine in half the time. They would be likely to work hard to find a way to implement it. So, in the next chapter, we'll delve into why adopting AI can be financially good, even for hourly billing lawyers.

If you're wondering about Simon, whom we introduced at the beginning of the chapter, his firm is slowly moving into a new way of practicing. Unfortunately, they are already behind their closest competitors in figuring out how to practice law more efficiently, and even in knowing how much it costs them to deliver individual pieces of legal work. To avoid facing similar or even longer odds in the future, they need to accelerate their evolution. Their more sophisticated competitors are certainly not slowing down, and Simon knows that there is no room for complacency. From talking to his peers, and by keeping his eye on the AmLaw 100 rankings, he knows that a firm's position is by no means secure. Firms in the top 10 tend to be stable, but a large share of the firms in the rest of the top 50 have moved up or down significantly over the past dozen years, in both revenue and profitability, and all of them are looking for competitive advantage.

The good news is that Simon's firm has recognized the need to change. We know firms where partners view doing the same work in less time as a silly exercise that leads them to earn less. One example we heard that has stuck with us involved a Biglaw partner asking a knowledge manager at his firm about the status of an automation project by asking how the "PRS" was doing. When the baffled staffer asked what he meant, the partner replied, "the profit reduction system."

The firms and legal teams that are pulling ahead are ones who understand that AI is creating new business models, new economies of scale, and new revenue opportunities that were never thought possible. This is our focus in Chapter 2. Let's dive in!

Notes

1. Should they worry about confidentiality breaches using unencrypted emails? Yes.
2. In this book, we will ignore tech-enabled services, which market themselves as AI but are really work done by people with the assistance of technology. They need to be considered separately, but since this is a book on AI, we will not do so here.
3. We are not aware of anyone seriously questioning whether using Excel or the like impacts confidentiality or privilege, and do not see any distinction with AI (apart, perhaps, from whether training a system raises issues, which we discuss in more detail in Chapter 5).
4. https://nysba.org/ethics-opinion-842/
5. There is a big difference between a vendor being, e.g., "SOC2 certified" and "hosting their application in a SOC2 certified data center." Large hosting providers like Amazon Web Services are usually certified themselves, so—while the latter is better than nothing—it is different than being certified yourselves.
6. NY CPLR § 4548 (2012) NY Civil Practice Law and Rules.
7. Dennis Kennedy's official ABA writeup states, "Actual usage might be higher than the reported usage. For example, many mobile apps are also essentially front-ends for cloud services. Many lawyers who do not think that they are using the cloud may, in fact, be using it every day, especially through mobile apps." https://www.americanbar.org/groups/law_practice/publications/techreport/abatechreport2019/cloudcomputing2019/

CHAPTER 2

#DoMoreLaw: How Doing Work More Efficiently Can Create More Legal Work, Not Less

Alyssa, a young lawyer, showed up at work a few minutes late on a rather ordinary Tuesday, after sitting in traffic en route to the Los Angeles law firm she has been working at for nearly two years. Emerging from the elevator on the 12th floor, the receptionist, Frannie, gave her a peculiar look, as did one of the firm's partners. It was as if they did not expect to see her at all. She was knee-deep in contract reviews, so she quickly headed for conference room 12-E, where she and five of her comrades had been poring over documents for a merger between two film studios. As Alyssa made her way down the corridor, she passed Monica, a paralegal who looked to be on the verge of crying. Alyssa opened the door to the conference room. There they were working diligently, just two of them, robot associates, one wearing Alyssa's identical outfit. It had finally happened. She was, as expected, replaced by a robot. Alyssa screamed in horror. And then she woke up.

Yes, AI can be scary, very scary, but it's not coming for you like *Ex Machina*. In fact, it's opening the door for legal professionals like Alyssa to do more, more interesting work.

Like other industries, there's an ongoing debate on the human impact of AI. However, more and more organizations have come to realize that AI augments lawyers rather than replacing them. It is changing how lawyers work.

By reducing the time-consuming and laborious aspects of their jobs, lawyers are now able to focus on more strategic high-value work.

While that sounds great, there is still a lot of skepticism. After all, if AI makes lawyers so efficient, saving time and money, how can there be more law to be done? We finished Chapter 1 with three unanswered questions:

- "If AI makes lawyers way more efficient, will we need fewer lawyers?"
- "How does being more efficient work out for me if I bill hourly?"
- "How do I justify the extra expense of the software?"

These questions are deeply interconnected, so, instead of treating them individually, we have made them the overall focus of this chapter.

Jevons Paradox: The More Efficient Legal Work Is, the More Legal Work to Do

People often think about how technology will dramatically shrink the amount of work they have to do. They focus on the negatives and zero in on the trimming. They worry about a 60% time savings in one place, 15% in another area, 5% somewhere else, and they see a bleak picture for their own careers. They don't recognize how greater efficiency can drive more work, for both current clients and potential new ones. The demand side for legal work is almost always steadily expanding, due to both (i) an economy that almost always grows over time, plus (ii) a world that is growing ever more complex and regulated. These two factors combine to drive a regularly growing need for legal guidance, whether that means additional research, reviewing details in contracts, having more leases or patents to review, or—even—more lawsuits.

A fallacy in much of the thinking about the growth or decline of legal work lies in equating "legal work" with "work that law firms do." It is true that law firm revenues have been relatively flat or only slightly growing since the Great Recession. Thomson Reuters' Peer Monitor service tracks demand for law firms over time. As you can see in Figure 2.1, after a deep contraction in 2009, growth has floated along just a percentage or two above or below 0% ever since.

But there is more legal work than what law firms do. Technology has enabled more in-house legal departments to retain work in-house, avoiding the premiums that law firms have charged for routine or process-oriented work. An entire Alternative Legal Services (or NewLaw or Law Company) sector has been built up to handle some forms of outsourced legal work, much of which is process-driven and lends itself to technology-enabled services.

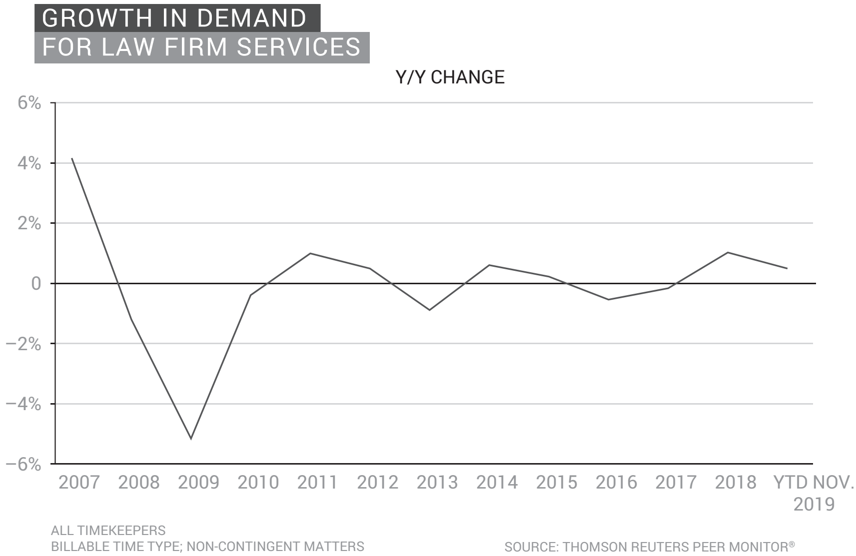

FIGURE 2.1 Growth in demand for law firm services.

Source: From Thomson Reuters, 2019 Report on the State of the Legal Market: Growing Competition Challenging Long-Held Law Firm Assumptions, Legal Executive Institute January 8, © 2019, Thomson Reuters.

The growth in the share of legal work that is retained by in-house legal is hard to measure in dollars, but one proxy for that growth lies in the number of lawyers employed by in-house compared to the number of them employed by law firms. In the United States, the difference is striking. In Figure 2.2, from a July 2018 analysis of the legal industry, it's clear that the growth in in-house employment is exceeding law firm employment by a large margin over 20 years. All the corporations that they work for are doing more and more legal work; it's just that law firms aren't being hired for all of it.

Throw in the legal work that's being performed by Alternative Legal Services Providers (ALSPs) and we see a widening gap between the total demand for legal services and the share of that work that's going to law firms. Much of that work is being retained by in-house departments or sent to ALSPs precisely because those organizations have been willing to apply technology to accomplish more with fewer resources. Law firms willing to make similar investments in technology might be able to claw back some of that gap.

So while this expanding legal market is one reason why we see more work in the future, the more exciting and bigger reason is that efficiency (and the lower unit prices that it brings) opens the door for greater overall demand for legal services.

Technology increases efficiency, resulting in less energy used for a specific task, which creates cost savings. However, demand can go *way* up as unit prices decrease. For example, Henry Ford was able to build a less expensive,

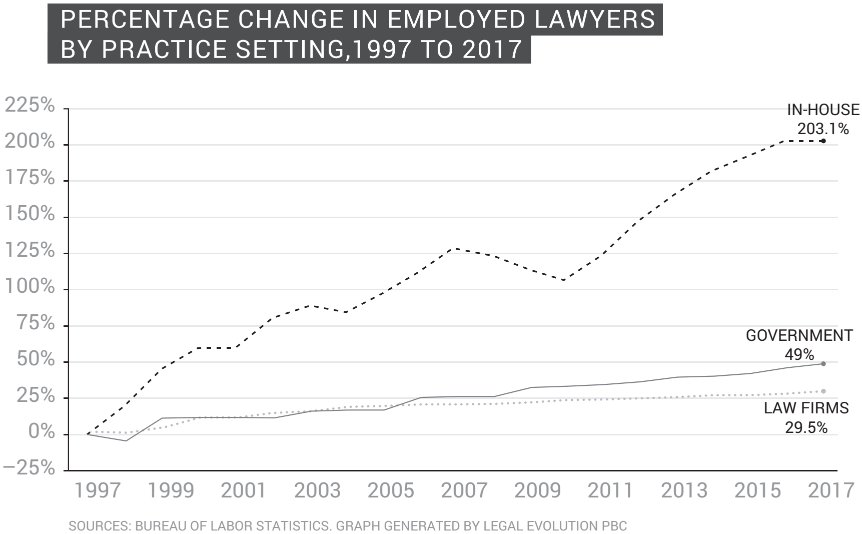

FIGURE 2.2 Percentage change in employed lawyers by practice setting, 1997 to 2017. Source: From Bill Henderson, Our journey to Big (067), September 16, © 2018, Legal Evolution PBC.

mass-produced automobile on an assembly line, which meant fewer jobs assembling each individual car. As aspects of production were performed on a moving assembly, the manufacturing time to produce one Model-T was reduced from 12 hours to 93 minutes, while using less manpower. This pushed the price of a Model T dramatically lower than other cars. In 1911, a Model T cost $700, where average competitor cars cost $1,100. As the Model T's price continued to drop, this drove huge demand. As Wikipedia tells:

> *In 1914, Ford produced more cars than all other automakers combined. The Model T was a great commercial success, and by the time Ford made its 10 millionth car, half of all cars in the world were Fords. It was so successful Ford did not purchase any advertising between 1917 and 1923.*

More efficient production meant far more automobile production jobs (first at Ford, then elsewhere) were created.

The refrigerator is another great example of how efficiency paradoxically drives more effort (see Figure 2.3). Back in the 1970s, it took some 800 kilowatt hours a year to run a single standard size refrigerator. Today it takes roughly 200 kilowatt hours per year to run a typical refrigerator. And, by the way, today's typical refrigerator is 20% bigger.

Based on refrigerators taking a quarter of the electricity to run, the world must use less energy on refrigeration today, right? Well, no.

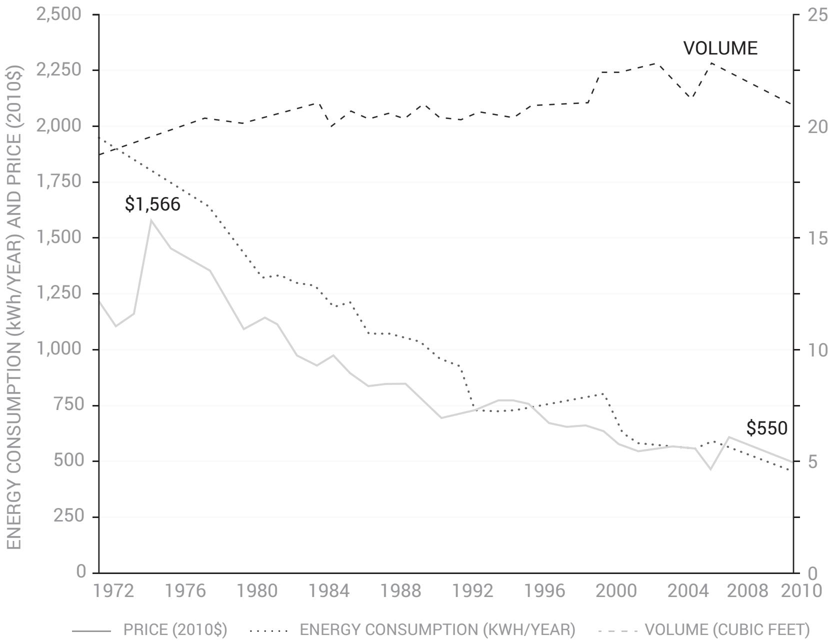

FIGURE 2.3 Refrigerator efficiency paradoxically drives more effort.
Source: From Meneghini La Cambusa Refrigerator, available at Robeys 窸

First, increased global standards of living mean people have refrigerators who wouldn't have had them before. For example, in 1995, only about 66% of households in China had a refrigerator. In 2018, 99% did.

The more interesting part of the refrigeration story is how inexpensive refrigeration made it possible for stores to expand their refrigerated foods sections and, as a result, offer a multitude of new products.

The milk section of a well-stocked grocery store once featured skim, half, whole, and chocolate milk. Today, even a random rural supermarket may now feature multiple brands of these classics (including organic and lactose-free versions), plus soy, almond, coconut, rice, oat, hemp, and cashew milk options. Like bake-at-home frozen croissants or baguettes? Okay. What about kale juice? You now have multiple choices! Literally, a convenience store today may have as much refrigeration as a supermarket did back in the 1970s. While it's easy to mock oat milk or kale juice, we are all richer for the choice. Noah remembers the lactose-intolerant kid in his early-1990s summer camp cabin who had to eat his Cheerios with juice, while everyone else got milk. Now, his ailment would be no issue. And—even without *needing* these new products—people happily buy them. Lots and lots of them. Almond milk is expected to reach a USD $6.77 billion market size in 2020. Meaning it must have value for some. Along the way, these new products are creating lots of new food production jobs that never would have existed before.

Affordability led people to buy more refrigerators. Today, it's not uncommon to have a two-fridge kitchen, an extra one in the basement or garage, a separate wine fridge in the kitchen, maybe a bar fridge, too. Plus an additional chest freezer, and perhaps one under the desk in your office. The demand for larger and, in some cases, multiple refrigerators is partially because they are so cheap to buy and run, and partially because there are so many more items to refrigerate today. Just as efficiency increased refrigeration and brought with it a demand for new products, technology-driven efficiency in legal should create demand for more lawyers.

These scenarios illustrate a phenomenon known as the Jevons paradox. Economist William Stanley Jevons saw how, during the First Industrial Revolution, technology was making coal usage more efficient over time. Surprisingly, this led to more coal usage. Efficient technology made coal effectively cheaper to run, leading to more possible uses, counterintuitively increasing overall coal consumption.

An article by author and researcher Darrin Qualman explores the paradox. Qualman discusses lighting and the cost of an hour of illumination. Adjusted for inflation, lighting in the United Kingdom was more than 100 times more affordable in 2000 than it was in 1900. This is because electric power plants are far more efficient, which has driven the price of lighting down. Therefore, the cost of running a single artificial light would be cheaper in 2000 than in 1900. Yet the Jevons paradox once again enters the equation when you look at the significant increase in the need for artificial lighting. As Qualman writes, "The average UK resident in the year 2000 consumed 75 times more artificial light than did his or her ancestor in 1900." Noah remembers his grandmother almost obsessively turning off lights in rooms not in use. Now we light heavily (for better or worse), because the unit cost is so low that it is no longer a barrier.

Law seems ripe for a Jevons paradox increase in usage. Legal services are generally very costly, and consumers have *lots* of unmet legal needs. Costs going down (combined with other delivery changes) could dramatically increase the volume of legal work done.

"Would You Like to Supersize Your Diligence?"

From 1992 to 2004, McDonald's ran a promotion that has remained seared in the public consciousness. For a relatively small amount more money, patrons could "supersize" their meal—getting an even larger fries and soda. After all, why stop at fries that contain 50% your daily recommended fat intake when you can have even more?! While supersizing has not stood the test of time at McDonald's (the 2004 Oscar-nominated documentary *Supersize Me* probably hastened its demise; supersize meals were cut six weeks after the film's premiere), it provides an important lesson in how lawyers who work more efficiently can do more law.

Today, reducing spend on outside counsel is a priority for in-house legal departments. Altman Weil's 2019 survey of chief legal officers is a good source of insight on this. This annual survey asks which "levers" CLOs are pulling to cut spending, and they also ask which of those tactics are most effective. The top two most successful strategies were "outsourcing to non–law firm vendors" (95% saying it drives "significant improvement to cost control") and "shifting law firm work to in-house lawyer staff" (93%). Others include "negotiating price reductions on portfolios of law firm work" (91%); "receiving discounts on law firm hourly rates" (86%); and "using alternative fee arrangements" (82%). Ben O'Halloran, former CLO of a large European private equity–backed company (and previously a senior lawyer at General Electric), says:

> *Law firm services are typically purchased flexibly, on-demand, resulting in a) a market-determined (and entirely valid) pricing premium, and b) limits on knowledge economies that can be achieved (because the on-demand law firm resources continuously vary and are less integrated into the client organisation, its operations and business priorities). Forward-looking legal teams are increasingly approaching legal work flows like process engineers, working to define categories of repeated legal work flows where quality improvements deliver meaningful impacts to the business, and then, where volume is sufficient, to in-source or sole source key parts of those work flows in order to improve quality and efficiency (often with cost savings as well).*

Shifting law firm work to in-house legal staff is an ongoing trend, one that goes against a more general trend toward outsourcing non-core work and services in the corporate sector. The Altman Weil data shows this as well: 36.3% of CLOs anticipated increasing in-house lawyer staff in the next 12 months vs. only 8.5% who were planning staff reductions. In a historical perspective, this study shows that on average, roughly four times as many corporations plan to increase legal staff each year vs. reducing staff since 2010. Thomas Barothy says, "I took over as COO of UBS's legal team in 2017. Since then, we achieved material annual reductions in our spend. Primarily, this has come through expanding our in-house legal team, doing work internally that we used to do externally, and implementing a dedicated outside counsel management team." In a very real sense, a law firm's biggest competitor is often its own clients, as those clients find ways to get better value by solving legal problems themselves with the help of their own staff and technology.

While some clients are after paying less, many seek something that sounds similar, but really is very different: better value legal work. Casey Flaherty (former in-house counsel; consultant to law departments, law firms, and other legal service providers; and author) has extensive experience with legal buy programs, which he says:

> *now almost always include large sections in their RFPs on the how, not just the who, of legal service delivery. Lawyer quality remains the threshold consideration. But, once that threshold is passed—once law departments are in the room deciding between the select firms they already deem excellent—lawyer quality stops being decisive. Demonstrable differentiators, including AI usage, can have a significant impact at the margins—and the margins are what matter at the final selection stage.*

Rosemary Martin, group general counsel and company secretary of Vodafone adds:

> *As a buyer of legal services, I look for value: not necessarily the cheapest option but the one that I think will deliver the outcome I am looking for, be that success in a case, speed in contract execution, or precision in defining the terms of a complex legal relationship.*

Ben O'Halloran concurs:

> *Legal departments are generally looking to maximise value from all of their spending, whether that be through choices between internal or external resources, process improvements, or deploying new technology. While law firm partners may sometimes interpret the client's objective as mere cost reduction, typically General Counsels are more focussed on improving the value they get from their law firm spending—and that involves both streamlining cost as well as expanding and maximising the potential benefits from outside counsel services.*

We know corporate law best, so will go there for an example of how lawyers can provide more value.

Let's consider a $400 million acquisition of a company. Typically, counsel reviews anywhere from 75 to 500 target company contracts during the due diligence process. However, a $400 million company might actually have 5,000–10,000 contracts. Why is such a low percentage reviewed? Is it because there isn't likely to be anything interesting in the unreviewed contracts? M&A lawyers *hope* so. The status quo approach is to review all "material" contracts. Is this an optimal approach? Let's explore further.

Material contracts generally come in two buckets:

1. *High-dollar value or otherwise strategically important contracts*. These tend to be easy to find. You ask the target or their investment bankers which contracts matter, and they give you the list. Then you review them.

2. *Contracts that say something that could be bad for the client*. Sometimes, contracts that aren't otherwise very important say something important. They have a badly drafted buried exclusivity or most favored customer provision in them, which brings in affiliates. Or an out-of-hand indemnity. Theoretically, deal lawyers would like to think that they catch these. But how? You're pretty unlikely to find a problematic provision in contracts you don't review.

In status quo review, only contracts in the first group are reliably reviewed. *Maybe* lawyers review a "sample" of other contracts as part of their review, but this tends not to be a scientifically drawn sample, at least not in a way that Alex and his PhD peers would recognize as a valid approach. Clients may be missing lots of dangerous information, but they and their lawyers would never know.

The only way to be sure is to review the contracts. Why doesn't this happen now? Generally, it's because doing more than a small review is simply not time- or cost-effective for most businesses. Happily, thanks to AI, total diligence is now possible. Rather than a 10% sample, you can now review 25%, 50%, or even 100% of contracts in question, in a manageable amount of time, for an acceptable (though not necessarily low) amount of money. Truth is, a nonrandom sample is often a poor research approach. Yet, many companies are willing to take a chance with it because they believe the odds are in their favor that they won't miss something problematic. In a lot of cases, it works out fine. If, for example, you trigger a change of control clause in a minor software license agreement, it's probably okay. If you breach, you may have to pay a small penalty. Oh well. We've seen this happen, where the client may happily incur a $20,000 penalty instead of spending an additional $300,000 on legal fees to find and avoid penalties like this. They consider it a pretty good risk/reward tradeoff. They're probably right when it comes to avoiding small penalties built into minor contracts. The problem is that contracts can have much worse things in them than this.

More fulsome contract review has different values for financial and strategic buyers. Typically, there are two types of company buyers: financial and strategic. Financial buyers (like private equity firms) are concerned with buying and reselling businesses at a gain. Since they tend to buy businesses and run them as is (in an isolated legal entity), financial buyers are less likely to have issues with a target company's contracts. However, they can still benefit from a faster and deeper review. A faster, light contract review early in the evaluation period can help financial buyers determine which deals to lean into. Also, findings from more thorough diligence can allow financial buyers to more accurately set a fair price for the asset. For example, if an exclusivity clause limits a target's scope of operations, its value may be impaired.

Strategic buyers, on the other hand, add companies they buy into their already-running businesses. This significantly raises the stakes on contract review. Contracts of an acquired entity are equally binding as those the company enters into in the ordinary course of business. Many companies put a lot of effort into ensuring that new contracts they enter into are properly approved. They should be equally careful about contracts they acquire in M&A. Sometimes, really bad things lurk in contracts, even seemingly inconsequential ones. Exclusivity. Non compete. Most favored customer pricing, indemnifications, uncapped liabilities, data transfer restrictions, and other clauses you might never find unless you dig down and look closely. These risks can compound when brought under a large acquirer's significant corporate umbrella. Imagine an emerging beverage company. If things go wrong as it grows, this could wipe out the company, but the company might not be that big so losses are naturally limited. The company is, to a certain extent, "judgment proof." If Coca-Cola or PepsiCo buy them, all of a sudden there is a whole lot more to lose.

The ROI of AI: Explained

Lawyers can add value by using AI to increase the number of contracts they review in transactions. Some clients might be happy to get a lower diligence bill thanks to faster AI-enhanced contract review. But many should be *very* interested in getting twice the diligence for the same price they paid the last time they did a deal, or three times for 30% (or 50%) more money. Figure 2.4 illustrates this. In this simplified example, we assume a 500 contract due diligence contract review, with a reviewer taking an average of 45 minutes a contract reviewing the traditional way, and 20 minutes per contract doing a thorough AI-enhanced review. Kira users consistently tell us they review contracts in 20–90% less time, so this 55% time savings is reasonable. While one review choice could be to use the extra time to reveal all additional contracts as thoroughly as the initial 500 "material" contracts, a better strategy might be to trust the software to spot issues in remaining contracts. While the software might make mistakes here that reviewers won't catch, you miss 100% of the dangerous provisions in contracts you don't review. These additional contracts wouldn't have otherwise gotten reviewed. In keeping with this strategy, we assume all contracts after the initial 500 will be reviewed in 5 minutes per contract. (In fact, a human reviewer would likely spend 1–2 minutes on many of these contracts, and a lot more on a few where the software or their intuition guided them to lean in. Averaging to five minutes per contract overall.) We include report preparation in this minutes-per-contract time assumption, but—likely—client reporting will be pretty slim on the non-material contracts (apart from where something gets found).

FOUR DILIGENCE REVIEW SCENARIOS

	TRADITIONAL	AI ENHANCED DILIGENCE CHEAPER	AI ENHANCED DILIGENCE NEAR SAME COST		AI ENHANCED DILIGENCE SUPERSIZED 35% UP-SELL	
CONTRACTS REVIEWED	REVIEW 500 CONTRACTS	REVIEW 500 CONTRACTS	REVIEW 2,108 CONTRACTS TOTAL INCLUDING 500 IN DETAIL		REVIEW 3,248 CONTRACTS TOTAL INCLUDING 500 IN DETAIL	
TIME PER CONTRACT	45 MINUTES PER CONTRACT	20 MINUTES PER CONTRACT	20 MINUTES PER CONTRACT FOR 500 CONTRACTS	5 MINUTES PER CONTRACT FOR ALL ADDITIONAL CONTRACTS	20 MINUTES PER CONTRACT FOR 500 CONTRACTS	5 MINUTES PER CONTRACT FOR ALL ADDITIONAL CONTRACTS
TOTAL REVIEW TIME	375 HOURS	166.7 HOURS	315 HOURS		410 HOURS	
HOURLY FEE	$350	$350	$350		$350	
TOTAL FEE (INCLUDING $10/CONTRACT FEE FOR USE OF SOFTWARE)	$131,250	$63,333	$131,336		$175,986	

FIGURE 2.4 Four diligence review scenarios.

We assume this is an hourly billing firm, and that the hourly rate for this work is $350 (low for the fanciest US Biglaw firms, but high for small law). We also assume the AI technology costs $10 per contract reviewed. In fact, many contract analysis software offerings are less expensive, but there's no need to shave the numbers too close in this example.

Here, we have four options. Traditional review, where a client gets 500 contracts reviewed for $131,250. AI-enhanced review, where all savings get passed on to the client, and they get the same 500 contracts, this time for $63,333. If the firm agreed with the client to keep roughly to the initial manual review budget but do the work AI-enhanced, they could get 2,108 contracts reviewed, including 500 thoroughly. That's more than 4× coverage. And, if the client was willing to "supersize" their diligence, paying some 35% more than the manual review cost for this work, they could get 3,248 contracts reviewed. Over 6× manual coverage, for 35% more money. Seems like a pretty good value to us, especially if they found anything dangerous in the contracts that would otherwise have been ignored.

Selling risk mitigation is a core skill of many top law firm partners. Upselling more thorough work—whether diligence, or precedent review, or something else—should fall right in their wheelhouse. Supersizing legal work seems unlikely to have a provocative, Oscar-nominated documentary created to criticize it. Instead, it will leave clients happy about paying their lawyers more money. Does upselling actually work? From 2017 to 2020, the average number of documents in a cloud project inside Kira has doubled. Though there are a few reasons this might be so, we think the most likely explanation is that people are doing bigger projects now because AI technology allows them to.

There's another piece to this puzzle; we didn't cover realization rates in the example above. We suspect the lawyer might have an easier time getting paid in full in the AI-enhanced situations. Improved realization rates are a core way hourly billing lawyers can do better financially through doing more efficient work. On average, US Biglaw firms have an 89% realization rate. That means that after discounting off their standard rate and reducing the hours billed to accommodate client demands, firms are leaving on average 11% of their potential billings on the table. Beyond that, many clients will write off additional charges, resulting in an even lower collected realization rate. In fact, this overall number hides important details. Clients often view partners—even very expensive ones—as good value. (As clients ourselves sometimes, we generally think they're right.) On the other hand, some clients refuse to pay for junior lawyers. Some practice areas have better realization rates than others. For example, in American bankruptcies, bills are approved by the US trustee, which is much less aggressive on law firm bills than a Fortune 500 legal ops or procurement team.

Figure 2.5 shows how AI use can impact realization rates. We will imagine a different AI-enhanced project from the one considered in Figure 2.4. Here, in scenario 1, a firm bills its client $200,000 for some junior lawyer work. In fact, the partner wrote off 20% of the amount their associates worked on this project before even sending the $200,000 bill, because they didn't think the juniors worked efficiently, and they worried about upsetting the client and damaging their relationship. These write-offs are common. Despite this preemptive write-off, the client only paid 65% of the diligence fee, still feeling that the work wasn't done efficiently. (The client is right!) Eventually, after lots of haggling, the firm got paid $130,000. Now, consider the AI-enhanced scenario 2. Here, the partner feels good about the efficiency of their team, so they bill all hours worked: $250,000. Throughout the matter, and in delivering the bill, the partner explains how their firm is focused on efficiency, and the client is happier about the value of the work they received. To be conservative, we assumed only a 10-point jump in realization rate, though—if the partner is good at selling value—this might be higher. Here, because the bill was higher (due to no preemptive write-off) and because of the

INCREASING REALIZATION RATES

SCENARIO 1 TRADITIONAL WORK		SCENARIO 2 AI ENHANCED	
DUE DILIGENCE FEE	$200,000	DUE DILIGENCE FEE	$250,000
WRITE-OFF RATE	20%	WRITE-OFF RATE	0%
REALIZATION RATE (OF JUNIOR LAWYER'S WORK)	65%	REALIZATION RATE (OF JUNIOR LAWYER'S WORK)	75%
		COST OF AI	$10,000
TOTAL PAID BY CLIENT	$130,000	TOTAL PAID BY CLIENT (LESS COST OF AI)	$177,500

FIGURE 2.5 Increasing realization rates.

higher realization (collection) rate, the firm makes an extra $45,500, despite us assuming that the AI cost $10,000. That's 35% more revenue! And—to keep the numbers simple—we didn't even look at matter profitability here. Suffice it to say that throwing out hours—either because you don't bill them or the client doesn't pay for them—is bad for profitability. Changes in realization rates can *really* make an impact. If a firm has an industry average 89% realization rate, and has over $1 billion (or $10 million, for that matter) in revenue, the money (and profit) it is leaving on the table can be pretty *immense*.

While the previous example was centered on hourly billing lawyers, fixed-fee work is getting more and more popular. In some jurisdictions, like the United Kingdom and Brazil, we believe the majority of transactional work (including due diligence contract review) is fixed fee. In a fixed-fee situation, lawyers who can generate the same amount of output with less effort are going to make more profit. Fixed-fee work can even be very profitable for firms, even at lower prices, if the firms get more efficient. Happily, there tends to be *lots* of room for more efficient work in law practice. If we remember the plight of poor Simon G. from Chapter 1 (whose firm lost its longstanding panel position with a key client), fixed fees (coupled with efficiency) are likely part of how Simon's equally prestigious competitors were able to so severely undercut his firm . . . and how Simon and his firm can fight their way back onto the panel next time.

Law firms can also realize ROI from AI by using their efficiency in their pitches. AI users tell us that AI has helped them win new work and retain work that might have otherwise gone to lower-cost providers. Law firms spend huge amounts to win new business. Noah's old firm, for example, threw (wonderful!) lavish parties for alumni. One was at the Central Park Zoo and included a seal feeding partway through. While Noah would like to think that they hosted him and his ex-colleagues just to catch up, he suspects the real purpose was to drive deals and litigation projects from alums who had moved in-house. While law firms don't all rigorously measure wins from using AI, others do. In 2016, Ragu Gurumurthy, Deloitte's chief innovation and chief digital officer, stated in a CNBC video, "[Kira] has tangibly enabled us to generate about $30 million of new revenue that we would not have generated otherwise." The impact of winning new business or retaining work that might have gone elsewhere can be very significant, easily covering all the time you spend reading this book, its cost, and the (much more significant) time and materials costs of implementing AI. Ultimately, delivering good value is good business. It's even better than putting on a fancy party.

Corporates who deploy legal AI tend to realize value in two ways. Many use AI to do work more efficiently. Since—for most businesses, apart from those that bill hourly—there's a strong view that doing the same work in less time is better, this is a pretty easy way to realize a return on investment. The more interesting way corporates find value is by being able to use AI to uncover information they wouldn't have been able to find without it. This can enable companies to make better, more informed business decisions and nimbly respond to environment changes. Here, the benefits can be hard to measure, but enormous.

Value Is in the Eyes of the Beholder

A satisfied client is someone who feels that their money was well spent. Remember the famous tracking shot through the kitchen in *Goodfellas*, as Henry and his date Karen enter the legendary Copacabana Club in New York City to The Crystals singing "Then He Kissed Me"? Doors are flying open ahead of them, as waiters scramble to grab a table and chairs and anchor them front and center in the best location in the house, before breaking out the finest wine. Now that's first-rate service and good value even at a price beyond what money can buy.

Many of us have encountered situations where we have spent more than usual and walked away saying, "It was well worth it." Value need not be expensive—maybe it's buying a box of Afeltra linguine for $6.90 (or even Barilla for $1.99), instead of Signature SELECT brand for $0.69. Value is all about feeling you did well for the money spent, relative to other choices. Danny Meyer (among New York's leading restaurateurs) believes that value can be had at any price; you just have to know how to find and package it. Meyer and his team have found great success at a range of price points, from fine dining at Gramercy Tavern and Union Square Cafe to the popular jazz clubs Blue Smoke and Jazz Standard, to 249 Shake Shack fast-food eateries. Meyer says, "Essentially, what's going to determine how you succeed in New York is how the people feel about the space, how good the food is, how they perceive the value, and most importantly how they feel treated."

Success in law, as in other service industries, is not just about price; it's largely about how clients feel about your experience (or that of the firm), how good the advice is, and how well they've been treated. Legal services are ripe for delivering clients better value work, because so many component parts are done slowly and not very well. This means that lawyers can potentially charge more and yet have clients walk away happier than they were when they were paying less.

MISSION (IM)POSSIBLE

By Dr. Thomas Laubert, Vice President and Group General Counsel, Daimler AG; Dr. Pietro Brambilla, Head of Digital Transformation Integrity and Legal, Daimler AG; and Dr. Jörg Hanke, Skadden, Arps, Slate, Meagher & Flom LLP

Picture this—you have been tasked with supporting the largest business reorganization project in years: the transformation of the company's two main operative business divisions into two legally independent entities by way of a hive-down. This means ensuring that all permits, authorizations, contracts, etc. required to conduct the respective businesses need to be validly transferred to the relevant new entity and/or amended or newly obtained. Any impediments such as change of control provisions, transfer restrictions, and obligations in connection with the consummation of the transaction (e.g., information or consent requirements) need to be identified. The framework conditions are demanding:

Time frame: one year to signing

Team: lean as possible

Documents to review: approximately 1 million active legal documents, such as contracts, certificates, and permits

By conservative procedure: mission impossible

Looking back to spring 2018, Daimler AG had decided to strengthen its customer focus and to increase the group's agility by separating the car and van and the truck and bus businesses into two new subsidiaries—internal project name: "Project Future." Upon consummation of the hive-downs, the new Mercedes-Benz AG would control the global business of Mercedes-Benz Cars & Mercedes-Benz Vans and the new Daimler Truck AG would be responsible for the global truck and bus business. Daimler AG, as the parent company, would be responsible for governance, strategy, and control functions, and would provide groupwide services. To achieve this, each of these entities would be required to be fully operational immediately upon effectiveness of the hive-down. It was an enormous effort in which the legal department also played a decisive role, especially regarding the necessary contract management.

Typically, a traditional approach to a large project like this might involve reviewing somewhere between 20,000 and 200,000 documents, which could take (at the upper end) more than 50 people working for approximately one year at the review and need a very large budget. How would we review one million documents and accomplish it in the expected time frame? Even with a super-heavy lift, a 200,000-contract review would only reveal what was in 20% of the documents. What about the other 80%? Should we decide to just look at samples and hope for the best?

This project demonstrated clearly that traditional ways of resourcing legal work might no longer be sufficient to deliver on business objectives. The traditional way also did not fit with the DNA of Daimler or the philosophy of its legal department. For more than 130 years, Daimler has been moving people and goods all over the world—safely, efficiently, comfortably, and with innovative technologies that have always kept the company a step ahead of the competition. It is this spirit that also drives the work of the legal department.

Innovation and technological change has been an integral part of the strategy of the Daimler legal function for many years. As such, we had the advantage that we already had created a specialized technology team within the legal department. Its aim is to promote the use of innovative technologies to drive automation, reduce complexity, increase speed, and improve efficiency in order to free up legal colleagues for more strategic and transformative work.

"Project Future" was a perfect opportunity to demonstrate the relevance of this transformative approach for the largest projects on hand—and to show that modern technologies such as artificial intelligence (AI) can take the importance of the legal department for the business colleagues to a new level. An important part of the better value equation when leveraging innovative technology is volume, and with AI you have the opportunity to accomplish what was never considered possible so far, such as creating a complete picture of a large document landscape.

Together with Skadden, our law firm commissioned for the project, we scanned the market and selected the contract review and analysis software from Kira Systems to do an AI-enhanced review of all active legal documents. To do this project, the legal team needed to train the software to find the information and legal concepts they sought in German and to verify if the prebuilt modules in English were sufficient for their purposes. Such training requires a certain number of documents providing for positive and negative samples.

IT infrastructure (Daimler preferred not to upload all data to an external cloud but to use an on-premise system) and all users had to be set up so everything was ready to go, all in short order. We did not always have sufficient samples to train the software for every legal concept (e.g., certain types of permits); nevertheless, the software still proved valuable for review purposes.

In the course of the review, it became clear that, worldwide, there were far more legal documents on file than initially expected. In sum, the 1 million that were initially expected to be reviewed was only 25% of the 4 million active documents. However, the project team kept their heads down and got it done.

Despite the massive volume, the review team consisted of less than 10 people. In order to ensure the best quality, the first-level review team checked approximately 80,000 of the most important legal documents manually with the help of the software. By that system, potential issues were highlighted that helped to focus on the relevant provisions and to speed up the review. All other documents were primarily analyzed by the software. The flagged provisions were just reviewed by a so-called first-level team, which also curtained related sections as well as a pre-agreed number of the other documents for quality check purposes. The team was supervised by senior lawyers who made decisions in cases of doubt. In addition, a second-level review team carried out quality checks throughout the whole volume of documents. Over the review period, trust in the capability of the software and the trained modules increased more and more, and consequently, the number of quality checks could be reduced.

In the process, the Daimler and Skadden team found meaningful information in contracts that never would have surfaced if only 20% of the documents had been reviewed (or, more realistically, 5%, given the emergence of a larger document universe than initially expected). Unsurprisingly, low-priority contracts were less likely to contain unexpected information. However, even these types of contracts provided for certain clauses requiring further action in order to consummate the hive-down (e.g., to inform a counterparty or to obtain a counterparty's consent).

In the end, the review team was able to finish the challenging task with a far more thorough picture of all of the contracts held by the company than would have been possible without the help of AI.

Mission accomplished.

This massive undertaking is a great illustration of the sheer volume of work that can be achieved by using AI. In an increasingly complex world facing an exponential growth of information, one needs exponential technology that is able to deal with these new challenges.

What else did we learn from the project?

Data is the foundation for AI. Having good qualitative and quantitative data sets is a prerequisite for running a successful AI project. This is why we have further strengthened our overall data and information strategy with a dedicated Data Officer for our organization. In addition, we are focusing our efforts on the targeted adoption of AI technology to have the greatest transformative impact with limited resources.

The use of AI clearly empowers people in the legal department to do higher-value work. However, innovation and transformation doesn't happen overnight. Driving digital transformation really takes a lot of commitment, because it is not just as simple as buying software and getting people on the team to open the application. It is equally important to foster behavioral change as well.

The close cooperation between internal and external lawyers enhanced by powerful modern technology turned this initial "Mission Impossible" to one of our most successful and efficient projects of the legal department in recent years. It clearly demonstrated the added value that a modern legal department can bring to the entire company.

Access to Justice

There are well over a million lawyers in the United States. More than 150,000 barristers and solicitors in England and Wales. More than 130,000 in Canada. Over 160,000 in Germany. Some 800,000 in Brazil. Yet, in all these places, many people go without a lawyer when they need one. A June 2017 *Washington Post* article noted that approximately 80% of low-income individuals in the United States cannot afford the legal assistance they need, while 40–60% of the legal needs of the middle class go unmet.

The American Bar Association cites access to justice as one of the fundamental principles of the rule of law. In a December 2017 article on ABA.com called "Access to Justice: Mitigating the Justice Gap," Leonard Wills wrote:

> *Access to justice consists of the "ability of individuals to seek and obtain a remedy through formal or informal institutions of justice for grievances." This process usually requires individuals to obtain legal representation—or at a minimum, legal advice. Without legal assistance, individuals can struggle to navigate through the complexity of court procedures. An individual's failure to understand court procedures and the substantive law-related issues of their case can lead to the loss of a home, children, job, income, and liberty.*

Access to justice is a problem, but it is also an opportunity. If lawyers could figure out how to package and price their services in a more appealing way, there is a vast latent market that could use much more legal service than they get today. AI (by driving more efficient work) can be part of the equation that enables lawyers to deliver legal services at a price current nonconsumers will be willing and able to pay.

Some Biglaw lawyers may have tuned out over the last few paragraphs. What does serving the poor and middle class have to do with them (apart from their pro bono program)? Well, even the biggest companies in the world let many potential legal problems go unsolved, finding the price to value equation not compelling. Think! Where do your clients face risk, or could use help, that they're currently ignoring because of cost, complexity, or speed? What if you could do the work for a third, tenth, or twentieth the cost? Or give an answer 10 times faster? Would that make clients pay attention to this area of concern? If so, then start to think about how. AI and other innovations may make executing on this opportunity possible.

ACCESS TO JUSTICE: A PRODUCT-MARKET FIT PROBLEM

By Jack Newton, CEO and Founder, Clio

Law and the legal system are an integral part of how our society operates. Yet we know that not all citizens of our society are able to access it. In fact, data from the World Justice Project shows that 77% of US citizens who encountered a legal issue did not have that issue resolved by a lawyer. Yet over 80% of lawyers tell us the number one thing they need in their law firm is more clients to grow revenue. Any economist would look at this with raised eyebrows—with such massive demand on the consumer side, lawyers should be complaining about having too many clients. Instead, legal is an inefficient market where supply is not effectively meeting demand.

There are many contributing factors to the state of the legal industry, many of which were heavily exposed during the COVID-19 pandemic. You have a court system that is slow moving and inaccessible to many people, at least in part because it has not evolved to stay in step with technological changes (an issue well-known long before the strain of a pandemic). On the consumer side, many people do not reach out to a lawyer because they perceive them to be inaccessible, expensive, and difficult to work with. When you consider that 40% of Americans would struggle to come up with $400 for an unexpected expense, it is not surprising that legal services would be out of reach. And, when you look at the data, legal services are not priced or packaged in a way that is financially feasible to most consumers.

While there is no silver bullet when it comes to solving the access to justice issue, there are many ways we can improve it. The biggest one is adopting technology. A technology-enabled lawyer can help more clients, without sacrificing their livelihood, by increasing accessibility to legal services and automating their administrative work so they can spend more time practicing law.

Using technology for the benefit of clients and legal professionals should be table stakes for any law firm. Yet there are still too many firms that view their address and the size of their boardrooms as the most important part of their client's experience. But when you consider how much more cost-effective it would be to deliver legal guidance through a Zoom chat as opposed to a meeting in a downtown office space with a marble-lined lobby, the numbers just make sense.

Technology can enable law firms and lawyers to deliver their legal services by increasingly working from their home offices, from a co-working space, or from remote locations. That dramatically changes the underlying cost structure of running a law firm. Combined with the productivity enhancements that technology can bring to a law firm through practice management software, document automation software, AI, and contract review software (to name just a few), lawyers can provide greater value to customers who would rather pay for results than ambiance. And lawyers can enjoy a more flexible working life that is not bogged down with pen-and-paper time tracking or wet signatures.

The access-to-justice problem does not exist in a vacuum—it is for all of us to solve. The good news is that it represents a huge opportunity for lawyers willing to adapt to a new way of thinking. The ability to deliver legal services in a new way, coupled with the productivity enhancements that technology can provide, allows lawyers to offer a completely different cost structure to the market and increase the accessibility of their legal services to everyone. This is simply a product-market fit problem, and one that the legal industry has the power to change.

Praise for AI FOR LAWYERS

"This book brings AI down to earth with an admirable lightness of touch."
—**Richard Susskind,** author of *Tomorrow's Lawyers*

"I learned a great deal from this engaging book. It was a quick, page-turning read. I highly recommend it."
—**Brad Karp,** Chairman, Paul Weiss

"Noah and Alexander have written a classic. This is everything you ever wanted to know about technology and the law but never dared ask. Through their lucid explanations of how the technology works, how it can be applied, and how it's fast-improving, your eyes will be opened to a future of legal practice that is more enjoyable, profitable, and just. As the pace of technological change accelerates exponentially, this book provides an accessible way to jump in and become part of the revolution."
—**David Morley,** Global Managing Partner, Global Senior Partner (2003–2016), Allen & Overy; Consultant & Portfolio Chair (2017–)

"This important new book provides a clear and concise guide to how artificial intelligence and related technologies are reshaping the market for legal services. Utilizing their unparalleled experience as the founders of the most important legal AI company and their extensive contacts in the industry, Waisberg and Hudek offer a thoughtful examination of the costs—and even more importantly, the benefits—of the application of technology to law. It should be required reading for anyone who cares about the future of our legal system."

—**David B. Wilkins,** Lester Kissel Professor of Law, Harvard Law School

AI For Lawyers is the definitive guide to artificial intelligence in the law. This isn't surprising given that Noah and Alexander created an AI-based business that's contributed to redefining the legal industry. Filled with insights and practical tips, every lawyer and legal professional should read this book."
—**James Goodnow,** CEO (Managing Partner), Fennemore Craig, P.C.

"Artificial intelligence is dramatically changing the practice of law and the global legal market. Lawyers have to learn not to live with it but to love it. What Waisberg and Hudek tell us about AI in law practice will at first be as disorienting, but then as reorienting, as was Richard Susskind's *The End of Lawyers?* a generation ago."
—**Stephen Gillers,** Elihu Root Professor of Law, NYU School of Law

BUSINESS & ECONOMICS General
$34.00 USA/$40.99 CAN

Cover Design: Wiley
Cover Image: © Edo Cuallo
Author Photos: © Joanne Son

Subscribe to our free Business eNewsletter
at wiley.com/enewsletters

Visit wiley.com/business

WILEY

Also available as an e-book

ISBN 978 1 119 72384 4
53400

9 781119 723844