PROBLEMS WITH PROBABILITY

Anthony J. Casey & Anthony Niblett*

Abstract

Some countries have explored the idea of using artificial intelligence (AI) systems to help triage the backlog of cases and facilitate the resolution of civil disputes. In theory, AI can accomplish this by establishing the facts of cases and predicting the outcomes of disputes. But the use of AI in the courtroom gives rise to new problems.

AI technologies help solve prediction problems. These solutions are typically expressed as probabilities. How should judges incorporate these predictions in their decision making? There is no obviously correct approach for converting probabilistic predictions of legal outcomes into binary legal decisions. Any approach that does so has benefits and drawbacks.

Importantly, a balance of probabilities approach – where liability is established if the AI predicts a likelihood of liability greater than 50 per cent and not otherwise – is not suitable when converting a predicted outcome into an actual outcome. Adopting this approach would significantly alter the outcomes of legal cases and have a dramatic and disruptive effect upon the law. The most notable disruption would be observed in settlement behaviour and outcomes.

* Anthony J. Casey is the Donald M. Ephraim Professor of Law and Economics and the Faculty Director of the Center on Law and Finance at the University of Chicago Law School, Chicago, United States. Anthony Niblett is an Associate Professor at the University of Toronto Faculty of Law, Toronto, Canada. Niblett is the Canada Research Chair in Law, Economics and Innovation, and a faculty affiliate at the Vector Institute of Artificial Intelligence. Casey acknowledges financial support from the Richard M. Weil Faculty Research Fund and the Paul H. Leffmann Fund. Niblett acknowledges financial support for this project from SSHRC and the Canada Research Chair program. In the interests of full disclosure, Niblett is the co-founder of Blue J, a start-up bringing machine learning to tax law and employment law.

We wish to thank Audrey-Anne Delage, Gordon Milne, and Arielle Ambra-Juarez for excellent research assistance. We thank participants in the *University of Toronto Law Journal* Symposium on Modern Challenges to the Judiciary in April 2022 for their helpful comments and suggestions.

I *Introduction*

Judges have always encountered problems with probability. In this article, we explore new problems with probability that may arise from the use of artificial intelligence (AI) technologies in judicial decision making. AI technologies help with prediction problems.[1] The outputs of the AI solutions to prediction problems are typically expressed as probabilities. A key challenge for the judiciary will be how to take these probabilistic predictions and convert them into clear, binary legal outcomes.

One potential advantage of using AI in the judicial system is reducing the backlog of cases in the civil court system. Civil disputes can drag on: nearly 15 per cent of civil court cases in Ontario take longer than two years to resolve.[2] Even for small claims of less than twenty-five thousand dollars in Ontario, disputes take an average of over nine months to resolve.[3] The situation is worse in other countries. In France and Italy, for example, the average civil law dispute takes nearly two years to resolve.[4] This backlog and delay can present a substantial barrier to access to justice.

Various countries have considered implementing AI systems to address this issue. For example, Brazil has turned to AI to help triage cases.[5] The use of AI to assist with the resolution of civil disputes has also been discussed and proposed in the Netherlands, the United Kingdom, Estonia, and China.[6] The topic has also received media attention in

---

[1] See Ajay Agrawal, Joshua Gans & Avi Goldfarb, *Power and Prediction: The Disruptive Economics of Artificial Intelligence* (Boston: Harvard Business Review Press, 2022) [Agrawal, Gans & Goldfarb, *Power and Prediction*]; Ajay Agrawal, Joshua Gans & Avi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence* (Boston: Harvard Business Review Press, 2018) [Agrawal, Gans & Goldfarb, *Prediction Machines*].

[2] Of the 270,266 civil disputes that were resolved in 2019–20 in Ontario courts, 40,061 had taken longer than two years to resolve. This percentage (14.82 per cent) is not atypical across the past five years. See Statistics Canada, 'Active Civil Court Cases, by Elapsed Time from Case Initiation to First Disposition, Canada and Selected Provinces and Territories' (10 March 2022), online: *Statistics Canada* <www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3510011601> [perma.cc/9A2W-GWT2].

[3] See Anthony Niblett & Albert H Yoon, 'Unintended Consequences: The Regressive Effects of Increased Access to Courts' (2017) 14:1 J Empirical Leg Stud 5 at 22.

[4] See European Commission for the Efficiency of Justice, 'CEPEJ Indicators on Efficiency' (last updated 27 September 2022), online: *Tableau Public* <public.tableau.com/app/profile/cepej/viz/EfficiencyEN/Efficiency> [perma.cc/HD6F-3U4G] (in 2020, the average disposition of a civil law dispute in France was 637 days, while, in Italy, it was 674 days).

[5] See Fausto Martin De Sanctis, 'Artificial Intelligence and Innovation in Brazilian Justice' (2021) 59:1 International Annals of Criminology 1 at 2–3.

[6] Monisha Pillai, 'China Now Has AI-Powered Judges' (16 August 2019), online: *RADII* <radiichina.com/china-now-has-ai-powered-robot-judges/> [perma.cc/FH2W-FBJZ]; Chris Young, 'China Has Unveiled an AI Judge That Will "Help" with Court Proceedings' (19 August 2019), online: *Interesting Engineering* <interestingengineering.com/innovation/china-has-unveiled-an-ai-judge-that-will-help-with-court-proceedings> [perma.cc/Q37R-ZRY7]. See also Jingting Deng, 'Should the Common Law System Welcome Artificial Intelligence: A Case Study of China's Same-Type Case Reference System' (2019) 3 Georgetown Law and Technology Review 223; Tom

North America, with some referring to the use of AI as 'robot judges.'[7] In theory, AI could help reduce this backlog. Consider two types of prediction problems where AI can streamline decision making:

> **1. Predicting the facts of cases:** AI can be used to determine facts. In the same way that radiologists can use AI to predict whether a bone is broken, AI can be used to determine whether a defendant caused damage with greater precision. For example, an AI algorithm may determine that there is an 80 per cent likelihood that the defendant caused the damage rather than just determining it is more likely than not that the defendant caused the damage.
>
> **2. Predicting the outcome of the disputes:** AI can also be used to help determine how a court would resolve a dispute based on the facts of the case and comparisons to previously decided cases on the same issue. For example, an AI algorithm might determine that there is an 80 per cent likelihood that the defendant would be held liable if the case were to go to a judgment. These predictions could potentially be used by courts to help decide cases more efficiently.

But the question of how the AI is implemented is fundamental. Judges need to incorporate these predictions into their decision making. We focus on this one specific question in this article: how judges should convert algorithmic predictions into legal outcomes? Legal decisions are typically binary in nature (for example, liable or not liable), but algorithmic predictions are typically probabilistic (for example, 80 per cent likelihood of liability).

It may seem that converting probabilistic prediction into binary outcomes is quite intuitive. Liability could be established if the AI predicts that the likelihood of liability is greater than 50 per cent. In many contexts, courts use a balance of probabilities or preponderance of evidence standard. But this view is too simplistic. There are important differences between the different types of prediction. Philosophical questions regarding the meaning of the probabilities must be addressed. There is no one-size-fits-all method for converting probabilities into legal decisions. Using the balance of probabilities approach for converting probabilistic predictions made by AI algorithms may be appropriate where the algorithm predicts facts. But we argue that it is not suitable for converting predictions of legal outcomes. Adopting the balance of probabilities approach

---

Fish, 'AI Shock: China Unveils "Cyber Court" Complete with AI Judges and Verdicts via Chat App,' *The Daily Express* (6 December 2019), online: <www.express.co.uk/news/science/1214019/ai-china-cyber-court-artificial-intelligence-judges-verdicts-chat-app> [perma.cc/MG39-E4PX].

[7] See e.g. Eric Niiler, 'Can AI Be a Fair Judge in Court? Estonia Thinks So,' *Wired* (25 March 2019), online: <www.wired.com/story/can-ai-be-fair-judge-court-estonia-thinks-so/> [perma.cc/H2L2-PAWA]. See also Dana Remus & Frank Levy, 'Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law' (2017) 30:3 Geo J Leg Ethics 501 (referring to algorithms performing legal work as robots).

to legal predictions would significantly alter the outcomes of cases and have a dramatic impact on the law.

Converting legal predictions into legal outcomes based solely on the balance of probabilities would also radically alter the settlement process.[8] Under economic theories of settlement, the plaintiff and defendant often reach an agreement based on the expected outcome of a case if the case were to go to court.[9] This expectation reflects the likelihood of plaintiff success. But under a balance of probabilities approach to outcomes, the likelihood of plaintiff success would be either 0 per cent or 100 per cent.

Suppose both parties agree that the plaintiff has an 80 per cent likelihood of winning one million dollars in damages. The parties would be wise to settle this dispute. They would likely settle in the neighbourhood of eight hundred thousand dollars. This saves time and expense, and reduces the risks of going to court. But if courts implemented a system where AI predictions are converted into legal outcomes and they use a balance of probabilities approach, the expected value of the suit changes. Under the same assumptions, the parties would likely to settle closer to one million dollars. There are some who benefit because of the conversion using the balance of probabilities. There are others who suffer. For example, plaintiffs with a predicted likelihood of, say, 20 per cent would lose out. The expected value of their dispute is now zero dollars. But, more importantly, the fundamental structure of the law in action has changed.

Our article proceeds as follows. In Part II, we discuss how AI algorithms can be used to predict facts and predict legal outcomes. In Part III, we discuss balance of probability in the context of AI prediction of facts. In Part IV, we explore why a balance of probabilities approach for AI prediction of outcomes will change the content of the law. Part V, we explore alternative ways to convert probabilistic predictions into binary legal outcomes. We also identify the significant trade-offs that come with each alternative. Part VI concludes.

[8] The argument here follows the analysis we initially presented in Anthony J Casey & Anthony Niblett, 'Will Robot Judges Change Litigation and Settlement Outcomes? A First Look at the Algorithmic Replication of Prior Cases' (2020) MIT Computational Law Report, Release 1.2 (August 14, 2020), online: <https://law.mit.edu/pub/willrobotjudgeschangelitigationandsettlementoutcomes/release/1>. There, we explored various alternatives and the effect on litigation outcomes and settlement outcomes in greater depth. Here, we focus on the ill-effects of the balance of probabilities approach when converting probabilities over outcomes into actual outcomes.

[9] See Lucian Bebchuk, 'Litigation and Settlement under Imperfect Information' (1984) 15:3 Rand Journal of Economics 404; David Rosenberg & Steven Shavell, 'A Model in Which Suits Are Brought for Their Nuisance Value' (1985) 5 Intl Rev L & Econ 3; Steven Shavell, 'The Level of Litigation: Private versus Social Optimality' (1999) 19:1 Intl Rev L & Econ 99; Steven N Durlauf & Lawrence E Blume, eds, *The New Palgrave Dictionary of Economics*, 2d ed (New York: Macmillan, 2008), sub verbo 'Economics of Litigation'; Andrew F Daughety & Jennifer F Reinganum, 'Settlement' in Chris William Sanchirico, ed, *Procedural Law and Economics* (Northampton, UK: Edward Elgar, 2012); Dru Stevenson & Nicholas J Wagoner, 'Bargaining in the Shadow of Big Data' (2018) 67:4 Fla L Rev 1337.

II *The use of AI to predict facts and outcomes*

Supervised machine-learning algorithms (one type of AI) makes predictions based on large datasets. With enough good data, these machine-learning algorithms generate predictions that are more precise and accurate than the predictions generated by traditional statistical techniques.[10] The success stories of AI are those that provide more accurate and faster predictions than humans are capable of. For example, when your email automatically redirects a message to the spam folder, a machine-learning algorithm has made a prediction that the message is spam. When facial recognition technology indicates that you are in a photograph, the machine-learning technology is predicting that you are in the photograph because the image's data is similar to the data in images of you. There is significant evidence illustrating when and how AI outperforms humans in making predictions.[11]

Here, we examine how judges can use AI to assist with their prediction problems.[12] For the purposes of this article, we differentiate between AI technology that 'predicts' facts and AI technology that predicts outcomes. An example may prove helpful. Suppose a plaintiff, Pattie, is suing a defendant company, Delivery Inc., for negligence. Pattie was injured when she swerved on her bike to avoid a car speeding the wrong way on a street. Pattie is seeking one million dollars in damages from Delivery. Pattie alleges that Xavier, a Delivery driver, was the negligent driver of the speeding car.

PREDICTION TECHNOLOGY CAN HELP ASCERTAIN THE FACTS

Suppose Pattie has no memory of the incident. A camera at the scene of the incident captured grainy images of the speeding car's licence plate. Visual recognition software can be used to predict the licence plate number. But this prediction of the licence plate number is not made with certainty. Let us say that the algorithm makes a prediction that it is 80 per cent likely that it was Xavier's car. The algorithm predicts that there is a 20 per cent chance that it was Yvette's car. Yvette has no relationship with Delivery.

---

[10] See Jerry Kaplan, *Artificial Intelligence: What Everyone Needs to Know* (Oxford: Oxford University Press, 2016); Sendhil Mullainathan & Jann Spiess, 'Machine Learning: An Applied Econometric Approach' (2017) 31:2 Journal of Economic Perspectives 87.

[11] See e.g. William M Grove & Paul E Meehl, 'Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy' (1996) 2:2 Psychology, Public Policy and Law 293; Claire Cain, 'Can an Algorithm Hire Better Than a Human?' *New York Times* (25 June 2015), online: <www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html> [perma.cc/4KU7-A98Y].

[12] See Jon Kleinberg et al, 'Human Decisions and Machine Predictions' (2017) 133:1 Quarterly Journal of Economics 237.

Prediction technology can be used to determine the facts of cases. The scope of 'prediction,' thus, goes beyond forecasting uncertain outcomes. This type of prediction problem is similar to those faced by radiologists. Radiologists use AI to predict whether a bone is broken or whether a tumour is malignant. This prediction is probabilistic. But it does not provide a forecast of an uncertain future event (for example, whether a tumour might become malignant.) Rather, this is a probabilistic interpretation of some ground truth. There is a correct answer.[13]

PREDICTION TECHNOLOGY CAN ALSO BE USED TO HELP ASCERTAIN ANSWERS TO LEGAL QUESTIONS

Now let us take the facts as given. Let us assert that Xavier was the driver of the car in question; Delivery's liability turns on whether Xavier is an employee or an independent contractor. If Xavier is Delivery's employee, then Delivery will be held to be vicariously liable for Pattie's injuries. The question of whether a worker is an employee or an independent contractor, under Canadian employment law, is not governed by a bright line rule. Instead, Delivery's liability will depend on the 'total relationship' between Delivery and Xavier.[14] The court will consider a variety of factors, such as the intention of the two parties, the control exerted by Delivery over Xavier, whether Xavier was driving his own car, Xavier's opportunity for profit and risk of loss, and whether Xavier was integrated in Delivery's business.[15]

AI technology analyses Pattie's case and compares it to a rich dataset of previously decided cases on whether a worker is an employee or an independent contractor. The AI then predicts the outcome based on judgments in similar cases.[16] But the predicted classification of the legal outcome is not made with certainty. The classification is probabilistic. Let us say that, given the facts of this case, the algorithm predicts that it is 80 per cent likely that Xavier would be found to be an employee of Delivery. To put this number in context, this means that the AI algorithm is predicting that if one hundred cases like Pattie's case were to come before the court, the court would hold that Xavier is an employee in eighty cases (that is, Delivery is vicariously liable) and would hold that Xavier is an independent contractor in twenty cases (not vicariously liable).

---

[13] See Agrawal, Gans & Goldfarb, *Power and Prediction*, supra note 1; Agrawal, Gans & Goldfarb, *Prediction Machines*, supra note 1.

[14] *671122 Ontario Ltd v Sagaz Industries Canada Inc*, 2001 SCC 59 at para 46.

[15] *Ibid* at paras 46–7.

[16] See e.g. Benjamin Alarie, Anthony Niblett & Albert H Yoon, 'Using Machine Learning to Predict Outcomes in Tax Law' (2016) 58 Can Bus LJ 231; Benjamin Alarie, Anthony Niblett & Albert H Yoon, 'How Artificial Intelligence Will Affect the Practice of Law' (2018) 68 UTLJ 106.

By using the algorithm to help make decisions, judges would essentially be replicating what judges have done in earlier cases. The algorithm predicts what the law would say about this particular case. The algorithm may be accurate, in the sense that it can predict how a judge would decide a case, but there is no 'correct' answer. Probabilistic predictions of outcomes can show the level of legal uncertainty in a case. This legal uncertainty could reflect differences in judicial attitudes and preferences. Some judges could be more pro plaintiff than others. They may be more willing to hold the wealthier defendants liable. Further, the uncertainty could reflect other variables that the machine-learning model does not consider.

* * *

The algorithm does not render a decision. A human judge is required to turn these probabilistic predictions into binary legal decisions. And there are good reasons to think that a one-size-fits-all approach to turning probability into outcomes is not optimal.

These two settings for predictive algorithms may appear similar. In both cases, algorithms make predictions that weigh in favour of finding Delivery liable with 80 per cent likelihood. But these two types of algorithms animate different problems in law. Philosophically, they are different types of 'prediction.' Algorithmic fact finding performs a very different function from algorithm legal prediction. We argue that judges should not convert the 80 per cent likelihood in these two different algorithms in the same way. Importantly, there are good reasons why a balance of probabilities approach to legal prediction should not be adopted.

Before we continue though, we address two concerns readers may have. First, these two types of predictive algorithms are clearly not the only two instances where AI can be used to assist judges. Indeed, much academic ink has been spilled analysing the use of AI to assess the risk of criminal defendants in bail assessments and sentencing.[17] The prediction problem that risk assessment algorithms seek to solve is different from the two prediction problems above. With risk assessment, the AI seeks to predict consequences of a legal decision. That is, if the judge grants bail to a criminal defendant, can the AI predict the consequences of this decision? Is the defendant a high flight risk? The prediction of consequences is not the focus of this article, nor is the use of AI in criminal law. We restrict our analysis to AI that predicts facts and outcomes in civil disputes.

Second, AI algorithms present a range of issues that have been explored by legal scholars. One concern is whether legal decisions should ever be influenced by algorithmic predictions because AI algorithms are distorted by bias.[18] Some argue that the predictions

[17] See e.g. Tracy L Fass et al, 'The LSI-R and the COMPAS: Validation Data on Two Risk-Needs Tools' (2008) 35:9 Criminal Justice and Behaviour 1095; Julia Dressel & Hany Farid, 'The Accuracy, Fairness, and Limits of Predicting Recidivism' (2018) 4:1 Science Advances, online: <https://www.science.org/doi/10.1126/sciadv.aao5580>.

[18] See e.g. Jason Millar & Ian Kerr, 'Delegation, Relinquishment, and Responsibility: The Prospect of Expert Robots' in Ryan Calo, A Michael Froomkin & Ian Kerr, eds, *Robot Law* (Cheltenham, UK: Edward Elgar, 2016) 102; Timothy D

these algorithms generate cannot be objective.[19] Others have argued that AI is not sufficiently transparent because some algorithmic predictions are not capable of being explained or justified.[20] The literature has also queried whether algorithms could ever be held accountable.[21] Further, some argue that algorithms trained on big data will misrepresent the law and entrench discrimination[22] and can erode privacy rights.[23] These are, of course, crucial issues. But they are not the focus of our inquiry here. We address a separate problem. Even if these algorithmic predictions are free from harmful bias, even if the algorithms are reliable, and even if the decisions can be explained and justified, important questions remain about how probabilistic predictions should be converted into legal decisions.

III *Predicting facts and the balance of probabilities*

AI can clarify the facts of civil disputes. Facial recognition software can help predict whether the defendant caused damage. Medical imaging technology can help determine what damage has been suffered. These predictions will become more precise and more accurate over time. The output of the predictive algorithm provides only a probabilistic prediction. But what should a judge do with this information? This is not a new problem

---

Robinson, 'A Normative Evaluation of Algorithmic Law' (2017) 23 Auckland UL Rev 293; Cary Coglianese & David Lehr, 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era' (2017) 105 Geo LJ 1147; Emily Berman, 'A Government of Laws and Not of Machines' (2018) 98 BUL Rev 1277; Frank Pasquale, 'A Rule of Persons, Not Machines: The Limits of Legal Automation' (2019) 87:1 Geo Wash L Rev 1; Elizabeth E Joh, 'Policing by Numbers: Big Data and the Fourth Amendment' (2014) 89 Wash L Rev 35; Julia Angwin et al, 'Machine Bias: There's Software Used across the Country to Predict Future Criminals and It's Biased against Blacks,' *ProPublica* (23 May 2016), online: <www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [perma.cc/5WGC-RXZY]; Solon Barocas & Andrew D Selbst, 'Big Data's Disparate Impact' (2016) 104 Cal L Rev 671; Kelly Hannah-Moffat, 'Algorithmic Risk Governance: Big Data Analytics, Race and Information Activism in Criminal Justice Debates' (2018) 23:4 Theoretical Criminology 453; Sandra G Mayson, 'Bias In, Bias Out' (2019) 128 Yale LJ 2218; Renata M O'Donnell, 'Challenging Racist Predictive Algorithms under the Equal Protection Clause (2019) 94 NYU L Rev 544; Bo Cowgill & Catherine Tucker, 'Economics, Fairness and Algorithmic Bias' (2019) [unpublished, archived at National Bureau of Economic Research].

[19] Caryn Devins et al, 'The Law and Big Data' (2017) 27:2 Cornell JL & Pub Pol'y 357.

[20] Maayan Perel & Niva Elkin-Koren, 'Black Box Tinkering: Beyond Disclosure in Algorithmic Enforcement' (2017) 69:1 Fla L Rev 181; Cary Coglianese & David Lehr, 'Transparency and Algorithmic Governance' (2019) 71:1 Admin L Rev 1. Cf Andrew D Selbst & Solon Barocas, 'The Intuitive Appeal of Explainable Machines' (2018) 87:3 Fordham L Rev 1085; Aziz Huq, 'A Right to a Human Decision' (2019) 106:3 Va L Rev 611; Lilian Edwards & Michael Veale, 'Slave to the Algorithm? Why a "Right to an Explanation" is Probably Not the Remedy You Are Looking For' (2017) 16:1 Duke L & Tech Rev 18.

[21] See e.g. Joshua A Kroll et al, 'Accountable Algorithms' (2017) 165 U Pa L Rev 633.

[22] Dan L Burk, 'Algorithmic Fair Use' (2019) 86 U Chicago L Rev 283.

[23] Arthur J Cockfield, 'Protecting the Social Value of Privacy in the Context of State Investigations Using New Technology' (2007) 40:1 UBC L Rev 41; Karen Eltis, 'The Judicial System in the Digital Age: Revisiting the Relationship between Privacy and Accessibility in the Cyber Context' (2011) 56:2 McGill LJ 291; Teresa Scassa & Amy Conroy, 'Strategies for Protecting Privacy in Open Data and Proactive Disclosure' (2016) 14:2 CJLT 215.

for the law or legal scholarship. This type of prediction problem is familiar to scholars of evidence and those who study the use of statistical reasoning in the law.[24] Scholars have debated hypothetical problems such as the 'blue bus problem' (where a plaintiff can establish that she was struck by a blue bus, and 80 per cent of the buses in town are owned and operated by the Blue Bus Company).[25] Some scholars have argued that law has a 'general aversion' to statistical and probabilistic evidence. There are good reasons, however, to think that these problems with probability – and the incorporation of statistical evidence into the law – are overstated.[26]

We suggest that there will be few 'new' problems of probability when it comes to AI and the prediction of facts in civil disputes. While the probabilistic predictions may be more precise and more accurate, there is little difference, philosophically, between the AI prediction and the non-AI prediction of facts. Put simply, there are good reasons why the typical preponderance of evidence rule should still be used. Under this standard, a particular fact is established if there is a greater than 50 per cent chance that it is true. In our example of Pattie and Delivery, the AI algorithm predicts that there is an 80 per cent likelihood that it was Xavier who was driving negligently. Under the preponderance of evidence standard, the weight of the evidence would counsel in favour of finding that Xavier drove the car. The 80 per cent is essentially treated as 100 per cent.

This type of prediction problem is backward looking. It is concerned with events that have already happened. The true fact is not probabilistic. The event either occurred with probability 0 per cent (that is, it did not happen) or with probability 100 per cent (that is, it did happen.) In our example above, either Xavier was driving the car, or he was not. Put simply, there is a 'correct' answer.

The preponderance of evidence rule reduces error costs. For example, imagine an urn with one hundred balls, eighty of which are black and twenty are white. Suppose you close your eyes and pull a ball at random, guessing its colour without looking. You repeat the exercise for each of the one hundred balls. A balance of probabilities approach suggests that you should guess 'black' every time, though you will get the correct answer 80 times out of one hundred. This is an error rate of 20 per cent. But this error rate is significantly less than if you guess according to likelihood (that is, guess black eighty times and guess white twenty times.) If you took this approach, you would be expected to get the correct answer only seventy-two times. Of the eighty times you select black, you would be expected to get sixty-four correct (80 per cent). Of the twenty times you

[24] See e.g. Frederick Schauer, *Probabilities, Profiles and Stereotypes* (Cambridge, MA: Harvard University Press, 2003) at 79–108.

[25] Ibid. See also Laurence H Tribe, 'Trial by Mathematics: Precision and Ritual in the Legal Process' (1971) 84:6 Harv L Rev 1329 at 1346–7; James Brook, 'The Use of Statistical Evidence of Identification in Civil Litigation: Well-Worn Hypotheticals, Real Cases, and Controversy' (1984–5) 29 St Louis LJ 293.

[26] See Ronald J Allen & Christopher K Smiciklas, 'The Law's Aversion to Naked Statistics and Other Mistakes' (2022) 28 Leg Theory 179.

select white, you would be expected to get four correct (20 per cent). The error rate would rise to 28 per cent.

IV *Using balance of probabilities to convert predicted outcomes into actual decisions will dramatically distort the content of the law*

The arguments for using a balance of probabilities approach are less persuasive when it comes to converting AI predictions over outcomes into legal decisions. There are important differences between 'predicting' facts (where there is an underlying truth) and predicting legal outcomes (where the prediction is essentially a forecast of the result in court). To illustrate the problem, suppose the facts of the case are agreed. In Pattie's case, the question is whether Xavier will be held to be an employee of Delivery (triggering vicarious liability) or an independent contractor. A machine-learning algorithm compares our facts to all previously litigated cases and predicts an 80 per cent likelihood that a court would find Xavier to be Delivery's employee. Correspondingly, it is 80 per cent likely that Delivery will be found liable for Pattie's injuries if the case were to go to court.

There is no obviously 'correct' way to convert this 80 per cent likelihood into a legal decision. Each potential approach to the problem comes with trade-offs. Under a balance of probabilities approach, judges would simply treat probabilities over outcomes in the same way as they treat probabilities over facts. If the predicted probability of liability is greater than 50 per cent, then the defendant is liable. If the probability is 50 per cent or less, then the defendant is not liable. Though a balance of probabilities approach may seem intuitive, applying it to decide outcomes would drastically affect the content of the law.

ADOPTING A BALANCE OF PROBABILITIES APPROACH IN CONVERTING PREDICTION INTO OUTCOMES WILL CHANGE THE RESULTS OF CASES THAT WOULD BE LITIGATED

Assume the predictions accurately reflect the law. The prediction in Pattie's case suggests that if one hundred workers in Xavier and Delivery's situation came before a court, the court would hold Xavier to be an employee in eighty of those cases. In the remaining twenty cases, Xavier would be an independent contractor. But letting the balance of probabilities determine outcomes would change the results of cases that would be litigated. Instead of eighty cases out of one hundred holding workers like Xavier to be an employee, the worker would be an employee in one hundred cases out of one hundred. Essentially, the probability of being an employee is no longer 80 per cent. It is 100 per

cent. If these cases were to go to court, the algorithm would change the outcome of 20 per cent of these cases.

MORE IMPORTANTLY, SETTLEMENT BEHAVIOUR WILL BE GREATLY IMPACTED

The vast majority of civil disputes settle before reaching court.[27] Under the current approach, settlement should be correlated with the likely outcome of court decisions. Settlement outcomes are a function of the probability of plaintiff success. If there is an 80 per cent likelihood of Pattie winning one million dollars in this tort dispute and both Pattie and Delivery know this, the parties should settle in the neighbourhood of eight hundred thousand dollars. But if courts take a balance of probabilities approach when converting predictions into outcomes, then settlement outcomes look very different. The shadow cast by the law in which bargaining takes place would change dramatically. Essentially, the grey area would be gone. The probability of Delivery being found liable in our example is now 100 per cent instead of 80 per cent. Any settlement outcome would reflect this change in the law, pushing it closer to the full one million dollars.

The impact of using predictive algorithms in this way would be starkest in cases where the probability of plaintiff success is around 50 per cent. In these 'close' cases, rational parties who largely agree on the likelihood of plaintiff success should settle for amounts around half the total value of the suit. In our example, if the disputed amount is one million dollars but now Pattie's likelihood of success is around 50 per cent, the parties should settle for around five hundred thousand dollars. This avoids the risk and cost of going to court. But if courts adopted the balance of probabilities approach to outcome determination, settlements would diverge to either one million dollars (if the probability is greater than 50 per cent) or zero dollars (if the probability is less than 50 per cent). The big winners would be plaintiffs who had probabilities just over 50 per cent. They would recover the full amount. The big losers would be plaintiffs who had probabilities just under 50 per cent. They would recover nothing. Cases where the plaintiff previously had a 49 per cent likelihood of success would be treated very differently to cases where the plaintiff had a 51 per cent likelihood. Currently, these cases will likely be settled for similar amounts. But if algorithmic assessments are used as the basis of legal decisions and a balance of probabilities employed, these two cases would have drastically different outcomes. Minute factual differences or even a narrow margin of error would snowball into major differences in outcomes under the law.

---

[27] See 'Civil and Criminal Cases' (last updated 1 September 2019), online: *Department of Justice Canada* <www.justice.gc.ca/eng/csj-sjc/just/08.html#:~:text=Lawyers per cent20often per cent20discuss per cent20the per cent20lawsuit,make per cent20it per cent20to per cent20the per cent20courts> [perma.cc/6AE3-PWFB]; Theodore Eisenberg & Charlotte Lanvers, 'What Is the Settlement Rate and Why Should We Care?' (2009) 6:2 J Empirical Leg Stud 111.

These changes would broadly affect the incentives of individuals and businesses to take precautions to avoid causing losses.[28] Currently, two different defendants facing predicted likelihood of liability of 49 per cent and 51 per cent would behave in similar ways. The two defendants likely take similar levels of precautions. This is because the likely legal outcomes from their respective actions are very similar. But under a balance of probabilities approach, the disruption in the law would have real effects.

Further, the nature of legal disputes would likely change. Litigants would likely dispute which algorithm should be used by the court to determine the extent of liability. Different classification models in statistics and machine learning yield different predicted probabilities even given exactly the same inputs. If one model predicted a 49 per cent chance of liability and another predicted a 51 per cent chance of liability, the choice of model – a question of procedure – would yield enormous differences in outcome. This would, no doubt, become a major source of dispute.

V *Alternative approaches to converting AI predictions on outcomes*

What alternatives are available to the (flawed) balance of probabilities approach for converting AI predictions into legal outcomes? In this part, we discuss three possible alternatives. Each alternative has its benefits and drawbacks.

ALTERNATIVE 1: 'RECOUPLE' LIABILITY AND DAMAGES

First, consider the possibility of recoupling liability and damages. The algorithm provides a prediction as to the likelihood of liability only. A balance of probabilities approach converts that likelihood into liability or no liability. But if damages and liability are recoupled, the outcome is a dollar figure that reflects the expected outcome. With our example of 80 per cent likelihood of plaintiff success on damages of one million dollars, this approach yields an award of eight hundred thousand dollars. Settlement outcomes are not as drastically affected as they would be under the balance of probabilities approach. If there is an 80 per cent likelihood of the plaintiff winning one million dollars, the parties are likely to settle in the neighbourhood of eight hundred thousand dollars. But this approach would necessitate a fundamental shift in how we think about liability

[28] See e.g. John Prather Brown, 'Toward an Economic Theory of Liability' (1973) 2:2 J Leg Stud 323; William M Landes & Richard A Posner, *The Economic Structure of Tort Law* (Cambridge, MA: Harvard University Press, 1987); Steven Shavell, *Economic Analysis of Accident Law* (Cambridge, MA: Harvard University Press, 1987); Robert D Cooter & Daniel L Rubenfeld 'Economic Analysis of Legal Disputes and Their Resolution' (1989) 27 Journal of Economic Literature 1067; Don Dewees, David Duff & Michael Trebilcock, *Exploring the Domain of Accident Law* (New York: Oxford University Press, 1996).

under the law. Defendants who are only 2 per cent or 5 per cent likely to be found liable would always be found liable to pay damages.

ALTERNATIVE 2: ASSIGN FULL LIABILITY ACCORDING TO LIKELIHOOD

A second alternative is to assign full liability, but only assign liability in line with the likelihood of liability. That is, if the algorithm predicts the likelihood of plaintiff success to be 80 per cent, the court would find eighty out of one hundred defendants liable and the other twenty would be held not liable. As with the first alternative approach, settlement tracks expected value: if there is an 80 per cent likelihood of plaintiff winning one million dollars, the parties are likely to settle in the range of eight hundred thousand dollars. Taking this approach would fundamentally change the way we think about legal uncertainty. In practice, this approach would convert unexplained variation in legal decisions into purely arbitrary variation. It makes winners out of those 20 per cent of defendants who are randomly selected to be held not liable; it makes losers out of the other 80 per cent. This approach might be unpalatable to many. Indeed, it appears to reflect 'randomized' justice.[29] Using this approach would generate inconsistent decisions, as statistically identical plaintiffs would be treated very differently under the law.

ALTERNATIVE 3: TRIAGE 'EASY' CASES ONLY

A third alternative is to limit the use the predictive algorithms to 'easy' cases only. For example, the algorithm could be used to determine the outcomes where the likelihood of liability is less than 5 per cent (that is, the outcome is no liability) or where the likelihood of liability is greater than 95 per cent (that is, the outcome is liability.) Other cases between these extremes can be deferred to human decision makers. This may seem like a sensible compromise. As we noted above, the impact of implementing a balance of probabilities approach is at its most dramatic when the likelihood of plaintiff success is in the neighbourhood of 50 per cent. But if courts restricted their conversion of probabilistic predictions into binary legal outcomes when the likelihoods approach 0 per cent or 100 per cent, they would leave many outcomes unaffected. The impact on settlement would not be overly disruptive either.

But, once again, there are important trade-offs to consider. Notably, there remain questions of where to draw the line. If we take a broader approach to save time and resources, the triage captures more and more cases. As more of the caseload is reduced, the 'triage' option begins to look like the pure balance of probabilities approach and reintroduces its attendant drawbacks. Irrespective of where the line is drawn, there will

[29] See Neil Duxbury, *Random Justice: On Lotteries and Legal Decision-Making* (Oxford: Oxford University Press, 1999).

always be cases that a human judge would have decided differently compared to the algorithm, even in 'low probability' cases.

* * *

We have explored three potential ways in which probabilities could be converted into legal outcomes. But there are many other potential options. Each alternative creates new trade-offs. The costs may be unpalatable for some even if the AI algorithm is reliable, unbiased, and transparent. These alternatives are premised on the idea on that at least some algorithmic predictions are converted into a legal decision. But the prediction may, of course, be used in other ways. The prediction can be used merely as an input – perhaps as a suggestion or recommendation – that the judge can consider, along with other relevant facts. For example, when a judge is deciding whether Xavier is an employee or an independent contractor, they would take into account the 80 per cent likelihood predicted by the AI algorithm in addition to all the other information available. This approach may do little to address the backlog of cases. Humans are still 'in the loop.' The delay and backlog of cases may not be reduced from this approach. That is, the courts may not be able to reap the putative benefits of the AI. Further, recent empirical work has explored how judges use risk assessment scores.[30] This research has shown that the biases of human judges can be reinforced and compounded when humans have the discretion to follow or ignore suggestions or recommendations of an AI algorithm.[31]

VI *Conclusion*

Judges face prediction problems. AI offers some promise for the modern judiciary in resolving some of these prediction problems. Litigants face enormous backlog, delay, and cost, which represents a barrier to access to justice. Judicial systems around the world are exploring the potential of AI to help lower these barriers and reduce the backlog of cases. In this article, we examined two distinct types of AI prediction: algorithms that predict facts and AI algorithms that predict outcomes. When AI predicts facts, little changes. There are strong arguments to use the balance of probabilities approach to convert probabilities into binary findings of fact. This approach reflects the current approach to fact finding. The structure of the law will not be dramatically affected.

---

[30] See e.g. Megan Stevenson, 'Assessing Risk Assessment in Action' (2018) 103 Minn L Rev 303; Megan Stevenson & Sandra G Mayson, 'Pretrial Detention and Bail' (2018) Faculty Scholarship at Penn Carey Law Working Paper No 2403; Bo Cowgill, 'The Impact of Algorithms on Judicial Discretion: Evidence from Regression Discontinuities' (2018) Columbia Business School Working Paper.

[31] Alex Albright, 'If You Give a Judge a Risk Score: Evidence from Kentucky Bail Decisions' (3 September 2019), online: *The Little Dataset* <thelittledataset.com/about_files/albright_judge_score.pdf> [perma.cc/C5R5-MQDQ].

But with predictions of legal outcomes, the calculus is different. The balance of probabilities approach that works smoothly for facts creates an enormous disruption to the structure of law when applied to predictions of the future. Its ripple effects would dramatically impact settlement outcomes. By extension, it would affect the legal risk of potential plaintiffs and defendants – and their behaviour in taking precautions.

The presence of algorithmic predictions in judicial decision making raises new problems with probability. It invites new and important questions about the relative value and costs of uncertainty in the legal system. Importantly, it fundamentally changes the nature of law. The conversion of *ex post* adjudication into *ex ante* rule making brings more than just expedient decision-making and a reduction of uncertainty.

Our main point, here, is that implementing AI predictions over outcomes into legal decisions is not easy. There is no obviously correct approach. Any approach to converting the probabilities into binary outcomes introduces trade-offs. Any discussion of the value of algorithmic decision making in law must acknowledge that the benefits and drawbacks of using AI will necessarily depend on how judges convert these probabilities into outcomes.