

This seminar was given by Cary Coglianese, Edward B. Shils Professor of Law and Professor of Political Science; Director, Penn Program on Regulation at the University of Pennsylvania Carey Law School. It was given on Thursday, March 25, 2021 as part of the Regulatory Policy Program's weekly webinar series.

Joe Aldy:

Welcome to the Regulatory Policy seminar. I'm Joe Aldy, the faculty chair of the regulatory policy program at the Mossavar-Rahmani Center for Business and Government at the Harvard Kennedy School. Let me open with a few reminders regarding the logistics of our online seminar. We are recording this seminar and will post it online. So if you know of someone who is interested, but could not attend the talk live, please let them know that they can access the seminar at the MRCBG YouTube channel. And I'll post a link in the chat in just a moment for you.

Joe Aldy:

[inaudible 00:00:29] remind friends and colleagues to register for Zoom links for each seminar on the Regulatory Policy Program webpage. And I'll post a link to that in our chat in a moment as well. We will take questions through the Q&A function in Zoom. Please click on Q&A in the bottom of your screen and type your questions. At the end of our presentation I will take the questions and pose them to our speaker. We're excited to have our old friend Cary Coglianese join us at the Regulatory Policy seminar today to present algorithm versus algorithm, a framework for government use of machine learning.

Joe Aldy:

Professor Coglianese is the Edward B. Shils professor of law and professor of clinical science and the director of the Penn Program on regulation at the University of Pennsylvania. Cary Coglianese specializes in the study of administrative law and regulatory processes with an emphasis on the empirical evaluation of alternative processes and strategies and the role of public participation technology and business government relations in policy making.

Joe Aldy:

He's also written on climate change policy, public participation and transparency in federal rule making, the use of artificial intelligence by government agencies in voluntary environmental programs. He's also the founding editor of the journal Regulation and Governance. And he founded and continues to serve as an advisor to the regulatory review. Cary, welcome back, at least virtually, to the Regulatory Policy seminar.

Cary Coglianese:

Well thank you very much Joe. It's great to be back and even if it is virtually, it's always nice. Although this was the easiest travel I've ever had to get up to Cambridge. Nice to see you and thank you for the opportunity to talk today about a project that grows out of a report that I prepared for the administrative conference of the United States, which is a government agency that operates a bit like an internal governmental think tank for good government practices in the administrative state. I should acknowledge that I have benefited enormously from a number of terrific Penn students as research assistants on this project.

Cary Coglianese:

One of whom, Alicia [Lay 00:02:42], is now my collaborator as we turn the report into an academic article and perhaps down the road this will spur us into other work too. So I'm very much eager to get your feedback and hear your questions and your comments about what I'll talk about today. I should also sort of situate this project at the outset in a larger academic and public discourse about the use of artificial intelligence generally and certainly with respect to its use by governmental entities. And I think it's probably fair to say that most of the work that one finds is much like these examples, emphasizing the negatives about artificial intelligence, highlighting the risks, the dangers, the ways that it can be used in unjust and improper ways.

Cary Coglianese:

Some of the titles of these books indicate that negative posture, weapons of mass destruction or if one thinks about comments that Elon Musk has made. He's talked about artificial intelligence as ultimately summoning the demons. Stephen Hawking has predicted that it will signal the end of the human race. Lots of concerns. And I think these are very serious concerns to be sure and there's value in research that's highlighting the risks of the use of artificial intelligence, though what I will say, my work differs in two respects. One here is that I'm willing to accept that there are downsides to the use of machine learning and artificial intelligence, but there also can be upsides as well.

Cary Coglianese:

The second difference is, of course, that I don't have anything near the kind of catchiness to the titles of this work. The report that I prepared and submitted for the administrative conference was titled A Framework For Governmental Use of Machine Learning. Not going to probably capture as many downloads or sales of then titles of some of the other work that's out there. But it has, I guess at least going for it, the claim that it has contributed to an actual statement, set of recommendations if you will, from the administrative conference that were adopted last December and published in the federal register this January, that recognizes that AI is an increasing interest to government agencies and offers agencies some advice and guidance about how to approach some of the dangers and realize the possibilities and positive potential of machine learning when it is used by government agencies.

Cary Coglianese:

I should say that at the outset that I am a public member of the administrative conference and I chair its rule making committee, the administrative conference is uniquely public private government agency that has certainly a wonderful full-time government employed staff, but also relies on the input from government members from other agencies as well as public members, principally academics. So what I'm going to say here today is simply my own views, not in any way the views of the administrative conference itself.

Cary Coglianese:

What I do want to talk about is what, in my view, is really the key question facing governments that are considering using this. When should they use it? When should artificial intelligence tools be used by government to compliment or substitute for human decision making processes. That is what I see as the question of algorithm versus algorithm. It's the human algorithms that make up the current decision making processes in government versus new digital algorithms, machine learning algorithms that power artificial intelligence systems. What I'm going to do today is make, I think, four points of first I'm going to emphasize how human algorithms have real limitations, which maybe will come as little surprise to

anyone, but I think it's a point that's often overlooked in this overall literature commenting on and analyzing the dangers of artificial intelligence.

Cary Coglianese:

There are also dangers and limitations. We're just more accustomed to them and I think we need to be reminded of them. Second then I will talk about the digital or machine learning algorithms and their possibilities, but I'll also recognize that they do have some pitfalls. I'm going to take a little bit of a, I guess you might say, detour in a way. I'm going to turn toward some legal analysis because if machine learning algorithms are intrinsically incompatible with our administrative law principles in the United States because they are so called black box algorithms, they have a degree of opacity to them, a degree of automaticity and self-learning properties that arguably could run up against various administrative law doctrines, then no need to really even consider them.

Cary Coglianese:

I will briefly highlight that legal analysis work that I've actually done in some other papers recently. And the bottom line of that though is that I think as a matter of legal analysis, there are no real inherent limitations on governments using them. They can use them wisely or not and they may well encounter legal issues in the way that they use them, but there's nothing intrinsic about the black box nature of these algorithms that I think would impede their adoption by government agencies. So then, at that point, if it's cleared legally then the question is when to use them and I'm going to offer a framework that I provided in my report to the administrative conference, a framework for government agencies to think about when to use them.

Cary Coglianese:

Ultimately that framework is going to be used by humans and the decisions about when to use algorithms will ultimately be in human hands so there's no escaping the possibility for human limitations, but there is, I think, a good possibility though that we can design processes that work better and when should the government go forward. That's where I want to end up at the end. Okay? So with that, I want to pause for a moment and if anybody wants to get into terminology, I'm not going to be overly precise or particular about the terminology. Obviously there are questions about what counts as artificial intelligence and different definitions.

Cary Coglianese:

There are an array of different kinds of machine learning algorithms and there's a whole host of other language that people use to refer to the kind of digital algorithms that I have in mind, whether it's, they call it predictive analytics, deep learning, reinforcement learning, neural networks, natural language processing, smart cities, and all sorts of other terms that are out there. I'm not going to be particular about that. What I'm really focusing on are a couple of general characteristics that I think make the kind of artificial intelligence and machine learning that makes my focal point here interesting and important.

Cary Coglianese:

Two characteristics. I mean one is the possibility of automating a good number of, maybe even more, tasks that governments handle today because of the power of some of the new artificial intelligence tools that can be developed, more things can be automated and that automation itself provides an opportunity to compliment or substitute even for human decision making. It also gives rise to what's

often thought of as algorithmic aversion and that's a characteristic that I think makes this choice important for government agencies to think through carefully.

Cary Coglianese:

The other is what might be thought of as a degree of autonomy to this, it's not just that some tasks are being automated, but with machine learning algorithms, they're being automated with the use of forecasting technologies or techniques that have a self-learning property to them that rather than having humans select the variables of interest and make a choice about the functional form to fit those variables together, that the algorithms are actually searching for the patterns in learning, in a sense, on their own.

Cary Coglianese:

Obviously humans are involved in specifying the objective that the algorithm is supposed to predict or optimize. Humans are involved in training the algorithms and so forth, but the underlying choices about the variables that matter and the functional forms are somewhat autonomous, and that also leads to a degree of opacity to these algorithms. Now that technology is advancing and we're being able to see some better ways of understanding how some of these algorithms are reaching the results they are, but it's not as intuitive at all as if we could put the results of these algorithms into some kind of causal language that might be more readily explainable and understandable by the public or even government officials themselves.

Cary Coglianese:

So anyway, I think these are general characteristics that much of this broader literature is focused on and that's how I'm going to be talking about artificial intelligence. I do also want to point out that there is a bit of a challenge here in talking about what is a highly heterogeneous set of algorithms and tools and techniques in such broad and general terms. And I recognize that. That's how I'm going to proceed, but I think when I get to my framework you'll see that ultimately I think the choices that government agencies need to make will be highly context specific, needing to consider the particular uses to which these algorithms are put, the particular properties that the algorithms have, and ultimately the kind of performance that they achieve.

Cary Coglianese:

So happy to talk further about that, but for now I'm going to kind of gloss over these many differences and speak broadly. The first point though now in the broader argument I want to make is to highlight how humans, which are making many of the decisions that governments are involved in today, have a tremendous number of well-known limitations. I mean, some of these limitations are physical ones. Neuroscientists estimate that working memory for humans can process about four variables, plus or minus, at a given time.

Cary Coglianese:

In part because of that we see in a number of different context human errors. About 12 million medical diagnostic errors a year. We see, in the medical area for example, the rise of the use of checklists to try to overcome those limitations. The World Health Organization has developed a set of checklists for surgeons, for example, in part to overcome those memory processing limitations. There are other, though, physical limitations that we know about, also for medicine, for example. Fatigue with medical

residents reporting their fatigued about 50% of the time and that fatigue is associated with an increase in medical errors.

Cary Coglianese:

There are analogs in the governmental setting. There's research that shows that parole decisions, at the beginning of the day when judges are making a series of parole choices in court, they tend to favor the petitioner in a majority of cases. But then as lunch time approaches the decisions tend to get demonstrably less favorable to the petitioners. After lunch, they systematically return to an overwhelming majority of the parole decisions favoring petitioners. So this is well documented in the legal context.

Cary Coglianese:

Aging is another physical limitation. There was a ProPublica report that maybe shouldn't surprise people too much about how we have a lot of older judges on the bench in the United States and there are documented instances of judges forgetting how to actually walk out of the court room to get back to their office or who have difficulty reading it out loud or who lack the recall of prior decisions or make decisions based upon non-existent evidence. These physical limitations are certainly real, they're combined with, at any age, a series of cognitive biases that behavioral economists and cognitive psychologists have uncovered.

Cary Coglianese:

One of which I just point out, confirmation bias or motivated reasoning. In some sense we have hardwired into our administrative process because when it comes to agencies creating new regulations they have to go through a notice and comment rule making process. That is they have to publish a notice of proposed rule making and take comments on it. Well how open minded can we really expect agency officials to be to what they receive in terms of the comments when they've already committed to what they want to do? I think that's highly prone to be affected by motivated reasoning and selectivity of evidence that comes from well documented cognitive biases.

Cary Coglianese:

In terms of other biases, implicit bias in racial or gender terms has been well studied. We know that just exposing people to very fast flashes of words that are related to crimes, for example. Just even a 75 millisecond flash in ways that people are not even able to perceive that they have seen these negative words will then tend to make them spend more time looking at black faces than white faces in experimental setting. So these individual physical limitations and cognitive biases are real and so too of course are group decision making problems and that is certainly applicable to decision making by government agencies and government officials and we have well-known instances whether it's the Bay of Pigs, the NASA Challenger explosion where group decision making has broken down, maybe becoming prone to group think, organizational behavior, specialist talk about garbage can decision making, we have errors in possibility theorem and we know how preferences in groups can cycle and lead to irrational outcomes.

Cary Coglianese:

And these kind of human limitations, I do think, result in many of the problems that people attribute to government these days, whether it's delays, unfairness, inconsistencies. Some research shows that among the thousands of social security administrative law judges who decide disability benefits, there

can be widely varying accounts and results from these judges, notwithstanding a more random assignment of cases. Some judges will grant disability benefits in 90% of the cases while others will grant these benefits in only 10% of the cases. The need for human decision making at the Board of Veterans Appeals that it can take more than a couple of years for veterans to get final rulings on whether they qualify for disability benefits.

Cary Coglianese:

These are just some examples of how these human limitations do contribute to real concerns about government and I think give us reason to be at least somewhat optimistic or at least give us reason to be open minded to the possibility that automation through machine learning algorithms could overcome some of these barriers because these algorithms are not prone to the same kind of limitations that humans have. And it has been shown that these algorithms have the potential for greater accuracy, ability to have a greater capacity to process data and process applications or other informational tasks, they could do it more quickly. And there's the possibility of them doing this more consistently if it's the same algorithm making the decision rather than 2000 distributed judges or officials around the country.

Cary Coglianese:

And we do see some evidence that this is well documented, can work. There's a study in Nature about how the Environmental Protection Agency could improve its targeting of violators of water pollution regulations, could improve the targeting of those violators by 600% even over say just a random selection of facilities to inspect. The Bureau of Labor statistics is using machine learning algorithms to help in coding workplace injury reports and finds that they can assign about 81% of these reports more accurately than humans and faster. The patent office is looking into using these tools for processing the two million trademark applications and one million patent applications that come in, in any given year.

Cary Coglianese:

The FCC received 22 million public comments on its net neutrality rule and machine learning natural language processing tools can help sort among these comments and find the discrete issues that the FCC or other agencies that receive a large number of comments need to focus on. So there's a great deal, I think, of potential and promise here, possibilities there are, potential pitfalls. Again, to return to that more negatively focused literature, which I take very seriously and again, they provide a very useful rule. It is important to think about the pitfalls. I mean, some of these range from just does government really have the data available and the processing capacity to use machine learning well and there's a lot of data certainly that the government has, but it isn't always linked up.

Cary Coglianese:

So there are efforts to create legal entity identifiers that can link these up, but that's a challenge and certainly a limitation. Another worry of course is cyber security. The more we automate with computerized systems the more prone government might be to hacks and we certainly have seen various cities around the country, Atlanta, New Orleans have their local government systems hacked. There are privacy concerns, both about just keeping all of the data safe and protected that are used in training algorithms and applying them, but there's also concerns that maybe you can use some of these algorithms to actually identify some personal properties or characteristics of people that they would prefer to keep private.

Cary Coglianese:

Lots of questions about the transparency and explainability, which I mentioned a little bit before about these being black box algorithms and certainly concern about bias in cases of the use of algorithms in the private sector where companies have found either high error rates for individuals who had darker complexion or Amazon, in a well-known case, found that its use of tools to screen potential employees or applicants for jobs led to some undesirable outcomes and biases against female applicants. And lastly, there's also, I think, an overarching concern about these tools making it easier and if it's easier for governments to perform their task it's also easier for governments to perform unjust or even say authoritarian tasks as well.

Cary Coglianese:

These concerns are, of course, I don't think totally distinctive to machine learning. I mean, for example, bias is not just a property of these machine learning algorithms. There's bias in humans and that's very difficult to detect. Humans tend to rationalize and evade and certainly wouldn't want to acknowledge often their biases. And often the kind of biases that do emerge in machine learning algorithms are coming about because the underlying data, which had been produced from human generated systems, that bias is coming from that original human source of bias. I think it may be easier to detect or certainly potentially easier to train an algorithm or to adjust it in ways that can address biases, at least probably easier than ferreting it out of and correcting it in human beings.

Cary Coglianese:

All right. So I think there's a real potential here for machine learning. It has to be used carefully of course. There are some concerns about it that are real and important to consider, but at the end of the day I guess we have to also ask, "Is this even an approach that is plausible under our current legal system?" After all, as Lincoln said, "Government of the people and by the people," not of the digital algorithms. And these algorithms, if they are really, truly black box, maybe they're not even compatible with our notion of open administrative government. Bottom line, I want to say, and I'll be brief about this is there's nothing intrinsically incompatible with the use of machine learning, even to substitute for human decision making and our existing standards of administrative law.

Cary Coglianese:

This is something in other work that I've done I've gone into much more depth on this question in an article called Regulating By Robot, Transparency and Algorithmic Governance. And in this work I've touched upon, for example, the non-delegation doctrine, which normally, in theory, inhibits Congress from delegating too much authority to agencies. There's a separate kind of private non-delegation doctrine that impedes the ability to delegate to private entities is the creation of algorithms to perform tasks that would otherwise be performed by human officials is this kind of impermissible delegation to robots or digital tools. I think the answer is clearly that it's not problematic. The non-delegation doctrine requires an intelligible principle to constrain the use of delegated discretion and machine learning algorithms don't work unless they have a very clearly defined mathematically precise object that's stated for them.

Cary Coglianese:

So much more precise than what the courts have approved under the non-delegation doctrine. In the private non-delegation doctrine I don't think necessarily it applies either because algorithms are not creating the kind of risks that that doctrine is concerned about, conflicts of interest and so forth. When it comes to procedural due process we do expect governments to be able to make decisions and explain

those decisions and provide reasons for them and go through a lot case law and discussion of legal principles and ultimately conclude that the kinds of explanations that are expected under this doctrine can be still providing even with a black box algorithm and in fact a Washington court has already, in the [Loomis 00:32:24] case, reached that conclusion.

Cary Coglianese:

With respect to transparency, yes these are black box algorithms, but we actually don't require total transparency of government as it is. Both due process and transparency principles tend to be pragmatic in the US. There are exceptions to them and I think there's more than sufficient ability to explain that an algorithm has been designed for a particular purpose, it's been tested to meet that purpose, and that is the reason an outcome has come about. There are some privacy laws, but those are not really anything that apply intrinsically or a barrier to machine learning.

Cary Coglianese:

Certainly other kinds of statistical tools or other uses of data would raise many of the same privacy concerns and have been managed well by government agencies and in fact some algorithmic tools, differential privacy, have actually helped solve some of the privacy concerns. What about the bias and equal protection concerns? Well I'll just briefly say that it's not likely that you'll see any intentional discrimination reflective of some kind of racial [anymous 00:33:41], even in an algorithm that is using variables that have racial characteristics within the data set because these are self-learning and no one is necessarily saying, "Key in on the racial variables."

Cary Coglianese:

The protected classes, the racial classes, are not the only data in the system and it's the combination of the whole data set that really is making the difference. You don't have the same kind of categorical treatment, a thumb on the scale of those racial factors that have troubled the courts. And even if the courts were to apply heightened scrutiny, I think there's probably a strong case in many of these instances where government is using these tools that there might be some compelling state interest that would meet even a heightened standard. So all this is to say that intrinsically I don't think there are any legal barriers to governmental use of algorithms, of digital algorithms.

Cary Coglianese:

There are possibilities that governments can use them unwisely, irresponsibly, and in ways that offend these legal doctrines, but ultimately I think it's up to the discretion of governments to use them, to choose when to use them, and we do see them using them in terms of regulatory enforcement and in criminal law enforcement, formal adjudication, the social security administration, for example, is using machine learning tool to check the quality and oversee the decisions made by all of those administrative law judges across the country.

Cary Coglianese:

Informal adjudication with decisions about patents or the like. In terms of performing analysis that's useful for understanding how agencies should adopt rules, it's being used there. And when they're analyzing these public comments. And of course there are autonomous technologies being used in various ways by government agencies including the US Postal Service looking into autonomous vehicles. A Stanford NYU study looked across the whole federal government and found about 150 or thereabouts use cases of algorithms, many of them what I would be considering these self-learning ones.

Cary Coglianese:

And this is pretty much how they broke them down. A lot of uses in research and analysis. The next most frequent use was enforcement. So far we don't see as much being used in adjudication and actually notwithstanding all of the attention to risk assessment by the criminal courts. The court system has been, I think, the slowest to adopt, or at least slower than the administrative part of government, to adopt artificial intelligence tools. Many of the risk assessment tools, if not all of them as far as I'm aware, are not really true learning algorithms, they're much more linear kinds of algorithms.

Cary Coglianese:

All right. So let me get to this final question, when to use machine learning. One way of answering that question could be to, and lawyers might be most inclined to look at this, is through a test that the Supreme Court articulated in the case of *Mathews V. Eldridge* when it articulated a balancing test for making decisions about procedural due process questions. And that test simply can be reflected in a simple equation where the judges are supposed to look at the private interests at stake and how a particular procedure leads to more accurate decisions by government and weigh that against the costs of governing.

Cary Coglianese:

This test is usually applied when litigants are coming in and saying, "The government didn't do enough." And so the court is saying, "Well, they have to do enough. Well, did they have to do more? Well we're going to take a look at these." I think this test is not a good one for governments to use in terms of deciding when to use machine learning because the private interests are going to be exogenous, that is they are what they are. The machine learning versus the humans is going to work on the improved accuracy and if machine learning algorithms tend to be more accurate than humans than they're going to fair well under that and they are going to lower the cost.

Cary Coglianese:

In some sense, this test is like made for machine learning algorithms. So I think that's one of the reasons, by the way, that there isn't an intrinsic barrier to the use of machine learning algorithms is because under this test they would pass pretty easily. And I don't think this test gives enough attention to other factors that may matter, including spill over effects for the society overall. One would then naturally kind of say, "Let's try to do cost analysis," and I'm sympathetic to that as a structural matter, but I don't think that, as a practical matter, the government will have the data ability or the ability to put all of the factors into some kind of commensurable value.

Cary Coglianese:

So I think you're really left with a framework to answer this question that's sometimes called a soft benefit cost analysis or some kind of multi-factor policy analysis that students are doing all the time at the Kennedy School, for example. So what are the factors that should come into account? Well I think first of all you have to think about the pre-conditions for the use of machine learning. You're not even in the ball game to use machine learning if you have a task that cannot be defined very precisely. You need some tasks that are very clear and not ones where, as [inaudible 00:40:13] has aptly put it, is often the case in law where we don't really have clarity of goals or precisions and we have to make what he calls incompletely theorized agreements in adopt language like reasonable care, which is not a very clear or precise goal and if that's what you're trying to produce is reasonable care or something along those lines, that's probably not going to be possible use machine learning.

Cary Coglianese:

You're going to have to have some data that's available. So I could see this being preconditioned being met more easily for government determining whether someone qualifies for black lung benefits than for determining whether someone qualifies for asylum where the standard is a well founded fear of persecution. I can see it being used, say in the criminal context for assessing the probability of DNA belonging to a defendant then because there's lots of DNA out there. We could study it. There's lots of data on that. Then the question of what was the probability of the defendant was at sixth and Main at 12:30 on March 12th in a red Corvette?

Cary Coglianese:

Those kind of very sui generis kinds of questions that come up in criminal law context, we wouldn't have the data available. And we should also think about external validity. In other words, if the world is changing, does the underlying data change at a rate that can keep up with the pace of change? Now this is a challenge for human decision making too if we have unanticipated events like a pandemic or stock market crash or something that changes the world and the variables. It's going to be difficult for humans too. There may be more predictable consequences or if we know something about the pace at which data is being refreshed and capable of being refreshed in a machine learning context.

Cary Coglianese:

Those are the preconditions just to use. Now I don't think that, in itself, is determinative. Then we have to look and see, "Well is machine learning really delivering improved outcomes? Is it achieving the goals that are at hand for the task being performed?" That's certainly a natural outcome to look for. But then also look for outcome in terms of the impact on direct users. Are they being discriminated against, for example? And then what about the broader societal impacts? There's a difference between using machine learning tools to sort US mail than using machine learning tools to determine who gets commercial pilot licenses, for example, in terms of the broader social impacts that are at stake.

Cary Coglianese:

Lastly, think about the risks to the agency. There are going to be some legal risk. There's also risks about public acceptability and we've seen this in instances like the UK with its application of algorithms for determination of school placements. I don't know how many of you are aware of what the city of Boston went through with re-designing its school bus schedules. That created a public backlash. I think these risks to the agency are a function first of the degree of determination that how much is the digital algorithm actually substituting for human decision making? Is it just providing an input to human decision making?

Cary Coglianese:

Is it creating a default rule where humans made tend to follow it, but they can't override it or is it actually replacing and taking the human out of the loop altogether? Another factor, of course, so the degree of stakes at issue and I offer, in the administrative conference report, this three by two matrix in which I, for heuristic purposes, say that when the stakes are low and the degree of determination is modest, those are going to be the areas, which would present the least amount of risks to agencies when you have high stakes and the algorithm is doing much more of the work you're going to have the greatest risks of someone contesting this in court or having some public backlash or hearings or so forth.

Cary Coglianese:

Putting it altogether, I mean this is sort of how I would envision the overall framework for governmental use, those preconditions, which obviously don't apply to the status quo, but only for machine learning. But then how well are each options, whether it's the use of machine learning as an input, as a default, or a decision. How well do those perform in terms of your outcomes and how much risk do they pose to the agency?

Cary Coglianese:

I would urge agencies to think about what the weights are to give to these various factors. I'm not suggesting that they all are equally worthy. In fact, I think agencies sometimes overestimate and overweigh the risks to themselves and they should be cautious about that. But that's where I kind of ended up. Let me offer four concluding recommendations. One is to have agencies really think these through, to fill out the kind of matrix that I was talking about, to do the planning and the hard work. Second, they need to think about whether they have the human expertise in house or whether they can get it to use these tools wisely.

Cary Coglianese:

Third, if they do go outside of government they need to think about how the procurement process is setting the agency up later if there are legal challenges or public concerns about the use of the algorithm. One of the situations that's come up in the handful of court cases so far is that contractors are claiming trade secret protection for their algorithm and then the government agencies are not able to do as much to explain to the litigants or the public. That's easily fixable if you think about the procurement issues in advance and put provisions in to provide for that kind of disclosure.

Cary Coglianese:

And lastly, I think that agencies should think about when to involve the public, especially if they think that these are algorithms that will have broader societal impacts, they would benefit, I think, could benefit from public input to help them think through some of the value choices and design features. So let me just return at the end to algorithm versus algorithm and say yeah, there's some dangers and real serious concerns that can accompany irresponsible and unthinking use of algorithms by government. But, on the other hand, there's real limitations to human decision making or human algorithms too.

Cary Coglianese:

And if machine learning tools, artificial intelligence, can be used wisely and well to achieve better results for the public and at comparable or even lower cost then, by all means, I think government agencies ought to be open to them. So I'll stop and invite your questions and look forward to the discussion.

Joe Aldy:

Great. Thank you Cary. That was fantastic. Even if you didn't have the most exciting title to jump off the screen, it was still a truly excellent presentation to get a sense of where we could see, I think, the use of innovative technologies to inform decision making and the processes of many aspects of the modern administrative state. So we have a number of questions in the queue. Let me remind everybody who's joining us today that you can use the Q&A button in the bottom of your Zoom screen and type a question out.

Joe Aldy:

I want to start with a question from my colleague Bill Clark. So Bill notes that one of the crises we face today is the collapse of trust in governance. More deliberation has been suggested as a response to this, but deliberation is about giving reasons. How does greater AI fit into a move to be more deliberative about giving reasons for government action as a means of building the trust on which effective governance depends?

Cary Coglianese:

Well I think reasons are not the only factor affecting trust. Performance also affects trust too. And consistency and fairness as well. So I think what you'd want to do is think about what's the context in which it's being considered to be used. And it may be that people just really would like to have their disputes resolved, their claims resolved in a fair or fast, efficient manner and that can actually build trust and I'll give you one concrete, very quick example. The company eBay has developed an automated digital dispute resolution technique. Actually it's a Harvard Kennedy School alum who designed this system.

Cary Coglianese:

And most of the disputes that people have on eBay are resolved without any human intervention. And yet eBay also reports that people who have a conflict, who have had a problem with a transaction with eBay, and go through their dispute resolution system, they come away and they're more willing to be return customers at eBay than people who never had a problem in the first place. So I do think that trust is important. I think trust can actually, in some instances, be promoted. And I think some of the algorithmic aversion also may be dissipate over time as members of the public become more aware of the, and familiar with, the use of these tools in other parts of their lives.

Cary Coglianese:

In fact, they may even come to demand them from government. "Why do I still have to sit on the phone for so long when I could have this resolved?" So it's a great question Bill and I think it is contextual, but it's also potentially something that machine learning could help with.

Joe Aldy:

Okay. Thanks Cary. So when we think about the potential for using AI in government, your example there, eBay, raises a question about how we can think about testing AI applications. That when we look in the private sector, a lot of these innovative technologies, they've done what we in academia would refer to as some variance of randomized control trials. So when you're trying to think about ways of demonstrating efficacy, perhaps building trust, especially given and reflecting on some of the work you've done in the past on retrospective review in analysis and regulations. What could we do here in terms of planning for and sort of the evaluation up front of these AI tools to be able to address a lot of these concerns that individuals or stakeholders have about the application of AI in government?

Cary Coglianese:

Right. And it's an excellent question. It builds on Bill's concern about trust too that the more we can do to validate these tools, and we have to do that. I mean, that's what I mean by better outcomes. And we can do that. We can start piloting these and we can collect data on how the human system currently is working and almost inherently we're going to have to have that data already available to work these algorithms. And then we can try out, not only train it on the retrospectively on some old data, but

before we launch it across the board, start running it side-by-side with human decisions and see if we can demonstrate and document that it's doing a better job.

Cary Coglianese:

And then I do think that over time if it does a better job and you then replace the humans, you won't be able to necessarily compare it with humans anymore because you'll lose that expertise. But you should, anytime you're really developing it further, you have additional versions of the system coming out, you should always, I think, be mindful of that comparison and that validation and then be open with the public about it. Maybe invite public or outside experts in to peer review those results as well.

Joe Aldy:

Great. Now we have a lot more data now on many different issues, in many different ways than we've ever had in the past and in fact one of the appeals of using machine learning is as a way to help manage these humongous data sets and extract the relevant signals there. And when we think about, for a number of regulatory agencies, they often have some discretion in their authorities to address what may be [emerging 00:53:57] risk. What might be the opportunities to use machine learning to identify sort of emerging risk as a way to motivate changes in the way we design and implement regulations?

Cary Coglianese:

I think that's actually a great question and a great and promising use for machine learning tools and in fact what we are seeing, if you recall my slide from the Stanford NYU study, we are seeing agencies using these tools a lot for that kind of analysis, trying to observe trends, see patterns that humans can't even detect. That's how I would see this being used. You have the Securities and Exchange Commission, for example, using a variety of machine learning algorithms, not just one, but they're using a lot of them to try to observe patterns in potential fraud behavior and then identifying places to intervene either with enforcement strategies or with redesign of rules. So lots of potential there.

Joe Aldy:

So we have time for one more question and we'd like to get a little bit into the process with the administrative conference of the United States. You produced this report to serve as the basis for some recommendations that the administrative conference adopted. A part of that process is typically the opportunity for you, as the expert, to receive feedback and to take questions from the government agencies who, at the end of the day, are already thought of as being at least the initial audience for the report and for these recommendations.

Joe Aldy:

What kind of questions or topics were raised most frequently by the agencies who at the end of the day are going to be the ones building on these recommendations?

Cary Coglianese:

Well that's a great question and I think the tendency among at least agency lawyers probably is naturally to be very cautious and resistant to change. And it's, in some measure I think, within the government those are the folks who I really have in mind and speaking to. And I want them to be aware that there really aren't inherent limitations or legal barriers to going forward. You do have to be careful about it, there's no question, but I think it's important for agency lawyers not to have an initial reaction, an instinctual resistance to change in automation, but instead to approach this as, "Well, let's recognize our

current system is not necessarily perfect either and if this can do better we as lawyers in agencies ought to be able to partner with the policy folks and the managers and develop systems that serve the public better, deliver more public value," as Mark Moore would say it.

Joe Aldy:

Well unfortunately it is now one o'clock and I apologize I did not get to all the questions that we received. But before we wrap up, let me remind everybody that we'll meet again next week, Thursday, April 1st at 12 PM for our presentation on cryptocurrencies, speculative bubble with a future of money by Timothy Massad, a senior fellow here at the Mossavar-Rahmani Center for Business and Government. And finally, please join me in thanking professor Cary Coglianese for his presentation and discussion today. Thank you so much Cary.

Cary Coglianese:

Thank you Joe.

Joe Aldy:

I hope you enjoy the rest of your day. Take care.