Children's Data Network

Los Angeles County Risk Stratification Model:

Methodology & Implementation Report

Model Version 1.0, Implemented as of August 2021

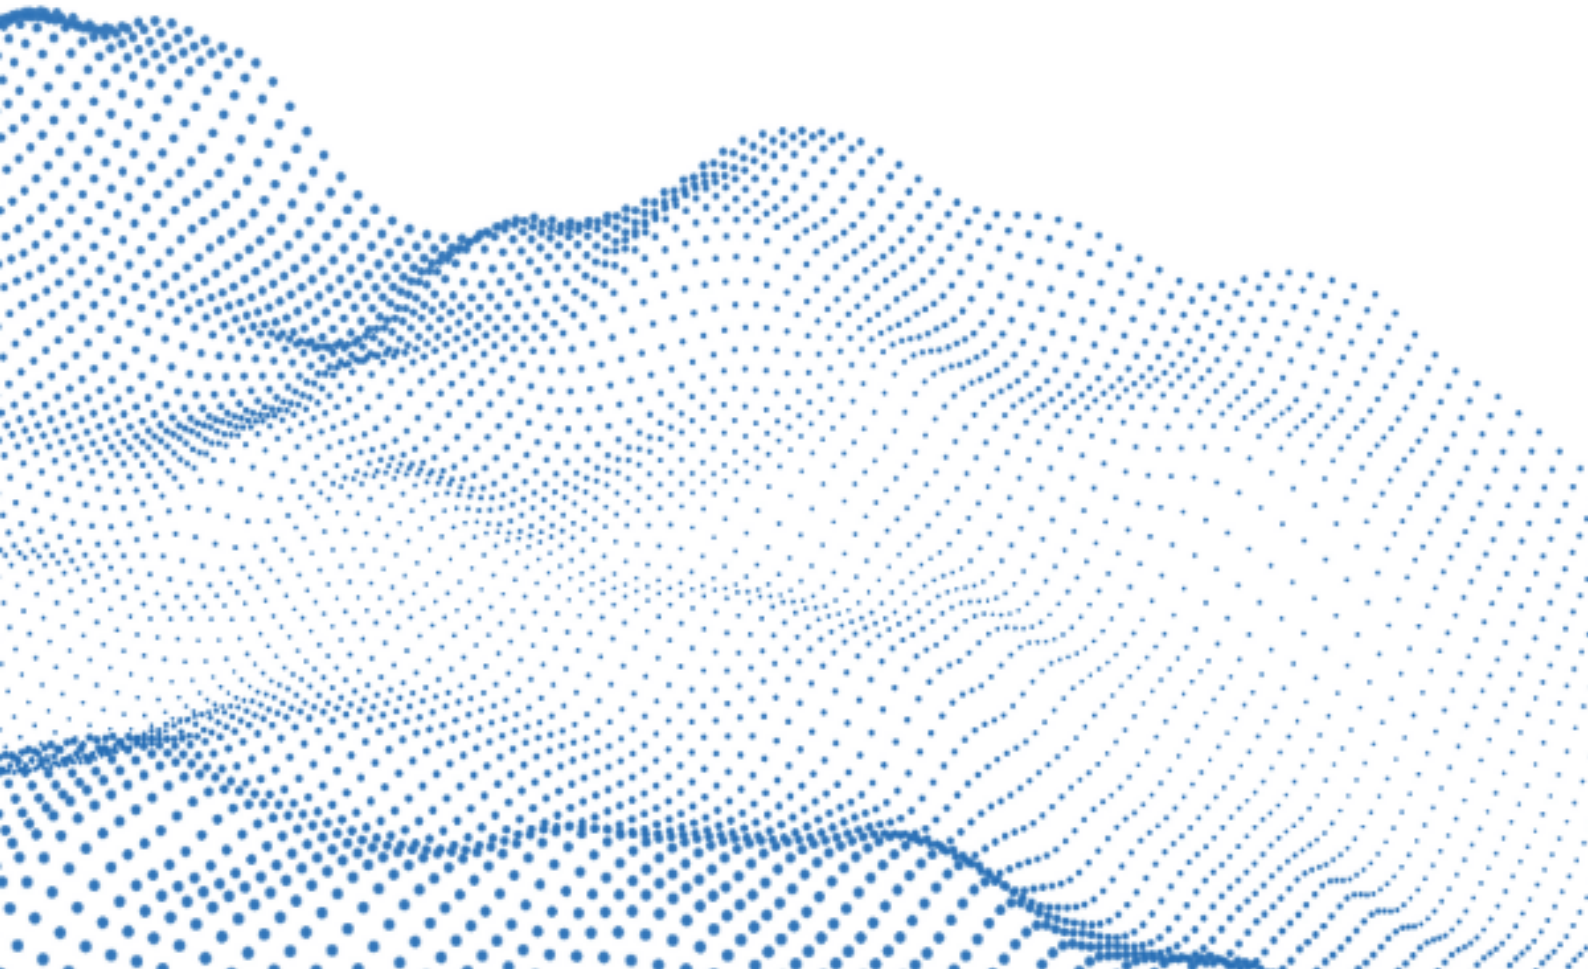

Investigators

Emily Putnam-Hornstein, PhD[a] | Children's Data Network, UNC Chapel Hill
Rhema Vaithianathan, PhD[a] | Centre for Social Data Analytics, Auckland University of Technology
Jacquelyn McCroskey, DSW | Children's Data Network, University of Southern California
Daniel Webster, PhD, MSW | California Child Welfare Indicators Project, UC Berkeley

Implementation Team
Emily Kulick, MBA
Stephanie Cuccaro-Alamin, PhD

LA Model Development & Deployment
Diana Benavides Prado, PhD *[Lead]*
Nina Anchugina, PhD
Kiran Chandradevan, MS

Base Model Development
John Prindle, PhD *[Lead]*
Huy Nghiem, MS
Tanya Gupta, MS
Michael N. Mitchell, PhD

Data and Pilot Implementation Support
Claire McNellan, MPH
Himal Suthar, MS

Project Management & Oversight
Larissa Lorimer
Jonathan Hoonhout

Supervision & Practice Consultation
Debra Waters-Roman, MSW
Ivy Hammond, MSW
Wendy Wiegmann, PhD

Supplemental Race and Ethnicity Analyses
Eunhye Ahn, PhD, MSW

Implementation Monitoring
Allon Kalisher, MSW *[Lead]*
Eliza Abendroth

Ethical Review
F. Brett Drake, PhD
Melissa Jonson-Reid, PhD

LAC DCFS Executive Leadership Team
Frank Ramos, MA *[Executive Sponsor]*
Jennie Feria, LCSW
Angela Parks-Pyles, MSW
Amoreena Jaffe, MPA

LAC DCFS Regional Office Leadership
Agustin D. Martinez, MSW *[Belvedere Office]*
Lisa Fleisher-Whitecrow, MA *[Lancaster Office]*
Jennifer Lopez *[Santa Fe Springs Office]*

LAC DCFS Technical Implementation
Rae Hahn, PhD *[Lead]*
Shu-Jiao Lin, PhD
John Booher
Vincent Chu
Kevin Phung
Sean Sun
My Trinh

LAC DCFS Racial Equity Review
Roxanna Flores-Aguilar, MSW
Lidia Manetta Escobar, PsyD
Luis Araujo, MSW
Danielle Rivera, MSW
Milena Abarca, MSW
Alen Khudaverdyan, MSA
Eboni Alexander, MA
Tiffany Jones, DSW, MSW
Kimala Lewis, MSW

LAC DCFS Project Management
Mary Parker, MPA
Nancy Salas-Silva, MSW

[a] As project principal investigators, Drs. Putnam-Hornstein and Vaithianathan assume full responsibility for the contents of this report and the accuracy of the data that are presented.

Acronyms:

ARA – Assistant Regional Administrator
AUC-ROC – Area Under the Receiver-Operating Characteristic Curve
AUT – Auckland University of Technology
BIS – Business Information Systems
CCWIP – California Child Welfare Indicators Project
CDN – Children's Data Network
CPS – Child Protection System
CWS – Child Welfare Services
CWS/CMS – Child Welfare Services/Case Management System
CSDA – Centre for Social Data Analytics
CSW – Children's Social Worker
CQID – Continuous Quality Improvement Division
ERDD – Eliminating Racial Disproportionality and Disparities
IR – Immediate Response
DPSS – Department of Public Social Services
DPH – Department of Public Health
DMH – Department of Mental Health
LAC DCFS – Los Angeles County Department of Children and Family Services
OOE – Office of Equity
PI – Pacific Islander
P&A – Prevention and Aftercare
PPV – Positive Predictive Value
RA – Regional Administrator
SPIRITT – Skills for Prevention, Intervention, Recovery, Individual Treatment and Training
RJC – Racial Justice Committee
SACWIS – State Automated Child Welfare Information System
SCSW – Supervising Children's Social Worker
SDM – Structured Decisions Making
SEIU – Service Employees International Union
SFS – Santa Fe Springs
TPR – True Positive Rate
UCB – University of California at Berkeley
UNC-CH – University of North Carolina at Chapel Hill
USC – University of Southern California

Appendices:

Report Citation:

Putnam-Hornstein E, Vaithianathan R, McCroskey J, & Webster D. Los Angeles County Risk Stratification Model: Methodology and Implementation Report (Model Version 1.0). Children's Data Network (August 2022): https://dcfs.lacounty.gov/resources/reports/

Context

Child welfare work – and investigative work in particular – is difficult and complex, and carries serious consequences. Child welfare supervisors need data-informed management tools to support the strategic direction of resources. This is especially important when frontline children's social workers are (a) gathering information concerning safety and risk during the course of an investigation of abuse or neglect and (b) ensuring proper community service connections have been made for children who remain in the home after an investigation. Delayed, incomplete, or inaccurate information, coupled with variability in investigation complexity, staff experience, and caseload size, can compromise children's safety and create an unequal landscape for families to receive the support they need.

Variability in complexity

Not all investigations are of the same seriousness or complexity. Not all families reported for maltreatment have the same service needs. A missed conversation, an unscheduled medical screening, or an incomplete risk or safety assessment by a social worker might amount to a minor administrative oversight for the majority of low-risk investigations and cases. Those same human errors, however, could have serious consequences for children and families in the context of the relatively small number of complex scenarios if not identified and addressed by supervisors and managers.

Variability in experience

Large public agencies often struggle with consistent staffing (i.e., experience, training, recruitment, retention) both within and across regional offices. This can lead to unevenness in the quality and completeness of investigations and practice with open cases, including the provision of support and services connections needed by families. Supervisors need to know when a more complex investigation has been assigned to a worker with less experience, because they may need more support.

Variability in caseloads and the context for decision making

Perhaps not surprisingly, offices that struggle the most with recruiting and retaining experienced staff members are also often those with (a) high volumes of child maltreatment investigations and correspondingly high numbers of cases opened for services, and (b) more complex cases due to above-average rates of poverty, substance abuse, and other factors that can compromise child safety. The resulting staffing and workload dynamics and increased demands on supervisors create a situation where not only are practice mistakes more likely to occur, but also the consequences of those slips are more likely to be serious. Supervisors in these situations, in particular, need tools that will help them manage their caseload of open investigations.

This report details the development of a risk stratification model and accompanying supervision and racial equity applications. These applications were designed to assist supervisors in Los Angeles County better manage information and as a result, support consistent high-quality practice with families and children during investigations.

1. Project Overview

In this section, we provide an overview of the development of a risk stratification model and accompanying applications by the Los Angeles County Department of Children and Family Services in partnership with the Children's Data Network, the Centre for Social Data Analytics, the California Child Welfare Indicators Project, and colleagues from Mathematica. This section is followed by a more detailed description of the methodology used to develop the model (Section 2); the three practice applications launched as extensions of this model (Section 3); a summary of the process for developing this project and next steps (Section 4); and an acknowledgement of the funders and partners who supported this effort (Section 5). Supplemental material has been organized into appendices that are listed with the references in Section 6.

1.1. Background

In 2021, the Los Angeles County Department of Children and Family Services (LAC DCFS) received and screened approximately 168,000 calls alleging abuse or neglect and carried out investigations involving 86,000 children.[1] During investigations, correctly assessing immediate concerns that may compromise a child's physical safety is critical. Also critical, and far more challenging, is assessing a child's risk of future harm. With the growth of predictive analytics and predictive risk modeling in other fields, child protection agencies are increasingly exploring how the increased availability of data and algorithms can support improved decision making, less biased determinations of risk, and more intentional delivery of services. The use of prediction models at key decision points has been received with both enthusiasm[2] and concern.[3,4] The simple fact is that overreliance on any one method for assessing immediate safety and future risk is problematic. At the same time, it is clear that the volume of child maltreatment referrals received each year, coupled with the massive amount of data and information frontline staffs are expected to review and incorporate into their assessment process, requires new technologies and approaches.

The current report details the development and piloting of a risk stratification model. The modeling and implementation decisions detailed here emerged from various sources: (a) feedback from community members collected during the last several years; (b) an earlier ethical review commissioned by the California Department of Social Services (CDSS)[5]; (c) conversations with and shadowing of LAC DCFS staff members; (d) meetings with the Eliminating Racial Disproportionality and Disparities (ERDD) workgroup; (e) reviews of historical[6] and more recent[7] audits that detailed areas where the system continues to struggle; (f) analyses of county and other data that point to preventable child deaths[8,9]; and (g) examinations of the shortcomings of existing risk assessment tools.[10,11]

Based on this collective feedback, several implementation choices were made for the pilot, including that the model would be introduced between, not attached to, specific system decision points (see Figure 1). This ensures that information from the model is used to support but not dictate any workforce decisions. The decision was also made to develop, from the outset, a "racial equity feedback loop" (discussed in Section 3) as an explicit acknowledgement that a combination of both structural racism and individual biases have led to system responses that have (and continue to be) unequal—and that unwarranted variation appears in both human decisions and the data generated as a byproduct of those decisions. By leading with an application of the model to examine decision making specific to Black children, we sought to both recognize that these biases exist and affirm that we have an obligation to reduce their impact on decision making by using data to help identify unwarranted variations in practice while also reviewing and revising policies that may have a disparate impact on families of color.

Figure 1. Pilot Implementation of the Risk Stratification Tool

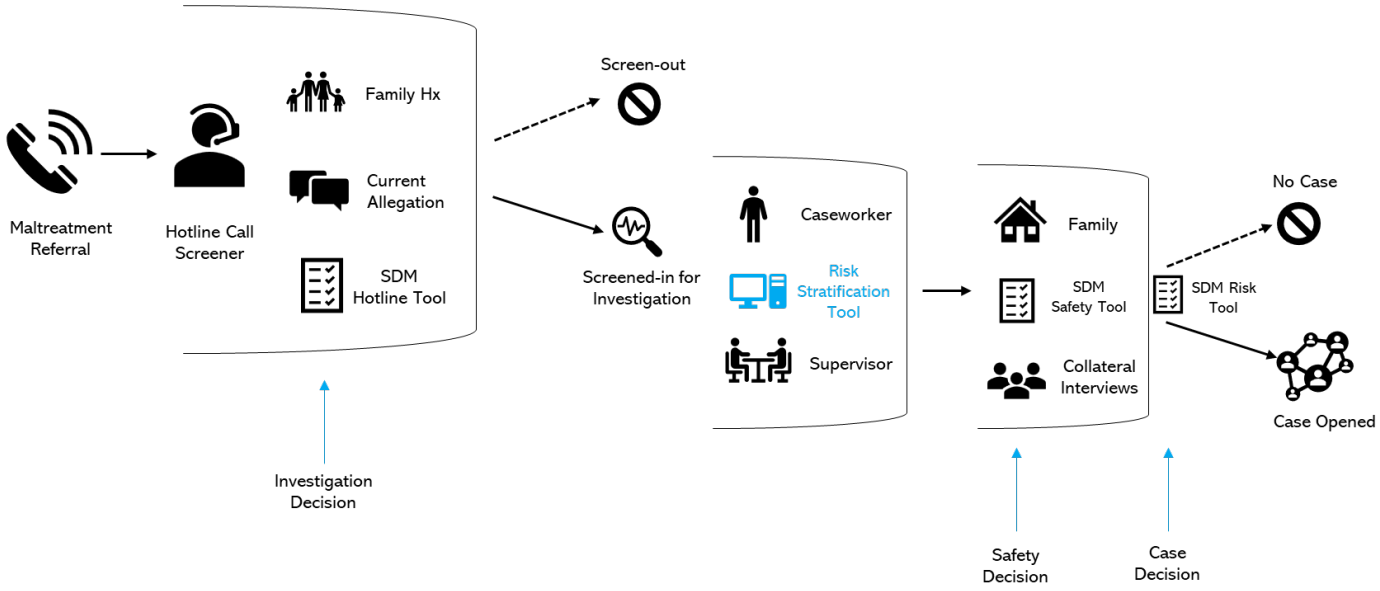

It is also important to note that the model is not being implemented as a standalone analytics tool. Rather, its implementation is part of ongoing efforts designed to elevate core practices and enhance support for investigations where the stakes are high by facilitating earlier and more consistent engagement with community partners. By adopting this approach, LAC DCFS will be able to better allocate staffing and supervision resources in an effort to guard against practice errors and human bias in cases where the consequences of those mistakes are more likely to be serious. It will provide supervisors with an additional tool for ensuring that investigations with more significant service needs—those that are often defined by inadequate or failed system responses in the past—receive the attention required to protect children and strengthen families. In this respect, we draw on arguments for "targeted universalism"[12]—or "setting universal goals pursued by targeted processes to achieve those goals"—as a framework for conceptualizing a risk stratification process in a public agency system. Child protection agencies operate with universal goals of child safety, well-being, and permanency and a commitment to family unity. But to achieve those goals, targeted approaches are often needed.

1.2. Research Team

Model development was led by the Children's Data Network (CDN) at the University of Southern California and the University of North Carolina at Chapel Hill, in partnership with the Centre for Social Data Analytics at Auckland University of Technology, the California Child Welfare Indicators Project at UC Berkeley, and the business information systems (BIS) team at LAC DCFS. Mathematica was engaged throughout the project for implementation monitoring.

Directed by Professors Emily Putnam-Hornstein at the University of North Carolina at Chapel Hill, Rhema Vaithianathan at Auckland University of Technology, and Jacquelyn McCroskey at the University of Southern California, the team featured university-based researchers, practitioners, and data scientists who collaborated closely with dedicated leaders and staff members from the LAC DCFS' regional offices, equity, BIS, policy and training, continuous quality improvement division (CQID), and risk management.

1.3. Problem

Nationwide, child protection agencies continue to struggle with high caseloads, decades-old technology and tools, and inconsistent worker and supervisory staffing (i.e., experience, training, recruitment, retention). Collectively, this can lead to unevenness in the quality and completeness of investigations and practice, compromising child safety and affecting the services and experiences of families.

In Los Angeles County, these challenges are no less acute.

"We have a very young workforce. We hired 3,000 workers within the last 4 years. Our supervisors are also very, very young."

LAC DCFS Supervising Children's Social Worker, October 2020

The county's Blue Ribbon Commission on Child Protection[6] found that workers "are overwhelmed by caseloads significantly above state and federal recommended levels. They often do not receive sufficient support, supervision, or training" (p. 3). The report further noted: "Often the least experienced social workers are assigned to assess complicated emergency situations without sufficient resources and support. Front-end investigation failures have consistently been found to be a major systemic weakness, causing many child fatalities and serious injuries" (p. 6).

Despite significant changes in the years since the commission released its recommendations, including the creation of the Office of Child Protection, improvements to cross-system and enhanced community prevention partnerships, and significant investments in hiring, a 2019 state audit[7] documented continued practice and supervision deficiencies, including systemic challenges in completing "safety and risk assessments, on time or accurately" (p. 1).

Child protection investigators and supervisors must make high-stakes determinations that affect child safety and family unity every day. Good practice and good decisions (that lead to good outcomes for children and their families) require a skilled workforce equipped with information. This project reflects an attempt to ensure supervisors have access to the data they need to effectively oversee investigations in a timely manner.

1.4. Objectives

Three primary objectives guided the development of the model, the applications that were developed, and the practices attached to those applications.

objective #1

Better align and deploy supervision and management resources to ensure that children are safe and families receive the services needed throughout (and following) a maltreatment investigation.

objective #2

Increase the use of information and data by supervisors to support quality and consistent casework that will reduce practice errors during investigations.

objective #3

Improve the use of data to identify screening practices and community reporting patterns that may result in unnecessary investigations disproportionately burdening Black and African American families.

1.5. Model

The core of this project is a risk stratification model trained using hundreds of coded data elements derived from hundreds of thousands of records maintained in the LAC DCFS Data Mart. Records contained in Data Mart are hosted on LAC servers maintained by the BIS team and are updated nightly based on information from the state's Child Welfare Services/Case Management System (CWS/CMS). A combined 341,428 Data Mart records were used for model training and testing, and 292 features were coded and used as predictors for model training purposes. This count reflects the total number of unique referred-child records generated between January 1, 2016, and December 31, 2017, with all associated history and information for individuals in those maltreatment reports. The model was trained using a LASSO regression method.[13] A 75% sample of referrals that were screened in for investigation during the selected period was used in the model training set, whereas the remaining 25% of screened-in referrals were held out during the training process and used for model testing purposes (test set). To ensure no crossover between the training and testing sets, we separated unique child-referral observations into each set. This prevented two children on the same referral from being split so that one child was included in the training data and a second in the testing data.

A child's removal and placement in foster care within 24 months of the maltreatment referral was used as the target outcome for training the model to classify risk. The decision to use this outcome was based on prior work in other jurisdictions,[14-16] previous modeling of California data in which training outcomes were compared,[17] and feedback from stakeholder groups. Approximately 13% of children investigated for abuse or neglect in LAC during this period experienced the target outcome within 24 months. As was true for the training outcome, model features were selected based on information included in models successfully deployed in other jurisdictions, previous coding work in California, and attempts to operationalize risk and safety concerns shared by staff members in the three pilot offices.

We evaluated the accuracy of different models using the area under the receiver-operating characteristic curve (AUC) on the test set (0.830 (95% CI: 0.826, 0.835)) – alongside other classification metrics. More than half (56.9%) of children that would have been designated for enhanced support were placed in foster care within 24-months, compared to 8% of other children investigated. Meanwhile, among all children who experienced a foster care placement during the 24-month follow-up period, 40.8% would have been designated by the model for enhanced support.

A full list of coded features used for model training can be found in Appendix A. Additional information about the construction of the data used to build the model and the methodologies used can be found in Section 2.

1.6. Applications

Three applications were developed for the pilot that drew on data from the risk stratification model. These applications included: (a) a notification to supervisors overseeing emergency response investigations alerting them to a new investigation that the model has been designated for "enhanced support"; (b) an investigation overview report that summarizes information that can be time consuming to assemble through existing case management system; and (c) a racial equity report that can be used by LAC DCFS to examine low-complexity referrals that were screened in for investigation and involved Black children.

The decisions made for delivering information from the model to staff members reflect our recognition that: (a) LAC DCFS staff members already have significant demands on their time, and the introduction of additional information will be most helpful if it is situated in existing workflows; (b) California's legacy case management system (i.e., CWS/CMS) makes the extraction and summarizing of historical information time consuming; and (c) data bias and human bias are both present in the child protection system, and new initiatives must acknowledge those biases and work to reduce their effect in measurable ways.[18]

Pilot applications are briefly described here. Additional details are provided in Section 4 and in Appendix B.

- **Enhanced Support Designation.** The model was used to identify new investigations with high levels of complexity and risk of future system involvement (i.e., top 10% of investigations based on future risk that a child will be removed and placed in foster care). By adopting a method for risk-stratifying investigations, LAC DCFS can explore ways to better allocate staffing and supervision resources to guard against practice errors in investigations where the consequences of those mistakes are more likely to be serious. The department also can have a more refined (risk-stratified) method for engaging multidisciplinary teams and community partners at the outset of investigations in an effort to connect families and coordinate services that prevent future system involvement.

- **Investigation Overview Report.** A subset of features coded from information in Data Mart and used to train and build the risk stratification model were presented in a report codesigned with LAC DCFS supervisors in the pilot offices. This overview report was made available for all open investigations in those offices, helping reinforce that the enhanced support designation is only one part of a larger effort to promote data-informed supervision in child maltreatment investigations. The report was updated nightly. Select features from the model were presented, along with clinical alerts identified by supervisors as pieces of history they frequently look for in CWS/CMS but that can be time consuming to gather.

- **Racial Equity Report.** Reducing subjectivity, bias, and unwarranted variation at the front end of the child protection system is key to good practice. To advance that work and lay a foundation for equity strategies that may emerge from the pilot, the model was also used as part of a racial equity feedback process. This effort was coordinated by the LAC DCFS Office of Equity in partnership with ERDD and CQI teams in each pilot office and was focused on hotline screening decisions for Black children. Specifically, data from the model were used to identify a sample of recently closed referrals classified as having a low likelihood of future system involvement, did not lead to any substantiated dispositions during the investigations, and yet were screened in for investigation. These referrals were qualitatively reviewed by CQI with the goal of better understanding how many might have been safely diverted to a community pathway and determining the policy changes that might be needed for that to happen in the future. The methodology used to select investigations for review is described in Appendix C.

1.7. Model Fairness

An examination of algorithmic fairness for the pilot model was conducted using several metrics. Because the applications attached to the model focused on investigations designated for enhanced support (i.e., investigations classified by the model to fall in the top 10% of those the most likely to lead to future system involvement and foster care placement), we focused our attention on that subset of investigations while also confirming that the model's overall classification accuracy was comparable for children from different racial and ethnic groups. Findings suggest that the model is well calibrated for children in each of the three largest racial and ethnic groups (i.e., Black, Hispanic, White). Additional details can be found in Section 2.

Please note that nothing in our descriptions of these findings should be understood as a commentary on whether there is bias in the underlying data, bias in the population of children who are reported for alleged maltreatment, or bias in historical decisions. Rather, these findings simply document the extent to which the risk stratification model makes accurate classifications across children with different racial and ethnic backgrounds.

"What does it mean for an algorithm to be fair? Surprisingly, there is a mathematical limit to how fair any algorithm—or human decision-maker—can ever be."
Corbett-Davies et al.[19]

1.8. Ethical Review

Building and testing a model using historical child protection records amounts to an empirical exercise. But the ethical use of any model requires a thoughtful use case designed to improve the "standardization, transparency,

and equity of practice."[5] Toward this end, we asked Professors Drake and Jonson-Reid, the authors of an earlier ethical review for the CDSS,[5] to review the specific applications and supervision use case implemented in Los Angeles County. They concluded that, "there is no indication that any aspect of the current efforts present ethical concerns beyond those of prior practice. Risk assessments (of a different type, and for different purposes) are common practice across the nation, and other, similar tools have been in use in Los Angeles County for years." Both the prior ethical review and the updated review for the current project can be found in Appendix D.

2. Methodology

In this section, we document the development of the risk stratification model implemented during the pilot. We outline the records used, features coded, machine learning methodologies explored, and the model's classification accuracy. The current report reflects the steps and procedures followed to develop and implement model version 1.0, the model used during the pilot period. This model, however, is not intended to be static. Nor should it be. The goal is to change and improve the model (and associated practices) over time. As modeling updates and modifications are made by the research team and LAC DCFS (and other partners), updated information will be released.

2.1. Data

Administrative records maintained in California's child protection and child welfare database, CWS/CMS, served as the source data used for training and testing the risk stratification model. CWS/CMS is California's statewide automated child welfare information system, designed to automate case management and data collection functions for California's child welfare services program.[20] The system includes records capturing everything from hotline referrals to case openings and out-of-home placements. To train a risk stratification model, we extracted all maltreatment referrals received by Los Angeles County between January 1, 2016, and December 31, 2017. This period was chosen because it allowed a 2-year follow-up window during which we could observe outcomes prior to any COVID-19 pandemic lockdowns.

2.2. Model Training Outcome

In earlier statewide analyses conducted by the research team, child outcomes observed at various points during the 24-month period following an initial investigation for abuse or neglect were coded. These included: (a) any subsequent referral, (b) three or more subsequent referrals, (c) a subsequent substantiated referral, (d) a subsequent inconclusive or substantiated referral, (e) case opening, and (f) out-of-home placement. For the current project, the decision was made to focus on training the model to classify risk of out-of-home placement within 24 months of the referral date.

This selection of an out-of-home placement within 24 months as the target outcome for LA DCFS modeling was based on prior work in other jurisdictions[14-16] and previous modeling of California data in which training outcomes were compared.[17] Feedback from stakeholder groups also informed the decision.[b,c] Specifically, LAC community members and other stakeholders felt there would be more consistency and accountability if we trained the model to predict an "outside of system" outcome that involved the courts. The placement of a child in foster care involves a court detention hearing and legal representation, a process with a notably different threshold than a decision to substantiate an allegation. Finally, the decision to train the model to classify the likelihood of foster care placement 24 months after a referral was also an attempt to identify risk trajectory differences that were proximate enough to the focal investigation to justify a different response but also far enough out that there would be a real opportunity to offer a preventive intervention.

[b] Stakeholder groups repeatedly indicated that they preferred training the model to predict an outcome that was somewhat "outside of system." Detention and removal to foster care was generally preferred because it was an outcome that would be observed following the involvement of lawyers and the courts. In contrast, while the decision to substantiate is guided by WIC 300 a-j statutes, it is largely at the discretion of the worker and child protection system.

[c] In the context of the follow-up period, the consensus was that the outcome needed to be proximate enough to the referral to be relevant and yet far enough out to ensure an opportunity for preventive interventions to be delivered and effect change. Shorter follow-up periods were cautioned against.

2.3. Model Features

All model input features coded and tested during the modeling process were derived from raw data tables originating in CWS/CMS and integrated nightly into the LAC DCFS Data Mart. To develop a list of model features (or predictors), a planning process was undertaken by the research team to operationalize a broad inventory of potential features that met one or more of the following criteria: (a) appeared in the peer-reviewed literature based on an association with serious harm (as either a risk or protective factor); (b) included in models implemented in Allegheny, PA, or Douglas, CO (sites that have implemented risk stratification models); or (c) emerged as clinically relevant history through conversations with the LAC DCFS staff. Based on feedback from community partners, we also sought to identify fields that could reflect family strengths and that would provide context.

Predictors were used as modeling inputs to train the model. For conceptual purposes, features were grouped into six "feature domains." These included characteristics of the current referral; demographics of clients on the referral; histories of prior referrals, allegations, cases, out-of-home placements; and other conditions and history. Within these domains specific features were coded for (a) the overall referral; (b) the child for whom allegations of maltreatment were made; (c) all children with allegations; (d) all siblings; (e) mother; (f) father; (g) alleged perpetrator; and (h) other adults. Although we coded race and ethnicity[d] and geography so that we could examine model performance for different subgroups, those fields were not included as features in the building the model based on feedback from community partners. In total, 292 features were coded and used as predictors for model training. A full list of these coded features is presented in Appendix A.

Table 1. Domains of Model Input Features

Domain	Description	# of Features
[1] Referral characteristics	Features describing the current referral, such as day and time, current allegations, and reporter source. (Office assignment and other geographic information were not included as model inputs.)	43
[2] Demographics	Features describing demographic characteristics of alleged child victims and adults on the referral, such as gender and age group. (Race and ethnicity were not included as model inputs.)	58
[3] Allegations	Features describing prior allegations for children on the referral and allegations for other children involving adults on this referral.	109
[4] Cases	Features describing the nature, timing, and counts of child welfare cases for individuals named on the referral.	25
[5] Placements	Features describing the nature, timing, and counts of placement histories for children named on the referral.	33
[6] Other information	Features describing other conditions, safety concerns, and history for children and adults on the referral.	24

2.4. Feature Coding

The coding scheme used a referral cohort approach. Specifically, we identified all individuals named on a given referral for maltreatment. These individuals were coded into different roles based on their designation in the underlying data system (e.g., victim, alleged perpetrator). When the same person had multiple roles—for

[d] Race and ethnicity is based on primary race/ethnicity status which is often collected from the reporting party at intake and may not be updated with the family's self-reported identity. Consistent with the CDSS convention, if Hispanic ethnicity is reported, the race and ethnicity feature is coded as Hispanic regardless of racial group.

example, if a mother was also the alleged perpetrator of maltreatment—we established a "role hierarchy rule" to ensure coding consistency (e.g., a person who was an alleged perpetrator would always be identified as such, even if they were also designated as a mother). The role hierarchy always classified a child as an alleged victim if they were named as such in a given referral. Following that, we classified individuals into roles based on the following hierarchy: (a) alleged perpetrator, (b) mother, (c) father, (d) sibling, and (e) other adults. The role hierarchy was based on earlier analyses documenting how features related to the alleged perpetrator are often the most valuable in classifying risk. Additional features were coded to provide the model with information about any duplication of roles (e.g., a feature to denote if the alleged perpetrator is also recorded in the data as a mother).

We then generated features to describe different characteristics of individuals recorded in a referral. We also generated additional features associated with the referral itself (e.g., time and day of the referral). After coding characteristics for each role on a referral and for the referral overall, we next constructed a training dataset in which each row represented a unique child-referral observation. The child with maltreatment allegations can be thought of (and is elsewhere referenced) as a "focal child" to denote the specific child who is the focus of other features coded in a given row of data. Model predictors included features of the focal child, alleged perpetrators, mom, dad, siblings, and other adults. We coded an indicator to document how many of the other ancillary roles were associated with the child-referral. For example, if the focal child had no siblings on a given referral, then a feature capturing the count of siblings would be coded as "0" and all sibling features would be set to "0." Finally, for each child-referral we coded their status on target outcome (whether they were placed in foster care within 24 months of the date the referral was received). This process left us with 341,428 child-referral observations representing 171,313 unique referrals.

2.5. Model Training and Testing

The model was trained using only referrals that were screened in for investigation. For model training purposes, we split child-referral observations that were screened in for investigation (*N* = 278,465) into a training set (75%) and a test (or "holdout") set (25%). Because children could have been referred multiple times, the data were split to ensure that a given child was exclusively assigned to either the training or test set. A system-generated referral ID was used as a blocking variable to guarantee that all children associated with a referral were exclusively assigned to either the training or test set, but not both.

A LASSO logistic regression model[13] was run on the training data with the model training outcome coded. The LASSO model was trained using weighted classification to account for the low prevalence of the outcome, using 10-fold cross-validation with three repetitions. The LASSO lambda parameter was selected using grid search over 100 lambda values in the range of 0.0001–100. The final selected model was the model for the lambda parameter that was one standard error from the lambda parameter of the optimal model. The hyperparameters were chosen to maximize the AUC.

The choice of the modeling method came down to ease of deployment and interpretation during the pilot period.[e] Earlier state-level modeling had experimented with machine-learning methodologies (i.e., random forest, XGBoost) and was used to establish a benchmark for performance. We compared the accuracy of the models generated by LASSO with earlier versions of the statewide random forest model,[21] which had emerged from that project with the best classification accuracy. We used the AUC on the test set for this comparison and determined that the LASSO performance was close to the earlier developed state-level random forest model after several likely forward-looking predictors from the state model had been dropped (AUC = 0.84, compared with 0.83 for

[e] It should be noted that LASSO is a type of regression designed to construct sparse linear models by shrinking some coefficients and making other coefficients equal to zero, while optimizing for prediction accuracy. Although LASSO produces a smaller set of nonzero weights than some alternative approaches, any attempt to interpret weights for individual features can be misleading, especially when the features are highly correlated (as in the current context). This means that in the presence of such correlation, individual weights cannot be interpreted as the unique contribution of that feature to the probability of the event being predicted.

the LASSO). We concluded that LASSO was suitable for an initial model deployment. Ongoing experimentation with alternative methods such as support vector machines or neural networks will be conducted in future refinements.

To further refine and finalize the LASSO model for deployment, we also assessed the PPV and TPR for the selected model after converting predicted probabilities (on the test set) to identify children with scores that fell in the top 10%.

2.6. Quality Assurance Reviews

During model training, data quality assurance checks were conducted on all model input features. A total of 42 referrals were selected from the research dataset. Eleven were selected randomly and the other 31 referrals were selected because they had positive values in domain areas involving rare characteristics or events (e.g., termination of parental rights termination, positive toxicity at birth). These criteria led to data reviews for 170 unique individuals on the selected referrals. For all individuals in the selected referrals, the corresponding CWS/CMS source tables and attributes were directly queried. Values on source attributes were then compared to raw and derived input features observed in the research dataset. All inconsistencies and coding issues identified were returned to the modeling team for review and if needed, correction. An additional round of review was conducted to confirm all identified issues were corrected.

A second round of data quality-assurance checks were completed following test deployment of the model in the LAC DCFS server environment prior to the pilot launch date. A total of 32 referrals were randomly selected for this process. These referrals yielded 152 individuals for review. For each selection, the corresponding Data Mart source tables and attributes were then queried for all referrals and individuals. Values on source attributes (from the production environment) were then compared to raw and derived input features in the research dataset. All identified inconsistencies were explored and reconciled in the code that was deployed. Following this recoding process, an additional round of quality-assurance checks was conducted to confirm that all identified issues had been corrected.

Finally, an ongoing quality-assurance report was developed as part of the deployment package provided to LAC DCFS. The code can be edited to examine input features and source data tables from Data Mart for specific referrals (or individuals). During the pilot, this report was used in conjunction with other methods to identify potential data issues that arose as part of ongoing model deployment and updates. This included a weekly management report made available to LAC DCFS leadership and regional office staff members during the pilot to help track the volume and flow of investigations designated for enhanced support in each office.

2.7. Model Performance

Table 2 presents overall performance metrics for all children in our test set. First, we present the base rate of foster care placements within 24 months among children investigated by LAC DCFS between 2016 and 2017 and selected into our test set.[f] Overall and conditional upon investigation, 13.8% of children experienced placement within 24 months of the referral. A larger share of Black children were placed in foster care than other children (19.7%), whereas a smaller share of Asian or Pacific Islander children (8.0%) were placed. Both findings are consistent with known disparities in California and Los Angeles County.

A fundamental question is whether a model can be applied "fairly" to subpopulations, especially individuals of different racial and ethnic backgrounds or who are characterized by different levels of socioeconomic status. Because there is no single definition of fairness,[22,23] we assessed the classifications produced by the model in several ways in the context of the proposed use case—designating a subset of already open investigations for enhanced support.

[f] Note that the data is analyzed at the child-referral level. The test set of 69,403 observations represents 50,906 unique children.

First, we examined the AUC for each racial and ethnic group using records that were held out as part of the test set. Even though our primary application of the risk stratification model focused on a very specific subset of classifications (i.e., the top 10%), a metric of the model's overall performance was useful given that LAC DCFS also chose to use the model to carry out a sampled review of recently closed investigations that had been investigated despite a low probability of future system involvement. The AUC for the model overall was considered "very good" to "excellent" (0.830; 95% CI: 0.826, 0.835) based on conventional heuristics.[24] Comparable model performance was observed by racial and ethnic groups, with greater variability (and larger confidence intervals) for smaller racial and ethnic groups (e.g., Asian or Pacific Islander children). Although not shown, we also examined the AUC for children at different ages and found that the model generalized well.

We next explored how well the model was calibrated. Specifically, we examined whether children in investigations classified as falling in the top 10% (i.e., those that would be designated for enhanced support) were equally likely to experience the outcome the model was trained to predict. Among children associated with investigations designated for enhanced support, more than half (56.9%) were placed in foster care within 24 months of referral, compared to 9% of children in the other 90% of investigations (not shown). Among the three largest racial and ethnic groups, the estimated placement prevalence within 24 months was similar. Meanwhile, among all children who were placed within 24 months, findings indicate that 40.8% would have had been designated for enhanced support by the model, potentially preventing the need for removal. Among children who were placed in foster care, enhanced support would have reached 48.4% of Black children, 37.2% of Hispanic children, and 41.8% of White children. Although approximately one third of children who were placed were removed as a result of the current investigation—meaning that there was little opportunity for prevention—nearly two thirds had placements that occurred following another referral for maltreatment.

Table 2. Overall Performance Metrics of the Selected LASSO Model (Test Set) [g]

		----------Placements in Foster Care within 24 months of Referral--------		
Group	**AUC (95% CI)**	Placement Prevalence (overall)	Placement Prevalence (among investigations designated for Enhanced Support)	Share of all Placements that would have been designated for Enhanced Support
Children investigated[h] (*n* = 69,403; 100% of test set)	**0.830 (0.826, 0.835)**	**13.8%**	**56.9%**	**40.8%**
Asian or PI children (n = 1,951; 2.8% of test set)	0.840 (0.802, 0.878)	8.0%	67.7%	38.4%
Black children (n = 12,497; 18.0% of test set)	0.805 (0.795, 0.815)	19.7%	58.0%	48.4%
Hispanic children (n = 41,950; 60.4% of test set)	0.824 (0.818, 0.830)	13.3%	56.2%	37.2%
White children (n = 7,274; 10.5% of test set)	0.845 (0.832, 0.858)	12.9%	53.3%	41.8%

[g] As part of our prelaunch quality assurance efforts, we asked colleagues at Carnegie Mellon University's School of Computer Science to separately code and independently confirm our findings regarding the model's performance by race and ethnicity. Their analysis was consistent with what is reported here.

[h] Native American children were included in the overall model, but we do not present separate statistics because the count in the test set was too small (*n* = 85). Likewise, children with missing race and ethnicity (*n* = 5,646) were included but not separately presented here.

We also documented the racial and ethnic distribution of children who would fall in an investigation designated for enhanced support using data from our test set (Figure 2). The distribution of predicted probabilities reflects group disparities present throughout levels of involvement with the child protection system. Specifically, Black (and although not shown, Native American) children were more likely to be in investigations designated for enhanced support, reflecting that the likelihood of future foster care placement is greater. Meanwhile, Asian or Pacific Islander children, who have notably lower rates of foster care placement following an investigation, present with a (right-skewed) risk distribution.

This is where the model use case—enhanced support and services—is important. 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?

Separate analyses of racial and ethnic differences in family engagement in community services in LAC suggest that although equally likely to be referred to community services, Black families may be less likely to accept and engage in services.[25] Lower uptake may occur for various reasons (e.g., service location, waitlists, service fit with family needs, trust in the provider). But as much as this dynamic contributes to heightened rates of future reports and foster care placements in the 24 months after an investigation, enhanced efforts to coordinate services that truly address a family's needs during the current investigation may be particularly impactful for historically overrepresented groups.

Figure 2. Distributions of Raw Probability Scores by Race and Ethnicity with Approximate Threshold (Dashed Line) for Enhanced Support (Test Set, binned by 100 increments)

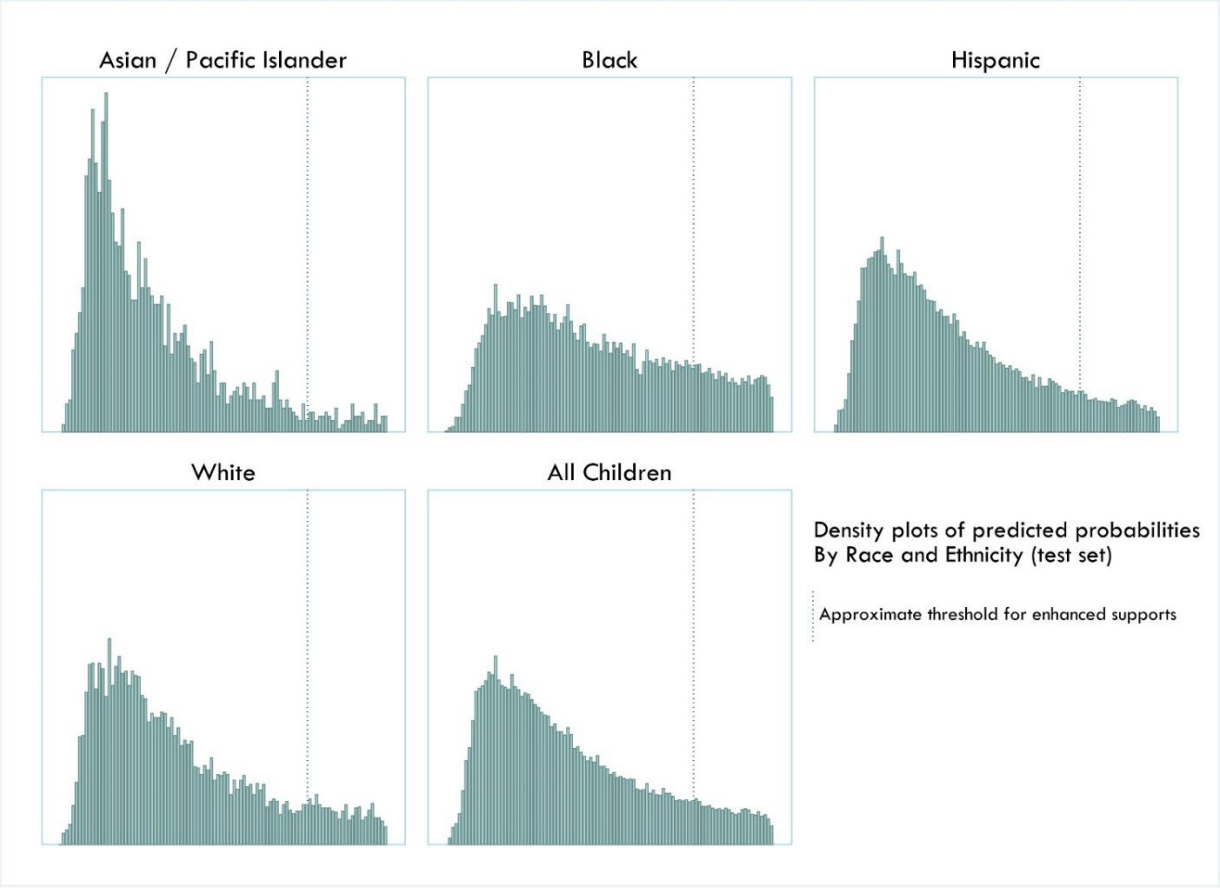

2.8. External Model Validation

The risk stratification model was trained to look for data patterns associated with a child's risk of future system involvement and placement in foster care. Not surprisingly, these investigations tend to be characterized by a high degree of complexity, as exhibited by the nature and degree of past system involvement. These investigations may benefit from different engagement strategies and more multidisciplinary efforts to ensure that any service shortcomings in past efforts to stabilize the family are well understood and addressed.

We must not forget, however, LAC DCFS's primary charge: to protect children from serious abuse that may lead to significant harm or even death. If a risk stratification model trained to predict future system involvement is poorly correlated with outside system measures of child harm (e.g., inflicted injuries, hospitalizations), then applications drawing on data from the model could potentially undermine child safety by directing supervision and other staffing attention and resources to the wrong investigations.

Consistent with external validations conducted in other jurisdictions implementing risk stratification models,[26] we examined whether near-fatalities and fatalities among children in investigations designated for enhanced support were different than what was observed for children in other investigations. We relied on two sources of data:

1. We used maltreatment near-fatalities recorded in accordance with WIC 10850.4, which requires that county welfare departments notify the CDSS of "every child fatality and near fatality that occurs within its jurisdiction and is the result of abuse and/or neglect" through the Statement of Finding and Information Form (SOC 826). This information, including the date of the injury event if applicable, is submitted to the state and recorded by the LAC DCFS in CWS/CMS.

2. In addition, we used child-level linkages between child protection records and statewide vital death records falling under the authority of the California Department of Public Health. Death records, available through 2020, were available through a data-sharing agreement and were anonymized for analysis by the research team. Death events were coded using the International Classification of Diseases, 10th revision.

We documented that the rates of maltreatment near-fatalities, maltreatment fatalities, and child homicides were consistently 2 to 5 times greater (depending on the criteria) among children in investigations the model designated for enhanced support (compared to all other children investigated for maltreatment). These findings indicate that although the model was trained to identify patterns correlated with maltreatment severity leading to a foster care placement, those same patterns were associated with a heightened risk of near-fatal maltreatment and inflicted-injury deaths. This means that as much as enhanced supports are directed to a relatively small number of investigations, those supports would reach children and families with a significantly heightened risk of experiencing a fatality.

In Figure 3 we present data for child deaths that were coded in vital records as homicides between 2016-2020 and involved a child in our sample. For each child who died, we examined the risk score generated by the model for all screened-in investigations that took place for that child prior to the date of death. We then took the maximum of those scores to examine how many children would have been designated for enhanced support – and further conditioned on the presence of a child in the investigation who was age 10 years or younger to align with the parameters of the pilot (n=19).

We observed that 52.6% – 1 in every 2 child homicide victims – would have been classified by the model for enhanced support during an investigation prior to death.[i] Meanwhile, during this overall period and using the

[i] Please note that for all external validations, we present data based on the highest risk score for the child before death rather than the highest risk score for any investigation the child was included in. The numbers were slightly higher (i.e., 57.9% of homicide victims would have been classified by enhanced support vs. 14.0% of children who did not die) using the other methodology.

same methodology, 9.9% of children who did not die would have found themselves in an enhanced support investigation. Please see Appendix E for additional fatality figures.

Figure 3. Child homicides by highest risk level estimated by the model prior to the date of death compared to the risk distribution of children who did not die during the study period

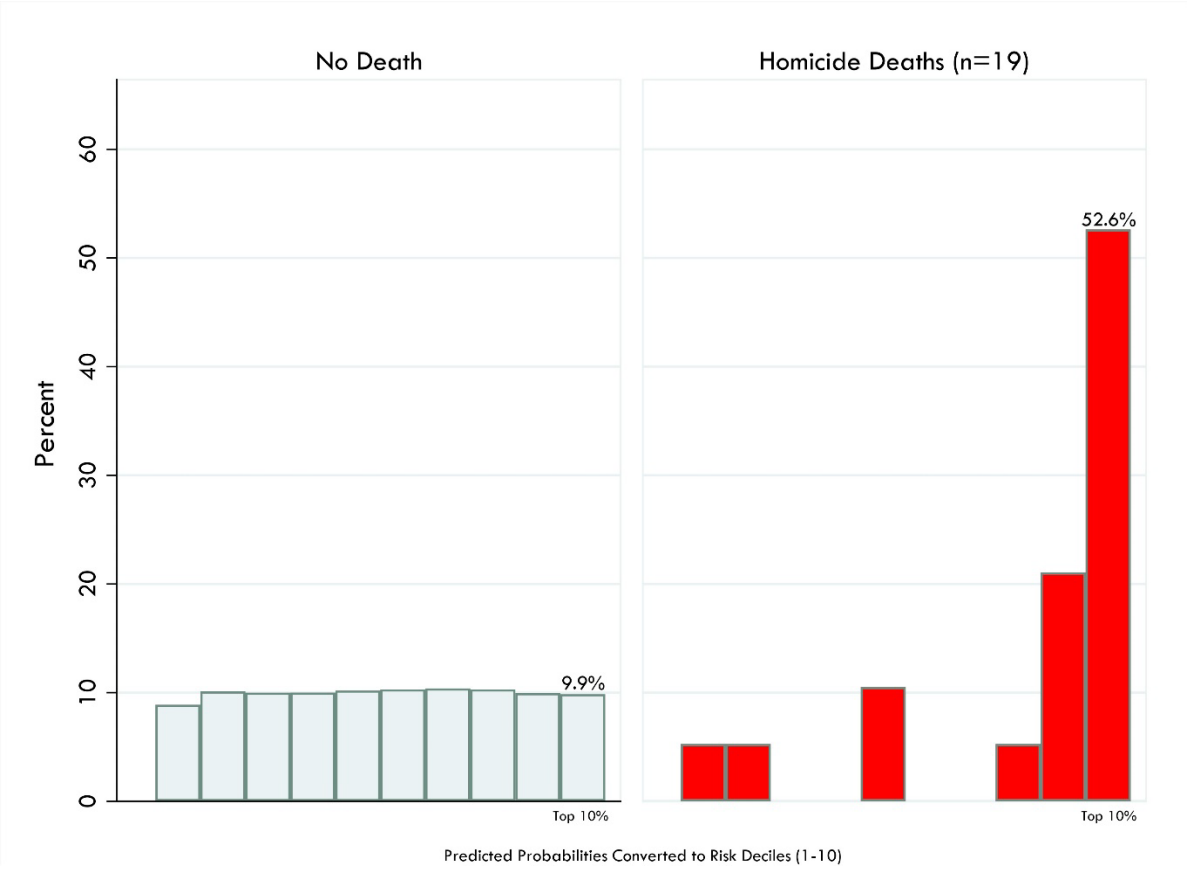

3. Applications and Practices

In this section, we outline three piloted applications that drew on data and information generated through the risk stratification model. It goes without saying that any model is only as good as the practices attached to the information that is provided. Each application launched during the pilot reflects an effort to support and augment the work of LAC DCFS supervisors and staff members. By extension, we hope these applications and uses of data will produce better outcomes for children and families in Los Angeles County.

3.1. Pilot Offices

Three regional offices were selected to participate in the pilot: Belvedere, Lancaster, and Santa Fe Springs. Several criteria guided the choice of these offices, including geography and concurrent participation in other pilot projects. Participation in the pilot meant that the model and its associated applications would be made available for open investigations assigned to the LAC DCFS staff in each office. The first 90 days were designated as a trial period during which feedback was gathered from the staff and refinements to the methods for delivering information were made. Of primary interest was how supervisors responded to investigations designated for enhanced support, whether they valued the data provided in the investigation overview report, what training was needed to use these applications, and which supervision practices were being explored and utilized. Post investigation outcomes were also monitored to confirm that there were no dramatic shifts in dispositions, case openings, or removals. Staff data and feedback during the initial launch were gathered by Mathematica and are summarized in a separate report.

3.2. Model Timing

The model was deployed to run nightly, timed to the receipt and integration of updated extracts of information from CWS/CMS into the LAC DCFS Data Mart. This means that applications that draw on coded features, or a specified risk designation, reflect the most updated information as of the last Data Mart refresh. Running the model to generate information in real time was not possible given that there is a lag in when LAC DCFS receives updated information from CWS/CMS that can be integrated into Data Mart. Likewise, building the model in the legacy case management system was not possible because it is maintained by an external vendor contracted by the CDSS. It should also be noted that LAC DCFS Data Mart is limited to records where Los Angeles County is the county responsible; the system does not include records from other counties where a client may have been served. Therefore, neither prior nor intervening child protection history in other counties is considered by the model.

3.3. Model Alignment with Other Tools

The risk stratification model does not replace any other required tools used for screening, assessment, or other decisions by LAC DCFS. The CDSS mandates the use of various Structured Decision Making (SDM) tools to guide specific decision points. As documented in Table 3 and depicted in Figure 3, these include: (a) the SDM Hotline Screening Tool (used to decide whether a referral will be screened in for investigation and the response time for initiating an investigation); (b) the SDM Safety Assessment Tool (used to assess a child's current safety in the home); and (c) the SDM Risk Assessment Tool (a manual algorithm used to classify the risk of future substantiated abuse). The risk stratification model was positioned to run after the SDM Hotline Screening Tool is completed but before the SDM safety and risk assessment tools are typically completed during an investigation. The goal was to provide information to supervisors and staff members at the outset of an investigation.

Table 3. Risk Stratification Model Versus Other Risk and Safety Tools

	Risk Stratification Tool	SDM Hotline Tool	SDM Safety Tool	SDM Risk Tool
Automated algorithm (no new data entry)	+	X	X	X
Considers hundreds of fields simultaneously	+	X	X	X
Available at the outset of investigations	+	+	X	X
Constructed and validated using local data	+	X	X	X
Nonproprietary, algorithm belongs to LAC DCFS	+	X	X	X
Used to make hotline screening decisions	X	+	X	X
Used to assess immediate safety concerns	X	+	+	X
Used to assess future risk of harm	+	X	X	+
Used to direct case opening or removal decisions	X	X	+	+
Required by the CDSS	X	+	+	+

3.4. Training

Prior to the launch of the risk stratification pilot, supervisors in the three pilot offices participated in voluntary training sessions designed to orient them to the risk stratification model and the relevant applications. Supervisors received training credit for attendance. Two trainings were offered and led by the pilot implementation team and LAC DCFS training staff. Each training included an overview of the risk stratification model, technical training on access and utilization of the model applications, a review of suggested strategies and practices for enhanced supports with complex-risk investigations, and a lengthy question-and-answer session. Please see Appendix B for some of the internal practice strategies that were outlined during staff member training.

To provide ongoing technical assistance, the implementation team created a downloadable navigation guide that could be accessed through a link from the investigation overview report. Additionally, a "help me" button embedded in the investigation overview report could be used to report problems or ask questions. The back-end log for this process was monitored on a weekly basis. Individual technical assistance issues were addressed directly with users, and global problems reported were addressed with regional administrators so that information could be shared with all supervisors.

3.5. Ethical Review of Applications

In 2018, Professors Drake and Jonson-Reid authored an ethical review of proposed use cases for a predictive risk model in California: Administrative Data and Predictive Risk Modeling in Public Child Welfare: Ethical Issues Relating to California.[5] As described in that report, "California is considering the following use cases: (1) to support Red Team reviews, (2) identification of candidates for providing prevention services under FFPSA, (3) explorative research on allocating limited intensive resources post Family Reunification, and (4) to assist with administrative CQI on referral screening practices" (p. 3). Although the CDSS did not proceed with implementation, it did issue a response to the review in which it advised against the use of a model at any specific decision point, but also noted other places where a predictive risk model might be usefully applied, including

"identifying which children and families would most benefit from supportive services or improving management of staff assignments" (p. 1).

The deployment of a risk stratification model between hotline screening and case opening (i.e., between decision points) and the development of applications to improve the quality of supervisory oversight during investigations emerged from this guidance provided by the CDSS. Likewise, to ensure continued attention to disparities, the model can be used to examine unwarranted variation across race and ethnicity through the racial equity feedback loop. As described earlier, an updated ethical review of the specific applications piloted with the LAC DCFS risk stratification model is available as Appendix D.

3.6. Piloted Applications

Three applications were developed for the pilot that drew on data from the risk stratification model. These applications were: (a) a notification to supervisors overseeing emergency response investigations, alerting them to a new investigation that the model has designated for "enhanced support"; (b) an investigation overview report that summarizes information that can be time consuming to assemble through the existing case management system; and (c) a racial equity report and feedback loop that can be used by LAC DCFS to examine low-complexity referrals that were screened in for investigation and involved Black children. Decisions that were made for these applications are detailed in the sections that follow.

APPLICATION 1: "Enhanced Support Designation"

Tied to Objective 1: Better align and deploy supervision and management resources to ensure children are safe and families receive the services needed during a maltreatment investigation.

Implementation	Description
[Open Investigations]	The model and risk stratification tool were implemented to examine only investigations that were already open (i.e., referrals already screened in for investigation). The model can review both historical information from CWS/CMS and information from the current referral, looking for patterns that distinguish investigations as uniquely complex and likely in need of services. *(The model runs overnight, meaning that for immediate response investigations, information will not be available until the following morning.)*
[Enhanced Support]	An enhanced support designation is generated if an investigation crosses the designated threshold (i.e., is in the top 10%). The language of "enhanced support" emerged through conversations with LAC DCFS and community partners. It was chosen because it does not conflict with language used in other risk and safety tools currently implemented in California and does not attach any potentially harmful label to a family. The language helps reinforce that these investigations may have significant service needs and that a family may already have a case opened for in-home services. No numbers or scores are attached to other open investigations. Likewise, for investigations designated for enhanced support, that designation is communicated in the form of a flagged label rather than a numeric risk probability. *(Approximately 10% of children are involved in investigations designated for enhanced support, reflecting a reasonable threshold from a supervision attention and workload perspective. The use of a labeled designation for a relatively small subset of investigations was devised to draw the appropriate level of supervisory attention.)*
[Existing Management Report]	The enhanced support designation was added to an existing supervision and management report produced by BIS. The chosen report was identified as one frequently consulted by supervisors and assistant regional administrators. *(This decision means that staff members do not have to go to a new location to find which investigations have been designated for enhanced support, helping reduce technology adaptation barriers to incorporating this additional information into supervision practices.)*
[Child Age]	An enhanced support designation is only generated if there is at least one child aged 10 or younger included in the investigation. *(Ninety percent of child maltreatment fatalities involve children aged 10 or younger. Reports that include only adolescent children (older than 10) are likely to present with different safety concerns and service needs than investigations involving younger children. Staff feedback indicated that separate applications or protocols for this population would be useful in the future.)*
[Information Access]	Information as to whether an investigation has been designated for enhanced support is available to all staff members involved with the investigation (supervising children's social worker, children's social worker, assistant regional administrator, etc.) and who would otherwise have permission to access information for a particular family or investigation. *(Shared access to the information is required to implement the desired engagement and practice interventions. Just as workers and supervisors can jointly review SDM risk levels and other designations, they are able to jointly review whether an investigation has been identified with more significant service needs.)*

APPLICATION 2: "Investigation Overview Report"

Tied to Objective 2: Increase the use of information and data by supervisors to support quality casework to reduce practice errors during investigations.

Implementation	Description
[Summary Data]	An investigation overview report summarizes the child protection history of individuals named on the referral and associated with the investigation. Feedback obtained during the codesign process indicated that it can be difficult and time consuming to review a family's history in CWS/CMS. This may be especially true for families in investigations designated for enhanced support, which may have extensive prior child protection history. The risk stratification model is run based on structured data elements of LAC DCFS involvement. Although the individual weights attached to each data element are of little use to the frontline staff, the structured history in a more easily accessed format seemed potentially useful to supervisors, especially if it could be accessed without additional navigation to other sites. *(Supervisors can access the investigation overview report during case conferences with staff members and when they are approving assessments and final dispositions to ensure relevant history is being considered in the context of the current report.)*
[All Investigations]	The investigation overview report was available for all open investigations in the three pilot offices. Staff members could click on any open investigation and see a consistent structured history of past system involvement for the family, regardless of whether the current investigation had been designated for enhanced support. *(Feedback received during the codesign process indicated that information proposed for inclusion in this report would be valuable to supervisors and others, regardless of whether the investigation had been designated for enhanced support. The creation of this report for all investigations reinforced that enhanced support designation is only one part of a larger effort to improve the consistency of front-end practice through the use of data.)*
[Service Relevant Indicators]	The investigation overview report was codesigned with the three pilot offices. During initial pilot planning meetings, the research team worked with the office staff to identify the information considered most relevant during investigations. The team also shadowed supervisors to observe how they currently navigate CWS/CMS to understand a family's history and document insights into both assessing risk and guiding the investigation process (e.g., records requests, discussions during family and collateral contacts). This process identified 16 clinically or service relevant indicators. These are automatically extracted and included in the investigation overview report, if present. *(To generate the investigation overview report, we were limited to structured data in CWS/CMS. We recognize that there is significant information that is only available in the free-text or narrative fields. The report does not replace a thorough review of records. Rather, it is a supplemental tool that provides a more consistent way to summarize structured history and ensure that it is easily accessible to the staff throughout the investigation process. Further, the report is intended to be iterative; it will be updated as we collect feedback from users and generate ideas for improvements.)*
[Simple]	Although there were requests to have dozens of data elements presented, the report is intentionally simple. It was designed to provide only high-level information that would provoke, and then help a supervisor conduct, a more careful review of case history in CWS/CMS. *(We recognize that it is important to not overwhelm users with more information than they could reasonably consume. To that end, the design and presentation is intentionally streamlined. The model is comprehensive in its use of data; the investigation overview report was designed to be simple.)*
[User Feedback]	In addition to ongoing implementation feedback, the investigation overview report includes a link so that users can report issues (e.g., incorrect data or problems with specific features, questions as to an investigation's enhanced support designation). *(This provided an opportunity to incorporate user feedback routinely so that improvements could be made during the pilot launch.)*

3.7. Enhanced Support Practices

Engagement with the regional offices during the planning phases revealed practices that could be usefully adopted and adapted for investigations designated for enhanced support. Rather than prescribe a uniform set of supervision practices or required management responses, LAC DCFS leadership outlined enhanced support practices, resources, and strategies that pilot office staff members could leverage to more strategically direct available resources toward families in an effort to advance early engagement and ensure the time needed to coordinate services. These practices include:

Augmented Supervision

Supervisors engage early with children's social workers to prioritize a complete review of data for an enhanced support investigation, ideally before the first family visit.

Investigation Teaming

Supervisors create an engagement "team" for an investigation designated for enhanced support. This could take several different forms, including early engagement with a community partner.

Internal Staffing and Multidisciplinary Team Meetings

If the enhanced support investigation involves a family with an existing open case, emergency response, continuing services, and any other assigned staff meet for clinical consultations.

Child and Family Team Meetings with Community Partners

Emergency response workers ensure that a child and family team meeting involves active engagement with trusted community partners.

Community Roundtable (Lancaster and Santa Fe Spring Offices)

Supervisors request a community roundtable for enhanced support investigations involving Black children. This permits more directed service planning earlier during the investigation.

Family Preservation Services

If the enhanced support investigation involves a family with an open family maintenance case, emergency response workers make a referral to family preservation.

Parent Partner

Parent partners' assignments are prioritized during the investigation process. Families receive support, information, and mentorship from parents who have successfully navigated the system.

Community and County Partner Linkages

Supervisors initiate preliminary and prioritized service requests to the department's contracted community providers and other county departments for families in investigations designated for enhanced support.

Assistant Regional Administrator Consultation

Supervisors request a prioritized assistant regional administrator consultation and review at the outset of a new enhanced support investigation, rather than at the end of investigations when certain risk factors are present (current practice).

APPLICATION 3: "Racial Equity Feedback Loop"

Tied to Objective 3: Improve the use of data to identify screening practices and community reporting patterns that may result in unnecessary investigations disproportionately burdening Black and African American families.

Implementation	Description
[Racial Equity Review]	The risk stratification model was used to generate a list of referrals involving Black children that the model designated as "low complexity" and with a low risk of future system involvement, and yet were: (a) screened in for investigation, (b) had no allegations of physical or sexual abuse, (c) did not emerge with any substantiated allegations at the conclusion of the investigation, and (d) did not have any associated case opened for services. *(Black children are disproportionately represented among those referred for maltreatment, both currently and historically. As such, equity review efforts focus on better understanding the dynamics specific to low-complexity reports for this population. The equity review process provides opportunities to review and examine upstream reporting [from the community] and screening [or hotline] decisions through a very intentional racial equity lens.)*
[Risk Threshold]	Although the risk stratification model is not attached to any decision making at the hotline, there are opportunities to identify "low-complexity" referrals with a low risk of future system involvement. Of interest is the extent to which some portion of these referrals might have been safely screened out and referred to community partners—and to then devise practice and policy changes that could ensure that is more likely to happen in the future. *(For the purposes of conducting racial equity reviews, a low-complexity threshold was set to include investigations meeting the selection criteria and where fewer than 5% of children are placed in foster care in the next 24 months. This translated into roughly 50% of investigations.)*
[CQI Team Review]	Per a special project directed by LAC DCFS executive leadership, the CQI staff co-located in each regional office systematically examined a sample of recently closed investigations. The goal was to better understand what led to the initial report from the community and to contextualize the screen-in decision. *(This focused review work was undertaken in collaboration with the CDN team and designated staff members from the Office of Equity. Findings will be summarized for an executive briefing and used to support community discussions. The methodology for selecting sample of investigations for review is included as Appendix C).*
[LAC DCFS Executive Team Briefings]	The CQI staff will report a summary of findings to the LAC DCFS executive leadership team in an effort to identify potential reporting or screening practice changes that may address bias.
[Community Engagement]	Findings from the initial review will be shared and feedback sought as to other ways the tool can be used to identify investigations that could be avoided in the future through practice or policy changes.

4. Phases of Work

Five phases of work took this project from conceptualization to implementation of the model, complete with the associated applications that were launched for operational use. These are outlined here. But it is important to reinforce that the work has not stopped—so we have included a Phase 5. The codesign process with frontline staff members and supervisors continues, with expected refinements to the nature of data presented to supervisors. Likewise, LAC DCFS decisions as to the appropriate team and supervision practices for complex-risk investigations will evolve with the help of the CQI team in each regional office. Finally, the Office of Equity is exploring ways that the racial equity feedback information generated through this project might serve as a model for other parts of the system, with the ability to review decisions in real time.

Phase 1: Project Planning and Design

- Project planning with LAC DCFS executive team
- Racial equity meetings with LAC DCFS Office of Equity
- Technical and environmental scan with BIS
- Supervisor shadowing in pilot offices
- Community and staff engagement (ERDD, Invest LA, SEIU)

Phase 2: Model Development and Application Design

- Development server configured by BIS
- Coding of features and other indicators for associated applications
- Initial model development (training and testing)
- Staff codesign and feedback on pilot use case for practice

Phase 3: Preimplementation Model Validation and Application Prototype Design

- External validations of model
- Fairness analyses
- Model quality assurance and refinement
- Finalization of use cases with LAC DCFS leadership
- Review of application prototypes with leadership team
- Pilot office orientation to applications and training
- Test implementation and quality assurance with BIS
- Community presentations of proposed applications (LAC DCFS townhall, provider meeting, SEIU membership)

Phase 4: 90-Day Trial Implementation in Pilot Offices

- Model implemented and applications published by BIS
- Research team site visits to pilot offices
- Mathematica focus groups and outreach to pilot staff
- Weekly meetings to review pilot office data
- Updates made to applications based on user feedback
- Initial data shared with county stakeholders and office staff

Phase 5: Pilot Office Monitoring and Learning

- Development of prototype for monthly management reports (including an inventory of closed investigations for office-specific racial equity feedback reviews)
- Weekly LAC DCFS and research team meetings to discuss data and practice
- Follow-up conversations with supervisors and staff members in pilot offices
- Cases sampled for racial equity review
- Examination of 6-month and 11-month outcomes

5. Funders and Partners

The work described in this report was possible thanks to the philanthropic support of the funders listed here. We are deeply appreciative of their longstanding efforts to improve the use of data in service of children, families, and communities by the dedicated staff of the LAC DCFS. And we are grateful for their collective confidence in our research team.

ballmer GROUP

CONRAD N. Hilton FOUNDATION

REISSA FOUNDATION

1ST LA first 5 la Giving kids the best start

THE RALPH M. PARSONS FOUNDATION

6. Project Team

Research Team	
Children's Data Network University of Southern California & University of North Carolina at Chapel Hill	The CDN is a data and research collaborative focused on the linkage and analysis of administrative records. In partnership with public agencies, philanthropic funders, and community stakeholders, the CDN seeks to generate knowledge and advance evidence-rich policies that will improve the health, safety, and well-being of children.
Centre for Social Data Analytics Auckland University of Technology	Improving the impact of social services, researchers at the Centre for Social Data Analytics work with local and international partners to explore how data can be used to improve the impact of social services.
California Child Welfare Indicators Project University of California, Berkeley	The California Child Welfare Indicators Project is a collaborative venture between the University of California at Berkeley and the CDSS. The project is housed in the School of Social Welfare and provides agency staff members, policymakers, researchers, and the public with access to critical outcome information on California's child welfare system.
Mathematica	Mathematica is an insight partner that illuminates the path to progress for public- and private-sector changemakers, applying expertise at the intersection of data, methods, policy, and practice and translating big questions into deep insights that weather the toughest tests. Driven by its mission to improve public well-being, Mathematica collaborates closely with clients to improve programs, refine strategies, and enhance understanding.
Agency Partners	
Los Angeles County Department of Children and Family Services Pilot Office Leadership & Staff Member Business Information Services Division Continuous Quality Improvement Division	LAC DCFS promotes child safety and well-being by collaborating with communities to strengthen families, keep children at home whenever possible, and connect them with stable, loving homes in times of need.
Office of Equity	The Office of Equity works to infuse equity into the day-to-day work of LAC DCFS and serves as an oversight body to support departmental accountability. The office is composed of sections focused on eliminating racial disproportionality and disparity, improving outcomes for LGBTQ+ youth, and the Women and Girls Initiative.
Invest LA Partnership	Led by LAC DCFS, the Invest LA Partnership includes more than 200 stakeholders from philanthropy, community-based organizations, service providers, and fellow county departments to join as partners in collective efforts to promote system improvements in safety, permanence, and well-being for children.

References

[1] Los Angeles County Department of Children and Family Services. Data and Monthly Fact Sheets. https://dcfs.lacounty.gov/resources/data-and-monthly-fact-sheets/

[2] Miller A. Want less biased decisions? Use algorithms. *Harvard Business Review*. Published July 26, 2018. https://hbr.org/2018/07/want-less-biased-decisions-use-algorithms

[3] Eubanks, V. A child prediction model fails poor families. Wired. Published January 15, 2018. https://www.wired.com/story/excerpt-from-automating-inequality/

[4] Keddell, E . Algorithmic justice in child protection: statistical fairness, social justice and the implications for practice. *Soc Sci*. 2019;8(10):281.

[5] Drake B, Jonson-Reid M. *Administrative Data and Predictive Risk Modeling in Public Child Welfare: Ethical Issues Relating to California*. St. Louis, MO: Washington University in St. Louis, Brown School of Social Work; 2018. https://www.datanetwork.org/wp-content/uploads/ethical-review-of-predictive-risk-modeling.pdf

[6] Los Angeles County Blue Ribbon Commission on Child Protection. *The Road to Safety for Our Children*. Published April 18, 2014. https://dcfs.lacounty.gov/wp-content/uploads/2020/03/Blue-Ribbon-Commission_Final_Report_April_18_2014.pdf

[7] Auditor of the State of California. *Los Angeles County Department of Children and Family Services: It Has Not Adequately Ensured the Health and Safety of All Children in Its Care*. Published May 2019. https://www.auditor.ca.gov/pdfs/reports/2018-126.pdf

[8] Inter-Agency Council on Child Abuse and Neglect. ICAN Reports. https://www.ican4kids.org/reports

[9] Pecora PJ, Chahine Z, Graham JC. Safety and risk assessment frameworks: overview and implications for child maltreatment fatalities. *Child Welfare*. 2013;92(2):143-160.

[10] Bosk E, Feely M. The Goldilocks problem: tensions between actuarially based and clinical judgment in child welfare decision making. *Soc Serv Rev*. 2020;94(4):659-692. https://doi.org/10.1086/712060

[11] Cuccaro-Alamin S, Foust R, Vaithianathan R, Putnam-Hornstein E. Risk assessment and decision making in child protective services: predictive risk modeling in context. *Child Youth Serv Rev*. 2017;79:291-298.

[12] Powell J, Menendian S, Ake W. *Targeted Universalism: Policy & Practice*. Berkeley, CA: University of California, Berkeley, Haas Institute for a Fair and Inclusive Society; 2019. https://belonging.berkeley.edu/sites/default/files/targeted_universalism_primer.pdf

[13] Tibshirani R. Regression shrinkage and selection via the lasso: a retrospective. *J R Stat Soc Series B Stat Methodol*. 2011;73(3):273-282.

[14] Vaithianathan R, Jiang N, Maloney T, Nand P, Putnam-Hornstein E. *Developing Predictive Risk Models to Support Child Maltreatment Hotline Screening Decisions: Allegheny County Methodology and Implementation*. Auckland, New Zealand: Centre for Social Data Analytics; 2017.

[15] Vaithianathan R, Kulick E, Putnam-Hornstein E, Benavides-Prado D. *Allegheny Family Screening Tool: Methodology, Version 2*. Auckland, New Zealand: Centre for Social Data Analytics; 2019.

[16] Vaithianathan R, Dinh H, Kalisher A, Kithulgoda CI, Kulick E, Mayur M, et al. *Implementing a Child Welfare Decision Aide in Douglas County: Methodology Report*. Auckland, New Zealand: Centre for Social Data Analytics; 2019.

[17] Children's Data Network. *Predictive Risk Modeling: Findings from California's Proof-of-Concept*. Published September 2018. https://www.datanetwork.org/wp-content/uploads/PRM_CWDAB_2018.pdf

[18] Tversky A, Kahneman D. Judgement under uncertainty: heuristics and biases. *Science*. 185(4157), 1124-1131. https://www2.psych.ubc.ca/~schaller/Psyc590Readings/TverskyKahneman1974.pdf

[19] Corbett-Davies S, Pierson E, Feller A, Goel S, Huq A. *Algorithmic Decision Making and the Cost of Fairness*. Published August 2017. https://dl.acm.org/doi/10.1145/3097983.3098095

[20] Child Welfare Digital Services. *CWDS Glossary*. https://cwds.ca.gov/project_glossary

[21] Chen C, Liaw A, Breiman L. *Using Random Forest to Learn Imbalanced Data*. Published July 1, 2004. https://statistics.berkeley.edu/tech-reports/666

[22] Kleinberg J, Mullainathan S, Raghavan M. *Inherent Trade-Offs in the Fair Determination of Risk Scores*. Published November 17, 2016. https://arxiv.org/abs/1609.05807

[23] Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 2017;5(2):153-163. https://arxiv.org/abs/1610.07524

[24] Hosmer DW Jr, Lemeshow S, Sturdivant RX. *Applied Logistic Regression*. 3rd ed. New York, NY: John Wiley & Sons.

[25] Palmer L, McCroskey J, Eastman AL, Rebbe R, Guo L, Foust R. *Los Angeles County Family-Centered Services: Using Administrative Data to Understand the Landscape of Community-Based Child Welfare Supports*. Published December 11, 2020. https://www.datanetwork.org/wp-content/uploads/FCS-final-report-12-11-2020.pdf

[26] Vaithianathan R, Putnam-Hornstein E, Chouldechova A, Benavides-Prado D, Berger R. Hospital injury encounters of children identified by a predictive risk model for screening child maltreatment referrals: evidence from the Allegheny Family Screening Tool. *JAMA Pediatr*. 2020;174(11):e202770. https://doi.org/10.1001/jamapediatrics.2020.2770

Contacts

Project & Report Questions

Emily Putnam-Hornstein, PhD
Children's Data Network, Faculty Co-Director
John A. Tate Distinguished Professor for Children in Need
eph@unc.edu
917-282-7861

LAC DCFS Media Inquiries

Media Inquiries:
Office of Public Affairs
Los Angeles County Department of Children and Family Services
publicaffairs@dcfs.lacounty.gov
213-371-6244

More Information

https://dcfs.lacounty.gov/resources/reports/

Report Citation

Putnam-Hornstein E, Vaithianathan R, McCroskey J, & Webster D. Los Angeles County Risk Stratification Model: Methodology and Implementation Report (Model Version 1.0). Children's Data Network (August 2022): https://dcfs.lacounty.gov/resources/reports/