SECTION 7

Allegheny Family Screening Tool: Methodology, Version 2

prepared by Rhema Vaithianathan, PhD (Center for Social Data Analytics, Auckland University of Technology), Emily Kulick (Center for Social Data Analytics), Emily Putnam-Hornstein, PhD (Children's Data Network, University of Southern California), Diana Benavides Prado (Center for Social Data Analytics)

CONTENTS

INTRODUCTION

This methodology report describes changes to the Allegheny Family Screening Tool (AFST), building upon and updating the original methodology report, Developing Predictive Risk Models to Support Child Maltreatment Hotline Screening Decisions (March 2017). The March 2017 report and accompanying documents include background information on the development of the AFST as well as an ethical analysis, impact evaluation and frequently-asked questions. As such, they provide context for this report about changes that have been made to the AFST since the original methodology report was written.

BACKGROUND

In August 2016, Allegheny County introduced a predictive risk model to support decision-making at the time that child abuse and/or neglect allegations are received. Version 1 (V1) of the AFST decision-support tool was in use from August 2016 through November 2018. Since then, a number of modifications have been made to the tool as part of the County's commitment to updating the model and related policies as source systems and variables are updated or policies are revisited. Modifications implemented in Version 2 (V2) of the AFST include changes to specific predictor fields used in the model itself, the modeling methodology, and County policies concerning the tool's use.

This Methodology V2 report provides information about changes made to the tool between the time the first report was written (April 2017) through April 2019. This report upholds Allegheny County's ongoing commitment to transparency by continuing to inform the community about changes to the tool and the County's policies. As this is a status report, details will likely change over time as the County continues to evaluate the impact of the tool and improve its accuracy.

MAJOR CHANGES TO THE AFST SINCE METHODOLOGY V1

Target Outcomes

AFST V1 consisted of two models. The first model, called the placement model, was trained to predict whether, within the two years following a referral, a child would experience a safety issue so significant that they would need to be removed from their home and placed in an out-of-home setting. The second model, called the re-referral model, was trained to predict whether within that same time period, a child who was initially referred and screened out would be re-referred as an alleged victim of maltreatment. Only a single score, the one that was the highest of the placement and re-referral models across all children on the referral, was shared with the call screener. For example, if there were two children on a referral and the older child scored 12 on the placement model and 15 on the re-referral model, and the younger child scored 7 on the placement model and 11 on the re-referral model, the score shared with the call screener would be 15.

The re-referral model (which predicted whether a child who was referred and screened out would be a re-referred within two years) was not as strongly linked to the primary outcome of concern, serious abuse and neglect. One of the reasons that the re-referral model did not have strong face validity is because high scores on that model could reflect children embroiled in custody disputes or other situations where there are frequent calls about the same issue. Additionally, initial incoming referral rates also represent the most racially disproportionate step of the referral pathway, and so a model predicting future referrals tends to overrepresent black children relative to white. Finally, the nature and characteristics of calls with higher scores using the re-referral model were resonating less strongly with screening staff as cases appropriate for investigation. An external validation that examined children's assigned risk scores against their medical encounters for injuries also suggested that the scores from the re-referral model did not create value above and beyond the placement model. In AFST V2, we have therefore restricted the model to predicting safety issues that are so significant that they lead to a court-ordered out-of-home placement outcome.

Predictors

Both V1 and V2 of the AFST use existing administrative data concerning children and adults named in a maltreatment referral to automatically generate a risk score. These integrated data are available to Allegheny County child protection staff through the County's data warehouse and reflect records originating from a wide range of sources. In the two years since AFST V1 was implemented, the characteristics of records and information in the data warehouse have changed as a result of changes made to fields in source data systems. This means some data included in the first release of the AFST may no longer be available or is now available in a different form, while other information is newly available. Changes in the source data systems and predictors used to build the AFST are outlined below.

County data sources *used in **both** the AFST V1 and AFST V2*:

- Child welfare records
- Jail records
- Juvenile probation records
- Behavioral health records

County data sources *used in AFST V2 that were **not** used in AFST V1*:

- Birth records

County data sources ***not** used in AFST V2 that were used in AFST V1*:

- Public benefit records (e.g., Temporary Aid to Needy Families [TANF], Supplemental Nutrition Assistance Program [SNAP])

In some cases, while a data source continued to be used to generate predictors in the AFST V2, the specific fields changed.

Despite the wide array of information about the history of referred individuals available in Allegheny County's integrated data warehouse, the need for call screeners and their supervisors to distill a large volume of information while making quick decisions meant that call screeners historically often relied heavily on the allegation (i.e., the nature of the maltreatment that was being alleged) as a main determinant of screening decisions. AFST V1 did not use allegations as predictors, which might have reduced its face validity with screening staff. AFST V2 includes allegations as additional predictor fields.

Public benefits data were excluded from V2 as these had changed over time and no longer aligned with the data used to develop V1. Additionally, a majority of the behavioral health fields used in AFST V1 were excluded in V2. In recent years, systematic changes occurred in how behavioral health diagnoses were defined and categorized. These changes meant that the behavioral health classifications in the research data used to build the model did not align with definitions currently “feeding” the algorithms. There was no information available that would allow these classifications to be harmonized across the time periods, and the team is working to restructure the behavioral health fields to reincorporate them into the model. The variables will likely focus on service type and severity, with additional predictors to identify if there were any prior services under each diagnostic category. The behavioral health variables that remain in the V2 model reflect aggregated indicators for whether each individual on the referral received any prior behavioral health service, as well as the number of days since the last behavioral health service.

A full list of predictor fields used in AFST V2 can be found in **Appendix B: Weighted Variables in AFST V2**.

Policy

AFST V2 is being implemented with a new visualization that signals to call screening staff that this is a new and improved model. Additional screen shots of the visualizations can be found in **Appendix C: AFST V2 Visualizations**.

FIGURE 1: **AFST V2 visualization provided to call screening staff**

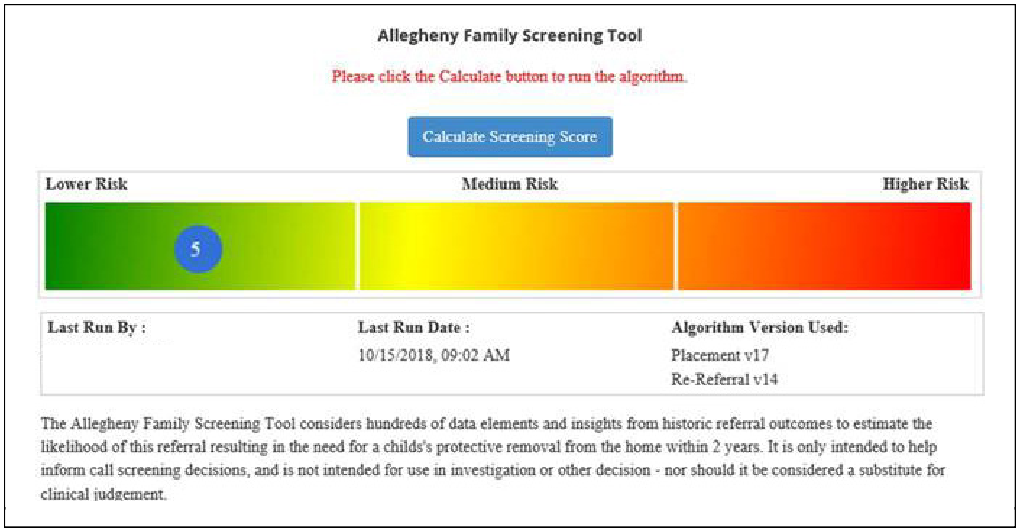

In addition to the new design, newly developed high-risk and low-risk protocols have been implemented. These protocols make use of "nudges," which default the highest-risk cases to be screened in and require supervisors to explicitly override the decision with written justification if they feel it should not be investigated; a similar default-based nudge with override capability was later added to the lowest-risk cases. The visualization displays the high- and low-risk protocol if the referral meets the criteria described in Table 1, below. If the referral does not meet the high- or low-risk protocol, then call screeners see the underlying risk score. The score for each referral continues to be the maximum score received among all children or victims on a referral. As noted above, the maximum score is now derived solely from the placement model, rather than both the placement and re-referral models.

TABLE 1: **Definitions and protocols for high- and low-risk referrals**

	DEFINITION	PROTOCOL	VISUALIZATION	PERCENTAGE OF ALL REFERRALS THAT FALL IN CATEGORY
High-Risk Protocol	Maximum score in a referral of greater than 17 and a victim child (or other child) age 16 years or younger	The referral is designated to be screened-in for investigation; however, supervisory discretion allows screen-out (override documentation required).	The following text is displayed: "High-Risk Protocol, High-Risk and Children Under Age 16 on Referral"	24%
Low-Risk Protocol	Maximum referral score in a referral of less than 11 and ALL victims and children are at least 12 years of age	Screen-out without investigation is recommended.	The following text is displayed: "Low-Risk Protocol, Low-Risk and All Children Age 12+ on Referral" and "recommended screen out".	4%
Other	All other referrals not defined as high-risk or low-risk	Full discretion and no policy recommendation	The categorical score is displayed on a horizontal bar with a gradient of green (1) to red (20)	72%

Modeling Methodology

The AFST V1 was developed using logistic regression; Methodology V1 utilized an Area Under the Receiver Operator Curve (AUC) to measure the probability that a (randomly chosen) referral that was a true positive had a higher risk score than a randomly chosen referral that was a true negative. A probability of higher than 50 percent indicates that the risk score was useful in guiding the screening decision. The Methodology V1 report found an AUC of between 76.9 percent and 78.3 percent.[1] As discussed in our research paper (Chouldechova et al., 2018), however, this reported AUC was over-stated because our split of records into training and testing sets failed to fully address sibling dynamics. Specifically, while referrals had been correctly split so that no unique child appeared in both sets, siblings with the same parent could have been inappropriately split between the test and research data sets. While this does not impair the performance of the previously deployed AFST V1, it does mean that original AUC was overstated.

1 See **Table 4** of Methodology V1 report.

For AFST V2, we explored a range of additional modeling methodologies to improve the AUC, including LASSO, XG-BOOST, Random Forest, and SVM. We discuss that process in detail in **Appendix A: Exploration of Modeling Methodologies for AFST V2**. In deciding which methodology we should adopt, we looked at 1) overall performance and accuracy for the specific high-risk group that serves as the focus of the County's policy, and 2) equivalent levels of accuracy for black children vs. non-black children.

We also gave due consideration to pragmatic questions of implementation and ongoing quality assurance. Given the large number of databases that are being linked in the AFST V2, quality checks and ongoing model maintenance are critical.

We ultimately decided to implement the LASSO model, and the remainder of this methodology report details the performance of that model. To assess whether the overall performance of the LASSO model across children of different racial and ethnic groups was similar, we computed the AUC by race. The overall AUC for the selected LASSO model is 75.97 percent,[2] the AUC for black children is 74.42[3] percent and the AUC for non-black children is 77.35[4] percent, suggesting that the tool was slightly better at predicting outcomes for non-black children than for black children.

2 95% Confidence interval (c.i.): 74.81%–77.13%

3 95% c.i.: 72.84%–75.99%

4 95% c.i,: 75.59%–79.11%

EXTERNAL VALIDATION OF AFST V2

External validation of the model is important to determine if the AFST V2 model, trained to predict the likelihood of a future child welfare out-of-home placement, is sensitive to more generalized and objective measures of child harm. Because true maltreatment rates are very difficult, if not impossible, to determine, we are left predicting measures of child maltreatment defined by the child protection system. As such, there are valid concerns that the AFST model, and other models trained to predict system outcomes like out-of-home placement, may be predicting the risk of institutionalized or system response rather than the true underlying risk of adverse events.

To address these concerns, we completed external validations of AFST V1 using medical records and critical events data. We have replicated those validations for AFST V2, as described below.

External Validation: Hospital Data

To externally validate the AFST V1 model using hospitalization records, we generated a probabilistic linkage between the County's maltreatment referral data and data from UPMC Children's Hospital of Pittsburgh. UPMC proved an ideal source of external data as it is the hospital that the majority of children in Allegheny County use. This means we had near universal medical encounter data (versus means-tested data) for children in the research dataset.

In our initial external validation, we documented that children who were identified in the highest risk groups by AFST V1 were the same children observed to have more generalized risk of relevant hospital events (see pages 19-23 of Methodology V1 for details on how the data was linked and what trends were observed for AFST V1).

We replicated this hospital validation for AFST V2, using the same linked dataset.

We examined hospital encounters (by cause) using four different approaches:

1. *Highest risk score and an injury encounter*: We looked at all unique children in our data, classified their risk based on the highest risk score assigned for any referral, and coded all associated injury encounters, regardless of whether the injury occurred before or after the child abuse and neglect referral.
2. *Randomly selected risk score and an injury encounter*: We looked at all unique children in the data, randomly selected a referral they were involved in and their risk score at that referral,

and coded their associated injury encounters, regardless of when the injury occurred relative to the selected child abuse and neglect referral.

3. *Highest risk score before an injury encounter*: We looked at all unique children in the data and coded the child's risk level based on the highest risk score assigned, but before a specific injury encounter.
4. *Randomly selected risk score before an injury encounter*: We looked at all unique children in the data, randomly selected a referral and associated risk score for each child, and coded a medical encounter as having occurred only if the selected referral date was before the injury encounter.

Figure 2 shows the pattern of medical (i.e., emergency department and hospital) encounters against each of the above different approaches. **Figures A** to **C** show a positive correlation between the AFST V2 risk scores and medical encounters for injury, abusive injuries and suicide. We also examined the association between cancer and the risk score as a "placebo" test; we do not see a strong correlation between cancer and risk scores, suggesting that the AFST is accurately identifying children at risk of abuse-related injuries only.

For more detail on how hospitalized injuries were classified see **Appendix D: Hospital Injury Classifications.**

FIGURE 2: **Children's medical encounters and risk scores**

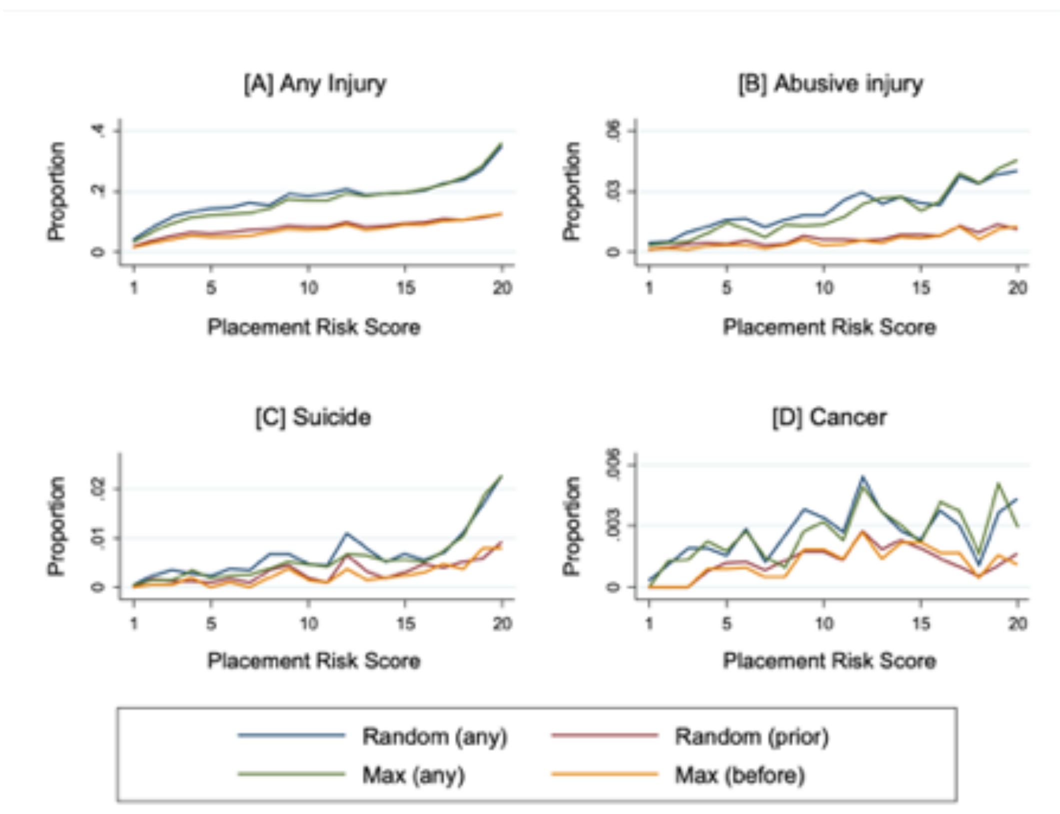

In **Table 2**, we report the odds-ratios for each type of medical encounter following a high-risk referral (as defined in **Table 1**). Note that the odds-ratio for non-black children is larger than for the black children, meaning that non-black children's risk scores at referral were more strongly correlated with later medical encounters.

TABLE 2: **Odds-ratio of medical encounter after referral (high-risk vs. non-high-risk)**

	ALL CHILDREN (N=82,211)	BLACK (N=36,302)	NON-BLACK (N=45,909)
Injury	1.73*** [1.67, 1.80]	1.41*** [1.35, 1.48]	1.89*** [1.79, 2.00]
Abusive Injury	1.46*** [1.34, 1.59]	1.23*** [1.10, 1.37]	1.60*** [1.41, 1.83]
Suicide	1.71*** [1.48, 1.97]	1.30* [1.05, 1.60]	2.23*** [1.83, 2.72]
Cancer	1.23 [0.95, 1.61]	.90 [.61, 1.32]	1.68** [1.16, 2.43]

*Note: 95% confidence interval under odds-ratio. *=p<.1;**=p<.05;***=p<.01.*

CONCLUSION

This report is part of an ongoing commitment to providing both Allegheny County and broader stakeholders with regular updates on how the AFST is evolving over time. We believe that the changes we have made improve the utility of the tool and increase the accuracy of screening decisions.

Evaluations of the impact of the model by independent evaluators are also underway and will be published as they become available.

REFERENCES

Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5-32.

Chen, T., & Guestrin, C. (2016, August). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785-794). ACM.

Chouldechova, Alexandra, et al. "A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions." Conference on Fairness, Accountability and Transparency. 2018.

Csardi G, Nepusz T. The igraph software package for complex network research. *InterJournal, Complex Systems*. 2006;1695(5):1-9.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267-288.

APPENDIX A: EXPLORATION OF MODELING METHODOLOGIES FOR AFST V2

A total of 82,211 unique child-referral observations were extracted for children referred for alleged neglect or abuse in Allegheny County between April 1, 2010 and July 31, 2014. Each observation reflected a unique child-referral record. Some children had more than one referral for a total of 46,507 unique children represented in the data. Outcomes for each child-referral record were observed until the end of the study window, July 31, 2016. To develop the predictive risk model, records were restricted to 45,801 observations in which the child was screened-in for an in-person investigation.

Each child-referral observation was attached to a set of 451 predictive variables describing the characteristics of the child, his/her family, the overall referral, and the alleged perpetrator of abuse. These variables included demographics of the family and alleged perpetrator, allegations associated with the referral, child and mother characteristics at the time of birth, as well as history of interactions with the child welfare system and with other social services such as jail, juvenile probation and mental health. The universe of screened-in referrals were partitioned into a 70/30 training (n=32,224) and validation set (n=13,577).

We used a graph-based method to partition the data into these two sets (Csardi G, Nepusz T., 2006). The method grouped all the children associated with a given referral into either the training or test partition. Because this method can lead to a lack of balance between the test and training partition based on the number of children on the call, balance in the count of children named on the referral was tested with a t-test to compare the average count between test and training set.

The model was trained to predict out-of-home placement within two years of the screened-in referral. Scores were generated at the child-referral level such that each score represents five percent of the referrals. For example, the child-referrals that score a 20 (the highest possible score) fall within the highest five percent of all child-referrals with respect to their predicted probability that the child will be placed in out-of-home care within two years of the scored referral.

Logistic regression method (LR)

This method was used to build an LR model on the training partition of the dataset and was used as the baseline for comparisons to other modeling alternatives.

LASSO regression method (LASSO)

The LASSO model (Tibshirani, 1996) was trained on the training partition using 10-fold cross-validation, with these folds selected randomly. The cross-validated model was trained to optimize for the AUC. The model selected 126 variables as weighted predictors of the target outcome along with the intercept term.

Random Forest (RF)

The Random Forest model (Breiman, 2000) was trained on the training partition with 500 trees and entropy as the splitting criterion. These parameters have been shown to provide the best results in terms of train and test performance in experiments with the Allegheny County dataset.

Appendix A (continued)

XG Boost (XGB)

The XGBoost model (Chen and Guestrin, 2016) was trained on the training partition with the following parameters: 1000 trees, learning rate of 0.01, maximum depth of individual regression estimators of 14, regularization lambda of 80, regularization alpha of 1e-05, minimum number of examples in a node of 1, a subsample ratio of columns per tree of 0.8, a ratio of number of examples of the negative class with respect to the positive class of 4.43, and a subsample percentage of 0.9. These parameters have been shown to provide the best results in terms of train and test performance in experiments with the Allegheny County dataset.

Modeling Results

Figure 3 shows the Receiver Operator Curve for the four modeling methods. As is clear from this Figure, LASSO, RF and XG Boost all perform similarly in terms of general predictive power: LASSO achieves an AUC of 75.97 (95% c.i. 74.81 – 77.16), RF achieves an AUC of 76.34 (95% c.i. 75.18 – 77.5), XGBoost achieves an AUC of 75.83 (95% c.i. 74.67 – 77.0). Logistic Regression achieves an AUC of 64.04 (95% c.i. 62.65 – 65.43), which is significantly lower than the other methods.

FIGURE 3: **Receiver Operator Curve for AFST V2 (test data only)**[5]

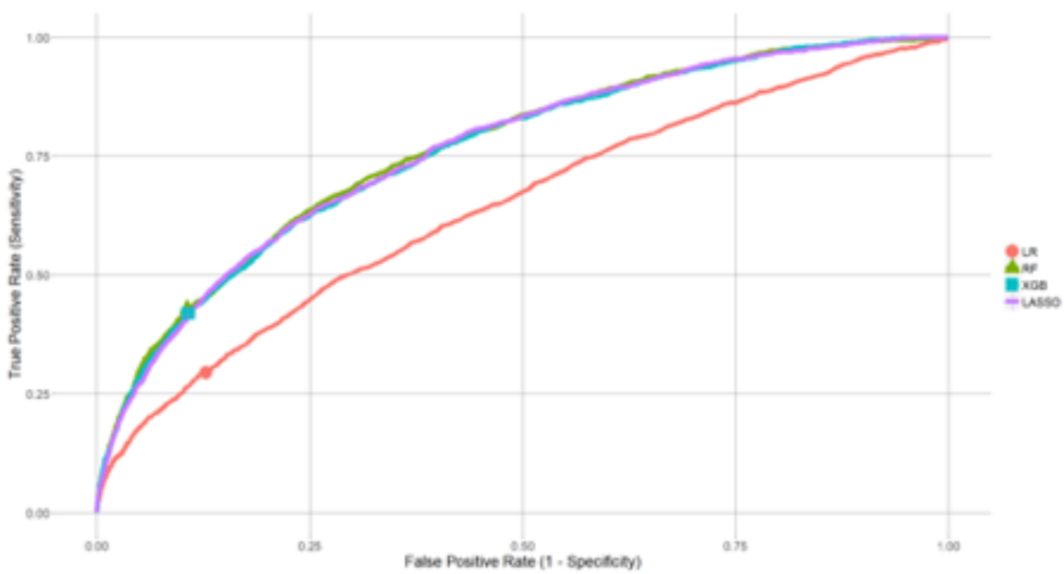

5 LR is a logistic regression model; RF is a random forest model trained with 500 trees; XGBoost is an XGBoost model trained with 1000 trees, a learning rate of 0.01, a subsample ratio of columns for each tree of 0.8, a maximum depth of 14, a minimum child tree node of 1, a regularization alpha of 1e-05, a regularization lambda of 80, a class weight for imbalance of 4.433, a subsample ratio of the training instances of 0.9, all these parameters selected by grid-search.

Figure 4 shows the outcomes by risk score for each of the models (for referrals in the validation partition only). The vertical axis shows the percentage of child-referrals that received a specific score (as noted in the horizontal axis) and in which a child was placed within two years in the test data and the 95 percent confidence intervals. Since Random Forest, XGBoost and LASSO produced the same AUCs, it is not surprising that they have very similar outcome rates by score.

Appendix A (continued)

For example, if we look at only those referrals where a child scored a 20, around 55 percent of those referrals will end up placed within two years – and that rate is the same across all models except for Logistic Regression.

FIGURE 4: **Rates of Placement Outcomes for Four Modeling Strategies (test data only)**

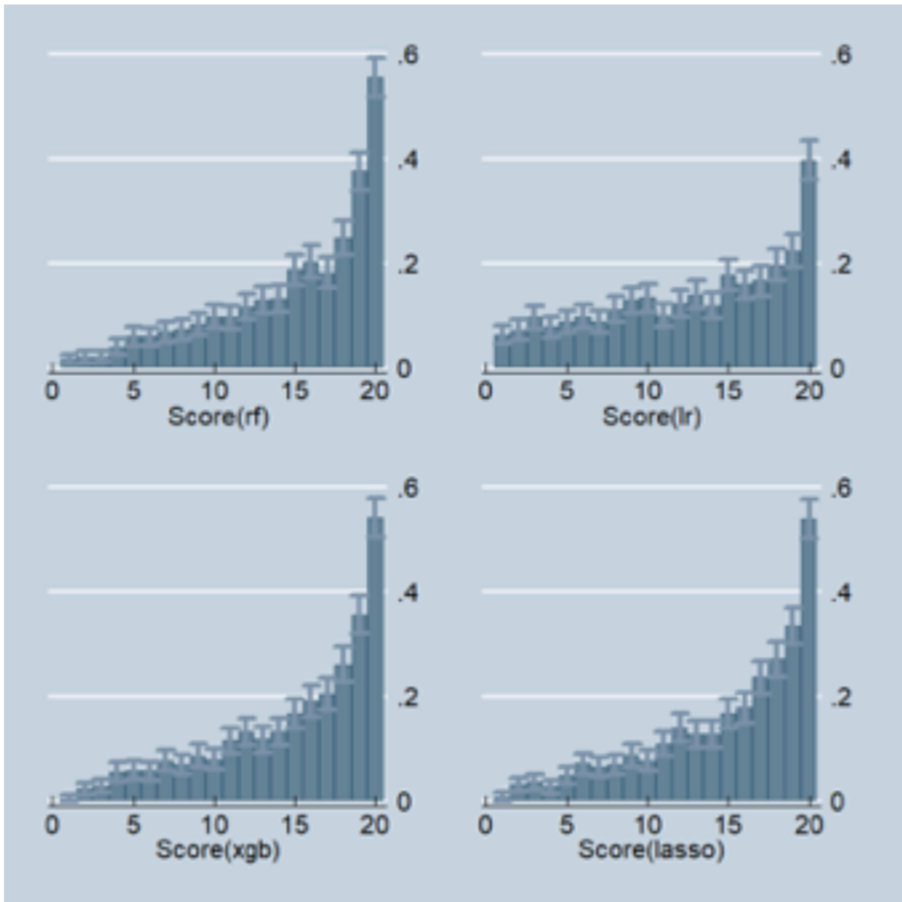

We also looked at fatalities and near fatalities to test whether there were any significant differences in the correlation between the scores and whether a child eventually experiences a fatality or near fatality. To do this, we estimated a logit regression where the dependent variable was equal to one if the child ended up having a fatality or near fatality that met the criteria for review under the provisions of Legislation Act 33 of Pennsylvania's Child Protective Services Law (CPSL) and zero otherwise. We restricted attention to children with a fatality event that occurred more than 50 days after the referral. Table 3, below, reports the results of this regression. The estimated marginal effects of a one unit increase in the predictive risk modeling (PRM) score (e.g., from 5 to 6) on the probability that the child will be a victim of a fatality or near fatality more than 50 days after the referral ranges from 0.074 per 1,000 to 0.059 per 1,000 for the

Appendix A (continued)

various models. All estimated effects are statistically significantly different from zero—but not statistically different from each other.

TABLE 3: **Marginal effect of a 1-unit increase in risk score on probability of a fatality or near fatality more than 50 days after the referral**

MODELING CHOICE	MARGINAL EFFECT (AVERAGE PER 1,000 AND 95% C.I)
LR	0.061 (0.025, 0.097)
LASSO	0.074 (0.040, 0.108)
RF	0.070 (0.035, 0.104)
XG Boost	0.059 (0.022, 0.095)

While the AUCs, outcome plots and mortality regressions provide general information about the accuracy of the algorithm across the range of scores, the more important metric for Allegheny County, given their protocols (as described above in Table 1), is to consider how well it serves to discriminate between high- and low-risk children.

Table 4 shows the positive predictive value (PPV) and true positive rate (TPR) for the four models with respect to the high- and low-risk protocols. The table is at the referral level and uses only the test data (that is, the referrals that were not used to build the models). The top part of the table shows the results for referrals which would have been flagged as high-risk by the tool, i.e., a referral where a child's score is greater than 17 and there is at least one child or victim on the referral who is aged 16 years or younger. The third row of the table shows the percentage of referrals that would be scored as high-risk by the protocol. Because the models identify different families as scoring greater than 17, the percentage of referrals that are identified as "high-risk" depends on the model. All models will flag around 25 percent of referrals as high-risk. The average placement rates for these referrals are between 35.4 percent for the Logistic Regression (LR) and 44.8 percent for the Random Forest. These rates are calculated at the referral level. For example, if the LASSO model were used to identify high-risk referrals, then 25 percent of referrals would be flagged; around 43 percent of all referrals would have at least one child in the referral who was placed within two years. These referrals would account for 55 percent of all referrals where at least one child is placed.

The low-risk protocol flags children as low-risk if the corresponding referral scores 10 or under, and all victims and children in the referral are more than 12 years old (i.e., intake date is after their 12th birthday). These referrals only account for between 2.9 and 9.2 percent of all referrals.

Appendix A (continued)

TABLE 4: **Comparison of Modeling Approaches for High- and Low-Risk Referrals**

	LR	LASSO	RF	XG BOOST
High-Risk Flag *(highest score on referral is greater than 17 and there is at least one child or victim on the referral who is aged 16 years or younger)*				
Proportion of referrals that receive the flag	23.41%	23.8%	25.1%	26.2%
Proportion of referrals flagged high-risk where child ends up placed within 2 years (PPV)	35.4%	47.6%	47.6%	46.2%
Proportion of all referrals where child ends up being placed, who are flagged	39.3%	53.7%	56.6%	57.4%
Low -Risk Flag *(highest score on referral is 10 or under, and all victims and children in the referral are more than 12 years old).*				
Proportion of referrals that receive the flag	9.2%	4.1%	2.8%	2.9%
Proportion of referrals flagged low- risk where child ends up placed within 2 years (PPV)	16.4%	7.6%	5.9 %	4.4%
Proportion of all referrals where child ends up being placed, who are flagged low risk	7.2%	1.5%	0.8%	0.6%

Table 5 shows the share of black children who are identified as high-risk and compares the performance across the models. Across all models, the share of black children flagged in high-risk referrals would be between 28 percent and 37 percent black (in the test sample). Rows 2 and 3 show the relative risk of being placed for black children vs. non-black children. This shows that conditional on race, the models are not miscalibrated in the sense that the relative risk of placement for black children is similar to that for non-black children. The relative risks are most similar for Lasso.

TABLE 5: **Comparison of Modeling Approaches for High-Risk by Race**

	LR	LASSO	RF	XG BOOST
Proportion of children flagged as high- risk who are black	28.1%	36.8%	36.1%	35.2%
Relative Risk of being placed if flagged as High- Risk and black vs. not flagged as High-Risk and black (95% c.i.)	1.95 [1.71, 2.22]	3.10 [2.73, 3.53]	2.72 [2.39, 3.08]	2.54 [2.24, 2.89]
Relative Risk of being placed if flagged as High-Risk and non-black vs. not flagged as High-Risk and non-black (95% c.i.)	1.70 [1.44, 2.02]	3.20 [2.70, 3.73]	3.30 [2.81, 3.88]	3.33 [2.84, 3.91]

Appendix A (continued)

Discussion of Model Choices

In deciding which methodology we should adopt, we looked at 1) overall performance and accuracy for the specific high-risk group that serves as the focus of the County's policy, and 2) equivalent levels of accuracy for black children vs. non-black children.

We also gave due consideration to pragmatic questions of implementation and ongoing quality assurance. Given the large number of databases that are being linked in the AFST V2, quality checks and ongoing model maintenance (e.g., to ensure that there is no feature drift) are critical.

We also analyzed whether the modeling methods resulted in differences in association between the fatalities/near fatalities and the AFST scores. We found that the scores generated by all models show positive correlation with the probability that a child was involved in an Act 33 fatality or near fatality more than 50 days after the score.

LASSO and Logistic Regression approaches, which consist of a simple set of weights, are easier to implement, while Random Forest and XG Boost, consisting of a sequence of linked trees, are hardest because of the difficulties with de-bugging the complex deployed algorithm.

The slight difference in PPVs by race suggests that non-black children are being given too high a score compared to black children. This phenomenon (that we first noted in Chouldecheva (2018)) is similar across all methods. However, when we consider the relative risk (conditional on race) with respect to the high risk protocols being implemented by the County, the models are choosing similarly risky groups.

APPENDIX B: WEIGHTED VARIABLES IN AFST V2

The weights of the model are available upon request from the Allegheny County Department of Human Services.

Definition of suffixes:

vict_othr	All other victim children named in this referral (other than the focal victim child who is being risk scored)
vict_self	The focal victim child being risk scored
prnt	The parent/guardian
perp	The alleged perpetrator. Please note, an individual on the referral could be included in multiple roles (e.g., an individual that is both the parent of the child and the alleged perpetrator).
chld	Other children named in the referral, but who are not identified as the victim

Weighted Variables LASSO:[6]

VARIABLE	DESCRIPTION
INFANT_VIC_NULL	=1 if the victim child <1 year of age at current referral; 0 otherwise
TOD_VIC_NULL	=1 if victim child is btw 1<=age<3; 0 otherwise
SC1_VIC_NULL	=1 if victim child btw 6<=age<9; 0 otherwise
VIC_AGE_SC2_NULL	=1 if victim child btw 9<=age<13; 0 otherwise
TEEN_VIC_NULL	=1 if victim child btw 13<=age<18; 0 otherwise
VIC_1_NULL	=1 if there's a single victim child in the referral; 0 otherwise
AGE_AT_RFRL_MISS_VICT_SELF	=1 if focal child has no age or invalid age; 0 otherwise
CHLD_3_NULL	=1 if there are 3 children involved in the referral who are not identified as victims of the referral; 0 otherwise
CHLD_AGE_INF	=1 if counts of the number of other involved children that are less than 1 year old at the time of referral; 0 otherwise
CHLD_VICTIM_VICT_SELF	=1 if focal child is specifically "Alleged Victim Child" (as opposed to just "Child"); 0 otherwise
FEMALE_NULL	= 1 if victim is female; 0 otherwise
BIO_DAD_NULL	=1 if victim in this referral has a bio dad identified in the relationship table; 0 otherwise
BIO_MOM_NULL	= 1 if victim in this referral has a bio mom identified in the relationship table; 0 otherwise
PERP_0_NULL	=1 if there is no perpetrator in the referral; 0 otherwise
PERP_3_NULL	=1 if there are 3 perpetrators in the referral; 0 otherwise
PERP_AGE_5564_NULL	counts of the number of perpetrators that are 55<=age<65
PERP_AGE_65_NULL	counts of the number of perpetrators that are more than 65
PERP_FEMALES_NULL	counts of the number of perpetrators that were female
PRNT_0_NULL	if there is no person identified as a parent in the referral
PRNT_AGE_19_NULL	counts of the number of parents that are 13<=age<20

6 The full set of variables is calculated for each child on the referral. The variable value is zero if the underlying data required to calculate the variable is missing. In many of the variable categories, an additional variable to indicate if data was missing was included.

Appendix B (continued)

VARIABLE	DESCRIPTION
PRNT_AGE_2024_NULL	counts of the number of parents that are 20<=age<25
PRNT_AGE_4554_NULL	counts of the number of parents that are 45<=age<55
PRNT_AGE_65_NULL	counts of the number of parents that are more than 65
PRNT_OVER2_NULL	if there are more than 2 individuals named on the referral identified as parents
IN_AJD_CHLD	= 1 if the child's MCI ID was created before the referral date; 0 otherwise
IN_AJD_OTH	= 1 if the person's MCI ID was created before the referral date; 0 otherwise
IN_AJD_VICT_SELF	= 1 if the focal child's MCI ID was created before the referral date; 0 otherwise
IN_HOUSEHOLD_NULL	= 1 if the victim is living in the mom's household; 0 otherwise (using InHousehold flag)
REF_PAST365_COUNT_OTH	aggregated no. of referrals in the past 365 days for all individuals involved with role of other (0 if missing)
REF_PAST365_COUNT_VICT_SELF	aggregated no. of referrals in the past 365 days for the focal child (0 if missing)
REF_PAST548_COUNT_VICT_SELF	aggregated no. of referrals in the past 548 days for the focal child (0 if missing)
REF_PAST90_COUNT_CHLD	aggregated no. of referrals in the past 90 days for all individuals involved with role of child (0 if missing)
REF_PAST90_COUNT_VICT_SELF	aggregated no. of referrals in the past 90 days for the focal child (0 if missing)
REFER_TIME_DAY_NULL	=1 if the intake time for the current referral is in the AM; 0 otherwise
PREVIOUS_RFRL_PERP	=1 if the perpetrator has prior referrals; 0 otherwise
PREVIOUS_RFRL_VICT_SELF	=1 if the focal child has prior referrals
FNDG_PAST365_COUNT_CHLD	aggregated no. of founded allegations in the past 365 days for all individuals involved with role of child
FNDG_PAST90_COUNT_CHLD	aggregated no. of founded allegations in the past 90 days for all individuals involved with role of child
SER_PAST180_COUNT_CHLD	aggregated no. of case openings in the past 180 days for all individuals involved with role of child
SER_PAST180_COUNT_PERP	aggregated no. of case openings in the past 180 days for all individuals involved with role of perpetrator
SER_PAST365_COUNT_CHLD	aggregated no. of case openings in the past 365 days for all individuals involved with role of child
SER_PAST548_COUNT_CHLD	aggregated no. of case openings in the past 548 days for all individuals involved with role of child
SER_PAST548_COUNT_OTH	aggregated no. of case openings in the past 548 days for all individuals involved with role of other
SER_PAST548_COUNT_PRNT	aggregated no. of case openings in the past 548 days for all individuals involved with role of parent
SER_PAST548_COUNT_VICT_OTHR	aggregated no. of case openings in the past 548 days for all individuals involved with role of other

Appendix B (continued)

VARIABLE	DESCRIPTION
SER_PAST548_COUNT_VICT_SELF	aggregated no. of case openings in the past 548 days for all individuals involved with role of victim
PLSM_NOW_NULL	= 1 if the referral was received during a placement episode; 0 otherwise
PLSM_PAST180_COUNT_NULL	victim's no. of placement episodes during the last 180 days
PLSM_PAST180_DUMMY_NULL	=1 if the victim was in placement during the last 180 days; 0 otherwise
PLSM_PAST365_COUNT_NULL	victim's no. of placement episodes during the last 365 days
PLSM_PAST365_DUMMY_NULL	=1 if the victim was in placement during the last 365 days; 0 otherwise
PLSM_PAST548_COUNT_NULL	victim's no. of placement episodes during the last 548 days
ALG_PR_12MONTHS_CNT_VICT_SELF	Count of number of total duplicated allegations (regardless of Allegation High Level Category) reported for child in prior 365 days to current referral.
ALGABS_CHLDBHVR_VICT_SELF	=1 if the focal child has an allegation in the Child Behaviors category on this referral; 0 otherwise
ALGABS_CRGVSUBABUSE_VICT_SELF	=1 if the focal child has an allegation in the Caregiver Substance Abuse category on this referral; 0 otherwise
ALGABS_IMMRISK_VICT_SELF	=1 if the focal child has an allegation in the Imminent Risks category on this referral; 0 otherwise
ALGABS_INADHOME_VICT_SELF	=1 if the focal child has an allegation in the No/Inadequate Home category on this referral; 0 otherwise
ALGABS_NEGLECT_VICT_SELF	=1 if the focal child has an allegation in the Neglect category on this referral; 0 otherwise
ALGABS_OTHER_VICT_SELF	=1 if the focal child has an allegation in the Other category on this referral; 0 otherwise
ALGABS_OTHREFSRC_VICT_SELF	=1 if the focal child has an allegation in the Other Referral Source category on this referral; 0 otherwise
ALGABS_PHYALT_VICT_SELF	=1 if the focal child has an allegation in the Physical Altercation category on this referral; 0 otherwise
ALGABS_PHYMALTRTMNT_VICT_SELF	=1 if the focal child has an allegation in the Physical Maltreatment category on this referral; 0 otherwise
ALGABS_PRNTCHLDCNFL_VICT_SELF	=1 if the focal child has an allegation in the Parent/Child Conflict category on this referral; 0 otherwise
ALGABS_SEXABUSE_VICT_SELF	=1 if the focal child has an allegation in the Sexual Abuse or Exploitation category on this referral; 0 otherwise
ALGABS_SEXCNTCTCHLD_VICT_SELF	=1 if the focal child has an allegation in the Sexual Contact Between Children category on this referral; 0 otherwise
ALGABS_TRUANCY_VICT_SELF	=1 if the focal child has an allegation in the Truancy category on this referral; 0 otherwise
ALGABS_UNWILLPRVDCR_VICT_SELF	=1 if the focal child has an allegation in the Unwilling or Unable to Provide Care category on this referral; 0 otherwise
ALGABSP_CHLDBHVR_VICT_SELF	total no. of prior referrals where the focal child had an allegation in the Child Behaviors category

Appendix B (continued)

VARIABLE	DESCRIPTION
ALGABSP_CRGVSUBABUSE_VICT_SELF	total no. of prior referrals where the focal child had an allegation in the Caregiver Substance Abuse category
ALGABSP_INADPHYSCARE_VICT_SELF	total no. of prior referrals where the focal child had an allegation in the Inadequate Physical Care category
ALGABSP_MEDNEGLECT_VICT_SELF	total no. of prior referrals where the focal child had an allegation in the Medical Neglect category
ALGABSP_OTHREFSRC_VICT_SELF	total no. of prior referrals where the focal child had an allegation in the Other Referral Source category
ALGABSP_PHYALT_VICT_SELF	total no. of prior referrals where the focal child had an allegation in the Physical Altercation category
ALGABSP_PRNTCHLDCNFL_VICT_SELF	total no. of prior referrals where the focal child had an allegation in the Parent/Child Conflict category
BC_FEMALE_OR_MISS_VICT_SELF	=1 if gender of child on birth certificate record is female; 0 otherwise
BD_AGE_18_19_VICT_SELF	=1 if father's age was 18-19 at time of child's birth; 0 otherwise
BD_AGE_20_24_VICT_SELF	=1 if father's age was 20-24 at time of child's birth; 0 otherwise
BD_AGE_25_29_VICT_SELF	=1 if father's age was 25-29 at time of child's birth; 0 otherwise
BD_AGE_40PLUS_VICT_SELF	=1 if father's age was 40 or greater at time of child's birth; 0 otherwise
BD_EDUC_MISS_VICT_SELF	=1 if father's education was "Unknown" or missing. This includes actual "Unknown" values, null values, and any invalid values; 0 otherwise
BM_AGE_30_34_VICT_SELF	=1 if mother's age was 30-34 at time of child's birth; 0 otherwise
BM_AGE_35_39_VICT_SELF	=1 if mother's age was 35-39 at time of child's birth; 0 otherwise
BM_AGE_40PLUS_VICT_SELF	=1 if mother's age was 40 or greater at time of child's birth; 0 otherwise
BM_EDUC_BA_OR_HIGHER_VICT_SELF	=1 if mother's education is "Associate degree", "Bachelor's degree", "Master's degree" OR "Doctorate or Professional degree"; 0 otherwise
BM_EDUC_LESS_HS_VICT_SELF	=1 if mother's education is "8th grade or less" OR "9th–12th grade; No diploma"; 0 otherwise
BM_MARRIED_VICT_SELF	=1 if mother is married; 0 otherwise
BM_PAY_MEDICAID_VICT_SELF	=1 if source of payment for delivery was Medicaid; 0 otherwise
BM_PAY_OTHER_VICT_SELF	=1 if source of payment for delivery was other; 0 otherwise
BM_PAY_PRIVATE_VICT_SELF	=1 if source of payment for delivery was private insurance; 0 otherwise
BM_PR_LV_BIRTHS_4PLS_VICT_SELF	=1 if there were 4 or more previous live births; 0 otherwise
BM_SMKD_3MTH_PRIOR_M_VICT_SELF	=1 if cigarette smoking before pregnancy is missing; 0 otherwise
BM_SMKD_3MTH_PRIOR_VICT_SELF	=1 if cigarette smoking before pregnancy; 0 otherwise
BR_MED_PREG_INF_YES_VICT_SELF	=1 if any of infections present/treated; 0 otherwise
BR_MED_PREG_RF_YES_VICT_SELF	=1 if any of risk factors are present; 0 otherwise

Appendix B (continued)

VARIABLE	DESCRIPTION
POVERTY_0_NULL	=1 if poverty rate = 0; 0 otherwise
POVERTY_30OVER_NULL	=1 if poverty rate is greater than 30; 0 otherwise
POVERTY_UNDER20_NULL	=1 if poverty rate is greater than 10 but under 20; 0 otherwise
POVERYRATE_NULL	=1 if no poverty rate available; 0 otherwise
ACJ_1_PER_PERP	% of months seen in Allegheny County Jail last 1 year
ACJ_1_PER_PRNT	% of months seen in Allegheny County Jail last 1 year
ACJ_1_PER_VICT_OTHR	% of months seen in Allegheny County Jail last 1 year
ACJ_1_PERP	total no. of months in Allegheny County Jail in the last year
ACJ_1_PRNT	total no. of months in Allegheny County Jail in the last year
ACJ_1_VICT_OTHR	total no. of months in Allegheny County Jail in the last year
ACJ_2_VICT_OTHR	total no. of months in Allegheny County Jail in the last 2 years
ACJ_3_PER_PRNT	% of months seen in Allegheny County Jail last 3 years
ACJ_3_PRNT	total no. of months in Allegheny County Jail in the last 3 years
ACJ_EVERIN_PERP	=1 if the person was in Allegheny County Jail before; 0 otherwise
ACJ_EVERIN_VICT_SELF	=1 if the person was in Allegheny County Jail before; 0 otherwise
ACJ_NOW_OTH	=1 if the person was in Allegheny County Jail at the time of the referral; 0 otherwise
ACJ_NOW_VICT_SELF	=1 if the person was in Allegheny County Jail at the time of the referral; 0 otherwise
JPO_1_CHLD	total no. of months in Juvenile Probation in the last year
JPO_1_PER_CHLD	% of months seen in Juvenile Probation in last 1 year
JPO_1_PER_VICT_SELF	% of months seen in Juvenile Probation in last 1 year
JPO_1_VICT_SELF	total no. of months in Juvenile Probation in the last year
JPO_3_PER_VICT_OTHR	% of months seen in Juvenile Probation in last 3 years
JPO_3_VICT_OTHR	total no. of months in Juvenile Probation in the last 3 years
JPO_EVERIN_OTH	=1 if the person was in Juvenile Probation before; 0 otherwise
JPO_EVERIN_PERP	=1 if the person was in Juvenile Probation before; 0 otherwise
JPO_EVERIN_PRNT	=1 if the person was in Juvenile Probation before; 0 otherwise
JPO_EVERIN_VICT_SELF	=1 if the person was in Juvenile Probation before; 0 otherwise
JPO_NOW_PERP	=1 if the person was in Juvenile Probation at the time of the referral; 0 otherwise
JPO_NOW_VICT_SELF	=1 if the person was in Juvenile Probation at the time of the referral; 0 otherwise
NO_BH_PERP	=1 if no behavioral health history for this person; 0 otherwise
NO_BH_PRNT	=1 if no behavioral health history for this person; 0 otherwise
NO_BH_VICT_SELF	=1 if no behavioral health history for this person; 0 otherwise

APPENDIX C: AFST V2 VISUALIZATIONS

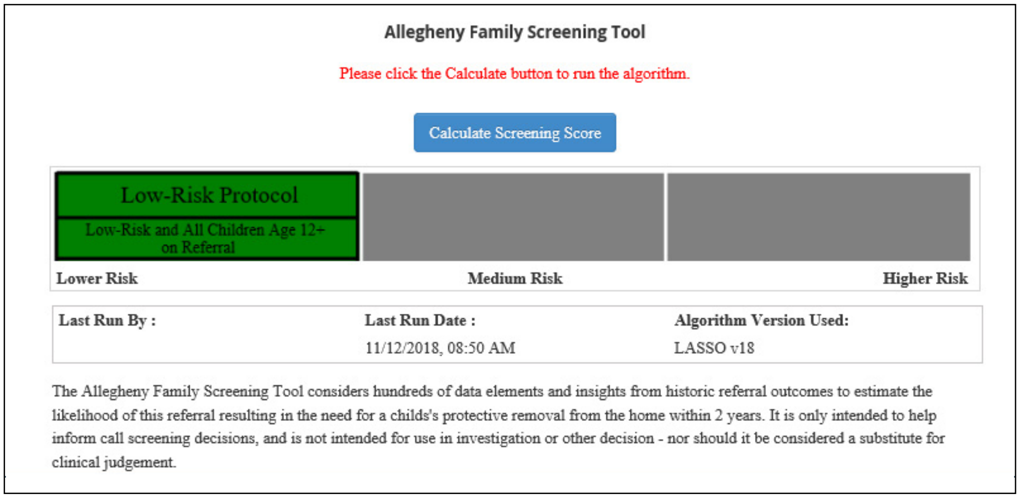

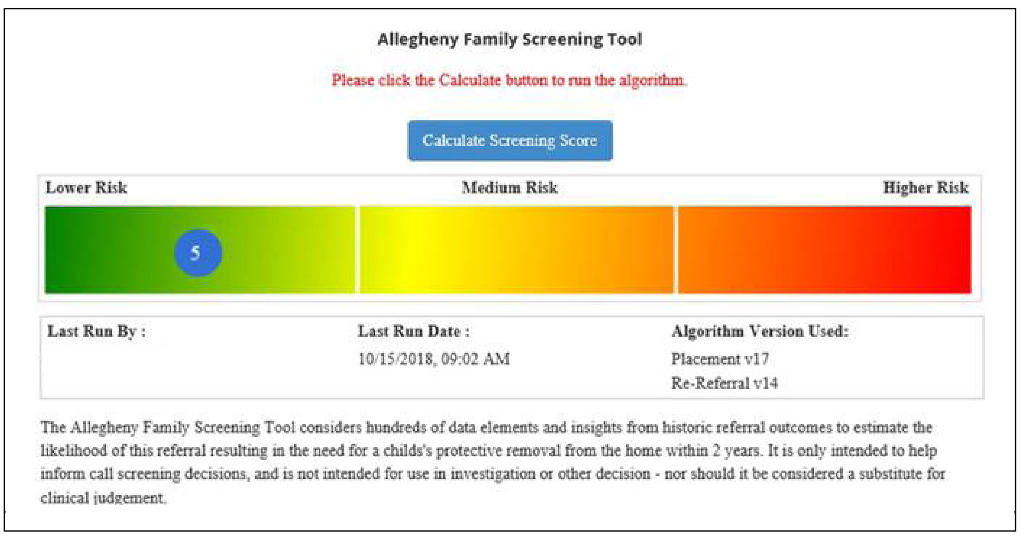

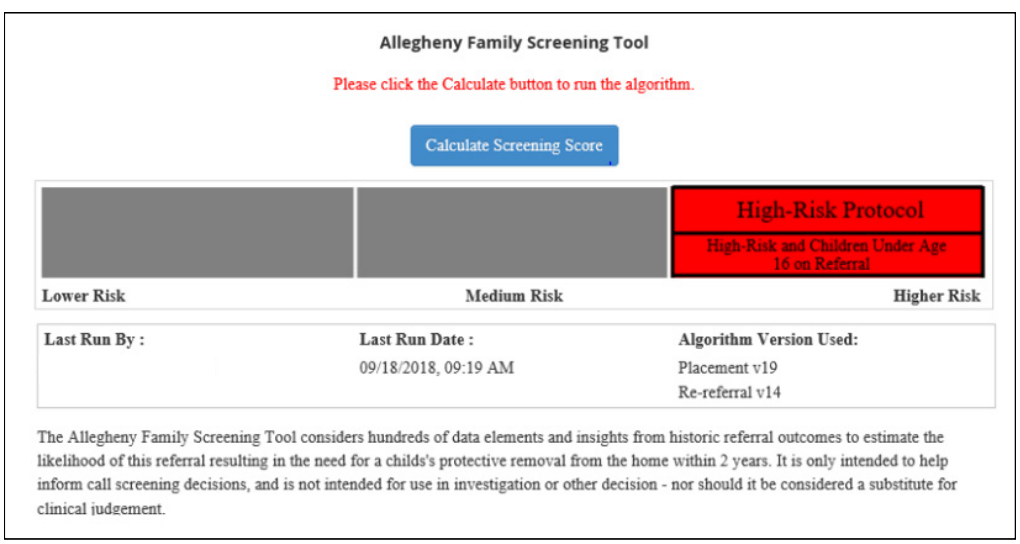

APPENDIX D: HOSPITAL INJURY CLASSIFICATIONS

Hospital Event Injury Type and ICD-9 Codes

INJURY TYPE	ICD9 CODES
Injury from physical activity	E0000-E030; E927-E9282
Injury from transportation	E8000-E848; E9290-E9291
Accidental poisoning drugs/pharms	E8500-E8699; E9292
Injury from medical procedure	E8700-E8799
Accidental fall	E8800-E8889; E9293
Injury from smoke/fire	E8900-E899
Accident climatic or natural disaster	E9000-E903; E9294-E9295
Accident due to abandonment/neglect	E9040-E9049
Toxic reaction from animal or plant	E9050-E9069
Accidental drowning	E9100-E9109
Accidental obstruction respiratory	E911-E9139
Accident struck by object/person	E914-E9269; E9283-E9289; E9298-E9299
Adverse effect therapeutic drug use	E9300-E9499
Self-inflicted injury	E9500-E959
Physical assault	E9600-E978
Injury on accident or purpose	E9800-E989