

Webinar recording



Can AI improve decision-making and outcomes when it comes to children's care and protection?

February 2026

A CELCIS EMERGING INSIGHT SERIES

CPD MEMBER  
The CPD Certification Service

CELCIS Centre for excellence for Children's Care and Protection

# Emerging Insights: Can AI improve decision-making and outcomes when it comes to children's care and protection?

## Featuring

**Chair: Claire Burns**, Director, CELCIS

## Speakers

**Professor Rhema Vaithianathan**, Professor of Health Economics and Director for the Centre for Social Data Analytics (CSDA), Auckland University of Technology (AUT).

**Professor Emily Putnam-Hornstein**, The John A. Tate Distinguished Professor for Children in Need, University of North Carolina and Faculty Co-Director of the Children's Data Network

**Joanna Redden**, Associate Professor, Western University and Co-Director, Data Justice Lab and Starling Centre

**Dr Shion Guha**, Assistant Professor, Faculty of Information, University of Toronto

**CB****Claire Burns**

Good afternoon, everyone and a very warm welcome to today's seminar, the second in the current series of our Emerging Insight webinars, this time with a focus on AI and children's social care. I'm Claire Burns, the director of CELCIS, and we do really appreciate everyone giving their time to step into this learning space with us. We are so aware of how busy everybody is, so the fact you've been able to find this time is something that we really appreciate. In our previous series, we explored [emerging issues in child protection](#), and we also shared key learning and themes underpinned by the findings from our own [Children's Services Reform Research Study](#).

In this series, we're exploring what AI might mean for children's social care, the services developed and practises used to respond to the care and protection needs of children, and what we all might want to learn about the impact of AI on their lives. In today's webinar, "Can AI improve decision making and outcomes when it comes to children's care and protection?" We'll hear about examples from across the world of where systems using AI have been built, the safeguards considered and put in place, and how these have been working. And we can learn from these international case studies, including from the US, the UK and Canada. We're joined today by an amazing panel, I have to say, four great contributors willing to share their insights and perspectives on this question. And just to say, I really appreciate it, I know the time that goes into preparing for these. We very much appreciate you taking the time to share your learning with us. So, we've got [Professor Emily Putnam-Hornstein](#), who is the John A. Tate Distinguished Professor for Children in Need at the University of North Carolina and is faculty co-director of the Children's Data Network. We've also got [Professor Rhema Vaithianathan](#), who is Professor of Health Economics and Director for the Centre for Social Data Analytics (CSDA), at Auckland University of Technology. [Joanna Redden](#), Associate Professor at Western University and Co-Director of the Data Justice Lab and Starling Centre, and who will also be joined by [Dr Shion Guha](#), Associate Professor in the Faculty of Information at the University of Toronto.

So, a really impressive panel that we're delighted to have with us. So at this point for the first presentation, I'm going to hand over to Professor Emily Putnam-Hornstein and Professor Rhema Vaithianathan. Thank you. Rhema, are you starting?

RV

**Rhema Vaithianathan**

Yes, thank you very much. Lovely to meet you and lovely to be invited. We're going to try something a little bit different. We thought Emily and I would have a conversation. I'm coming at it more from sort of tech point of view. I'm an economist. I got into this work almost 20 years ago, but I started in health, where obviously using data, predictive analytics, machine learning has a much greater history. I started around 2007 in the NHS actually, in the UK, working with teams building and deploying predictive analytic tools in hospitals for re-admission risk. Emily and I have been working together for 10 years, so I thought, why don't we start by me asking Emily a bunch of questions so that we get a little bit more of a casual conversation. We're all a bit over these huge kind of webinars with lots of presentation slides. So Emily, thank you. Do you want to talk a little bit about yourself and then we'll crack into some case studies?

PE

**Putnam-Hornstein, Emily**

Thank you, Rhema. Everyone should know that because Rhema and I have been partnering so long, we often joke that we're like the old married couple, talking over each other. So, this will be perfect for the format. Really nice to be here. Claire, thank you for the introduction. Thanks for the opportunity to join. By way of background, I am in a School of Social Work. I have worked in child protection or child welfare in the United States for 20 plus years. And my entrance to work around algorithms and AI and predictive risk modelling was quite by chance. I had always been interested in how front-end decisions are being made. I never thought of myself as a researcher focused on implementation science or tool development, and then came to appreciate that rather than just writing papers about the ways that data could be used to identify children who are at greatest risk, there were opportunities to partner with folks like Rhema to actually develop and operationally deploy those tools for frontline workers. So that's what I think we will spend our time talking about today in the context of some different case studies.

RV

**Rhema Vaithianathan**

Thanks, Emily. I know we are really busy right now rolling this out across

LA (Los Angeles County). So, LA for those on the call who may not know to contextualise it, the largest county in the US, they've been working to deploy a predictive risk model, and we've been in a trial and now we are scaling it out this week, I think. So, Emily, do you want to share some of that as a case study, a little bit of why they wanted to take this option and what it, how it works? Those are the first two questions.

PE

### **Putnam-Hornstein, Emily**

So, as Rhema mentioned, Los Angeles County is huge. It's larger than most states in the United States. There are more than 3 million children in the county, and their child protection agency is charged with the very difficult task of receiving calls from community members and mandated reporters, having to make decisions about what does or does not get investigated, and then having to make decisions about how to staff and support those investigations. The work in Los Angeles County has been in some ways under discussion going back to 2014. I'm not sure how many of you have seen 'The trials of Gabriel Fernandez', which was a kind of a Netflix documentary, but that is focused on Los Angeles County. And so, as many of you know, what often happens is if there is a high-profile child death or child tragedy, we get leadership and legislators, and public outrage and people are trying to figure out how can we do something different moving forward. It's my opinion that often those efforts are not very thoughtful. They're too overly reactive to one single isolated tragedy. But I will say that it was a child death that sparked the work that Rhema and I have been engaged in in Los Angeles for the last seven years. And I remember it very distinctly because I received a call from the then Child Welfare Lead from Los Angeles County. And I had spoken with him on occasion, but he was not typically calling me up on a weekend. And he said, Emily, we've had another really bad case where a child died. And it hasn't hit the media yet, but it will. And what I'm tired of doing is, I'm tired of having cases where something bad happens and then we look at all the things we missed in terms of information and decisions and people we should have called, doctors records we should have gotten that we just didn't do when the child was alive. And we only discover that after the child has died. And he said, is there a way that we can use some of the algorithmic work and predictive risk modelling work that Rhema and I had actually already started in another state to help the workforce and help supervisors who are overseeing open investigations for alleged child abuse or neglect to identify a small subset where the complexity of the work, the risk to the child, the potential for harm is really elevated so that

we're making sure that we do absolutely everything to connect those families to services, to prevent future harm and, that as much as we are all human and mistakes happen, that we're also making sure that we are bringing forth our absolute best social work practise for the cases where the stakes are highest and greatest. And so that's what actually sparked the work in LA. And Rhema, go ahead.

**RV Rhema Vaithianathan**

Yeah, so can you just tell me what a typical supervisor's situation is today versus before the tool came in? What's changed as a result of that?

**PE Putnam-Hornstein, Emily**

Yep, absolutely. So, in LA, a supervisor overseeing an open investigation would have anywhere from 20-40 open investigations at any given time. Those are investigations where caseworkers are going out gathering information. And typically, the supervisor will meet with the worker throughout the 30 days that it's open. But the supervisor is heavily dependent upon information that the caseworker brings to them about how serious the situation is or what the service needs and concerns are. With the development of an algorithm, what we are doing is on a nightly basis, we're looking at all the new investigations that have come in in Los Angeles County, and we are pushing an e-mail alert to a supervisor if their unit has been assigned what we call an Enhanced Support Investigation. And what that means is that when the algorithm crawls through all of the existing administrative case management data, and in California, we have data around reports of abuse and neglect, foster care placements, you name it, going all the way back to 1998. And so what we're doing is we're using that data, the model's looking at hundreds of features, and it's identifying a small number of investigations, and that gets proactively pushed to a supervisor so that the supervisor knows that these are the investigations where additional time and attention on my part will be very helpful to the worker and to the family.

**RV Rhema Vaithianathan**

Thanks, Emily. So, you know, I have, and I'm sure you've heard this too, Child Welfare Leaders saying to us that they don't think severe fatalities, severe maltreatment and fatalities are preventable. And I think for many years, we sort of thought that as well. But in LA, the findings have been quite surprisingly acute about this issue. So, it'd be really interesting for people to hear a bit about what that initial trial showed in terms of actual

mortality. And more interestingly, why you think that effect was - what was the supervisor doing differently as a result of this e-mail that arrived in their inbox to say, hey, you know, one in 10 of your investigations opened yesterday and the algorithm says it's got a lot of risk factors.

**PE Putnam-Hornstein, Emily**

Yeah, so we deployed this algorithm and this supervision tool in three of the 19 regional offices in Los Angeles County. And what we were able to do is we were able to look at the outcomes for children in offices where we gave this additional enhanced data and information to supervisors versus the 16 offices where we as researchers could look and say which of the investigations did the model think were really high risk, but that information was not passed along to the supervisor because it was just a pilot and a trial. And Rhema, as you mentioned, we found some very significant safety gains. We saw a 20% plus reduction in near fatalities and fatalities in the offices where we gave the supervisors that information. And what I would attribute that to is, yes, it was the algorithm, but what it speaks to is actually when we equip our frontline workforce with the right data and the right information at the right time, we are able to make better decisions about safety and risk. And so, by making sure that we were surfacing salient history and contextual information for these investigations and equipping the supervisors with that so that they could work more closely on those investigations, we were able to implement interventions that were protective and that resulted in the reduction in child harm.

The other thing I just want to point out is a lot of people think "tool" and they think it's something where they have to fill out a form. That is what Los Angeles County had been using before. That's generally how we have delivered data and information to our workforce. And what I think is novel with these new AI applications is that we're able to pull information out of data systems and deliver it in a way that is more useful and more actionable. And that sounds so simple, but at least in the United States, we have not done a good job equipping our critical frontline child protection staff and supervisors with good technology. And so I view this as, on the one hand, a really innovative step, but also one that arguably the private sector and others took 10 plus years ago. And we are just now catching up in making sure that we're giving our public sector workers tools that allow them to do a good job.

**RV Rhema Vaithianathan**

So it's just an intervention that curated information to supervisors and gave them some risk flags and that just precipitated their attention and that attention was enough to prevent, huh?

**RV Putnam-Hornstein, Emily**

And I want to add one thing to that, Rhema, which is that I also think that in child protection, we always have that tension. Are we focused on the current allegation of abuse and neglect in isolation, or are we contextualising it in known clinical family history? And that clinical family history can be very difficult to unearth. A supervisor told me that what we could give them through our tool in a click and a matter of seconds took them 41 clicks in the data system to go back and read, right? So even if we're thinking about it purely as there are massive efficiency gains, and most child harm, I believe, happens not because we have a poorly trained workforce, but because the workforce is incapable of mining through all of the data we expect them to hold in their decision-making process.

**RV Rhema Vaithianathan**

Interesting. I'm sure that's something hopefully that resonates with a lot of people. Do you want to turn the tables and ask me questions that you think my social workers always ask?

**RV Putnam-Hornstein, Emily**

I will, so, Rhema, I think that for social workers such as myself and colleagues, we hear a lot of different terms thrown around. There's a lot of hype, AI, AGI, machine learning. So, can you walk me through at a high level, what do these terms mean and how do you think about the kind of applications in child protection settings?

**RV Rhema Vaithianathan**

That's such a good question because I don't think even the experts know exactly how to pass what's machine learning, AI and AGI. So, AI, as you know, has a lot of hype right now, but what it is, is just a continuation of the prediction problem. So, what we're using in LA is kind of a standard machine learning tool. It's been around for decades. The idea is we take structured data, so facts like was there any findings of abuse or neglect? Was mom herself removed when she was young? Things like that, structured data, and we ask a model to predict something bad happening

to that child. What's called AI now is large language models. Those sort of take text as their input, which in most cases, for those of you who have interacted with ChatGPT, or Claude or other things is your question and ask what's the predicted next paragraph or page that I can give you that will best satisfy you when that's the text you've given me. So, it's essentially the same prediction exercise but using text. So, it's not that different to what we're doing in a predictive risk model. It's just taking text as an input. AGI is this kind of more sort of conceptual difference, which is people consider AI and machine learning in the same class of prediction problems - you think of them as just parroting what they hear. So, they don't have real intelligence, they just think of them as a very smart parrot that says, you know, 'Polly wants a cracker', because they realise that that gives them a cracker. They're not really conceptualising it. AGI is an idea of a machine that knows about the world in the same way a child grows up to know about the world. So, I don't know if that helps too.

PE

**Putnam-Hornstein, Emily**

It does help. Can you talk, Rhema, a little bit about, for the LA model, what kind of data were being used there? Was it both, do you need both the textual data and the kind of structured spreadsheet data? How did you approach that?

RV

**Rhema Vaithianathan**

Yeah, so we just use structured data typically in all our applications because before these recent large language models, the way that text was able to be ingested by those tools was not very good. And we have very good predictive power, just understanding structured data. To be honest, the written clinical notes that we find in child protection systems are sometimes quite incoherent. For example, when you compare what we call face-to-face notes, notes from workers who've gone to visit families, they will sometimes be contradictory to each other because different workers went out. One worker did see some evidence of mental health challenges. Another worker will say, no, there was no mental health challenges. One worker will see substance use in the house. Another worker will say, no evidence of substance use in the house. So, one of our challenges, to be honest, is to improve the quality of some of those caseworker textual notes rather than to use it right now. We're trying to help them improve quality of those.

PE

**Putnam-Hornstein, Emily**

And I guess we have had some success. We built models, deployed models in several different states. Rhema, what do you think is the biggest barrier to child protection agencies adopting these tools? Like what more do we need to know to feel confident that this is the right path as opposed to these of paper and pencil risk assessments tools that are still being so widely used?

RV

**Rhema Vaithianathan**

Yeah, child protection systems are very conservative when it comes to technology, mainly because they've been burnt a lot and they sometimes see this as a tech problem. The other thing is a lot of the vendors out there are used to working with systems where they can take the data off site and take it into their own system, process it, and send it back. We cannot do that in child protection because of the high confidentiality nature of their data. So, data privacy can often mean that most of the kinds of AI systems out there are not available to our child protection system. Thirdly, I think we have a lot of educating to do to help our child protection system leaders discern what's a good use case, what's a good vendor, what does it mean to purchase and pay for high quality predictive analytic tools and what does it mean to be sort of just sold a little bit of, you know, there's a few sort of snake oil salesman, of course, in AI, as in all other areas, and how do you make sure you're not getting caught into one of those?

PE

**Putnam-Hornstein, Emily**

Perfect. And then I guess I have one or two final quick questions if we have time. So, you know, one of the things I think is great about you, Rhema, is that you are super-duper technical, but you also spend a tonne of time shadowing and talking to frontline staff. I would say I think very fairly that all of the tools have been co-designed. And so, I'm just kind of wondering, how should young social workers think about these tools? And what have you heard after some of these tools have been deployed in places like Colorado and Pennsylvania and Los Angeles from the workers themselves?

RV

**Rhema Vaithianathan**

Yeah, we get some great insights. I drive with them. I do work with them. I walk with them. And so, I think we've really learned a lot from that. One

of the things that I find workers are learning to do is to be complementary skills to these tools. And I think that's really important, especially for our young social workers that are graduating today. They are digital natives, and they get frustrated with the systems we give them. They're like, 'Why can I order my pizza in less time than it takes for me to figure out what the family history is? This is mad'. So when they see our tools, which sort of curates and gives them information at the tips of their fingers, that's almost what they expect. And so now they understand how to work as complementary with those. Like, I know it gives me information in this context, but I know it's really weak in giving me more narrative, more contextualised, more clinical judgment. So, I'm going to take what it gives me and then I'm going to do the additional work to unearth more context, more judgement and so on. So that's one of the things I've really found very interesting. They've also said that they have become much more sympathetic to their workers because they can see at a glance how complex, because we give them a complexity flag. So, they've now started saying, 'Oh, you're dealing with a lot of complex cases'. And they started now seeing objective data to make them more supportive.

I think it's really important for folks in the audience to realise that this is not a general, it's not good all the time. It's not good for all purposes. It's a technology we have to curate and use for very specific solutions, like LA said they had a problem. We worked together to identify the solution where this could be a technological answer. This is not a one answer, you know, one ring to heal the world, as it were.

**CB** **Claire Burns**

Great, thanks so much. I mean, I've got about 5 questions already, so I'm sure everybody else who's listening in will have, but I will try and resist asking you my questions at the moment and go on to the next contributor. But lots to come back on, Rhema and Emily, that was such a good input, so interesting. So, going to pass over now to Professor Joanna Redden and Dr Shion Guha.

**JR** **Joanna Redden**

Hi there. Thanks everyone for having us. We decided to share some slides with you so that we can provide an overview of some relevant work. Shion and I have prepared to talk for about 10 minutes each, and maybe we'll just do a quick intro. So just to let you know, I am Joanna Redden. I

am presenting research that is really built on social science research methods and a social science perspective.

**SG Shion Guha**

I'm Shion Guha with the University of Toronto. I am not a social worker. I am a computer scientist. So, I just want to start by saying that I know absolutely nothing about social work, but I have worked with social workers very closely over the past decade or so.

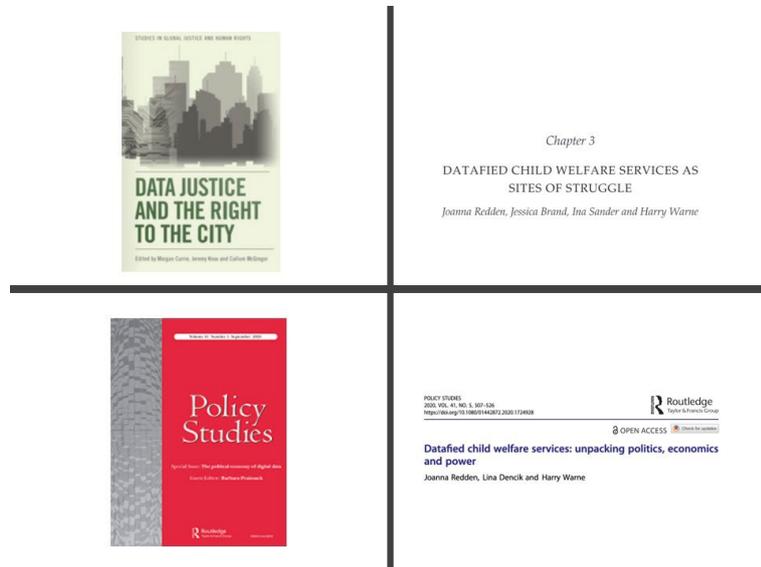
**JR Joanna Redden**

So, in the time that we have with you, before we move into a discussion, and thanks so much to Emily and Rhema for getting this conversation going, we're going to provide a really brief overview of some very specific research that's relevant to you. I wanted to share some findings from some recent research that I have done in collaboration with many others. Then Shion is going to present some research findings from work that he's been working on and then we'll close and hopefully we'll have lots of time for discussion. So, I wanted to start off today just by noting for those in the room who may not be aware, is that the [What Works for Children's Social Care Centre](#) did a really fantastic large-scale project a number of years ago. [They provided a technical summary, a high-level research summary](#), as well as shorter condensed material presenting their research findings. They worked with four local authorities to pilot the use of predictive analytics to see if these tools worked for social workers and if they led to better outcomes.

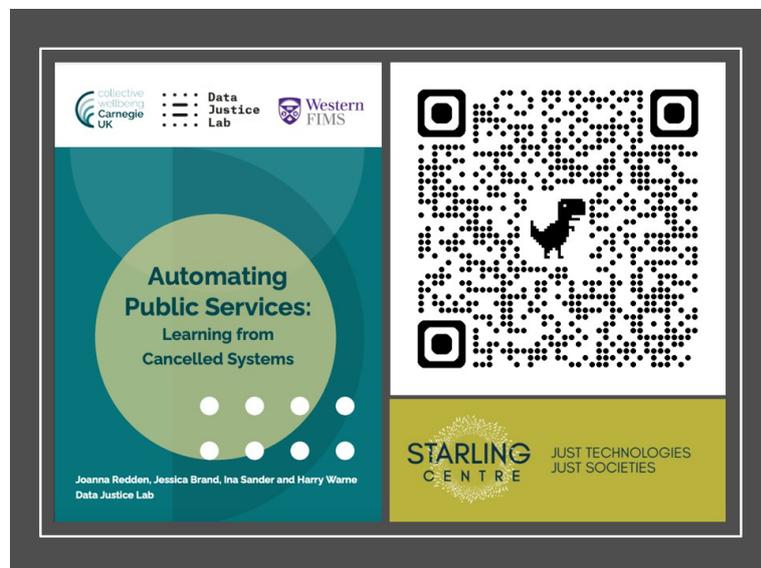


They also did polling with 800 social workers in the UK, and they did community consultations. This is a great resource for anyone out there

who's wondering and looking to make use of these tools in the UK at this moment. And in talking with the organisers, I know that there are a number of local authorities who are doing this. I also wanted to share with you and just leave with you that the research that I'll be summarising, has been published in different locations.



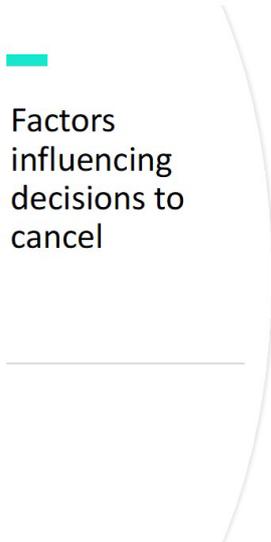
So, there's an open access book chapter. There's also an open access article that you can check out or you can [e-mail me](#) for copies of if you want more detailed accounts of what I'm going to be talking about today.



But for the most part, what I'll be presenting today is a research project that we did that was published in 2022. It was supported by Carnegie UK. And this project involved looking at where and why different government agencies in different countries were making the decision to pause or cancel their uses of automated decision-making systems, which included predictive analytics applications in a number of cases. The idea behind

this research is that we argue it's really important for us to be paying attention to what's working, but also what's not working. As you all know, uses of AI in any child welfare context is really high stakes. There's a lot of research that has been done to date, but one of the key pieces of information that's hard to get is what doesn't work and why. And we argue that learning from those who have made decisions to pause or cancel their uses of AI is really, really important, especially if you want to prevent the kind of harms that have occurred previously. And so, in connection with this work and what this work involved, we did a scoping analysis.

Just to note, it's really difficult to find out where AI is being used by government agencies, let alone where it's being cancelled. So we, in various research organisations that I've been part of, are part of a widespread chorus calling for algorithmic registers - for governments to publish registries of where and how AI is being used. And very few are doing that at this point, even if they have committed to doing so, as the UK Government has done, but is not fully doing at this point. So, you can imagine how hard it is to find out where cancelled systems have occurred. So, what we set out to do was to identify cancelled systems. We did a high-level scoping review, which involved media scans, government document review, and reaching out to experts in Australia, New Zealand, Canada, the United Kingdom, Europe, and the United States. As part of this research, we identified 61 systems that had been cancelled. Our research ended in 2021. I'm sure we'd find a lot more now. And the first part of this report is this high-level overview of the cancelled systems. We also did case study investigations. And so, for our case study investigations, we did interviews with community organisations, with government workers, with lawyers, with activists, and with affected communities, and we compared. We did case study investigations, we had three sets of case studies. We looked at systems that had been cancelled in the area of policing and justice, in the area of child welfare and in the area of benefits administration. And those much more detailed case study investigations are available in the second section of this report. Overall, we found that there were a wide range of factors that led systems to be cancelled and that a number of these factors often happen in concert.



## Factors influencing decisions to cancel

Civil society critique or protest	26
Critical media investigation	24
Legal action	19
Government concern - privacy, fairness, bias, discrimination	13
Critical government review	12
Political intervention	8
Government decision - procurement, ownership	6
Other	5
Corporate decision to cancel availability of system	3

And for our purposes today, I really just want to highlight the fact that we found a lot of civil society critique and protest in connection with systems that had been cancelled, as well as critical media investigation and legal action. One of our most significant findings is that often, when there is information available about how a system is being used, whether it's through media coverage or whether it's through civil society actions, this then raises concern about the implications, which then can lead to, in a number of cases we found, legal action on different legal basis because often the legal basis upon which these systems are used is not entirely justified and there are different reasons for systems being cancelled because they have been successfully legally challenged. We also found systems cancelled as a result of government concern related to a range of factors as well as critical government review.

So as mentioned, these are the case studies that we looked at. I want to focus on the findings from our comparative analysis of child welfare systems that were cancelled. And one of the things I wanted to highlight here is that across the systems that were cancelled, we found key areas of concern that relate to the uses of predictive analytics for risk assessment in child welfare. One area is a lack of transparency surrounding these applications, which makes it difficult for them to be interrogated and challenged by users and people who are affected, concerns about the ways in which these systems can reinforce and exacerbate discrimination, that these systems are often inaccurate and unreliable, that these inaccuracies can lead to stigmatisation of parents, and that the implications of this are not being considered enough and that the data informing the development of these systems is often limited and

problematic. The research that has been done on these systems demonstrate that there are a range of actors who have been involved in problematising the use of predictive analytics for child welfare. The range of people involved in challenging or problematising their use include government workers as well as those within civil society. And that there are particular political economic factors at play when a decision is made to implement the uses of predictive analytics for child welfare to risk assess families. And these political, economic, contextual factors really matter and they're not talked about enough. For example, in the UK context and the system that we looked at, we identified that the austerity context was a major driving factor in the decision to implement predictive analytics, which is that we saw local authorities in the United Kingdom at the time that we were investigating, seeing their budgets reduced by a third, which led to increasing pressure on those on the front line to do more with less. We saw pressures to combine data and in order to identify families who are troubled in order to try to find ways of getting those families more services while at the same time we saw vendors marketing the development and use of their tools as a way to address the significant constraints that local authorities are playing. But at the same time, the public sector - local agencies who are being targeted with tools for sale, are not adequately being resourced in order to have the time and research capabilities required to really challenge the marketing of these tools.

We see differing levels of transparency involved. And we see ongoing concerns being raised by communities that aren't being responded to adequately. And there is in fact very limited opportunity for the community subjected to these tools to raise concerns about them and to challenge them. In short, our cancelled systems research pointed to a range of things that could be done differently going forward.

## Recommendations

- ✓ Create and maintain **Public Registries**
- ✓ Resource public organisations including regulators to support **greater transparency and accountability**
- ✓ Enhance **procurement support**
- ✓ Require **Impact Assessments** and recognise the need to **address systemic injustice**
- ✓ Review the **legality of uses** of automated systems
- ✓ Shift the **burden of proof** required to implement a ADS
- ✓ Engage the **public**
- ✓ Understand the **"No Go"** areas
- ✓ Take **responsibility in accounting for ADS history**
- ✓ Ensure a **politics of care** approach

One would be to create and maintain public registries. Another would be to resource public organisations, including regulators, to support greater transparency and accountability about whether, if, and how these tools should be used. This resourcing would include procurement support to ensure that local authorities have the capability to do the interrogation that they need to make sure that tools are safe and ethical. We argue that impact assessments need to address systemic injustice. A lot of the data that exists within some of these systems, whether we're talking about the criminal justice system, we're talking about child welfare, even if we're talking about the education system, there are long-standing histories of systemic injustice that exists within those data sets. And so, we are arguing that any sort of attempt to use these data sets to develop models to predict the future need to employ and kind of address systemic injustice and also we need to ensure that there is auditing across the lifespan of these systems that includes post-deployment auditing. There needs to be a review of the legality of these systems, the burden of proof - and this is something that Michael Sanders and Vicki Clayton found in their [piloting of predictive analytics](#). And one of the things that they argue, which is really important, which is that any organisation that wants to make use of predictive analytics in such a high-stake context has to prove that their systems are both effective and that they will not cause harm. There needs to be widespread and more meaningful engagement of the public through such things as extended peer review. There needs to be a responsibility to recognise that there are no-go areas where the risk of harm is so great that communities may decide that they don't want uses of these tools. And there needs to be, particularly in this case, where there has been so much research that has been done on this area, there

needs to be a responsibility. If you are going to develop and make use of a tool in a new context, then there should be a requirement to take into account the history of the similar systems that have caused harm and have been challenged as unjust in other contexts. So I will stop there and I will hand it over to Shion.

SG

## Shion Guha

Thank you, Joanna. So, as I pointed out, I'm a computer scientist, so I know very little about social work. But in the past decade or so, I've worked with the [Wisconsin Department of Children and Families](#) in great detail, studying some of the systems that they had put into place. And then in the past couple of years after I moved here, where I'm currently at in Toronto, Ontario, Canada. I've been kind of getting pretty deeply embedded with the various children's aid societies in Ontario that are the arm of the provincial government that administer and provide child welfare services. So, I'm going to talk about a couple of things today. Number one is that I am sure all of you have faced or are increasingly facing pressures from the powers that be to implement different types of predictive systems, or possibly in the last couple of years, AI systems into your day-to-day decision making. The problem is we lack a fundamental understanding of whether or not these systems actually help us make better decisions.

The image shows a screenshot of a research article page. At the top left, there are logos for ACM Digital Library, Association for Computing Machinery, and ACM Open. Below these is the DOI link: <https://dl.acm.org/doi/10.1145/3476089>. On the right, there is a PDF download icon and statistics: PDF Download 3476089.pdf, 10 February 2026, Total Citations: %, Total Downloads: 4314. The article title is "A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare". The authors listed are Devansh Saxena (Marquette University, Milwaukee, WI, United States), Karla A. Badillo-Urquiola (University of Central Florida, Orlando, FL, United States), Pamela J. Wisniewski (University of Central Florida, Orlando, FL, United States), and Shion Guha (University of Toronto, Toronto, ON, Canada). The article was published on 18 October 2021. There is a QR code on the right side of the page. At the bottom left, it says "Open Access Support provided by: Marquette University, University of Central Florida, University of Toronto".

So this slide that you see here is actually the result of a very deep long-term ethnography that I conducted with my collaborators in Wisconsin, where we aimed to study exactly how people make decisions if they are confronted with algorithmic predictions in front of them. And it's largely based on a few things. What we realised is that, and that's changing a

little bit with AI, but we realised that largely in North America, the various algorithmic systems that are present aim to predict things like risk. And the problem is that risk is something that different people have different understandings of, and we do not often agree on what exactly risk is and how we should measure risk and whether or not the algorithms that are claiming to measure risk actually measure risk. And this is true even within child welfare departments with experienced social workers. I'm pretty sure that if I took a straw poll around the room and asked folks about, what do you think is risk of neglect, we would get like 15 different definitions. Or if I said, hey, what is risk of placement instability? I'm sure many of you would give me very different definitions. And so, the issue is that if we have predictive systems that are trying to predict risk, and then we take downstream decisions based on that understanding of risk, then we need to understand what exactly this risk situation is. So, this is kind of like the two things that I want to touch on a little bit.

I wanted to foreground a little bit of my students here, but the Human Rights Commission in Ontario actually criticised risk assessment standards and tools in child welfare, both the vanilla paper ones that we've had for the longest time, as well as the newest algorithmic systems that they have in many children's aid societies in Ontario that are already in deployment and they don't really work very well. Within Ontario and Canada, we've had significant criticism from all sides about how we've been failing our children. Obviously, a lot of these issues are structural, and we could do with fixing the structure itself, but who cares about that? Let's tinker with technologies instead.

The image shows a screenshot of a research article page. At the top left, there are logos for ACM Digital Library, Crossmark, and ACM Open. Below these, it says 'Latest updates: <https://dl.acm.org/doi/10.1145/3544548.3581308>'. The article title is 'Rethinking "Risk" in Algorithmic Systems Through A Computational Narrative Analysis of Casenotes in Child-Welfare'. The authors listed are Devansh Saxena, Erina Seh-Young Moon, Aryan Chaurasia, Yixin Guan, and Shion Guha, with their respective affiliations. On the right side, there is a PDF download icon and statistics: 'PDF Download 3544548.3581308.pdf 10 February 2023 Total Citations: 23 Total Downloads: 1,184'. Below that is a trophy icon and 'Published: 19 April 2023'. At the bottom right, there is a QR code and citation information: 'CHI '23: CHI Conference on Human Factors in Computing Systems April 23–28, 2023 Hamburg, Germany'. The page is framed by a green border.

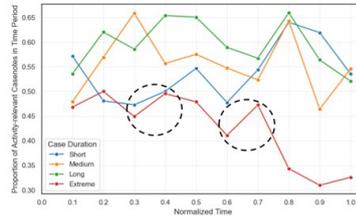
So, this is where we kind of come to the modern day, right? So, we've kind of understood how people make decisions from these risk assessments, but this is the actual reality of what people are being asked to do. So many children's aid societies in Canada have been asked to use AI. And by AI, what they really mean is, could you please use some enterprise version of ChatGPT or Co-pilot in order to make decisions about children? Yes, it might sound horrifying to this particular audience, or maybe not. Maybe many of you have been confronted with this as well. So, we wanted to do a study to see if there are algorithmic systems that can summarise information from both administrative data as well as case narrative data and then present them to caseworkers. Does that actually work? Are those AI systems actually doing the things that they're doing? And I'm obviously abstracting a lot of the technical stuff that's going on behind the scenes, but I'll just make one note; that we have discovered various mathematical limits to predictions that actually prevent us from making strong claims like this predictive system will solve this problem. Actually, there are mathematical limits to that, and we should not be making such claims, even though there are many startups and other companies nowadays who are trying to make such claims.

So, we did this study, and here's the thing. Child welfare agencies want to ensure that workers are providing effective services for families, right? But what happens is, and this might be very familiar to all of you, in child welfare, we often deal with different kinds of cases. And a lot of the times we care about case duration, and case complexity and case duration are often correlated with each other. So, the shortest cases are often the cases that are easiest for workers to deal with, whereas the long and extreme length cases are often the most complicated. But most of the time, social workers do a fantastic job with short cases, medium cases, even most long cases. Most of the times, you all have developed amazing heuristics, lived knowledge and experience to deal with all of this. So, if AI systems are really good, shouldn't they be trying to help us solve the problems that we cannot solve right now with our own heuristics and lived experiences? Well, that's what we really wanted to study.

# Supporting CW workers in Canada

(Under review)

- Divided casenotes by case duration
- Extremely long term cases make fewer references to Service Plan Goals compared to shorter term cases
- Black circles denote when new Service Plans are drawn up for a case in Extreme cases



Please read, reuse and cite: <https://tinyurl.com/simons-ucberkeley> | [shion.guha@utoronto.ca](mailto:shion.guha@utoronto.ca) | <https://www.hcds-uoft.ca>



I will abstract this figure for you. I just want you to take a look at that red line. The red line is the most important and interesting line here. The red line points to the most extreme length and complex cases, and this is actual data from Ontario. And we're studying the performance of AI tools as they exist now. And what we really see, what is the big takeaway here that I'm hoping we could discuss and that you could take away from this, is that current AI tools are absolutely not good at solving the cases that we really want them to solve. We don't want them to give us amazing recommendations and suggestions for short cases or easy complexity cases, because those are the ones that we already know how to deal with. Shouldn't AI be trying to solve the problems that we don't know?

# Supporting CW workers in Canada

(Under review)

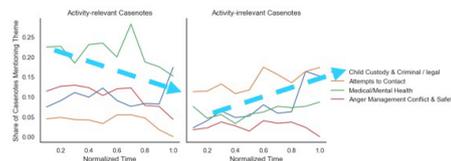
## Manual labeling vs. LLM

- LLMs struggle at identifying Service Plan goal relevant Regular Casenotes as cases become more complex and uncertainty increases

## Tracking thematic Service Plan goal relevance using BERTopic

- Caseworkers are working on new child welfare concerns that emerge in Extremely long cases that are not outlined Service Plans (right plot)

Case Duration	Agreement	Cohen's $\kappa$	FPR	FNIR
Short	0.804	0.604 [0.535, 0.670]	0.265	0.135
Medium	0.782	0.550 [0.493, 0.602]	0.324	0.135
Long	0.727	0.402 [0.352, 0.447]	0.470	0.146
Extreme	0.730	0.470 [0.440, 0.500]	0.382	0.134



Please read, reuse and cite: <https://tinyurl.com/simons-ucberkeley> | [shion.guha@utoronto.ca](mailto:shion.guha@utoronto.ca) | <https://www.hcds-uoft.ca>



So, here's some scary figures. But essentially, what I want to point out over here is that we have established some computational limits, some actual technical limits, to what current AI systems can or cannot do. And

this is published research. So, if you are getting pressure from other folks, 'Hey, please implement AI, it'll help you solve the cases that you can't currently solve'. Please do point them towards this particular result because we have established that current AI tools can't do that. And unless we significantly rethink how AI tools work, which we aren't really doing right now, this limit, this computational limit will continue to be in case.

## MAIN TAKEAWAYS

Risk models often learn the intervention system. not latent harm.

- Administrative outcomes are often **system contact labels**, produced by human discretion and bureaucracy.
- Narrative records surface **intervention work, constraints and power** which predictive risk models typically erase
- The meaning of risk is **temporal and process dependant**; static risk scoring collapses actual human dynamics into a number.
- LLMs do not solve this in long, complex cases; **reliability declines as uncertainty and case entropy rise**.
- When prediction ignores intervention, it becomes punishment.

Please read, reuse and cite: <https://tinyurl.com/simons-ucberkeley> | [shion.guha@utoronto.ca](mailto:shion.guha@utoronto.ca) | <https://www.hcds-uoft.ca>



My point here is very simple. Hopefully, these are some easy takeaways. Right now, we have two different types of algorithmic tools. We have predictive risk models that are trying to make predictions, things like risk of neglect, maltreatment, placement instability, so on and so forth. The things that social workers really care about. And then you also have AI tools that are saying, okay, feed us data and we are going to give you all kinds of summarised information that will hopefully help you take the right decisions. The problem is that these types of models often learn intervention systems and not really the latent harms, so they don't actually surface the latent harms; they surface something else, and when you combine things like administrative records and narrative records, what you really get are these tools often surface things like constraints and power relationships which are very useful. I'm not saying that we shouldn't be using these kinds of systems. I'm not a Luddite. But I think that the ways in which these current systems are designed and implemented don't really work for the ways in which we want them to within Social Work. And finally, one big takeaway is that risk is not a static construct. It's not a singular understanding at one period of time. It's temporal and process dependent, right? Risk changes over time. So, if

you have a predictive algorithm or some type of AI tool that is flagging something as high risk, medium risk, low risk, it's only doing that at particular snapshots. And that's not very accurate, because what it does is it gives you an average outcome, and most of the times in child welfare, the ways in which we've developed our heuristics, we don't care about average outcomes, and that's a problem. And finally, LLMs, large language models, ChatGPT, Copilot, Gemini, Perplexity, all of these tools that we have, they do not solve this. Particularly in long, complex cases, their reliability significantly declines. This is the computational and mathematical limit, which currently we do not have a solution for. So, I will stop here. Hopefully I have kind of presented a more technical kind of understanding about how these types of tools, while very useful, currently don't really solve the things that we want them to. Thank you.

**CB** **Claire Burns**

That's a bit of a mike down moment, really. So thank you so much to you and Joanna. Just such interesting inputs about AI, but so much of what you said as well, I think, reflects some of the real challenges around implementation of any sort of activity in social care, about, are we doing the readiness work? Are leaders getting the time to think about how they're stepping in, what they're procuring? So, if I could just put some questions to you. I suppose what I've picked up listening to both of the presentations is that there is some risk in here and we have to be careful about it. But I think, Rhema, from you, what I also heard was, it's not that we are taking every element of a social worker's skill and wanting to replace it, but actually how can we have elements of AI that help the assessment process. So, the assessment process, as you're saying, Shion, is very nuanced, it has to be contextualised, it has to be seen within what we think is acceptable at any time. But are there elements of this that actually can help us do our job better that doesn't replace all of our assessment skills, I think, is what I'm hearing.

So, suppose just coming back to Emily and Rhema first, and then I'll come back to Joanna and Shion, you can understand, predict that some of what people are saying is, what is the experience of challenging bias? What is the issue of bias within this that that may take workers to a place where they would look back and say, well, actually, that influenced my assessment skills in an inappropriate way. So, I wonder if I could come to you two firstly, what elements of bias have you found in this? And are there ways of challenging it?

**RV****Rhema Vaithianathan**

So in terms of bias, for sure, there are some particular areas where workers are realising and one is sexual abuse. The tool is a little bit - there's a lot more nuance around the assessment of sexual abuse because of the lack of history often that people who have that as an allegation have. In terms of race, so we have had randomised controlled trials of our tools done by independent researchers, so not our research team, and they have consistently found a reduction in the racial disparities of decision making. Essentially, what we find is that when people are told, when workers are told that black families or families of colour are also low risk in the tool, which you know the majority are, they then change their decisions because without the tool, they tend to conflate race with risk. That's the evidence from the randomised control trials that essentially works off of reversing some kinds of human biases that people might have in their heads, with more structured information that they have at their fingertips that helps them reduce particularly racial bias. I think we probably gave you links to that randomised control trial which happened in Colorado. Emily, I don't know if you have any tips.

**RV****Claire Burns**

Yeah, thanks. And it's interesting because in our previous webinar, that was one of the things that was flagged as well, was that sometimes when we're talking about the use of AI and stigma or discrimination, we actually can forget that that happens between individuals as well. So that those lenses are often there with workers as well, you know, so it's not to say that it's not a polarised thing. So, Emily, is there anything you'd want to add in there?

**PE****Putnam-Hornstein, Emily**

No, that's exactly what I was going to just mention. I think Rhema and I are always careful to acknowledge that there's of course going to be bias that's baked into the data that the model's looking at. And then there's also bias that continues to operate in our human decision makers. One way that I have found it helpful to think about this is what are we trying to do with these predictive risk models? We're trying to hopefully change the trajectory of the family, moving forward. So, if the model is helping us see which children we have historically either overserved or underserved, the idea is that we're doing something different and preventative for them. So, Shion, I would agree with you completely. A lot of the details

here have to do with how are these models being used and applied in a way that serve the workforce and serve the families they're trying to serve. And that's where I think that kind of human algorithm interaction is very critical to all deployments.

**CB** **Claire Burns**

Joanna, anything you want to add here?

**CB** **Joanna Redden**

Yeah, I think there's a few thoughts to add. So, in terms of bias and discrimination, we know that and, I think as Rhema and Emily mentioned, you know, humans, of course, can be biased too, right? But I think one of the things that we've seen in previous research and concerns expressed around uses of AI is that because essentially, one of the goals is to make systems faster and more efficient, there's a real potential, and this has happened in different cases, of accelerating and speeding up bias and discrimination. So, I'm thinking in particular of applications in different areas like policing and fraud detection. There's concern too, I think, about the way in which systems are deployed in that some of the social workers that we were talking to in the research that we did in one particular local authority who were being asked and consulted in the development of uses of predictive analytics, one of the things that the social workers in this particular local authority said was that the use of predictive analytics for risk assessment didn't really help them identify what would be helpful in terms of strength-based development. And so, they were finding and they were raising concerns about how the tool that they were being asked to trial wasn't really helping them with positive reinforcement. And then I think one of the things I always want to talk about in these kinds of situations now, because it's not talked about enough, is that we have to remember that we're not talking anymore about one type of system in use, right? So, we're talking now about uses of AI across areas that are going to affect families in different ways. So, we have AI systems being used in hiring. We have AI systems being used for benefits administration. We have AI systems being used in healthcare. And we have a quite extensive body of research now documenting that marginalised communities are disproportionately negatively affected by all these kinds of systems, right? And so, we have to be thinking about compounding effects. And so, any investigation of the impact of these systems should also be thinking about and doing impact assessments that take into consideration that this kind of this aspect of compounding effect.

**JR** **Joanna Redden**

Shion, did you have anything?

**SG** **Shion Guha**

Yeah, I'll add one thing. So, in these conversations, we almost always talk about bias originating from data and bias originating from human beings. There's one more critical point of bias that nobody talks about because it's inconvenient to do so. And that is that every single model has a particular bias associated with it. So, the choice of what the underlying model, statistical or otherwise is, is also deeply biased. Different models do different things. So, you know, if I make a flippant comment like, oh, I have a model that predicts risk of neglect in this particular jurisdiction. And you have a particular model that does that. And both these two models are statistically different, but they're trying to accomplish the same thing. There's a particular bias in the mathematics of the model. This is something nobody really wants to talk about, particularly big tech never wants to talk about it because it's extremely inconvenient to do so. Then you'll almost always find the conversation veering towards, 'Oh no, our data is so terrible', which is true. Our data is terrible, but it doesn't mean that our models are also not terrible.

**CB** **Claire Burns**

Yeah, thanks very much. And again, I think it comes back to... As again, I think you've given us an appropriate level of caution around how we use these things. But again, Rhema just coming back, it was interesting what you were saying around, how does it help us to curate information, things that can take workers a long time to find out actually how can we pull information together that then enables our assessment, you know, there's still that assessment element in it as well. I was also struck by the parallels of both your contributions about the level of readiness and buy-in that's often required and the kind of pre-work that's required. Joanna, I was just thinking about your presentation around if leaders were deciding to step in, all of the things that they would need to consider, around procurement. I think it was really interesting, seeing how many organisations have got their fingers burnt already about decisions being made really quickly about an organisation and this is going to be the silver bullet, and it never is. So, I'm going to ask you to come and certainly just to talk about again, what should our leaders be thinking about and considering and what does that buy-in process look like? But

Rhema and Emily, I was also just thinking about what was the process of enabling buy-in with the workforces? Did you feel that that's always sufficiently done? And also with communities, to what extent are communities aware that you're using this tool or do you think it's important? So, I'm going to start with you around readiness and buy-in from workforces and then go back to Joanna in terms of what would be her lessons around strategic leaders.

**CB Putnam-Hornstein, Emily**

I'm happy to talk a little bit about LA. And just briefly, Claire, I think that one of the reasons our tools have been embraced and useful and impactful, which we can clearly see in the data, is that they truly were co-designed with the agencies we worked with. And so there was a significant... I mean, we must have had more than two dozen different meetings with community stakeholders to surface potential concerns to clarify very explicitly how this model was going to be used and how it was not going to be used. We met with the union, we met with workers, and then what actually happened is once the tool was deployed, many of the staff in the three pilot offices that were using it became the biggest champions. And so, they were the ones who were helping us, as we rolled it out into additional offices, to explain how it was supporting their practice, their supervision of cases, their work with children and families in a very positive way. So, you know, I know that there was a mention of the Eckerd College and other tools that have been dropped. I will say that Rhema and I are very proud that every jurisdiction we have partnered with, because I think it's been a true partnership, are continuing to use and expand the use of the tools that we have built for them and with them.

**RV Rhema Vaithianathan**

Just to be aware that the systems we're going into, they already have things, often paper and pen, that they call risk and safety assessments. So, this is not a new thing and the workers hate those because essentially I have shadowed workers doing structured decision-making risk assessments, they have to go onto their screen and count things, count how many neglect calls have come in and put that onto a piece of paper. So, the fact that we just go in and grab it and give it to them in two seconds, they love it. So given that those, they've seen a legacy of tools and systems that no one has for 20 years ever questioned with the biases in these, the efficacy of these, the accuracy of these, not a single

question. And so we found a lot more engagement because of that. Whereas I guess in some of the counties you're talking about in the UK, maybe they're not already using those tools. So, there might not be the same readiness for something easier.

**CB** **Claire Burns**

Yeah, thanks, Rhema. Joanna, I don't know if there's anything else you want to say, that was really comprehensive view you gave us about all of the things that needed to be considered before you step in. And I think a lot of our workforces will have experienced things just happening to them, things being introduced, you're now using this system with very little kind of co-production or buy-in. So, I thought just maybe if you could expand a little bit about, you know, if we've got leaders who are thinking about stepping into, well, how can we use some of these, can we use AI? What are the range of things you'd want them to consider before stepping in to choose something?

**JR** **Joanna Redden**

Yeah, I think one thing would be the histories of these systems. I think that's really important. So for anyone thinking about pursuing the use of these systems, they would have opportunities that we haven't had. So, it can be very difficult to get insider views from those who have tried and applied these systems in different places, particularly if they've been controversial, if they've been cancelled. But, I think probably that anyone working within government who wanted to really consider whether or not they wanted to pursue these systems, they could call up other government agencies in different countries and find out what those experiences have been. I think there's been a lot of research done on this topic. One of the first things I would do if I knew that my local authority was being pushed to look into some of this, I would forward on that work by the What Works Centre for Children's Social Care. That was really fantastic, detailed piloting work that was done there. And it's all documented. I think that one of the areas that I think requires consideration is how and in what ways the public is going to be informed and engaged. What methods will be provided for public engagement? Because we know this is going to be controversial. And there needs to be a means for meaningful public consultation. There needs to be a means for independent review. So, in the cancelled systems research that we found, some of the best practises we identified had to do with government agencies pre-emptively engaging independent reviewers and

then making those reviews public. And then I think one of the things that we didn't talk about in the report, but that is now really front of mind for us, is function creep. So, one of the things that we've seen, for example, in some of the systems that we looked at that were cancelled around the use of predictive analytics to identify children who might be likely of exploitation, as we mentioned in our kind of pre-conversation setup this morning. Now, that same system, it was cancelled to identify children at risk of sexual exploitation, and it's now been repurposed to identify children who are likely to become criminalised, which is leading to children being excluded from school wrongly. And so, I think we have to really be thinking about it, and one of the things I would encourage are what kind of protective measures are you going to put in place to ensure the future of any sort of data sharing arrangements that you enable. And one final thing I'll say is that data sharing and giving people access to information in a timely fashion is hugely important. And you can do that without the predictive analytics piece, right? So you can give social workers the ability through different kinds of tools to share information and access information in real time. That doesn't necessarily mean you have to attach a risk assessment to it. And we know, for example, through research that was done in Manchester, that that has proven really helpful for social workers in Manchester. Shion, did you have anything to add?

**SG** **Shion Guha**

Yeah, I think one of the things that I have learned about social workers, not being a social worker myself, is that social workers get extremely annoyed if they feel like their decision-making agency has been taken away from them. And I think that creates a divide between top level bureaucrats, administrators and politicians who look at predictive systems or algorithmic systems in general as a means of doing two things. Number one, achieving more efficiency in providing services. Number two, reducing costs. And I will point out that there has been absolutely no evidence anywhere in the history of public administration research that any of these two premises that leadership cares about is correct. So that's like a top-level understanding, right? Like why are these tools coming down? And then when they hit frontline workers, frontline workers are like, but hang on, we would like some tools, as Rhema and Emily pointed out, we would like some tools, but we don't want these tools. Why these tools? And what we're seeing with the current AI summer that's going on is that whole cycle once again, where leadership wants some type of AI

usage because there is a feeling that AI will solve most of our problems, many of our problems. But again, there's absolutely no evidence that it does or it can. And in fact, in some of the work that I talked about, which is, to my knowledge, some of that first steps in that research process, we've seen that it doesn't exactly do what caseworkers want. And there was a question in chat. Can it realise this and reduce critical thinking on the behalf of social workers? And the answer is absolutely yes. And you know, the four of us are faculty. We see this in our students all the time. Students, undergraduate students constantly use AI. And I feel like we stand on the same front when we say that. Yes, it is, you know, we're seeing this in students, that over-reliance on AI actually reduces those critical decision-making capabilities. So, there are some dangers in how the collision of algorithmic systems kind of affects how caseworkers actually do their job. So once again, I want to stand on the side and say that I absolutely do not advocate for saying things like, oh, we should never have any systems. That's not what I'm saying at all. What I'm saying is that we need to carefully understand exactly where to deploy the system. And currently, what happens from a top-down perspective is leadership sees AI as this big hammer, and then everything else is a nail, and then we're going to use the AI hammer to beat all of the nails. And that's not the right way of doing it, right? So as a takeaway, so if it helps anybody in the audience make arguments to whoever is trying to compel you to use these systems. There are ways, there are things that you can do, but maybe that's going to clash with administration's viewpoint of how these tools ought to be used. So, there is that participatory co-design process that I think Emily talked about, that these conflicting value tensions between what different stakeholders want and need out of tools can be resolved, should be resolved. So I'll just stop there.

SG

### **Claire Burns**

That's great. You've answered about four questions there, Shion, so that's brilliant. Four questions that I was going to put, because people have talked a lot and I just want to acknowledge that, the questions around in what way it can erode core social work, core social work skills. And again, Joanna, I'm just really struck by you saying the amount of work that's required to upfront this to make it work right. And I think, and your question about your note about what has the driver here been really, been really important as well. I mean, we're all in a context where, you know, we've got a workforce crisis, we've got increased levels of need, and about how that might drive some of what drives some of this. I just

want to come back to Emily and Rhema because there is quite a lot of questions as well just around the benefits of the tool that you're using and I think particularly people are saying one of the big challenges that we have, but our evidence tells us this is an international challenge, is that collating information between services. So, can you say a little bit about what might some of the challenges there have been about pulling information together across education and health? You talked a bit about and how the tools help to overcome some of that.

PE

**Putnam-Hornstein, Emily**

Yeah, so the way that we have approached that in the jurisdictions in which we have worked is we've always started with the rule that the algorithm should not see any cross-sector data that the community has not already decided frontline child welfare caseworkers can access and see. So, what that means is in a jurisdiction in Pennsylvania where they had an integrated data warehouse and the community had already sanctioned that that information should feature into child protection decisions, the algorithm is built from that data warehouse. In a jurisdiction like Los Angeles, they did not have existing integrated data so the algorithm is built purely off of their child welfare or child protection data. So, I just wanted to kind of lead with that as a place to start. I think that people often assume that you need cross-sector data to build good, accurate models. That's an empirical question. We have certainly found that it is not always the case that more data from other systems improves the accuracy or the nature of the decisions that our child protection workers need to make. But I also just wanted to respond to something that Joanna said. I think, Joanna, you said, we can have data that's linked and comes over, and we don't need to use an algorithm or a risk tool to kind of support the analysis of that. And I will say that I have spoken directly with states and child welfare workers who are truly resistant to having all sorts of cross-sector data come over if they are not also given the tools to help them manage and know how to use it, because it can actually become a liability if I, as a child protection worker, am presented with information from a criminal justice system, without any parameters for how that historical data should feature into my present day thinking. So, I would actually argue that these tools can bring structure and transparency and broader community conversations around the use of cross-sector data that do serve children and families. So, I will see if Rhema has anything to add there, but...

**RV****Rhema Vaithianathan**

Yeah, I mean, I don't know if people know this phenomena of a day of information overload. So, if you plot decision quality on the vertical axis and the amount of information a person has on the horizontal, it's an inverted U shape in the sense that sometimes the more information you give people, the worse their decisions are. Because when you give people too much information, what they start doing is using heuristics of mental shortcuts. Different people choose different five pieces of information they always look at because otherwise they just get overwhelmed by the quantity of information. That's why predictive analytics sitting alongside a curated set of information, I think is the sweet spot. The predictive analytic tool takes, ingests hundreds of pieces of information and then curates those hundreds of pieces into a few pieces that a person can start their questioning and then go into the case system to see more. So, I understand where you're coming from when you say you should just give them the data. But I think it's just not humanly possible to process so much information.

**CB****Claire Burns**

Joanna, is there anything you want to respond to there or do you just?

**JR****Joanna Redden**

Yeah, sure, yeah, no, I was speaking about a very particular database. It was used in Manchester when we were doing our case study investigations. I'm not sure if it's still in use, but it was an integrated database. It was very controlled, so it wasn't, Emily, pulling in everything. I mean there was very specific legal constraints on what information can and cannot be shared. It was specifically designed to, at the time, to be operating within the legal limits and to enable caseworkers to have access in real time to particular kinds of data, not all the data, but particular kinds of data that they were finding they were needing in high crisis situations, right? And so in this case, it was an information sharing tool. And so also, to your point, Rhema, it wasn't a tool that was meant to provide them with tonnes of information. And you know, it was a kind of a real time tool meant to capture what was happening with someone who was presenting or there were concerns about high risk in a crisis situation. And before what was happening is that case workers were having to get on the phone to call the different bodies in order to get access to information that might have been documented the day before

or, you know, two days before or something like this. So, this was preventing some of that real time consuming work of getting on the phone and calling different, different busy people to get access to information. So, it was a specific kind of data sharing tool.

**PE Claire Burns**

Thanks very much. Emily, I wonder if I could just - there's a couple of really technical things that people have asked about the tool and then I'll just ask everybody to have just a kind of last message that they would want people to go away with before they've finished. But Emily, so just questions around who owns the product? And do the tools require all notes to be electronic records as opposed to written notes? And someone's also asked are you also able therefore to identify that children and young people who you haven't intervened, you haven't needed to intervene with and that, you know, and that's been legitimate. So sorry, those are about three different questions.

**PE Putnam-Hornstein, Emily**

So, I gave the example of Los Angeles using their model to alert supervisors when they have very kind of complex investigations. So, there's more support and supervision provided. We are now also trialling a use of the model to help identify when calls have made it through the screening process. But if we slow down and review the history, the model thinks that actually this is not a family we are likely to see in the future. And so, there is a team meeting and there's explorations of ways that we could safely reduce the footprint of the public agency by diverting those to community partners. So that is something that is being tested. There was a question around kind of paper versus digitised. Yes, we need digitised data. But right now, in LA, we're just using their structured data elements because the county does not have access to all of the case narrative data. So that's where we started. And then the question about who owns the model. When we come in, we feel strongly that these data were generated through a public system in support of child safety goals. So, we come in, we build the model. We do not claim intellectual property rights to that model itself - it belongs to the agency with which we are working. So, Rhema, I don't know if you want to add anything, but those are my quick three answers, Claire.

**RV Rhema Vaithianathan**

No, I've got nothing to add.

**CB** **Claire Burns**

Yeah, okay, thanks very much. But all just all really interested about the tool that you're using. So, I suppose you've got one minute each if there's any other and the summing up any other message that you would like to give us. So, Shion, would you like to start - a final point that you would like to emphasise?

**SG** **Shion Guha**

No, thank you very much for inviting me to this. This was fantastic. And if anybody in the call would like to reach out, please feel free to do so. I'm always very happy to talk about these things. Thank you.

**CB** **Claire Burns**

Thanks, thanks so much, Joanna.

**JR** **Joanna Redden**

No, nothing for me. I want to echo Shion's thanks. And I also just want to give a shout out to you as the organisers for these webinars. This is such a great idea. It's so important to get different perspectives and to have a chance to talk and share ideas on these really high stakes issues. So, thanks so much to you as the organisers for organising this event.

**CB** **Claire Burns**

Yeah, thanks so much, Joanna. I'll come back to that. Emily, anything that you would like to?

**PE** **Putnam-Hornstein, Emily**

I will just briefly echo the thanks and say I think the most important thing to kind of consider when exploring these tools is always the compared to what. Rhema and I are not positioning these as crystal balls. They will not solve all agency problems. But I think that even a cursory look at tools and decision-making processes as they exist today compared to what's possible with these analytic tools. In my mind, it's very clearly a no-brainer that agencies should be moving there and workers should be supportive. Thank you.

**CB** **Claire Burns**

Okay, thanks, Emily. And I think Joanna's point that we need to learn a lot from what has happened in the past, but we need to continue to learn

from projects like yourself where you are actually stepping into what is a risky place to step in. And you know, we've got so much to learn from that as well. So, Rhema, anything from you?

RV

**Rhema Vaithianathan**

I just would encourage people to think of use cases where you improve the quality of services we give. You as leaders are going to come under pressure to adopt tools just to reduce costs. And I would really urge people to look and ask, if that's all we get out of this technology, I'd quit today. What we can do with this technology is really help improve quality of the experience for our families and the quality of casework that our workers do. And that's where I think we need to lean into this technology.

CB

**Claire Burns**

I think that's a superb place to finish because I think a lot of that aligns with what Joanna was saying, but I think where we all are, we've had a lot of differences of opinion today, which is great because it really makes us think about everything that we need to consider. But you've finished us on that ultimately this has to be about better support for children and families. That has to be the ultimate goal if, we can all hold on to that. So can I just say what a fantastic hour and a half listening to you all. Thank you so much. Again, and we hope we'll keep the connections going. Thanks.