Do Risk Assessment Tools Make Racial Disparities in Child Welfare Better or Worse?

Los Angeles County and a team of scholars are collaborating on a "Risk Stratification Model" – an example of 'predictive analytics', which uses complex algorithms to prioritize risk in child maltreatment reports.

Algorithms are often viewed skeptically by Progressives and leaders of BIPOC communities as potential contributors to "heightened surveillance", meaning practices that disproportionately screen families of color into child protection.

The LA County project addresses this concern with an oversight group focused on bias, and by piloting the approach in predominantly Black communities. This Sounding Board article describes a related project in Allegheny County, Pennsylvania which positively impacted racial disparities by eliminating bias in 98% of screening decisions.

This tool may also help optimize limited resources, compensate for inexperienced workforces, and prevent families from getting deep into the system.

The challenge is that this is being implemented without additional resources.

Transcript of podcast on LA County Risk Stratification Model

Predictive analytics, broadly defined, is the use of data to stratify child welfare caseloads, including as they enter the system, into the highest risk cases which should get priority for casework and services, and lower risk cases which should be theoretically addressed with less urgency and potentially referred to community agencies for services. The risk stratification model was developed by scholars Emily Putnam-Hornstein and Rhema Vaithianathan, and in its current iteration is being implemented through an extensive and close collaboration with Los Angeles County, community leaders, Mathematica and other partners. There is also a current blog on this topic in *Sounding Board*, a monthly article by Dee Wilson which is usually published in the online newspaper The Imprint. In that article Wilson focuses on the use of the algorithm in Allegheny County using a model also developed by Putnam-Hornstein and Vaithianathan.

As mentioned in the blog, Progressives and activists, particularly in Black, Indigenous and Persons of Color, or BIPOC communities, have been wary of predictive analytics because they see it as potentially just one more way to identify BIPOC families and drag them into child protection. You can find one critique of this in an article called "Algorithmic Justice in Child Protection: Statistical Fairness, Social Justice and the Implications for Practice" by Emily Keddell in the journal Social Science, which I again put a link to it in the script for this podcast or you can just look at up on the Internet.

A key point that is made in this and other critiques of predictive analytics is that it is very difficult to do well. People who collect the data have to be well-trained so they are consistent in gathering the information that becomes the input for the system. Reporting needs to be consistent across racial and ethnic groups so that the algorithm doesn't reproduce under- and over-reporting. Algorithms need to be understandable to communities to increase their trust of the data, and to the people who have to use them so workers can detect when the output may not make sense.

These of course are typical management problems, and our review of articles opposing the use of predictive analytics in child welfare show that they ` tend to be fuzzy on the analysis and appear to be trying to make an ideological point rather than looking at the issue fairly. So the open-minded question in our view is not whether predictive analytics somehow inherently exacerbate racial disparities, but whether they are practical to implement in real world situations, and if they actually help address the many complex issues in child welfare, including racial disparities, and actually help managers accomplish their mission. Looking at this approach more broadly, it is typical of countless efforts to reduce the need for human resources in organizations – business, nonprofit, or government – through the use of technology

This question turns not on looking at the issue through a political lens, but on whether the developers did a good enough job in addressing the complex technical challenges of the model and whether working with LA County they make their risk stratification model accurate for that community, whether they have implemented it well so far, and if it will improve the overall efficiency of the system enough to make the model viable over time. Parenthetically, in my experience managing public operations, the implementation part is vastly undervalued. It is probably equal in importance to designing an excellent program or model.

Returning to LA County, the information for this podcast is derived from their "Methodology & Implementation Report", which is from August 2021. This report describes their process in great detail and if you read it I think you will agree that the multi-disciplinary project leaders appear to have done everything possible to be careful and precise in designing the algorithm, in including staff, managers and the community in its design as well as the rollout of the implementation, and to diligently review how the algorithm was working at each stage of the process. In short, they seem to have done as good a job of managing this implementation as is possible for anything in the real and often chaotic world of local government. So unless someone wants to make the case that these algorithms are inherently too complex to ever implement, or inherently biased for some reason, it makes sense to assume that the results are compelling, and worthy of fair consideration, as well as further refinement and development.

The way that the report demonstrates the effectiveness of the algorithm is to look back and see if it would have predicted such cases accurately. To do this they used an outcome of placement in foster care within 24 months. This was a consensus metric arrived at with managers in the system and community representatives who wanted an outcome measure that was outside the system, that is not in the control of caseworkers and supervisors to influence or manipulate. Using this measure helps reduce skepticism that the results could be consciously or unconsciously affected by bias, and as a result increased staff and community confidence in the results. Using this outcome measure, the algorithm accurately predicted 57% percent of children who ended up in foster care during that period of time, which again was the consensus view of a measure that indicated the algorithm accurately identified high-risk situations. The algorithm also predicted accurately that child fatalities and near fatalities for maltreatment were two to five times greater for the group designated as highest risk.

The effectiveness of the model was reinforced by results in Allegheny County using the second generation of this model, known as the Allegheny Family Screening Tool, or AFST-2. In this setting, high risk cases were defined as the 25% of cases with the highest risk scores. The model reduced racial disparities for these cases by 9.6%, which accounted for 98% of the disparities gap. In addition the model reduced the percentage of Black children removed to an out of home placement within three months by 2.9%, which accounted for 76% of the racial disparities gap.

Returning to LA County, the term used for the highest risk cases was “enhanced support”. This was considered the least value-laden term to describe high risk situations, and was applied arbitrarily to the 10% of cases who were most at risk according to the model rather than 25% in Allegheny County.

To address racial disparities, the algorithm built in a “racial equity feedback loop”. This took a close look at the opposite end of this continuum, namely Black children in the lowest 5% of the assessment model. Cases were individually reviewed to ask the question of why cases with the lowest risk profiles were screened in and investigated. This and similar information was fed back to an ongoing Eliminating Racial Disproportionality and Disparities (ERDD) workgroup to help tweak the algorithm and modify the model such as through updated training.

The designers and implementers of this model used a term called “targeted universalism” which means taking the universal values and goals of the system - which include child safety and well-being, permanency for the child, and striving for family preservation – and using them to target enhanced support families.

The overall goal of the data produced by the model’s algorithm is, to quote,

> “These data can help supervisors identify a small number of investigations where children are significantly more likely to have future system involvement that leads to placement. The question then becomes: What can be done now, during the current investigation and in partnership with community agencies, to ensure the appropriate services are put in place to reduce the likelihood a removal will occur?”

This gets to a critical issue. In several places the report quotes audits and statements from stakeholders to the effect that workers are on the whole relatively inexperienced and overwhelmed by their caseloads. The question is whether reprioritizing existing resources will enable an otherwise overstressed and under-resourced system to get better results. This turns primarily on the concept of enhanced support for families, and assumes that supervisors and caseworkers will be able to focus resources on the 10% of families that fall into this group.

The enhanced support practices include:

- supervisors working closely with caseworkers before the first visit to review the data
- putting together an investigative team which could include a community partner
- multidisciplinary team meetings

- meetings between county staff and community partners
- a community roundtable for investigations involving Black children
- family preservation service
- partnering up families being investigated with parents who have successfully gotten through the system
- priority service requests to community providers, and
- consultation on a priority basis with assistant regional administrators.

So, enhanced support will clearly require a lot of resources in a resource strapped program, including a lot of meetings and consultations for caseworkers that are described throughout this report as being already pressed well beyond their capacity.

I did some research for a client a number of years ago on what factors enable states or counties to reduce turnover among child protection workers. The key finding was that there were two inflection points where workers decided that this career was or was not for them. One was at six months in the other at one year. The chances that the worker would stick with child protection and make a career of it went up significantly if supervisors could help workers get through those first difficult six months and then again over the hump at the one year point. The factor that made a difference was whether the county or state unit was adequately staffed to the point that a supervisor or senior worker could take one third of their time or more to mentor and coach new workers. From the description in the report, it doesn't sound like Los Angeles County is at or near that point.

The algorithm produced by this admirable project is a public-sector instance of trying to improve efficiency and accuracy through technology. It was clearly able to accurately identify high risk cases, i.e. those that need enhanced support. The big question to me is whether the service delivery end of this model has the resources to be implemented successfully. I hope so, and I will look forward with anticipation to the next update on this project to see if a better allocation of scarce resources is in fact leading to better outcomes.

Rich Gehrman

Executive Director, Safe Passage for Children of Minnesota

3/3/2023