

INTERVIEWER: So we are so excited for today's panel. We are very lucky to have a number of experts and legal scholars on this panel who will help us understand the complex issues involved in the use of artificial intelligence risk assessment tools in sentencing.

First we have the Honorable Judge Jed Rakoff, who will be moderating today's panel. Judge Rakoff has served as a federal district judge for the Southern District of New York since 1996. Since going on the bench in 1996, Judge Rakoff has authored over 1,500 judicial opinions and has also frequently sat by designation on the Second, Third, and Ninth US Circuit Courts of Appeal.

We also have Professor Deborah Hellman, who is the David Lurton Masee Professor of Law and the director for the Center of Law and Philosophy here at UVA Law. Professor Hellman's work focuses on equal protection law and its philosophical justifications as well as the relationship between money and legal rights.

Next, we have Professor Jessica Eaglin from Indiana University Maurer School of Law. She is an expert in criminal law, evidence, and federal sentencing law.

We also welcome Mr. Alex Chohlas-Wood, the executive director of the Stanford Computational Policy Lab. Alex is a current PhD student at Stanford University. He has led the development of data driven tools in both the private and public sectors, including as the director of analytics for the New York Police Department.

And finally, we welcome Ms. Julia Dressel. Julia is a software engineer at Recidiviz, a technology nonprofit building data analytics tools to power decarceration efforts.

So first I would like to invite each of our panelists to offer some brief opening remarks. And maybe we can start in the order that they appear on my screen. So Professor Hellman, would you like to go first?

DEBORAH HELLMAN: Sure. And I'm so thrilled to be here. Thank you both journals for putting on this fabulous event. And I feel very honored to be up here with such great panelists. I'm sure it's going to be a super interesting discussion. So within this issue or in this space, what I've mostly focused on are issues related to the fairness of algorithmic decision making or questions of discrimination. I'm actually becoming interested in other areas, but that's really what I'm going to focus on here.

So the issue of algorithmic fairness in the context of criminal justice got a lot of attention through expose, I guess you'd say, in *ProPublica* magazine that emphasized a particular algorithmic tool called Compass and the way that it affected-- the way that it was used in the criminal justice context.

And essentially, the dispute about that tool created some interesting questions because the tool treated different groups of people the same in one respect and differently in another respect. And it wasn't possible to treat the two groups of people, let's say blacks and whites, but this would be true with regard to other groups as well-- it isn't possible to treat them the same in both dimensions, thereby giving rise to the question, which of these dimensions, are we to care about?

So let me tell you a little bit about the dimensions and some thoughts I have on it. And here I'm mostly drawing from the paper "Measuring Algorithmic Fairness" that Shweta referred to a moment ago. So one of the ways that you could treat the two groups, let's say blacks and whites, the same, you could call equal predictive accuracy.

And here the idea is that a black person and a white person who are scored by the tool, if they both get the same score, that score should be equally predictive for blacks and whites or for each of those two people. They should be equally likely to recidivate if we're using it to predict recidivism. So that's one way in which we can treat the two groups differently.

Another way you could call error rate balance. And here the idea is that a black person and a white person who get the same-- excuse me-- who do not go on to recidivate, who over time don't recidivate, are equally likely to have been characterized as low risk by the tool. So the first metric that I call people predictive accuracy begins with a score and asks about its ability to predict reality. And the second measure begins with reality and asks about how well it's captured by the score.

So in that paper I make, I guess you could say, three contributions that I just want to describe briefly what they are. And I'm not going to say that much about them. But I'm happy to talk about them in Q&A. So the first point that I make is that the first measure, that is that the score should be equally predictive for each of the two groups, is a measure that relates to what we ought to believe.

It tells us what we ought to believe about whether the people of those two groups are likely to recidivate and, as such, as a measure about what we ought to believe, is less apt or maybe inapt as a measure of fairness, because fairness, generally speaking, relates to what we ought to do. And a measure that's focused on what we ought to believe isn't a naturally good tool to tell us something about fairness.

The second measure, the one that is about whether a black person and a white person who don't go on to recidivate are equally likely to be characterized well or poorly by the tool, does relate to issues that are in the kind of fairness space. So I think it is relevant to fairness though it's more suggestive of unfairness rather than being absolutely a measure of unfairness.

And the third point is that if the tool were to explicitly take account of race, which the tool does not, that is actually had a racial classification within it, it might be improved both as to accuracy and as to fairness. But most people who develop and work on these tools think that there are legal prohibitions that would prohibit explicit racial classification within the tool.

And the third point I want to make is that that's not as clearly so as one might think. That is one can make arguments for why that isn't problematic from the perspective of constitutional law. Obviously, there are arguments you could make against that. But it isn't the slam dunk that people take it to be. So I'll stop there. And happy to talk about any of that in the Q&A.

INTERVIEWER: Thank you, Professor Hellman. Alex, do you want to go next?

ALEX CHOHLAS-WOOD: Sure. So I'll also maybe just briefly-- I don't know how well people are familiar with this idea what risk assessments are but maybe just briefly introduce the idea of what these tools do. So risk assessments are generally meant to inform these high stakes judicial decisions like I'm most familiar with the pre-trial context where you're thinking about whether to detain somebody before their trial, but also I know we're talking today about sentencing decisions that are made by a judge as well. So these tools are meant to inform those decisions.

And really the key strategy is that these tools are meant to predict misconduct, whether it's future recidivism or, in the pre-trial context, you might also be concerned about whether somebody is going to appear at their court date. Most of these tools are actually relatively simple statistical instruments.

They're something that you could probably calculate on the back of a napkin if you really came to it, even though a lot of these end up being proprietary and not visible to the public as tools themselves. And these instruments or these statistical algorithms usually bring in factors such as age and criminal history. There are other factors like race and gender that Professor Hellman mentioned that maybe are much more controversial in terms of inclusion of these algorithms.

So I'll talk a little bit about the advantages. I'm guessing a lot of the discussion today will be about the concerns, which I share some concerns myself. One advantage of using these tools is that it's possible that they'll bring consistency to judicial decisions. One context that our lab examined looked at a pre-trial detention decisions in one jurisdiction and found that if you looked at two different judges, one judge was detaining 50% of the people that they saw who came to their court and another judge was releasing 90% of the people that came to the court. And that was even after adjusting for case factors.

So essentially, a defendant's outcome could really vary very widely depending on which judge they were assigned to. And I don't think that any of us want that outcome in the criminal justice system. And so the idea here is that these tools could, in theory, help bring those different detention rates or other important decisions in line across judges and make them more consistent for everybody.

The other very related advantage is that, in theory, these tools also can provide a great amount of transparency to these very important decisions in the criminal justice process. So for example, if you can examine exactly what's going on under the hood, you know exactly what factors are being considered, the court essentially is making a very explicit statement about how it constitutes or understands risk.

And in that sense, defendants and anybody else who's interested in the process can very clearly understand how risk is being measured for any given defendant. And I think that's a pretty marked improvement over the existing process where decisions are essentially made by a judge or really this much more opaque process that happens right now where we don't really have a clear sense of how that risk decision is made.

The other thing, backing away from the really nitty gritty here about risk assessments themselves, is I think there's a larger question about how or really what the policy is that algorithms are being used to support in general. And so I think there's a normative question here that's really relevant for this context for sentencing, for example.

We have to answer the question first of whether we actually want to use prison or sentencing in general to avoid future recidivism. And if we agree on that premise, then I think risk assessments are generally a very good tool to help inform that decision, because if we're trying to prevent future recidivism, these are tools, like I mentioned, that can help bring a consistency to that estimation and transparency to that process.

But if we imagine that there might be other uses for sentencing like rehabilitation, I would argue that risk assessments are a pretty poor tool in that sense. And so I think there's much bigger normative questions that we have to answer about what the policies are the algorithms are supporting, the AI is supporting, before we can really answer questions about whether the algorithms are appropriate or the best tool for the decision. And I'll stop there.

JED RAKOFF: So Julia, do you want to go next?

JULIA DRESSEL: Yeah. Thank you. I'll keep this brief so we can get into the discussion. But I wanted to first talk about how important the racial context of the specific United States criminal justice system is in any conversation about algorithmic or statistical tools that are used within that system.

You can hear the word machine learning. But any statistical model that's making individual level prediction within the criminal justice system is going to be necessarily built on historical data about what has happened in that system. And you're going to feed the model with historical data about who and what kinds of people have been trapped over and over again in that system and then ask that machine or whatever you're building to make a prediction about the potential future recidivism of an individual person.

And understanding the patterns of what has happened in that system historically is super important in understanding how those tools are going to be influenced by those patterns in the decisions that they make. And so I just wanted to give a couple quick statistics on the racial discrepancies that we currently have in the United States criminal justice system, because I think it's just super important context to have.

And a couple high level stats that we have are that Black adults in our system-- Black adults in this country are almost six times more likely to be incarcerated than a white adult. And Hispanic adults are around three times more likely to be incarcerated than a non-Hispanic white adult and that Black and Hispanic-- excuse me-- Hispanic individuals in our country comprise around 29% of the United States population but they make up about 57% of the prison population.

And so overall, we've got this system that has historically not treated people of different races and ethnicities equally and has historically, as it emerged from the Jim Crow South, has disproportionately policed and incarcerated people of color, predominantly Black people, in the country. And so as we're thinking about how an algorithmic tool is used in sentencing, like Alex said used in pre-trial, it's used all over the system-- actually, these tools are used in pretty much all decisions that are made within the criminal justice system-- it's important to know these are the patterns that we've had in our system since its inception. And it's not surprising that any tool that we've made is aware of those patterns and is using those patterns that it learns about who gets trapped in the system over and over again to inform its decision about who it thinks will continue to get trapped in that same system.

And specifically about the work that I do, I in my role at Recidiviz work with a lot of state data on what types of decisions are used, using risk assessments. And it can be anything from parole decisions to what kind of supervision level somebody is on, if they're on probation or parole. They're used to determine whether someone's eligible for early release from probation.

They're used to determine the programming that somebody might be qualified for when they enter prison-- so a lot of things. It's not only sentencing, and it's not only kind of at the beginning entrance of the system, which-- people sometimes assume this is only used to kind of make beginning decisions about how somebody will be either detained pre-trial or sentenced once they are in a sentencing situation. And I think I'll leave it at that, and I'm really looking forward to the conversation.

**SHWETA
KUMAR:**

Thank you, Julia. I think Professor Eaglin is connected to audio, so Professor, you do want to give your opening remarks now?

**JESSICA
EAGLIN:**

Thank you for the opportunity to participate in this conversation. So I became interested in actuarial risk assessments in the sentencing context as they were proliferating in response to or as a potential partial solution to the phenomenon, the historical phenomenon, of mass incarceration and the exponential increase in the size of our prison population in the United States. And there are so many issues with the tools, some of which I know have already been mentioned. And I apologize to Professor Dressel, I missed a few of her comments.

But I want to spend my few moments talking a little bit about expanding beyond the challenges within the tools, and understanding and problematizing the proliferation of these tools as an innovation to potentially reduce incarceration. So one of the things that particularly interests me is the way that the tools actually shape how we think about the phenomenon of mass incarceration itself. And I have written in the past about what are the conceptual bases on which we legitimate or justify the expansion of these tools in response to mass incarceration.

It's supposed to save us money. It's supposed to identify those individuals that are potentially the most dangerous or the most likely to recidivate-- however we define that, and I'm happy to talk a little bit more about that in the question and answer-- into the future. But it also changes these other soft concepts that we think particularly are meant to constrain the expansion of the carceral state.

So I've written quite a bit about the line between incapacitation and rehabilitation. And if we think that these tools are actually facilitating our ability to identify individuals for effective rehabilitation, which they're often offered as a justification to do, and to the extent that what they actually are doing are identifying individuals who are subject to more social surveillance through the carceral state, sometimes through things like drug courts and treatment, which are supposed to be benevolent and to connect individuals with much needed social services but often can expand to justify the offering of treatment services within the carceral state, and as we at the same time are reducing or eliminating services to communities most in need outside of the carceral state.

Similarly, I've talked a little bit in my research about our ideas of racial justice and the ways that actuarial risk assessments can actually legitimate a notion of racial justice as what I've called a formal idea of racial justice. When we are more concerned with treating people the same, when actually the reality is that people come to the carceral state and face sentencing from very different backgrounds-- and those different backgrounds are captured and standardized in actual risk assessments and in deeply problematic ways.

So there are the traditional controversial factors-- things like age, things like gender-- and while most actuarial risk assessments now no longer consider race as an explicit, perfect predictive factor, they did up until about 1974. But even without explicitly considering race, we know that it has-- that actual risk assessments have a negative impact in terms of tending to identify individuals of color disproportionately, however we want to define that, and obviously that is a question about fairness that Professor Hellman has already talked a bit about.

But part of the reason why we have this issue is because the factors that are used to predict recidivism risk are not actually objective in our deeply racially and economically stratified society. So things like age of first arrest are actually disproportionately occurring in communities of color and poor communities in particular due to policing practices that we know are deeply problematic. But even things that we might-- so things that are not criminal history are also deeply problematic.

So actuarial risk assessments are continually being offered, that take into consideration factors like zip codes. The Pennsylvania Sentencing Commission considered whether to include zip codes as a predictive factor and actuarial risk assessment. That makes a lot of sense, again, if we understand the history of policing and disparate policing, where we know communities of color are subjected to overpolicing and under-policing and other ways that are ultimately reflected in the actuarial risk assessment itself.

And so if we think about not just the design of the tool but contextualizing this pursuit of the tool as a response to mass incarceration, it's a far less-- I guess it's a far less uncontroversial-- it's a much more controversial solution or partial solution that is being implemented in the criminal justice system. And in a recent paper-- and I'm happy to talk about this in Q&A as well. A recent paper that I wrote also talks about how the epistemological implications of actuarial risk assessment specifically for judges, and the sense in which we can't understand actuarial risk assessments without understanding the long history of transformation in this space for judges that were triggered in part by the expansion of sentencing guidelines-- I write quite a bit about the Federal Sentencing Guidelines in particular-- which in some ways inverted what judges are supposed to be doing or have long conceptualized themselves as doing at sentencing, from engaging in questions about individual justice-- how did this person end up in the carceral state?

And what kind of punishment would be appropriate for this person? To something that is far more in the vein of what David Garland has called individuation, which is a focus on how we can generalize the process of individualizing in a way that I think can actually obscure the expanding carceral state and the implications that has and that are also reflected in the actual risk assessments themselves. So I'll leave the rest of it for Q&A, but those are my-- a different way to think about the conversation about actuarial risk assessments in the sentencing context in particular.

SHWETA KUMAR: Great. Thank you so much, Professor Eaglin, and thank you, everyone, for your preliminary remarks. So without further ado, I will let Judge Rakoff take it away.

JED RAKOFF: So thank you very much. I'm intimidated. Everyone on this panel knows a lot more about artificial intelligence than I do. In fact, until about maybe 10 years ago, I thought artificial intelligence was what the Court of Appeals used when they were reversing me, which always seemed to be highly artificial, but now I know better, of course.

I do want to put some questions. And each of my questions will be initially directed at one member of the panel, but I encourage other members of the panel to join in as well. And my first question is directed to Professor Eaglin.

At least the cases I've read involving the use of artificial intelligence in sentencing have focused on the fact that it measures recidivism, which is usually used, in those cases at least, as a basis for increasing the sentence beyond what it otherwise would be. And I wonder whether that's a proper form of sentencing. This goes more generally by the notion of incapacitation.

And looked at from one angle, what you're saying is, we're going to increase your sentence not because of what you did but because of what we predict you will do in the future. And I wonder whether that is really an appropriate way for sentencing to proceed. So starting with Professor Eaglin-- but I encourage other members of the panel to respond as well.

**JESSICA
EAGLIN:**

Thank you for that question. So what's interesting in the context of actual risk assessments is that proponents of the actual risk assessments would say that this is, in fact, an incorrect use of the tools as well, that they are being offered in particular as a way to divert individuals. And they're less concerned with what happens to individuals who are identified as high-risk and much more focused on what happens to individuals that are identified as low-risk.

Now, we can understand that as problematic. I think that it is, and part of the reason I think it's problematic to focus on it in this way is because that in itself has taken us away from thinking about why this individual is actually being identified as high-risk. And in previous work, I've done quite a bit of exploration of what it is. What is actuarial risk? What does recidivism risk even mean?

Most of the actuarial risk assessment tools that are being used today are predicting risk of re-arrest in, like-- usually, I think it's about two years, two to three years' time period. In part, we can't understand why somebody would be identified as being high-risk of rearrest without thinking about problematic policing practices that disproportionately focus on individuals that have already been connected in some way or touched by the criminal justice system, which isn't to suggest that every single person that is identified as high-risk is a person that wouldn't be-- whether a court was thinking about it from an incapacitation or a retribution perspective, somebody that is appropriately receiving some kind of punishment in prison as a potential sentence.

But instead, the dilemma really lies in, to my mind, this policy decision to use this as a way to think about addressing and reducing incarceration. We shouldn't be surprised that the courts are actually giving longer sentences to individuals that have been identified as high-risk. It's basically what the tools are sending as a message, even if they're not being advocated to be used in that exact way.

This is deeply controversial. But it's also something that we can only understand in the context of a sort of shift that has been happening in the sentencing context and really in criminal administration for decades, away from retribution as the guiding theory of punishment, if we ever really think that actually came to fruition, which I question. Because our conception of what incapacitation is, is quite broad and has influenced our ideas of moral desert as well.

So I think that part of the dilemma, as I stated in my comments, is about narrowing the scope of what judges should be doing to-- first, starting with which of the theories of punishment should be guiding them, instead of passing judgment on this question of how often we are using incarceration as our primary mode of addressing social problems that are falling at their doorstep and that actual risk assessments do nothing to resolve.

JED RAKOFF: Any other member of the panel want to comment on that? Let me move then to a very different kind of question. The algorithmic scales that are used for the purpose of predicting recidivism have been the subject of some criticism, that they are not really that reliable. I read recently that some studies think that they are wrong 25% to 30% of the time, which seems like a pretty high error rate. So I wonder, Julia Dressel, if you wanted to comment on that.

JULIA DRESSEL: Yeah, absolutely. So I have seen in my research-- specifically it was on the tool Compass, which is the tool that sparked the loudest kind of wake-up call around 2016, I think, with the ProPublica study. And that tool specifically has been found to have around a 65% accuracy, which is even worse than-- you're saying 25% to 30% wrong. That's 40% to 35% wrong.

And when you're thinking about the point of view of somebody whose life is in the balance-- you know, this is either a question of whether they will be released before their trial, which could be weeks, months before their trial, or if it's influencing sentencing, there are a lot of people where an inaccurate prediction in this context has very, very disastrous consequences. But I think actually more prevalent-- or a more important question to be asking than what is the minute accuracy rates of these tools is the fact that there is no current widespread regulation of what the performance of these tools is.

And they have grown so prolifically throughout the criminal justice system because they have been acquired in the same way that Microsoft Office gets implemented in a judicial district or something like that. And so there are processes of procurement where different systems within the corrections landscape-- they are being sold technology, and they're procuring this technology and using it with a lot of promise that they are going to really help either judges or a parole board make decisions.

And until recently, there wasn't a lot of awareness that these tools were even being used, and until recently also, there hasn't been much of a demand for regulation, and for the tools to have to prove that they are actually doing what they say they're doing. And so there isn't a lot of knowledge on what the standard accuracy is for a lot of these tools, because there hasn't been a lot of standard regulation for what the accuracy is for, what standard measures of fairness are and should be held accountable for the implementation of the tools.

JED RAKOFF: Yes, Professor Hellman?

DEBORAH HELLMAN: I just wanted to ask a question about that research and to call attention to a little ambiguity in what we might mean when we talk about whether they're accurate or inaccurate. Because if a tool makes the prediction that it's 75% likely that so-and-so will recidivate, and so-and-so doesn't recidivate, that doesn't mean that it wasn't 75% likely. It's like the polls that say-- you get the weather report that says there's a 75% chance of rain, and it doesn't rain, it doesn't mean that the weather report was inaccurate.

So when we're talking about prediction, it's not so easy to say-- I mean, we have to look in the aggregate. So I'm just curious about the studies you're referring to about whether they're making a-- giving a report about the percentage of people who don't go on to recidivate or whether they consider that kind of a case to be an error. Because I think it's a little bit different when we're talking about prediction.

And the other thing that complicates it, of course, is, if you release people and they recidivate, you have data on that. If you don't release people, you actually don't have data about the accuracy. And that's a limitation that's in the nature of things. So these questions, I think, are more complicated.

JULIA DRESSEL: Yeah. Thank you for bringing that up. I think those are all really good points. Another thing I'd like to say is that the implementation of the tool affects the outcome that it's predicting. Right?

And so the fact that if you're rated as high-risk in a pre-trial setting, you're more likely to be detained pre-trial-- that is impacting the data on whether people released pre-trial are going to commit a new crime before their trial. And the studies that I was particularly talking about are looking at very, very binary classification. And so saying-- given the people that we have usually two years of data for after they were assessed-- is looking at a very binary, did they end up rearrested or not, and making that 65% accuracy claim on whether or not the risk prediction was towards recidivism or away from recidivism, and then whether or not that binary event happened.

But even measuring what the accuracy of a tool is, is super complicated. Because like Professor Eaglin said, the number of people that-- or, sorry, the trends of people that are more likely to be policed and are more likely to be arrested are highly, highly impacted by race, by neighborhood. And those mean that the measurement of the actual thing that you're trying to identify from a prediction standpoint is super complicated and is also influenced by the prediction that happened to that person in whatever context.

JED RAKOFF: I wasn't totally sure I understood the practical impact of Professor Hellman's question. Because what is the difference between saying, I predict Mr. Jones will recidivate, but my measurement is only right 75% of the time, or my saying, I think there's a 75% chance that he will recidivate. Doesn't it really sort of come down to the same thing? It's a 25% erroneous measure.

DEBORAH HELLMAN: Except that what you did is, you gave Mr. Jones a score of, let's say, seven, that you think it's 70% likely that he'll recidivate. And then, you get all the people who got a seven, and 70% of them did recidivate, and 30% didn't. Then we might say that the tool was perfectly accurate in that 70% of the people who got a seven recidivated and 30 didn't.

Because the tool wasn't saying in that instance, will recidivate. The tool was saying it's 70% likely that they'll recidivate. So you can't really answer the question, is it accurate as to the individual case. Because no matter whether he recidivates or doesn't recidivate, if the tool said it was 70% likely that he would, there's no way to answer whether it's right or wrong as to him. There's only a way to answer in the aggregate.

JED RAKOFF: It's interesting that because this is used by the judges, it's not subject in the federal system and in most states systems, I believe, to Daubert analysis. If the prosecutor introduced one of these measures, it would be subject to a Daubert challenge. And one of the most important factors that a court would have to look at is error rate before even admitting it into evidence.

But since it doesn't come in that way-- it comes in through use by the judges-- that kind of critique is not authorized. Let me ask a different question that's sort of related. Some of the companies who manufacture these instruments don't allow lots of the underlying key information about their design to be assessed, because they're trade secrets, and they're in competition with other companies.

But doesn't that really present a serious problem in assessing whether they're [INAUDIBLE]? Let me ask Alex Chohlas-Wood to comment first on that.

**ALEX
CHOHLAS-
WOOD:**

Yeah. Sure, I definitely agree with the characterization that it's problematic. One way that I find that's actually really useful to think about these algorithms is that really, we're making some false distinction with algorithms that they're sort of an entirely new thing. But in reality, they're just policies, like everything else that happens in government.

And I think there are reasons that bills are available to the public. Actually, one thing I think that we're all very intimately familiar with that's been coming up recently a lot is this idea of vaccine prioritization. And that system in itself that we all know-- you know, the eldest people got vaccinated first or were eligible first-- that whole process, you could absolutely call that an algorithm.

And I think it would be absurd for us to imagine a scenario in which a private company had some algorithm that was not inspectable by the public that prioritized people in the public to get vaccinated. I think that's pretty obvious to us why we would see that as being very problematic. And in some sense, like Julia mentioned earlier, these are really high-stakes decisions, just like a vaccine is. And so we should really insist that these algorithms, if they're used in practice, are transparent.

And I think it also gets to this idea that Mr. Eaglin was mentioning earlier, that there are definitely problematic factors that could be incorporated in these algorithms. And so the idea of transparency really is super important. This did come up in a 2016 Supreme Court case in Wisconsin. I believe it's called Loomis versus Wisconsin.

And actually, one theme of that case from Loomis, who was contesting the results of the use of the Compass on his sentencing decision was that in fact, he thought that the inclusion of gender in the algorithm was improper. And that highlights, aside from the inclusion of gender, which I think is a separate discussion, the idea that he knew that gender was used and could contest its use in the calculation of his risk score-- points to this promise of risk assessments as being, in their best form, a transparent way for risk to identify risk and to measure risk. Otherwise, we get in the situation that I think we can all pretty easily understand that this is a problematic situation for risk assessments if they opaque.

JED RAKOFF: For those who are interested, the citation for the Loomis case that was just referred to is a 81 Northwest 2nd 749 decision of the Supreme Court of Wisconsin in 2016. The decision also says something that I found interesting. That's the nicest word I could put on it-- which is that the judge would have reached the same conclusion even without the algorithm.

Well, fine. So why did we have the algorithm, right? But let me ask a question that's already come up in some of the discussions so far. And that is whether these scales reflect racial and gender bias, and whether that in itself makes them highly problematic. And let me ask Professor Hellman to start off the discussion about that.

**DEBORAH
HELLMAN:**

Thanks. That's a great question, and I want to pick up something that both Professor Eaglin said and also Julia Dressel. So I think actually, we should worry less about whether there's race or gender traits in the algorithm than the fact that-- and I think this was Julia Dressel who emphasized this, that we have a past that is filled with injustice, and data is just information about the past.

And so to the extent that you're making decisions based on past data, it's likely going to carry forward some of those injustices. But I think one-- and that's true about algorithms that are fancy AI, and that's true about the algorithms that are just decisions in our head. If the judge looked at the defendant and is trying to make a decision about whether to release him or her on bail, that judge is-- we don't know the little black box there, too. We don't know what's going on, and that's also likely to be infected by things from the past that we would characterize as injustice.

One thing I want to emphasize, though, is that there are two ways in which injustice infects the past, and therefore can make its way into these algorithmic decisions. And I think they are different, and it's important to emphasize their difference. One gets a lot of attention, and the other one gets less attention.

So the one that gets a lot of attention, I think, is the idea that biased policing practices-- let's say if the police are policing African-American or other minority communities more, then there's going to be more data about arrests of those people in those communities, and that if you're looking at that past data to make decisions to make predictions about the future, that's going to import the bias policing practices into the predictions going forward. So I would call that accuracy effecting injustice.

But there's also the fact that we've had all sorts of injustices in the past. And because of those past injustices, members of minority groups, let's say, are poorer. And being poor often is associated with higher crime rates or less employment or other things that get factored in. And so that past injustice is baked into the data, too, but not necessarily in a way that makes it less accurate. So both factors are in there.

And I think we have to think about how we want to use these tools, because they do reproduce both forms of injustice. Sometimes they make it less accurate. Sometimes they don't make it less accurate, but we still should be concerned about hearing that injustice forward.

But I think coming back to then something Alex said in the beginning, we also have to ask the comparative question. It's not as if our minds or the minds of judges are unbiased and transparent, so it's a comparative question about, are we making things worse or making things better. Because the alternative isn't a system that's free of those problems.

JED RAKOFF: So Alex, you wanted to-- well, several people wanted to comment. So well, let's start with Alex, and then Jessica, and then Julia.

ALEX CHOHLAS-WOOD: I'll make this quick. I completely agree, Professor Hellman. And then, just to extend that, I think it's really important to not only think about these algorithms in isolation but to really evaluate them in context. And I'll sort of just give these two examples that illustrate why this is really important.

One is that, let's imagine that we have a super unfair biased algorithm that's implemented somewhere. But as you mentioned, maybe the judges in that jurisdiction, they completely ignore it, it's never used, and actually outcomes don't change there at all. So even though this algorithm itself might be really problematic, if the outcomes aren't really changing, that's like a situation in which we might really be focused on the algorithm being problematic, but in fact there's not really anything going on.

Conversely, it's possible that we could create an algorithm or maybe even a-- a pipe dream, but let's imagine we can create an algorithm that's like perfectly unbiased. It's totally fair, we put into practice, and in that jurisdiction, judges only use that algorithm selectively to incarcerate Black defendants. And so I think there's this myopic focus on the algorithm itself.

I think there's a huge importance to actually evaluate these things in practice and see how they influence judicial behavior over time. And unfortunately, that research has been somewhat lacking, but that's really where I see the big gap happening-- how are these changing judicial decisions, and how is it impacting all the communities we care about, including Black communities?

JED RAKOFF: Jessica? Now you're muted.

JESSICA

EAGLIN:

This is not my day. I'm sorry. It's Friday afternoon. My mind is elsewhere.

OK, so I have several things I want to say about this. You know, my first comment-- and I think-- is that we can't understand inequities that-- I pause to characterize them as things of our past, and they are very much things of our present-- but inequities that generate from, for example, disparate policing, which I think we are all very much aware of, or things like poverty, which actually themselves can also be built out of policing practices. And that doesn't mean we can fix all of that in a tool, and I think that's very much part of what Professor Hellman is getting it.

But it's also tricky when we start to think of these things as completely independent when we go into designing a tool. And to then perceive of the tool as being itself objective because it treats these things as different when actually in our social reality these things are all building on one another-- so I take pause with that. But I do think that it highlights a point that I was actually going to make earlier in response to a comment Alex made, which is, it's not just about what judges do with a tool.

Because part of the problem is our perception that the tools are in fact objective when they're not. They're very much constructed. You know, those big questions like, what are the kinds of things we are predicting? When we say recidivism, what is it that we're measuring? Well, we tend to be using the kind of data that is easily accessible.

So things like low-level arrests, misdemeanor arrests-- those kinds of things, they're just very easy for us to get our hands on as developers. I shouldn't say, as, because I'm not one. But what we've seen is that these are kind of the easy metrics to use.

But it's not necessarily that those are the metrics that we should be using. Now, to this point, you know, what else could we be measuring? Well, we know, again, going back to things like selective policing practices, that we also have this deep history of things, for example, like white-collar crime.

We're not measuring that. That's not in our actuarial assessments. Why isn't that in our actuarial risk assessments?

Those are normative decisions about what are the things that we choose to measure, that exist within the tool, and then if we perceive of those tools as being objective, are going to influence judges. And I appreciated Judge Rakoff's point-- are we surprised that judges are using them? No, like this is what they were designed to do. And for a judge to say that this is what we would have done otherwise doesn't negate the fact that the tools themselves have these problems.

Now, Alex also added earlier on saying, this is a problem of transparency. And I think that something I've written about in the past is, what is the problem of transparency? Is the problem of transparency a question of-- that this was a privately developed tool as opposed to a publicly developed tool? That is a problem.

But the deeper problem is that we don't understand what the tools are doing, that we don't even understand what the data is doing. What is the data being selected? So there are so many necessary points of intervention, not just through judicial critique, but literally through public engagement about, what does it mean to adopt a tool, that I think up until this point have been largely eliminated from the discourse around adopting these actuarial risk assessments.

It's not just a question of what tool we should use, but like what are the tools supposed to be telling us? We need more intervention into that, and particularly from the communities, in my opinion, that are going to be most affected by this. And we already know that they are the communities that are already being disproportionately pushed into the carceral state.

So if we are going to start with questions about transparency, I think that there's an obligation for meaningful transparency, not technical transparency. But that up until this point is missing from the conversation.

JED RAKOFF: And Julia, you wanted to comment as well.

JULIA DRESSEL: Yeah. A lot of what I wanted to say has already been said. But what I think is important is thinking about-- if you are going to build any actuarial risk assessment-- that word's always hard-- or algorithmic tool in this context, you're thinking, OK, what are historical features that are predictive of recidivism or predictive of finding yourself incarcerated in our country? And you're going to land on employment status, age-- past criminal convictions is usually the most predictive factor in ending up in the system again.

You're going to land in all of these categories that very descriptively show how our country has criminalized different groups of people. And so you have the criminalization of poverty in our country. You have criminalization of mental health issues, and you have the criminalization very explicitly of people of color, very historically in the system.

And so no matter kind of which categories you latch on to and have enough data to be able to build a predictive tool, you are always going to be reinforcing historical categorizations of criminality going forward. And so anything that you are using that has any form of reliability in a predictive way-- excuse me-- is going to be picking up on the ways in which we've disproportionately criminalized certain types of people or certain situations in the past and reinforcing and making that even worse going into the future. And then, on top of it, when you have those inside of a black-- maybe a black box technology or something that is hidden behind trade secrets, those categories of criminalization get solidified as objective truth. And we start relying on them as more objective categories instead of looking at the system that we have, identifying where there are discrepancies, and trying to ameliorate those discrepancies.

And so going back to something I think Professor Eaglin brought up before, which is, we need to start asking questions of, why are we relying on these tools? In which contexts are we relying on these tools? And what questions are we trying to ask?

And if we are using all of this information about someone's need for housing and somebody's past criminal history and the fact that they don't have a job, et cetera, et cetera, all these kind of risk factors for ending up in the system again, and we're asking a question of, oh, what's this person's risk because we need to mitigate risk in a way that we need to mitigate their risk by punishing them by detaining them, et cetera. Instead, can we start asking questions of what are these people's needs does this person have a housing need? Let's get them housing.

Does this person have a mental health need? Let's divert them from prison, and we need to be investing in resources that help in that way. So it's asking different kinds of questions and realizing that if we rely on the patterns of criminality that we have built over a century in order to make any kind of prediction of what will happen in the future, we will never get away from this system of mass incarceration that disproportionately affects, negatively affects, people of color in our country. So that's what I wanted to say.

JED RAKOFF: So before I get to my next question, the last comment, I think, deserves a further elaboration. As I'm sure everyone--

JULIA DRESSEL: Professor--

[INTERPOSING VOICES]

--for that last comment.

JED RAKOFF: Everyone in this audience knows that this country has by far the highest incarceration rate of any country in the world, averaging over the last two decades two million and more people in jail and prison, 60% of whom are either young Black males or young Hispanic males, a 500% increase over what our incarceration rate was in 1960. It's a terrible situation.

It is driven, in my view, in substantial part by mandatory minimums and career offender statutes, and sentencing guidelines. And I want to focus, leading to the next question-- mandatory minimums, career offenders, statuses-- there's nothing the judge could do about that. Federal judges have been on record for at least 15 years, saying mandatory minimums are terrible and create harsh and inhuman sentences in many situations. But we're stuck with them until the legislature does something about it.

Guidelines are a different matter. They were originally mandatory, the first 15 years in the federal system, but now they're discretionary. And a question that's really hard to get at is whether they still impact sentencing or not.

And I want to put that same question in the context of these algorithm measures of recidivism. Do-- is it your view-- it's really hard to know-- but is it your view that judges are really using this, or really just using this as a tool to camouflage a sentence they may have already arrived at for other reasons? And that's such a hard question that I'll start off with Professor Eaglin. (LAUGHING) As a criminal law professor, you got to know the answer, right?

**JESSICA
EAGLIN:**

All right. So thank you so much. I actually-- I want to respond to your question, and I also just want to add a quick point in conjunction with what Julia just said. So in response to what Julia said earlier, I think also just to add to her point, part of the problem is about thinking about solutions outside the carceral state, which we are very reticent to do.

And so part of the way the actuarial risk assessments have revived themselves in criminal sentencing has been through-- they're being offered now as risk and needs assessments as opposed to just risk assessments. And this is supposed to be a new way to think about what this technology is doing, but it also reflects a willingness to identify what people need and give it to them with sort of the criminal arm hanging over their head. And that's the only way they can get it.

So they've got to pay for the services in the carceral state, and they can't get it, and they're continually-- we're moving away from giving them the opportunity to get such services that they need any other way. And so we really do have to understand this in a more conceptually large framework, and I think that that's a lot of what Julia was getting at.

So to this question that Judge Rakoff is raising, do judges use this? Of course they use it. Do they think they use it? Maybe not.

And I think that the analogy to sentencing guidelines is spot on. I mean, if we look at the history of federal sentencing guidelines in particular, judges were righteously angry when the Sentencing Guidelines were rolled out, and said, this is not what sentencing is, and this reflection-- this mechanical, numerical reflection of what matters at sentencing is not a one-to-one to what we do.

And as we have seen, in the 15 years of mandatory sentencing guidelines, suddenly by the time the guidelines were rendered advisory, judges forgot or didn't ever learn how to sentence without that mechanical technology. Right? And I think that we cannot understand why judges are not up in arms about actuarial risk assessments without understanding that they have already been conditioned to think in a completely different way about what sentencing is supposed to be and what kind of factors matter at sentencing.

We have moved as a society in a direction toward saying that the things that matter are the things that we can measure, and actuarial risk assessments are a thing that we can measure, and so of course it is going to be a thing that very much influences judges in their sentencing. I don't think it's the only thing they take into consideration, but I do think, in line with the way that I and I think many others have understood sentencing guidelines, that in a number of circumstances, they're making judges' job harder.

They're making it harder for judges to be able to see and to distinguish why different people end up in the criminal justice system, what is the punishment for this crime but also for this specific person. Actuarial risk assessments aren't actually giving us a lot of tools to help us do a better job of that. They're instead shifting the orientation of what it is we think matters at sentencing to the things that we can count, and then expecting and finding that judges are going to be influenced by that in many different ways.

JED RAKOFF:

So following up on that, just to take the liberty of being moderator to chime in, I thought the point you just made is so terrific. I was a criminal defense lawyer before all these laws came into play. And so I have some feel for how sentencing was in the, in my view, good old days or bad old days, depending how you look at it.

But most of my colleagues have never known anything other than mandatory minimums, career offender statutes, and sentencing guidelines. And so even though the sentencing guidelines are no longer mandatory, they play, in my view, a big role, because they don't know better, so to speak. They never experienced anything, because they're a lot younger than me.

So this is a serious problem. I don't know if anyone else wanted to comment on that before we move on. I should say that we're now at the point where if there are questions from the audience, we'd be very happy to take them. And you can direct your questions through the chat facility to Shweta, and she'll repeat them, then, for the benefit of us. And do we have one already, Shweta?

**SHWETA
KUMAR:**

We do not yet. So if anyone has any questions, please feel free to put them in the general chat. Yeah. Sorry, Professor Eaglin, go ahead.

**JESSICA
EAGLIN:**

I have a comment if we're waiting for questions. I guess it's a comment and a question turning it to the moderator. Because before the call, before the talks started, we actually were discussing very briefly the question of Loomis, and I know that Alex raised it as well, the Wisconsin decision from 2016.

And I think that there are so many things to critique about that decision. Although I think it was a well-intentioned decision, there are many issues with it. But I wondered what your thoughts were, in particular in light of this question that you raised, Judge Rakoff, about the way that we understand what sentencing is and how that may have evolved, if you have any thoughts on how that might have been reflected in the Loomis decision.

JED RAKOFF:

So first, I want to say that I love Wisconsin. It is a beautiful state. I always enjoy whenever I have a chance to go there.

But I respectfully think that the decision was wrong, wrong, wrong, in many, many respects. But most in the sense that I don't see how a court can allow an algorithm to be used that you can't assess. This is the point that Alex was making before.

Transparency is central to the judicial process, and certainly to the sentencing process. And to say, oh, I'm going to make my sentence on the basis of a conclusion that was reached through an instrument, but I don't want to tell you how that instrument works, and I don't even want you to find out how that instrument works-- just take it as a given, and bam, go to prison. So that does not strike me as the way to proceed.

The great-- this is a little bit off the subject, but not totally. The great problem with sentencing in this country has historically been that it's perceived, not without reason, that different judges give different sentences for essentially the same crime and the same individual. I think that's often greatly overstated, because in any case, there are often 100 factors operating that a sensitive judge needs to take account of.

Nevertheless, certainly when I was a criminal defense lawyer, we all knew who were the harsh judges and who were the soft judges, and we always tried to manipulate our cases before the soft judges. But curing that problem has, I think, led to worse problems. And this is just, I think, another example. Here is an attempt to give some scientific veneer to assessing just one factor, namely recidivism.

And the very fact that it can't be measured, although query whether those measures really are good-- but assuming even that they are good, by the very fact that it can be measured, it takes on a life, and importance-- that I think is unfortunate. The same thing happens with the guidelines. The guidelines are more extreme in illustrating the point.

So the original guidelines for many years said that one ounce of crack cocaine was to be treated as the equivalent of 100 ounces of powder cocaine. There was no scientific basis for coming out with those figures. The Sentencing Commission just thought they sounded good, and they were reacting to the so-called crack cocaine epidemic of the late 1980s.

The result of that was that crack cocaine sentences, which were predominantly of African-Americans, were vastly disproportionate to the sentences imposed on white cocaine powder cocaine users. And this, in my view, clear racism was justified as being an objective attempt to eliminate disparities through the sentencing guidelines and their arithmetic approach.

Now, that's an extreme example, and what we're dealing with here in this panel doesn't fall into that extreme. And I want to be clear on that. But it illustrates to me the point that this mumbo jumbo about numbers can often be very misleading. Now, do we have any questions?

SHWETA Yes. We do have a few questions.

KUMAR:

JED RAKOFF: That's great. Because I'm so embarrassed by my last answer. I wanted to get out of here. But OK.

SHWETA So first up is Denny. You can just unmute and ask a question.

KUMAR:

AUDIENCE 1: OK. Thank you. Hi, everyone. Thanks for the discussion. It's very informative.

I kind of have two questions, and they both sort of relate to the data question that Professor Hellman had, which is-- so for Compass, it sounds like they're using data to predict recidivism. But I can also see how changes in laws and social programs, like defund the police and all these new projects, could affect the dynamics and the cause and effect of that prediction tool.

So is there sort of like a statutory re-analysis or basically checking the newest set of data isn't on these changes for accuracy? Or does it just perpetually go forward?

And the second question is related to the historical data aspect. I imagine that at some point, somebody creating the tool is also inserting their own normative views and deciding specifically which algorithm to use or what data to curtail or which outliers to cut out. So if we're seeing judges are a black box, human-driven process, why would a human-driven process at arm's length be better, writ large?

JED RAKOFF: So Professor Hellman.

DEBORAH Sounds like that was directed at me. So I don't mean to be the-- I actually have, I guess-- I don't mean to be the
HELLMAN: contrary one on the panel, but maybe that's healthy for dialogue. Because I think I have a more mixed view of these tools than some folks on the panel.

But I will say, I have a mixed-- not-- I'm just saying the more positive things sometimes, because I think it's too simplistic to think that they are this black box and the other thing is not a black box. Because when the judges that Judge Rakoff referenced-- the hard one and the soft one-- make their sentencing decision, we actually have no way to figure out what the factors are that went into those decisions.

And in fact, even if the person him or herself was trying to figure out, how much weight did I give to this factor versus that factor-- personally, I don't know about you, but when I examine my decision making, I don't know how much weight I gave to various factors in making decisions. It's very, very hard to know. So we have to keep that comparative in mind, and it may be the case that we actually have more information about the factors affecting some of these algorithmic decisions than we do-- I mean, the AI technological ones than we do the human ones.

Though I do think that-- as several comments referenced-- we tend to kind of fetishize the algorithmic tools in a way that we are more skeptical of in human decision-making. We more readily recognize that it can be erroneous or biased or something like that. But it's not as if it's perfectly transparent and updated for new data as it comes in all the time either. So let's just remember that.

But I think-- and this relates to the question also that David [? Luban ?] put in the chat about what tools-- I'd like to know what other people think about this-- what tools we have for requiring things of these tools. And so now, I'm just going to spitball here, because I really don't know.

And I think some other people have more information than me about-- but I don't see why the legislature could not put whatever qualifications or requirements it wants. That is the state legislature of whatever state, I think, could say that any tool that's used to make sentencing decisions has to be validated in certain ways, has to be updated in certain ways. For your question, Denny, I think they could. Right?

I mean, I don't know whether there are any such laws in place right now, but I don't see any reason that they couldn't. Looks like Julia is leaning forward. She maybe knows about whether there are some laws like that. So this is outside of my area of expertise, and I'm going to pass the ball.

JED RAKOFF: Other comments?

SHWETA KUMAR: Professor Eaglin, you can go ahead and--

JESSICA EAGLIN: No, no, no. I don't-- if you were already-- you were leaning forward, you go first. I'll follow up after you.

JULIA DRESSEL: I was saying, to my point I made a while ago, that the awareness of the misuse or the disparate impact of these tools I feel has grown a lot over the last couple of years, and I do not think that the legislation has caught up to it yet. And there have been some effort. The biggest one that I'm thinking of, specifically because I live in California, was SB-10 recently, which was the Senate bill to eliminate cash bail with the direct replacement of cash bail in the whole state with algorithmic risk assessment.

And within the wording of that bill, there was some kind of attempt to-- actually, not really. There was poor attempt at what the regulations would be of those tools, of the implementation of those tools, of the validation of those tools, and of the consistent-- yeah, basically, the consistent validation that those tools are actually working. And so there hasn't been much legislative regulation of how transparent or what features can be used, et cetera.

That does not mean that there shouldn't be. I think that there absolutely should be. But there hasn't really been, that I'm aware of. Maybe you know more than I do about that. [INAUDIBLE] elsewhere.

**JESSICA
EAGLIN:**

If I could just hop in-- so yeah, they're not regulating it. And I think part of the problem is like, why aren't they regulating it? And I've written about this a bit.

I mean, you know, I think that we need like tons of regulation in this space. Like questions about what kind of data we should be able to use. We don't have that. People are-- it's just a normative thing.

And it's crazy. Like sometimes-- for example, I think in Wisconsin, the actuarial risk assessment that they adopted was adopted by the Department of Corrections. Like, it's not even-- and then they adopted another one that was advanced by the Judicial Council, and we don't really know who picked which one or why.

You know, but it's not just like what tool it is. It's like what's in the tool is not really being regulated, down to simple questions like, what is the recidivism risk that we are measuring? They're not asking that question.

Is the data coming from the state in which this tool is being applied? Not asking that question. You know, and the deeper question is, why not?

And I think that this goes to David's question, David [? Luban's ?] question. They definitely can. And there have been suggestions that they should be making some conditioning.

I mean, the argument that you get from the developers as well-- we can't be too transparent in our tools. Like, this is what makes them competitive-- is that each of them is slightly different. And while we're happy to work with individual decision makers, we have our own incentive to keep this private for our own financial reasons, et cetera, et cetera, trade secret.

But the legislators have to actually have the will to do this, and they don't. And until recently, they have not. I mean, in New York, there's been an initiative for accountable algorithms across the public sector. So not just in the context of sentencing, for example, but also in the context of like algorithms that are used to determine which schools high school students are getting into, and just that whole complicated process that they have out there.

And this dilemma that I highlighted in response to Alex's point has come up in that space, where it's like, what does transparency mean? Like, transparency can be a deeply procedural thing. Like, do we know what the data is?

To me, that's not good enough. It's got to be substantive. Do we know why we have chosen this data? Like, what are the reasonings for it?

We really aren't there yet, and I think that the reason that hasn't happened is because of the perceived objectivity of the tools, which has been and is continuing to be critiqued and attacked, not just in the legal space, but obviously of course in the space of public discourse, which is so deeply important for-- if we're going to use them, which I am deeply skeptical that they are doing what we benevolently think that they can be doing. But even if we are going to use them, then we need to change our own framework about, what does it mean to be adopting some sort of technology into this decision-making process?

And it's not enough to say that judges would be judges would be transparent so-- or lack transparency, and the tools have the potential to be transparent. That's not good enough. That's like very far from good enough, because you can embed all kinds of decisions into a technology and not ever have to ask meaningful questions about it that don't exist in the language of technology or technologists themselves.

And so it's very important, I think, that sort of law-- I mean, this is why we're having conversations like this-- that individuals, law students, law professors, legal minds start to be critical of the tools. And that has everything to do with our own perspective about the technology first and foremost.

JED RAKOFF: So I think, Shweta, we have time for one more question if there is one.

SHWETA KUMAR: Yeah. So one question from Rachel Martin is, do you think there's any way to use algorithms, not necessarily the ones we have now, in a way that can provide some of the promised benefits like reducing incarceration while still avoiding the pitfalls?

JED RAKOFF: So who wants to take a crack at that? Alex.

ALEX CHOHLAS-WOOD: Yeah. I think I definitely-- maybe not super common viewpoint here, but I think that there is a lot of promise in using algorithms to allocate benefits. One example is an intervention that I'm working on right now, where we're helping people get to court by sending automated text message reminders to them to try to help them remember that they have a court date. And this is something that's been shown to be very effective at increasing appearance rates. And as a consequence, people can avoid having arrest warrants put up on their head for missing a court date.

And so basically, you can do this very simple, very cheap thing of sending basically a text message that costs a couple of cents to somebody. You can avoid this whole horrible outcome where somebody has a warrant placed for their arrest and they go to jail because they missed a court date, just by doing this very simple thing. And the way that we're using algorithms in that context is by personalizing the content of the reminder so it's most effective for that individual.

And of course, there's still equity concerns there about what the personalization means and what kind of data we're using. That's something that we think about all the time, but certainly, there are many contexts in which algorithms can really be used to reduce incarceration and improve outcomes across the board.

JED RAKOFF: Julia and then Jessica.

JULIA DRESSEL: Yeah. I'll say it super briefly. Something that we're working on at Recidiviz is modeling the potential impact of a policy. And that is using more predictive analytics, which-- the words get all mumbled together when we're talking about machine learning, artificial intelligence, everything like that. But being able to use the trends of how people have moved through the system historically and how certain maybe crime classifications, et cetera, have impacted the flow through a state system, and then saying, OK, given this policy, what do we predict the system-level impact will be?

And you can actually look, if you have enough data, on how race is related to certain classifications that are related to that policy. You can say, oh, we predict that this policy will actually have this positive impact on reducing disparities in the state. And so you can use data and an understanding of what has historically happened to make a prediction at a system-wide level, which is pretty different than making a prediction of an individual person. But a system level-- this is how we predict the system will respond to this state-wide change-- and try to model also how that will have an impact on the racial disparities in that state and push for changes that are going to be reducing those disparities in that system.

JED RAKOFF: And Jessica?

JESSICA EAGLIN: So I just want to say that, in response to Alex's comment, I think he said it spot on, that this is a way. This is not the only way, and it might not be normatively the way that we want to reduce incarceration. It has tons of additional costs, costs that somehow get lost in the conversation.

For example, the cost of designing and maintaining an actuarial risk assessment, many of which the states are paying forward to these private companies-- if the judges are saying they would have done the same thing anyway, which I question, do we want to spend our money this way, or could we spend it in another way? I am familiar with the project that Alex is talking about, in terms of identifying and notifying individuals of court dates, and that sounds important. It is important.

But I also think that there are ways that we can address and reduce incarceration that are more structural. And we can think about whether we want to put the onus on individuals to individually reduce the incidence of incarceration, or if we also want to change the way that we as a society think about offering services and things like that, that would also be more significant in my mind. But alternative ways to think about reducing incarceration-- and the thing that really gets my goat-- there seems to be many things on this panel that I have many things to say about.

But the thing that really gets me is when people suggest that we can't reduce incarceration without this, that we need the tools. We don't need the tools. We choose to use the tools. We choose this route, but that doesn't mean that it has to be the only way or the preferred way. And we need to find a way to find a foundation for that conversation.

JED RAKOFF: I think we are pretty much at the end of our time. I did want to express my appreciation both to all the fine panelists and also to our two sponsors, for putting this, I thought, very informative, panel together. So Shweta, anything further that you wanted to say?

SHWETA KUMAR: I think that's it. I just wanted to say, thank you again so much, everyone. This was such a helpful, insightful panel. So thank you again for your time.