



Allegheny County Department of Human Services Response to Independent Reviews of the Hello Baby Predictive Risk Model and Planned Intervention and Implementation

September 2020

The Allegheny County Department of Human Services (DHS) solicited the feedback of two independent ethicists regarding the Hello Baby predictive risk model and planned design of the intervention and its implementation. Michael Veale of University College London and Deborah Daro of Chapin Hall at the University of Chicago reviewed the planned design for Hello Baby and explored general ethical considerations. The following outlines DHS's response to the analysis and describes how DHS incorporated the ethical findings into the intervention's design and implementation. We have also responded to each of the reviewers' specific recommendations, which can be found as an appendix to this review. Overall, DHS is in agreement that Hello Baby is an innovative prevention strategy that "could result in families avoiding the most costly and intrusive child welfare services and prevent children from serious physical and emotional injury."

EXPLAINABLE

The ethicists emphasized the importance of understandability for the end user when it comes to the program's eligibility and key tenants. DHS has created a series of communications to support and deepen understanding for end users of the Hello Baby program and the data sources that support the model.

Additionally, DHS has published all of the weighted variables included in the final Hello Baby model and will provide the weights of the variables upon request. It is important to note that this is a prediction model and not a causal model. Therefore, even researchers and experienced data scientists cannot interpret the final list of variables and their corresponding weights. Variables that may independently be strong predictors may have been omitted if they were highly correlated with other variables included in the model.

AUDITABILITY

The ethicists acknowledge the importance of auditability. They emphasized the need for DHS to enable interested third parties to probe, understand, and review the behavior of the algorithm through disclosure of information that enables monitoring, checking, or criticism through the provision of detailed design and technical documentation.

All Hello Baby reports (e.g., methodology, ethics and evaluation) will be publicly available on our website. Previously, DHS has openly shared information about predictive risk models (PRMs) with any interested parties (including public sector entities, researchers, journalists and graduate students) through a variety of mechanisms including community meetings, site visits, webinars and conference presentations, and intends to do so with Hello Baby. Additionally, DHS's practice is to openly invite researchers to use our data and like the Allegheny Family Screening Tool (AFST, a previously implemented PRM), the weights of the model are available upon request from the Allegheny County Department of Human Services.

FAIRNESS

The ethicists suggested that DHS should examine eligibility fairness and bias. The methodology report comprehensively explores these issues and additional reports will be published that provide a more technical evaluation of the PRM tool across racial subgroups, looking at a range of different measures of fairness. We took steps to ensure that the eligibility criteria are fair and establish that we are not failing to offer priority services to children from a specific racial subgroup. These steps to counter racial bias include testing the model against standard fairness metrics, validating the tool's predictive accuracy using a range of objective measures of harm, and excluding predictors that are only observed for a sub-section of the population such as Medicaid-funded services.

CONSENT

Ethicists recommended providing clear language on the opt-out process for the use of data. DHS has put together an [FAQ document](#) with specific language for families and providers. Hospital staff, as the first point of contact with families, will share a handout that introduces the Hello Baby program and outlines the opt-out process. Families will then be sent a mailer to the address provided to the hospital outlining this process a second time. If a family does not elect to opt-out using one of the provided methods before their data is run and scored, they will still be able to opt-out of services as program participation is voluntary. Information about the opt-out process will also be available at any time on the [Hello Baby website](#). Traditional referral pathways into the Hello Baby priority program have been retained so that families who do opt out have alternative ways of getting matched to the services they need.

INTERSECTION WITH CHILD WELFARE

The ethicists acknowledged the importance of fully detailing the intersection of Hello Baby participation and Child Welfare. Parents on an active child welfare case or referral have the same opportunities as all parents to opt out of either the Hello Baby Predictive Risk Modeling process or choose not to receive Hello Baby services despite being eligible. In addition, opting out, refusing, or electing to stop participation in the Hello Baby program will not impact a family's child welfare case or status. Hello Baby services are not mandated, and like all services provided by child welfare, parents are able to refuse or accept services.

DHS USE OF DATA

The ethicists had specific recommendations for the use, storing, and access of the Hello Baby model score. DHS has created policies and procedures to address each of these concerns, which are summarized below.

Use

Scores will not be shared or retained; the only documentation will be the referral to the service for which the family is eligible (i.e., Priority or Support level). Referrals will not have scores attached. There will be no documentation for families falling below the service eligibility threshold (i.e., those who are eligible for Universal services only). Clarification of these standards will be included in the forthcoming practice guidelines and policies.

Storing

The predictive model is run in a database in DHS's Data Warehouse. The data warehouse combines data from many sources and is used specifically for analytics. Scores will not be shared or retained outside this system, which is available only for analytics. These databases will be maintained permanently for all clients for whom scores are calculated and will be used for research and validation purposes as noted.

The system that will be used to provide and track services is the Synergy Family Support Center (FSC) case management tool. The only documentation passed to this system will be the referral to the service for which the family is eligible. This application will contain data only for clients who are eligible for services with a score of 15 or higher. Referrals will not have scores attached, and there will be no documentation for families falling below the service eligibility threshold. Clarification of these standards will be included in the forthcoming practice guidelines and policies. Since only eligible clients are loaded, there is no need to redact this data.

Access

Scores are not accessible to end users outside the analyst teams who perform analytic activities, such as research, validation, and monitoring of performance.

TOOL EVALUATION, QUALITY ASSURANCE, AND POLICY CONSIDERATIONS

Tool evaluation

The ethical analyses found ongoing monitoring, both of the tool and the intervention, to be essential. DHS is committed to ensuring that the model is working properly and will seek feedback from stakeholders, throughout implementation and maintenance. Additionally, given the large number of databases that are integrated to support the Hello Baby data model, quality checks and ongoing model maintenance are critical and will be implemented. The [Allegheny Family Screening Tool](#) was modified several times as part of our commitment to updating the model. See Allegheny Family Screening Tool for more information about this process. Related policies are revisited and updated as source systems and variables change. DHS will publish regular data about the model and the program.

Quality assurance

DHS solicited proposals for both a process and an impact evaluation of Hello Baby and intends to contract with Urban Institute and Chapin Hall who will work in partnership on the process and impact evaluation. DHS and the evaluators will select an evaluation methodology that balances rigor with our responsibility to serve as many families as funding allows. The impact evaluation will examine, among other things, whether the program reduces serious abuse and neglect and improves child and family well-being. The process evaluation will inform engagement and service delivery and will be critical to our understanding of the mechanisms by which Hello Baby works. It will also inform our ongoing quality improvement efforts and provide important information for the program manual we will use during the program's implementation. The evaluation will be in place during the planning and implementation phase in order to guide program development.

Policy considerations

Several policies and procedures were conceived within the context of the ethical considerations:

- a. Policy on the use of, access, and storing of the PRM score
- b. Procedure for community partners to refer to Hello Baby services
- c. Policy about the intersection between Child Welfare and Hello Baby and whether refusing Hello Baby services affects a family's child welfare case.

APPENDIX A

APPENDIX A: DHS'S RESPONSE TO HELLO BABY ETHICIST REPORT BY MICHAEL VEALE

Recommendation 1

Allegheny County should continue to vigorously look for and invest in other ways to reach out to families in need of support services that do not rely on algorithmic systems, to complement the strategies that do.

Response: *Hello Baby isn't replacing existing services; rather, it is adding additional services and an additional pathway for families with young children. Consistent with DHS practice, a full range of supports — separate from the predictive risk model (PRM) — will continue to be available to any family in Allegheny County. These include Family Support Centers and other resources listed on the [Hello Baby website](#).*

Hello Baby's most intensive services will also be available to high-need families identified by community partners (e.g., hospitals, pediatrician, ob/gyn, community agencies) through a referral mechanism that is currently being developed.

Recommendation 2

The County should pledge that this predictive system be used to provide only voluntary supportive services, rather than to start investigations or to directly inform coercive powers.

Response: *Information generated by the algorithm will be used only to determine a family's service eligibility and will never be used in the child welfare intake or investigative decision-making process. Although a child protective services caseworker may infer a family's level of risk by their participation in certain Hello Baby services, the fact that there are other referral pathways to Hello Baby mitigates that risk.*

Clear definitions and restrictions on the use of and access to the data are key elements of the model and will be reflected in the forthcoming practice guidelines and policies.

Recommendation 3

The County should monitor how scoring and service targeting affects the volume of data captured on this subset of individuals and on communities and take appropriate measures to avoid these groups being disproportionately over-surveilled.

Response: *This is a valid concern, not only for Hello Baby but for any prevention service. Various remedies (e.g., excluding prevention-focused services from the data) are being considered. The newly-formed [Pittsburgh Task Force on Public Algorithms](#), of which DHS is a part, will be examining such issues and will inform our guidelines and practice. Additionally, the evidence suggests that if there are surveillance effects, they are empirically trivial.¹*

1 Drake, B., Jonson-Reid, M., & Kim, H. (2017). Surveillance bias in child maltreatment: a tempest in a teapot. *International journal of environmental research and public health*, 14(9), 971.

APPENDIX A

Recommendation 4

Allegheny County should ensure that at least one other screening-in pathway for offering these targeted services exists. Such a pathway should have a primarily clinical element rather than solely relying on an actuarial model and should aim to capture families whose cases and experiences might not be well quantified for procedural or substantive reasons.

Response: *A referral process for community partners is being developed and will be included in the forthcoming practice guidelines and policies.*

Recommendation 5

Scores should be redacted on production to “Ineligible” if they fall below determined thresholds for service ineligibility, rather than unnecessarily keeping them in numeric form. A proportionate exemption to this could be determined if they were deemed necessary for validation or research, in which case these should be kept separately and redacted in the operational system.

Response: *Scores will not be shared with service providers; the only documentation will be the referral to the service for which the family is eligible. There will be no documentation for families falling below the service eligibility threshold (i.e., those eligible for Universal services only). Clarification of these standards will be included in the forthcoming practice guidelines and policies.*

Recommendation 6

At the point at which there are no plans to make use of the value of a score through a standard operating procedure, the score should be purged.

Response: *See above.*

Recommendation 7

If there are justified cases for the score being retained for downstream service provision, tailoring or as a flag for re-engagement, the score should be made as coarse as is possible and compatible with those purposes.

Response: *There are no situations in which the score will be used for downstream service provision or as a flag for re-engagement.*

Recommendation 8

The County should examine options for creating a pledge in a binding and public manner to ensure that the intended purposes of the system do not expand without a clear and accountable process and that the safeguards do not diminish.

Response: *The newly-formed [Pittsburgh Task Force on Public Algorithms](#), of which DHS is a part, will be examining such issues and will inform our guidelines and practice.*

APPENDIX A

Recommendation 9

The County should produce and make public a strategy for ensuring that the system and DHS's expertise is disseminated in the most responsible manner possible to other interested public sector bodies. With this strategy, DHS may wish to consider legal instruments such as licensing schemes and institutional arrangements such as secondment and collaboration agreements.

Response: *Consistent with our practice, all Hello Baby reports (e.g., methodology, ethics and evaluation) will be publicly available on our website. Previously, DHS has openly shared information about PRMs with any interested parties (including public sector entities, researchers, journalists and graduate students) through a variety of mechanisms including community meetings, site visits, webinars and conference presentations, and intends to do so with Hello Baby.*

Recommendation 10

Allegheny County should ensure that easy-to-understand explanations of why an individual was offered targeted services are provided upon request. In documentation or during outreach, the opportunity to request this information should be actively presented.

Response: *We are working with the Camden Coalition to develop appropriate communication strategies and are currently user-testing suggested language. Practice guidelines will reflect the results of that testing. With support from the Camden Coalition, all relevant staff will be trained in the forthcoming practice guidelines.*

Recommendation 11

Allegheny County should ensure that individuals can request all information that relates to them used by the County in the creation of this score and that individuals are actively informed about their ability to do this.

Response: *ClientView is an existing DHS tool that allows clients to access their own service records (as well as service plans and assessments). We are creating a "client portal" (AccessMyInfo), a webpage that will provide residents with a place to find local resources and also provide easy access to ClientView (which is being redesigned and will also be renamed). In order to access their private records, clients will use a pin number they will receive upon enrollment in DHS services. The pin number will be automatically generated and included in a letter sent to clients; it will also be available by submitting a form online or requesting it from a caseworker or provider.*

The forthcoming practice guidelines and policies will include additional ways in which to inform clients of these resources.

APPENDIX A

Recommendation 12

Allegheny County should seek to publish online versions of the models it produces after ensuring sensitive data is secured.

Response: *This will be included in the methodology report that will be published on our website.*

Recommendation 13

Allegheny County should publish and maintain “model-centric” explanations online: metadata about the models including at least their main inputs, logics and optimization targets; performance in practice, including on sensitive and salient subgroups; and practices and processes around model maintenance and use.

Response: *The methodology report, process and impact evaluations, and outcomes of the [Pittsburgh Task Force on Public Algorithms](#) will all address this recommendation.*

Recommendation 14

Allegheny County should invite researchers to examine and audit their systems and give them heightened access to software, data and frontline workers. A list of such collaborations should be published.

Response: *It is, and will continue to be, DHS practice to openly invite researchers to use our data, following the process outlined [here](#). The process and impact evaluations will address practice issues. The results of these evaluations will be published on our website.*

Recommendation 15

Allegheny County should seek to work with researchers to create and make publicly available peer-reviewed research on the models’ characteristics and their use in situ.

Response: *See above*

Recommendation 16

The County should begin a program of work on discrimination and inequality in information use and service provision more broadly which involves, but is not limited to, statistical understandings of fairness. Where deficits in modelling are discovered, particular attention should be paid to how different interventions might seek to mitigate these, rather than solely attempting to adjust the model to compensate.

Response: *DHS extensively explored the statistical fairness of the model. The methodology report describes these findings. This work will also be informed by the results of the [Pittsburgh Task Force on Public Algorithms](#).*

Recommendation 17

The County’s work in building the above strategy should be deeply participatory and open to comment.

Response: *See above*

APPENDIX A

Recommendation 18

Maintenance of the model over time should have a strong qualitative dimension, seeking feedback from different sources of on-the-ground knowledge, such as frontline workers, to understand how the model performs on different groups, and how the phenomena being modelled might be changing over time.

Response: DHS is committed to ensuring that the model is working properly and will engage in quality assurance, including seeking feedback from stakeholders, throughout implementation and maintenance. Additionally, given the large number of databases that are integrated to support the Hello Baby data model, quality checks and ongoing model maintenance are critical and will be implemented, including testing for feature drift, and other sources of model degradation. The Allegheny Family Screening Tool was modified several times as part of our commitment to updating the model. See [Allegheny Family Screening Tool](#) for more information about this process. Related policies are revisited and updated as source systems and variables change. DHS will publish regular data about the model and the program.

Recommendation 19

The County should develop methods to monitor oversight, critical examination and the use of these scores by all those who have access to them or are tasked with approaching families for voluntary targeted services.

Response: See above.

Recommendation 20

The County should ensure that investment decisions do not create undesirable over-reliance on scores, such as through under-investment in alternative sources of information or through reduction of time available for each task due to efficiencies of automated systems.

Response: As has been our practice, PRMs will continue to serve as a decision support tool, used in conjunction with clinical judgment and practice. Our commitment to clients' access to their own information; our plan to encourage referrals from community providers; and our evaluation and quality improvement processes that monitor outcomes based on both PRM scores and clinical judgment all confirm our commitment to the appropriate use of technology in a human services environment.

APPENDIX B

APPENDIX B: DHS'S RESPONSE TO HELLO BABY ETHICIST REPORT BY DEBORAH DARO

Issue/Recommendation

Little information is provided as to how the opt-out process will be explained to families. Using such language as “we will be determining what other needs you might have by reviewing your prior experiences using information maintained by state agencies” may well raise red flags for families, particularly those who are already involved in the child welfare system or who have a less favorable view of public agencies

Response: *Language is currently being user-tested and will be informed by the responses. Appropriate language will be included in the forthcoming practice guidelines and policies.*

Issue/Recommendation

One would not be able to obtain participant-level data on those opting out of the program, but aggregate numbers would provide some indication if those opting out include a higher proportion of parents presenting some of the common characteristics associated with future parenting difficulties.

Response: *The outcome evaluation will examine group differences between those who participate and those who choose not to do so.*

Issue/Recommendation

The County is electing to use “passive consent” presumably to maximize the number of births available to be screened using the PRM. This approach is considered appropriate only if the intervention or strategy involves minimal risk to the participant and if obtaining written approval for the procedure is not practical or feasible. It is not clear if this approach has already been approved by the County’s Institutional Review Board. If it has, then the approach has been judged appropriate in this instance. If it has not, the County will need to make the case as to why it is not asking parents to “opt in” for the screen.

Response: *The “passive consent” is only for running the PRM, which commits clients to nothing. If they are determined to be eligible for services beyond the universal supports, they will be contacted and still may decline to participate with no adverse effects. Even after agreeing to participate, families may decide to opt out of services at any time.*

Allegheny County does not have an institutional review board.

Issue/Recommendation

Program material did not make clear what the County will do if an active CPS (child welfare-involved) parent “opts out” of the PRM process. Formal procedure needs to be put in place as to how these cases will be addressed. You will need to review with them how refusing services or dropping out of the program will impact their CPS status.

Response: *The current thinking is that refusing or dropping out of the program will not impact a family’s CPS status. Hello Baby services are not mandated, and like all services provided by CPS, parents are able to refuse or accept.*

Detailed information on this topic will be included in the practice guidelines and policies.

APPENDIX B

Issue/Recommendation

While not all outcomes associated with an intervention can be predicted, the burden on the provider is to be as transparent and direct as possible at the onset of services and address any questions participants raise in a complete and honest manner. DHS should explain how the family was selected (they may or may not recall any conversation in the hospital at the time their baby was born) and the risks and benefits associated with accepting these services.

Response: *Healthy Start and DHS will develop scripts/discussion points, currently being user tested, prior to implementation. This will be included in the forthcoming practice guidelines and policies. Research into community perceptions of PRM tools tells us that families do not necessarily want to be approached using the lens of risk but to focus on strengths.² As a study participant in research undertaken on algorithm use in the public sector stated, “The computer tool is the ‘why’ to approach a family not the ‘way’ to approach a family.” We are cognizant of the need to engage using a strength-based approach.*

Issue/Recommendation

Offering “priority” to these families when allocating available service slots could result in fewer prevention services being made available to families who, while challenged, fail to score at the highest end of the PRM scale.

Response: *With limited resources, it’s important to identify those most in need so that they can receive the appropriate services. We maintain that Hello Baby will help us to more effectively deploy services and, rather than decrease services, will actually free up resources for those who need them most.*

Issue/Recommendation

May alter not only the aggregate success these programs have achieved in other domains but also create an assumption that such services exist primarily to avoid child abuse and neglect.

Response: *Although we are trying to avoid child abuse and neglect, we are not promoting Hello Baby as such but rather as a support for new parents. Because it is not affiliated with child welfare/child protective services, there is no reason to link the two.*

2 Brown, Anna, et al. “Toward algorithmic accountability in public services: A qualitative study of affected community perspectives on algorithmic decision-making in child welfare services.” *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 2019.

APPENDIX B

Issue/Recommendation

Very few of the variables in the model address proximate challenges. No variables in the model reflect a parent's current access to health care, informal social supports and current psychosocial health, characteristics that, when lacking, are often associated with poor health outcomes for both the child and mother.

Response: *That's true, we are looking at a specific subset of data that has been correlated with poor outcomes (maltreatment) by age three — and we have specifically focused on not making poverty a factor. The providers will use other family factors to tailor the approach for each individual family.*

Issue/Recommendation

Beyond its potential to alter the overall composition of the participant pool enrolling in prevention services, the PRM also may impact provider perceptions and behaviors toward their clients. The PRM gives service providers additional information on a family's history that may alter the way workers interpret the conditions they do observe.

Response: *The sole purpose of the Healthy Start partnership is to support families at highest risk. Healthy Start has been engaging with families for decades and the staff is trained to avoid stigmatization. Staff will know that the referrals they receive are high-need families, but they won't know any details and won't know if the family was referred by the PRM or by a community partner; they will learn about the family's individual needs as the relationship develops and will develop a plan in partnership with the family.*

Issue/Recommendation

Heightened awareness of a family's circumstances may create surveillance bias, resulting in a higher probability of a family being reported.

Response: *Yes, surveillance bias is a danger of all home visiting programs and indeed any service that connects with families and children. It's true that more people in the house increases the possibility that they will observe maltreatment. However, a key element of our training will be on how to differentiate between true maltreatment and indicators of poverty, which can masquerade as maltreatment for untrained staff. Additionally, the evidence suggests that if there are these effects, they are trivial, and certainly not of a size to outweigh the real benefits of prevention programs.³*

3 Drake, B., Jonson-Reid, M., & Kim, H. (2017). Surveillance bias in child maltreatment: a tempest in a teapot. *International journal of environmental research and public health*, 14(9), 971.

APPENDIX B

Issue/Recommendation

Implementation should be carefully assessed and the link between program characteristics and participant outcomes monitored.

Response: *These elements will be included in the process and outcome evaluation design.*