



**CENTRE FOR
SOCIAL DATA ANALYTICS**

Implementing the Hello Baby Prevention Program in Allegheny County

Methodology Report

Prepared by the Centre for Social Data Analytics at the Auckland
University of Technology

September 2020



Note:

This report discusses Version I of the Hello Baby predictive risk modeling (PRM) County Tool, which was in use as of September 2020. Therefore, the information contained in this report is preliminary and based on a model that has been and continues to be revised as we learn more about the impact of the PRM. An outside evaluation will also inform modifications of the model. This paper will be updated and revised to reflect changes.

Please contact the authors for up to date information about the PRM tool.

Authors:

This report was written by Rhema Vaithianathan, PhD (Auckland University of Technology), Emily Putnam-Hornstein, PhD (University of North Carolina) and Diana Benavides-Prado, PhD (Auckland University of Technology).

The project was led by Rhema Vaithianathan with the contribution of the following staff and contractors at the Centre for Social Data Analytics: Emily Kulick, Nina Anchugina, Huyen Haley Dinh and Katerina Taskova (now at the University of Auckland).

Contact: rvaithia@aut.ac.nz

Acknowledgments:

The research team would like to acknowledge the leadership of Marc Cherna and Erin Dalton at the Allegheny County Department of Human Services.

Citation: Vaithianathan, Rhema; Diana Benavides-Prado and Emily Putnam-Hornstein. *Implementing the Hello Baby Prevention Program in Allegheny County*. Centre for Social Data Analytics. Auckland, New Zealand. September 2020.

Executive Summary

While Allegheny County (Pennsylvania) has a rich array of supports for families, many of these services are not accessed by families with the most complex needs. We find that the vast majority of families in programs such as home visiting do not fall into the categories of “high” or “complex” needs. This lack of connection with the families most in need means that the County is missing opportunities to prevent harm and protect children.

The Hello Baby program was launched in September 2020 as a better way to engage and support new parents in Allegheny County. The program includes matching a family’s needs and preferences to existing services, alongside wholly new services designed to more directly support families. The objective is to work actively with parents to help children thrive and to prevent adverse outcomes.

The Hello Baby program has a graduated approach to services, implementing the following tiered structure:

- **Universal Tier:** Universal services will be offered to all Allegheny County families with a baby born at any birthing hospital in the County. These services include access to resources on the Hello Baby website and the Hello Baby support-line, which will offer families a single resource to identify and access information on available programs. This may include referrals to [Family Centers](#), determining eligibility for state and federal programs that may have a means-tested component to them, and access to Nurture PA, a text-based mentoring program for new mothers, in addition to other County programs. (While all families can access services from the Universal Tier, when we refer to families in the Universal Tier in this methodology we mean those who have not been identified as eligible for the Family Support or Priority Tiers.)
- **Family Support Tier:** Families identified with moderate needs will be prioritized for support from outreach workers from an existing network of community-based resource hubs called Family Centers. Families are able to self-refer to access these resources. Under the Hello Baby program, Family Center staff do not wait to see which families find their way to services. Instead, staff actively contact eligible families, based on a predictive risk modeling tool, helping them to navigate the array of services offered in the County (e.g., home visiting, family support, childcare subsidies).
- **Priority Tier:** Families with the most complex needs will be contacted by a family engagement specialist and a social worker as part of the *Hello Baby Priority* tier offered through the community-based agency, Healthy Start, Inc. After obtaining a family’s consent, the family support team will undertake an initial needs assessment, ensure the family’s prioritization into existing services, and provide ongoing support and case coordination. Families can disengage, or re-engage, with services at any point.

While all families with newborns will have the opportunity to participate in the Universal tier, referrals to the Support and Priority tiers can originate through three different approaches: (1) self-referral through a Family Center or by calling the Hello Baby support-line; (2) provider referral from a birthing hospital, midwife, physician, or other clinical professional; and (3) the Hello Baby predictive risk model

(PRM), which uses universal County records to prioritize higher needs families with newborns for the Family Support and Priority tiers.

By adopting the predictive risk modeling (PRM) approach, the County is able to augment traditional self-referral and provider-referral pathways, with proactive outreach. This overcomes the limitations of traditional pathways which rely on families or their providers knowing about the program, referring families appropriately and doing so early enough to be effective.

The Hello Baby PRM tool uses universal data about the family held in existing administrative data systems falling under the authority of the County. The PRM tool draws upon a set of weights applied to 59 different existing universal data features for a newborn and their family. These features emerged from historical Allegheny County data and were chosen based on their relationship to a child's future risk of experiencing severe abuse and neglect and a removal from the home by age three. A family's service tier ranking is also strongly correlated to a variety of other childhood adversities including the mother being booked in jail, homelessness and maternal mortality, post-neonatal infant mortality, and preventable injury death. Compared to other ways of determining eligibility for services, which rely on a few criteria such as Medicaid eligibility or maternal age, the PRM is a much more accurate way to identify children who are at risk for suffering harm.

The ability of a program like Hello Baby to address and reverse racial disparities depends, ultimately, on its efficacy in serving and truly meeting the needs of Black communities. The PRM tool functions as a small part of the overall program, helping to ensure families who may not traditionally self-refer, or who may not interact with providers who actively refer, are still given an opportunity for eligibility in the Family Support and Priority tiers. That said, we have taken steps to ensure that the eligibility criteria are fair and to establish that we are minimizing the unwarranted selection of children from specific racial sub-groups for Family Support and Priority services. This has included testing the model against standard fairness metrics, validating the tool's stratification accuracy using a range of external measures of harm, and excluding predictors that may be available for only a subset of newborns (e.g., Medicaid-funded services). Additionally, ongoing monitoring and evaluation of the program will occur to ensure that services are effective for all subgroups and lead to help and support, not more surveillance.

The County's experience with another project where a PRM tool was used, the [Allegheny Family Screening Tool](#), has been instructive, suggesting that such tools can reduce racial disparities.

The key to success lies in implementation by the providers who will need to effectively engage and then serve families with culturally appropriate and responsive services. Allegheny County is fortunate in being able to partner with Healthy Start, Inc. as the provider for families on the Priority Tier. They have a strong track record in engagement – and a commitment to reducing racial disparities.

All providers of prevention programs – including Healthy Start workers – are mandated reporters. This means that if they observe abuse or neglect in the families they serve, they are required to report this to Child Welfare. An anxiety amongst some communities is that families that agree to services might end up with increased referrals to Child Welfare. This phenomenon is referred to as “surveillance bias” and is sometimes raised in opposition to prevention programs. However, the evidence suggests that if these effects exist, they are small, and certainly not of a size to outweigh the real benefits of prevention

programs.¹ Additionally, the County has undertaken extensive training of Healthy Start staff to ensure that they understand the signs of abuse and neglect so that they are not unnecessarily referring families to Child Welfare.

The County 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 these evaluations. DHS and the evaluators will select an evaluation methodology that balances rigor with the County's 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 The County's 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 the County 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.

The Pilot: Year One

Hello Baby Priority referrals will be limited to a subset of births in the Allegheny County region during the initial implementation year. For this pilot, there are nine Family Centers and Healthy Start locations that will receive referrals based on PRM eligibility for services. Evaluation of program implementation will guide expansion of services to make Priority tier services available to selected families with complex needs countywide.

Overview

There is little doubt that the prevention of maltreatment, rather than its treatment, is critical to reducing the individual and public health burden of abuse, neglect, and their sequelae. Despite many notable advances in our understanding of child maltreatment, however, primary prevention is an area in which the field's knowledge of what interventions work, for which children and families is incredibly underdeveloped.

There is limited knowledge as to, exactly *how*, and with *what supports*, health and human service systems can most effectively move upstream to engage families in voluntary programs that prevent maltreatment and reduce the need for child protective services.

It is clear that efforts to prevent maltreatment require innovative new approaches.

The Hello Baby program reflects a universal, county-wide effort to redesign and re-focus family support services and prevention dollars. This effort involves a universal, hospital-based program established to engage and enroll families in early intervention programs tailored to the service needs of a given newborn and their family. This strategy aligns with the policy framework outlined by Daro and Dodge (2010) which "would embed high-quality targeted interventions within a universal system of support that begins with an assessment of all newborns and their families". Although much of what defines

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.

Hello Baby is grounded in decades of research concerning family engagement, childhood adversities, and child development, selection into a pathway of voluntary Hello Baby services is unique in that it is not based exclusively on a clinician-dependent referral system or a diagnostic screening. Rather, this program will include a population-level risk stratification model to ensure children and families most in need of services are appropriately identified and prioritized for services.

In this methodology document, we outline the fundamental challenges of developing eligibility rules that ensure that all families have access to the supports and services of adequate dosage and intensity.

Section 1: Background

As federal investments in programs and evaluations reflect (Boller et al., 2014), there is a widespread appreciation that the success of various early intervention programs (e.g., home visiting) are dependent upon model fidelity (Daro, 2010), and that successful implementation must take into account the systems in which children and families are served (Foster-Fishman, Nowell, & Yang, 2007; Kilburn, 2014). Importantly, this strategy reflects a pragmatic recognition that ensuring fidelity to evidence-based models is critical and more effective than, for example, prescribing features that home visiting programs should have. Unfortunately, despite these efforts, the field is not yet in a position to generalize the relationships between specific program features and observed outcomes (Daro, 2010; Kilburn, 2014).

It is also not clear how systems operating upstream of child protective services (CPS) can most efficiently and effectively serve a specific case-finding function for engaging high-risk children and families—providing a means of early identification and referrals to home visiting and other prevention programs. This case-finding function is important because it has implications for how screening and referral programs may be most successfully embedded in current systems that touch families and young children, and because effective case finding ensures that limited resources and home visiting services slots are effectively targeted to families with children who are at greatest risk of maltreatment.

Despite evidence suggesting that the greatest costs and benefits of high-intensity home visiting programs occur for the highest risk children and families (Kilburn, 2014), to date there have been no population-level examinations of different approaches that could be employed to ensure children at greatest risk of maltreatment, specifically, are being identified and triaged into these very services. Indeed, research from other countries suggest that the families accessing these home visiting services often fall in lower-risk groups (Vaithianathan et al., 2019).

Given the absence of universal home visiting or intensive family support programs in the United States, knowledge of which children and families are identified and then referred to prevention programs is critical. This information is needed because the County has an ethical obligation to ensure that when limited service slots are available, children and families with the greatest concentration of risks are appropriately prioritized for services. Similarly, the County has a fiscal imperative to ensure that when programmatic effects vary across population risk levels, efforts are made to thoughtfully triage clients to maximize benefits relative to costs.

As jurisdictions move deliberately forward with an agenda of research to build the evidence base concerning the prevention of child maltreatment, it is crucial to develop a better understanding of

opportunities to be increasingly strategic in the selection of children at high risk of abuse or neglect for referral to prevention programs.

Equally challenging is determining how to most effectively harness existing early intervention services to strengthen families and prevent the conditions that lead to CPS involvement during the first few years of a child's life. Hello Baby offers an opportunity to examine the extent to which a multi-tier family support program, with a risk and needs stratified prioritization model effectively identifies, engages, and serves children at the highest risk of future adversities, including child abuse and neglect.

Section 2: Local Context

Over the past two decades, the Allegheny County Department of Human Services (DHS) has collaborated closely with the Allegheny County Health Department, the community, and faith-based organizations to deliver prevention services. These services aim to strengthen families, improve children's outcomes, and reduce involvement with CPS. Most of these services have been focused on reaching families in higher-poverty neighborhoods of the County. Providers operate with a high level of respect for family and community strengths; efforts have succeeded in greatly increasing resources to families in Allegheny County. DHS contracts with several providers to operate a network of 26 [Family Centers](#) (FCs) across the county. Annually, close to 6,000 children and their families are enrolled in one of the neighborhood-based centers. FCs are community hubs where parents can attend programs, access resources, and connect with other parents. Centers are staffed by individuals who often live in the neighborhoods they serve. FCs are governed, designed, and improved by participants and community members. A [recent evaluation](#) found that areas within Allegheny County served by FCs had fewer maltreatment investigations once the level of social disadvantage and population size were considered (Wulczyn & Lery, 2018).

Robust collaborations in Allegheny County have made more services available to families, yet historically low parental uptake and difficulty with ongoing engagement mean many families have either never participated in preventive community services or have only received very light touch interventions. Using matched data from DHS's [Data Warehouse](#) and records from encounters with early childhood prevention services, analysis shows that among children born in Allegheny County only 15% (2,014 children of approximately 13,141 born per year), receive any prevention service. And, despite 26 FCs across the County, as few as 3% receive services from FCs. Additionally, while infant mortality has been falling for the County as a whole, for Black infants the rates remain high and persistent mortality disparities remain.

While family support and maltreatment prevention are county priorities, so too is child protection in the more traditional sense (i.e., after allegations of abuse or neglect have been made). Yet, the limitations of a system that must wait to provide services or interventions until *after* allegations are received becomes clear when one considers that in over half of critical incidents in Allegheny County (i.e., cases where a child died or nearly died as a result of abuse and/or neglect), there had not been a child protection referral prior to the event. This means the County had no opportunity to support the family before the child was injured. The unfortunate reality is that a system that provides support services only *after* a

child comes to the attention of CPS is unlikely to reduce the prevalence of child maltreatment in the population.

It is certainly possible that referrals for maltreatment lead to protective interventions and therefore children who do not have contact with child protection face a higher risk of severe abuse and neglect. Yet, an equally plausible conclusion is that many of these children come from families with quite apparent indicators of high and complex needs, but these risks go undetected and unrealized up until the point that a child is harmed. In other words, these may have been families who were simply not offered services of sufficient intensity early enough in a child's life. DHS explored some characteristics of high need families to understand their system involvement rates prior to birth and interaction with prevention services post birth. While the data available in the DHS Data Warehouse and included in this analysis is only a partial view of early childhood prevention services recipients in Allegheny County, we found that the vast majority of children in families enrolled in home visiting programs were not classified as high or complex needs. Thus, with limited resources generally, and a restricted number of service slots for intensive in-home services specifically, it is essential to better target and service families with children who are at greatest risk of maltreatment.

To focus support and services on families who could most benefit from them, DHS sought a proactive mechanism to better identify families with children at risk of adverse outcomes, including serious maltreatment. The goals were to:

1. Reach and engage more families who might benefit from support.
2. Connect families to services *before* incidents of abuse and neglect.
3. Ensure that families with the most complex needs have priority access to the most intensive supports the County offers.

Section 3: The Hello Baby Program and Eligibility Levels

Graduated Services as Universal Prevention

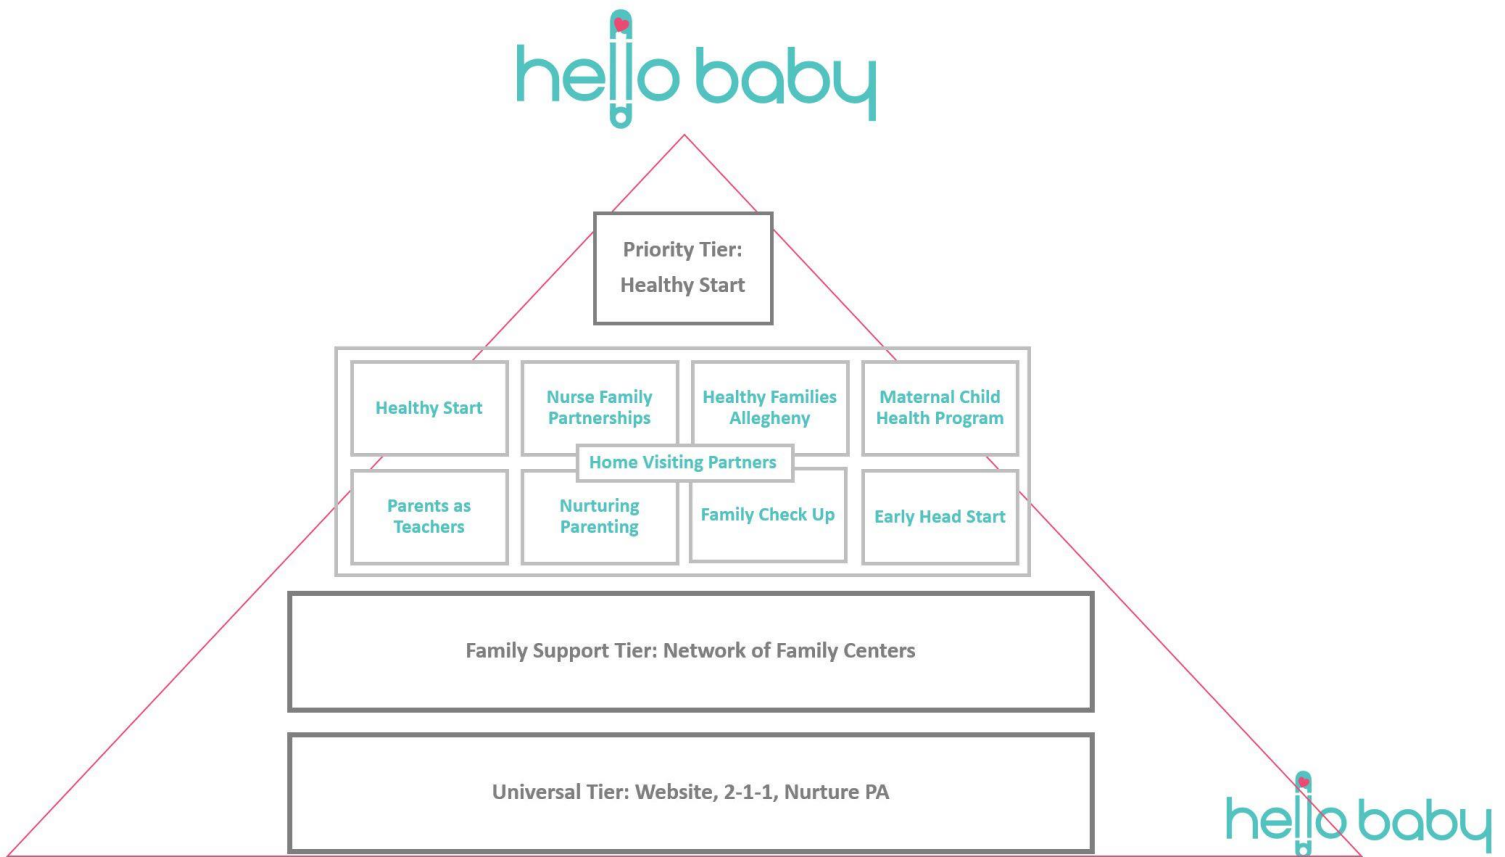
In the context of a universal prevention program everyone is, by definition, eligible for the service. While there are many benefits associated with universal services (e.g., widest reach, no eligibility rules, reduced stigma associated with service usage), the downside is that services are stretched thinly across many families. This means that often those families with the greatest needs may not receive the intensity of services or supports required to truly change outcomes, while other families attract support disproportionate to their needs.

Hello Baby was designed to provide “graduated services” within a universal prevention program, with progressively more intensive engagement for families determined to have higher needs. Under this

program, some families will be given prioritized access to more intensive service offerings, and therefore eligibility criteria are required. The approach is described below and depicted in Figure 1:

- **Universal outreach** begins either at one of Allegheny County’s primary birthing hospitals, or one of the obstetric or pediatric practices throughout the County. Every new Allegheny County mother will be given information about the Hello Baby program and may receive an initial visit from a nurse to talk about the program in more detail. A few weeks later, a postcard will be sent to reiterate the introductory information and let families know about the Hello Baby support line (available 24/7), the Hello Baby website (an inventory of parenting support information), and other available community resources for parents such as NurturePA, a texting service staffed by volunteers who can answer questions about newborns and parenting issues. These services are available for all families in the County, regardless of need.
- **Hello Baby will reach families with moderate/higher levels of need** through the *Hello Baby Family Support* tier which uses the existing network of Family Centers (FCs). FC outreach workers will visit families in their neighborhoods, welcoming the new babies and attempting to engage and connect parents with the rich support and services they offer (e.g., home visiting, family support, childcare subsidies, Head Start, basic needs). In addition, any family can request services from an FC through the Hello Baby website, warmline, or texting service.
- **Families with greater and/or more complex needs** will be eligible for the *Hello Baby Priority* tier. This constitutes a two-person team established to engage families, assess their needs, and to jointly, and with the family, develop a plan that leverages parental strengths. This tier has been designed to clear barriers to service connections for as long as families wish or until their child turns three years of age. This assistance can take the form of concrete goods (e.g., diaper, formula, food), assistance with immediate needs such as transportation to the pediatrician and eviction prevention, and connections to community resources such as evidence-based parenting programs, mental health and/or substance use treatment, child care and emergency housing.

Figure 1: Hello Baby Tiered Services



Program Selection Approaches

There have typically been three ways in which families are selected for different levels of graduated services: (1) professional judgement; (2) threshold criteria; or (3) a structured assessment tool. Eligibility criteria that depend solely on the professional judgement of a service provider (e.g., family physician, social worker) are increasingly rare as it has become clear that professional judgement alone is poor at establishing long-arc risk, compared with structured tools that rely on objective criteria (see Schwalbe, C. [2004]. Re-visioning risk assessment for human service decision making. *Children and Youth Services Review*, 26(6), 561-576. and papers cited therein). There are also concerns about equity of access when relying solely on professional judgment to choose families. This approach is likely to favor families who are already engaged in services and known to providers, leaving other families with unmet service needs to be overlooked.

Another way to determine eligibility is to use threshold (or rules-based) criteria. This typically involves setting easily verifiable administrative criteria (e.g., Medicaid eligibility, parental age, or marital status) to determine program eligibility. Since a single factor is often too inclusive and would leave too many families eligible for limited slots, often there are a set of eligibility conditions that must be met. For example, the well-known Nurse Family Partnership (NFP) requires that participants be first time mothers, enrolled no later than 28 weeks, and meet income eligibility requirements. The limitation of this approach is that these eligibility requirements are not always associated with a relative risk of

adversities, so these programs may not serve those who are in the most need. As reported in Section 7 of this report, we found that using Medicaid eligibility and other traditional methods of selecting “high risk” families resulted in a group of eligible families that have lower relative risk of adversities in a range of outcomes as compared to the group identified by a PRM.

Finally, structured assessment tools function as a more complicated threshold model, requiring risk factors to be disclosed by families, often through lengthy questionnaires with items that are aggregated. Families are determined eligible by meeting a certain number of positive risk factor responses. These questionnaires can often be highly invasive, and many families might be reluctant to disclose the full extent of their needs for fear of stigmatization. Indeed, certain groups such as those with a history of criminal justice involvement might be reluctant to share their history because they fear it might act against their interests.

Risk Stratification Model

All families of newborns residing in Allegheny County will be welcomed to, and can participate in, the Hello Baby program. To determine prioritized eligibility for more intensive levels of Hello Baby services (i.e., Hello Baby Family Support tier, Hello Baby Priority tier), the County will use a predictive risk model (PRM), coupled with complementary professional and self-referral pathways.

As with structured assessment tools, PRMs can be thought of as a type of threshold tool, but PRM’s overcome numerous limitations of current approaches to determining program eligibility. The primary strength of a PRM is that it is able to use verifiable data retained in the County systems about a family without any additional data collection. This means that rather than using one or two observable criteria, or subjecting families to long and invasive questionnaires, PRM tools can draw upon a more complex and nuanced set of criteria. Additionally, the attributes can be chosen and weighted to help identify families at highest risk of the types of adverse events that the program is trying to prevent, a step which is not typically taken with structured assessments. An advantage of PRM tools that are based on existing data in existing systems is that families can be proactively informed about their eligibility – and this can occur as soon after the baby is born. This reduces the information barriers that families face. Additional advantages include the ability to screen the entire population of interested families with newborns in the county, on average 13,000 per year, and help identify families who are in the highest need groups. It also counters the natural inclination of frontline services to fill their service slots with families who are easily engaged – rather than those in greatest need. For example, evidence from a New Zealand predictive modeling study showed that the majority of children who were enrolled in an intensive home visiting program did not in fact come from the riskiest families as identified by a predictive risk model (Vaithianathan, 2019).

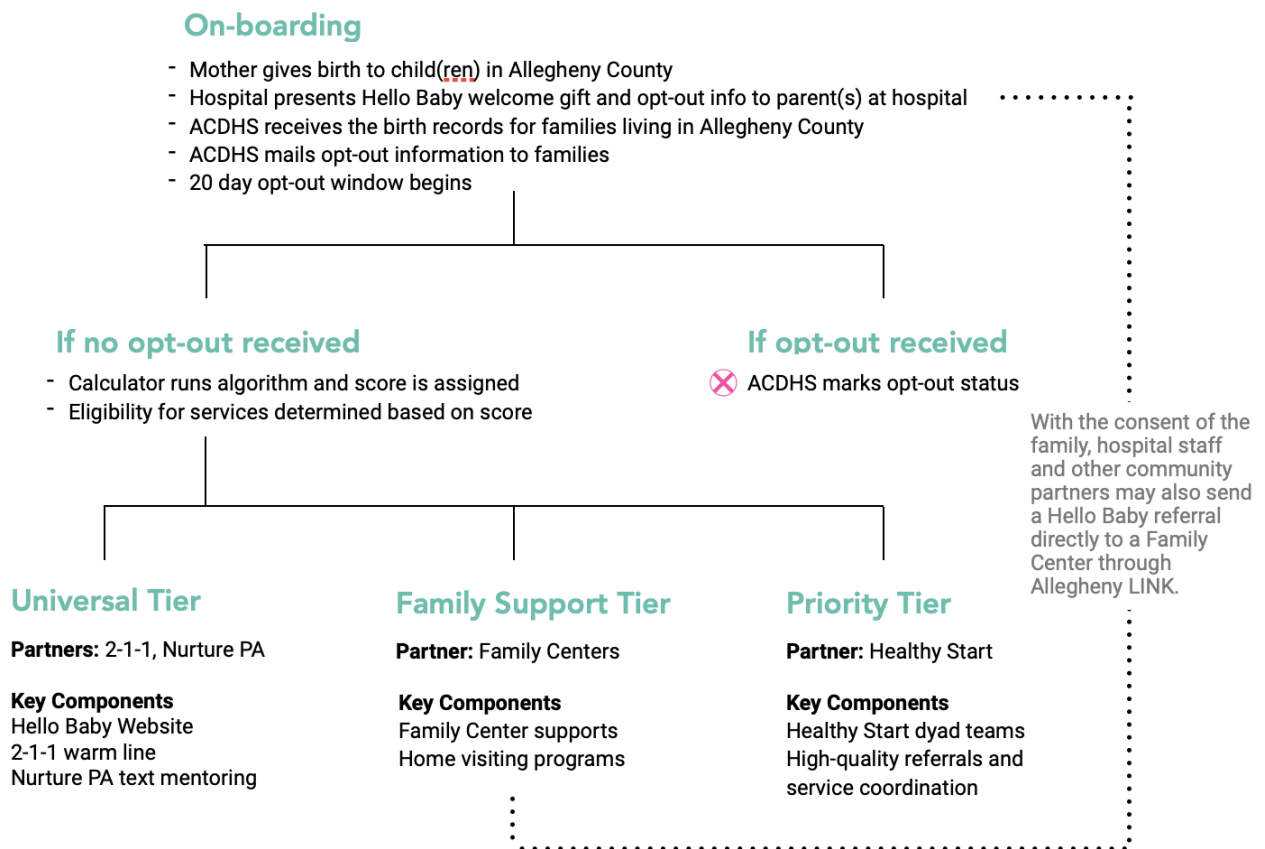
Service Pathways

Beginning with the launch of the program in September 2020, every family will learn about Hello Baby at the hospital and/or their obstetrician/gynecologist or pediatrician’s office after giving birth in Allegheny County and will receive a postcard in the mail a few weeks later. The postcard will provide general information about Hello Baby and how to access support as well as explaining how families can opt-out. If the family doesn’t opt out, DHS will use the Hello Baby PRM to determine if they are eligible for additional supports.

Universal services will be offered to all Allegheny County families with a baby born at any birthing hospital in the County. Families with moderate needs will be contacted by outreach workers at the Family Center in their geographic region. Families with the most complex needs will be offered more intensive services by Healthy Start, Inc. as the community partner to provide this service. Healthy Start teams (family engagement specialist and social worker) will reach out to these families to better determine their needs and to connect them to the best resources the County has to offer.

While the Hello Baby PRM will be the primary mechanism to identify families eligible for more intensive services, community partners can also make referrals for families that they feel would benefit from additional support. These community partners may include Family Center staff and other home visiting providers. In addition, any family that feels the need for additional support can contact Hello Baby staff to request services. The figure below outlines the pathways to service across program tiers, beginning at the time of birth. Additionally, all families can reach out to any named partner organization at any time to request services.

HELLO BABY SERVICE PATHWAYS



Section 4: Model Methodology

In this section we describe how we used historical data to train and validate the Hello Baby PRM tool. The tool is designed to be used shortly after the birth of a child and calculates the risk that a newborn child is removed from home and placed in foster care within the first three years of life. The tool is then used to allocate children who have the highest risk to the Priority group, middle risk to the Family Support group and the rest to Universal group.

Background

All PRM tools must be “trained” using historical data. Training is a statistical method for building a PRM model. This model is a set of rules that summarizes the correlation in historical data between predictors or features of the child and family and the training outcome (removal from home within three years). This means that when a child is born, the set of rules can be used together with the predictors of the child and families to calculate the chance that the child will have the removal outcome.

The Hello Baby PRM was built using anonymized records for children born in Allegheny County between 2012 and 2015. Birth records were used to define a cohort of children who were both born during this four-year period and where maternal residence fell within the county (N=52,520).

Each birth record contained information concerning not only the child born, but also the mother and father associated with that birth. Information about the father is missing for 12.6% (6,636 of 52,520) of the records and information about the mother is missing in only four of the records. This allowed for a range of self-reported health, sociodemographic, and service history information to be coded for the child and their family.

An initial set of 459 features or predictors were extracted and coded for newborns and their parents. All feature data was housed within the County’s integrated Data Warehouse. Only features captured within “universal data systems” – or systems that are not explicitly means-tested – were included. This means that past encounters with the criminal justice system were coded as features because there is only a single criminal justice system (vs. a public and private system). In contrast, features related to a parent’s mental health history were not included as only individuals receiving public health insurance (vs. private health insurance) have records contained in the County’s Data Warehouse. Race was tested as a predictor and found not to be useful, and therefore excluded from the final candidate set of predictors. However, this is not to say that race might not play a role in the modeling. In Section 7 we go into greater detail on how the model performs across racial sub-groups and we also provide more technical analysis in forthcoming reports.

Table 1 summarizes the data sources with features that were tested in the final model. This table has been organized to present the county data source from which records were extracted, the individual or unit at which features were coded, examples of features derived from a given data source, and the total number of features from this particular source and unit of analysis. Although most features were coded to reflect the characteristics or history for a given individual (e.g., mother, child, father), child protection and child welfare records from the Office of Children, Youth, and Families were coded at the family level. Column four provides the full count of the candidate set of features that we experimented with.

The Appendix provides a detailed inventory of all features included in the model and their coding. Weights attached to features in the final model are available on request.

Table 1: Predictors utilized for the Hello Baby PRM tool

Data source	Individual	Features (sample)	Training features used in initial modeling (count)
Vital Birth Records	Child	Gender; Birthweight; Preterm birth	15
	Mother	Age; Prenatal care; Prior births; Education level; Smoking in last 3 months	56
	Father	Presence of father on the birth certificate; Age; Education level	15
Children, Youth and Families	Mother	Count of prior referrals; Count of prior episodes of involvement (active cases) *	53
	Father		51
Allegheny County Jail	Mother	Count of months with at least one jail booking*	5
	Father		5
Allegheny Courts	Mother	Count of months active with Common Pleas Criminal court; Count of months active in Magisterial District Court*	60
	Father		60
Homeless and Housing Support Services	Mother	Count of days in an emergency shelter; Count of episodes in transitional housing / rapid re-housing; Count of episodes in permanent supportive housing*	41
	Father		38
Assisted Housing	Mother	Count of months receiving housing support through the Housing Authority City of Pittsburgh; Count of months	20

	Father	receiving housing support through the Allegheny County Housing Authority*	20
Juvenile Probation	Mother	Count of months active in juvenile probation, Count of prior juvenile probation placements*	10
	Father		10

* Features are coded across time periods including within the last 1 year, 2 years, 3 years, ever, and active at the time of the birth.

Outcome(s)

Predictive risk models are trained to predict a target outcome or outcomes. Since we are looking to use the PRM to stratify birth cohorts for different levels of Hello Baby services, the ideal outcomes to train the model are adversities that children experience that with appropriate services can be reduced or averted. A variety of outcomes were tested for training the Hello Baby model (e.g., child welfare referrals, case openings, removals, removals for physical abuse alone). Likewise, different lengths of postnatal follow-up time were used to observe each outcome (e.g., one year, two years, five years).

Ultimately, one measure of future child welfare involvement was selected and defined as the target for model training: a removal and placement in foster care between a child’s birth and their 3rd birthday. The decision to focus on this outcome was guided by considerations of: (a) modeling performance (i.e., the outcome that led to the greatest risk differentiation), (b) face validity (i.e., removal and placement in foster care reflects interaction with the child protection system we would hope to prevent through earlier services and supports), and (c) external validity (i.e., placement can be defined as an “outside of system” indicator in that a removal to foster care involves court actions).

One reason that children’s early years were chosen as the target for model training is that early experience of adversities has long term consequences; research has shown that toxic stress in early childhood leads to epigenetic changes and behavioral alterations resulting in poor outcomes in multiple domains, including physical and mental health, learning and behavior. The risk factors studied in the Adverse Childhood Experiences Study² suggest that child abuse or neglect (along with maternal depression and substance use) is a major source of toxic stress. By training the model on an outcome that captures an indicator of severe abuse and neglect experienced during these early years, we focus on a crucial period in childhood development.³

For the cohort of children born between 2012 and 2015, 18.1 per 1,000 children (952/52,520) were removed and placed in foster care at some point between birth and age three.

² Felitti, V. J., Anda, R. F., Nordenberg, D., Williamson, D. F., Spitz, A. M., Edwards, V., & Marks, J. S. (1998). Relationship of childhood abuse and household dysfunction to many of the leading causes of death in adults: The Adverse Childhood Experiences (ACE) Study. *American journal of preventive medicine*, 14(4), 245-258.

³ Garner, A. S., 2013. Home visiting and the biology of toxic stress: opportunities to address early childhood adversity. *Pediatrics*, 132(Supplement 2), S65-S73.

Model Training

To build an eligibility model for the Hello Baby program, features were coded for each birth record from historical records and used for model training purposes. This training process allowed for an examination of the association between the initial slate of features and our defined outcome. The performance of a risk stratification model is assessed based on the combination of feature weights that best predict the outcome serving as the training target.

The likelihood of experiencing a removal by age three (target outcome) was modelled using the LASSO (Least Absolute Shrinkage and Selection Operator) Regularized Logistic Regression algorithm. The LASSO regularized form of a Logistic Regression model ensures certain predictor weights are set to zero while minimizing prediction error, given the sum of the absolute value of the weights is less than a constant. Thus, it is capable of both predictor selection and regularization, which results in more easily interpretable and accurate models.

Firstly, we partition our research data of 52,520 set into a training set containing 39,365 records and a testing set of hold-out records of 13,155. Siblings and step-siblings were forced to be in either the train or the test splits to ensure that the test data set was independent of the training data set. Furthermore, individuals who died within the outcome period were excluded from the training data.

Then, the risk model was instantiated through the R package named 'glmnet' (Friedman, 2010). In the model training phase, the aforementioned LASSO constant, which is often symbolized by lambda, was optimized within the range of $1e-04$ to 100. The lambda parameter was tuned through a 10-fold cross-validation procedure that was repeated three times, yielding three random partitions of the original training set. Those results were again averaged to produce a single estimation. Among 100 different lambda values we tested, the model that corresponds to the best lambda was considered as the final model.

We should note that before arriving at the final modeling approach outlined above, multiple models were trained and tested, including those that captured features that were ultimately dropped: (1) those originating in public behavioral health systems and (2) measures of race/ethnicity. Race/ethnicity was excluded because when we met with members of the community, it was immediately apparent that its inclusion raised concerns that bias would be directly introduced into the model. Additionally, race added no predictive power to the models. Forthcoming technical reports will provide more details.

Likewise, features coded from the County's behavioral and physical health data holdings were excluded because we only had these records for the subset of people for whom these services were paid through the public Medicaid system. Privately funded physical and behavioral health records are not available and there were concerns that if the model was trained based on only a subset of all behavioral health history, it might bias the stratifications and classifications produced by the model.

Alternative modeling approaches including Support Vector Machines, Random Forest, RiskSlim and Boosted Trees were also tried, but there was minimal additional accuracy and the challenges of deploying some of these methods meant that we went with LASSO. Forthcoming technical reports will provide more details.

Classification Errors

The use case for a PRM to establish the level of service eligibility for the Hello Baby program is fairly simple: to provide a mechanism for universally screening the population of newborns in the County to identify those born into families with the greatest need and then prioritize those families for supports and services. By stratifying the newborn population, and offering graduated services, families with the greatest need can be given priority access to higher intensity services, which have limited service slots due to their associated cost.

In this context, there are two primary “classification mistakes” that can be made by the PRM tool. The first is to identify a newborn and their family as having a high need for services, even though this family would not have had adverse outcomes regardless of services or supports offered. This type of error (called a “false positive”) may lead to the potential provision of case management and other support to a family that may not need them. Whilst unlikely to be harmful to these families, there are real inefficiencies with these types of errors as it means that children with higher needs go unserved.

The other type of classification mistake that can be made occurs when the model fails to identify a newborn and their family as high need (called a “false negative”). In this case, families receive a lower tier of services than is required, and because those services were inadequate, the child experiences preventable adversities. The failure to flag a family with high needs means that the County misses an opportunity to prevent harm and is of particular concern.

These two concepts are measured by the *True Positive Rate (TPR)* (which tells us the proportion of families who experience the adversity that are flagged as eligible) and the *Positive Predictive Value (PPV)* (which tells us the proportion of families who are eligible that experience the adversity).

A high TPR means significant numbers of newborns in families with high needs are correctly classified by the model as eligible for the greatest support level through Hello Baby. TPR is, of course, dependent on how many newborns are flagged as eligible for services and what the baseline prevalence of the measured adversity is.

A high PPV means that most of the newborns who are selected by the PRM as eligible for the highest levels of Hello Baby services are shown to have adversities the County hopes to prevent. PPV also depends on prevalence, because even a perfect prediction of a low prevalence event such as mortality would still have low PPV when it comes to capturing 5% of the newborn population. To understand this, note that the baseline rate of removals is 1.8 per 100, so when we choose 5% of the population for the program, even if we captured every child who was removed, and therefore the TPR was 100%, the PPV would still only be $1.8/5$ or 36%.

As a generalized measure of model performance and accuracy, we also used *Area under the Receiver Operator Characteristic Curve (AUC)*. A PRM with an AUC of 0.50 is considered to be a PRM that is no more accurate at identifying a child who will suffer adversity than a totally random guess. A PRM which has an AUC of 1.0 is perfectly able to classify children who will be removed by age three.

Performance of the Selected Model

Our selected PRM model proved highly predictive (or highly correlated with future outcomes of foster care placement by age three) and has an AUC of 92.4% (95% C.I. of 91.2, 93.6) on the test data.

Table 2. TPR and PPV (calculated on Test Data only)

	Hello Baby Priority	Hello Baby Family Support	Hello Baby Universal
Share of the population	5.0%	24.3%	70.1%
Share of all children removed who are eligible (TPR)	54.1%	42.6%	3.3%
Prevalence of removals amongst those eligible (PPV)	20.0%	3.2%	0.0%

Notes: The total sample size in test set is 13,155, with Priority Group consisting of 662 births, Support of 3,204 births and Universal of 9,289 births.

Row 2 of Table 2 shows the share of children who end up removed by age 3 that would be eligible for the various levels of the Hello Baby program. Recall that the baseline rate is 1.8 per 100. The high level of predictive power of the model means that even though Hello Baby Priority only comprises 5% of the birth cohort, it accounts for over half (54.1%) of the children in the test data who end up being removed. For Hello Baby Family Support, where 24.3% of the cohort are eligible, we find that 42.6% of all removals are captured in this group. Only 3.3% of children removed would be ineligible for Hello Baby Priority or Family Support.

Row 3 of Table 2 shows the expected prevalence of removals within the eligible population. Almost 1 in 5 children eligible for Hello Baby Priority (20.0%) in the test population are removed by age 3; compared with 3.2% in the Hello Baby Family Support and less than 1 per 1,000 in Hello Baby Universal.

Section 5: Model Validation

One of the challenges with training a model to predict a system-based outcome such as a child's home removal and placement in foster care is that it is not necessarily an objective measure of severe abuse and neglect, or high and complex needs. That is, if the child protection system is poor at identifying

children at risk of harm, then training a model on children who are removed might not actually identify children at true risk of harm.

The purpose of Hello Baby Priority is to match needs with services and to provide the most intensive services to those families with the highest and most complex set of needs. In this section we undertake “model validation” by using a range of external adverse outcomes that the model was not trained to detect, including infant and maternal mortality (see Table 3). The purpose is to establish the complex set of challenges facing families flagged as eligible for the most intensive Hello Baby services using a PRM tool.

Table 3 describes the range of outcomes that we used to validate the model, and the prevalence of these outcomes. The top section of the table lists a set of measures we refer to as “universal adversity indicators”. These are measures of abuse and neglect, maternal homelessness, maternal jail booking, and maternal mortality. They are outcomes that reflect adversities being experienced by children and recorded in the administrative data sets. They were developed to capture Adversities of Childhood Experience (ACE) which have been shown to be harmful to later adult outcomes.⁴ They are also restricted to measures that we can observe for all families. Note that none of these measures were used to train the model, which is what makes them valuable as a validation exercise.

A second set of indicators we use are three different child mortality rates: any mortality, post-neonatal infant deaths and mortality with a code of accidental, violent or maltreatment death. In this latter class we also include near fatalities that occurred due to maltreatment and were investigated as part of PA state requirements (sometimes called Act 33 fatalities and near fatalities). These events are rare – which means that we must be careful about interpreting the exact numbers – but they are the best “ground truth” measures of harm that we have in the administrative data systems.

The bottom part of Table 3 reports on outcomes that are only observed for a fraction of the population who are Medicaid eligible. For these measures, we restrict attention to those 14,562 children who are reported in the birth certificate information as being born on Medicaid.

⁴ Felitti, V. J., Anda, R. F., Nordenberg, D., Williamson, D. F., Spitz, A. M., Edwards, V. and Marks, J. S. (1998). Relationship of childhood abuse and household dysfunction to many of the leading causes of death in adults: The Adverse Childhood Experiences (ACE) Study. *American journal of preventive medicine*, 14(4), 245-258.

Table 3: Rates of adversities used to validate model

Validation Outcome	Definition	Source Data	Birth Cohort Prevalence cohort prevalence (rate and population count)
Universally measured adversity indicators			
Abuse and neglect (case opening)	Case opened for child welfare service by age 3	DHS Children, Youth, and Family Service Records	37.5 per 1,000 1,968
Homelessness Services assessment	Any interaction by Mother recorded with the LINK call center that assesses people for homeless services; and where they met eligibility within 3 years of birth.	DHS Homelessness and Housing Services Records	7.5 per 1,000 395
Maternal jail booking	Had a booking in Allegheny County Jail within 3 years of birth.	Allegheny County Jail data.	9.7 per 1,000 511
Maternal mortality	Maternal death observed to the end of the period for which death records were linked.	Vital Death Records	0.42 per 1,000 22
Any of the adversities listed above	At least one of the adversities listed above (maltreatment, homelessness, maternal jail booking, maternal mortality) used for validation purposes	As above	46.2 per 1,000 2,429
Child mortality			
Child mortality (any cause)	Any cause child's death up to end of the period for which death records were linked.	Vital Death Records	6.2 per 1,000 323
Violent, accidental and maltreatment related child mortality/near mortality	A combination of child death with a coding of injury, drowning and other preventable reasons and referrals to State for fatalities or near fatalities caused by maltreatment with 3 years.	Vital Death Records and County referrals under Act 33.	1.8 per 1,000 93
Child post-neonatal mortality	Child's death due to any cause during the post-neonatal period (28 days to 1 year)	Vital Death Records	1.6 per 1,000 84
Adversity indicators measured only for Medicaid eligible			
Maternal behavioral health crisis (Medicaid-funded only)	Medicaid funded mental health crisis services used within 3 years of birth.	Medicaid data	56.9 per 1,000 828

Maternal ER visits for medical care (Medicaid-funded only)	Four or more Medicaid funded emergency room visits for maternal medical care within 3 years of birth	Medicaid data	226.3 per 1,000 3,295
--	--	---------------	--------------------------

Notes: (1) Mortality is established using vital death records derived from the County’s Department of Public Health. Due to lags in certifying official death records, we cannot establish that all child or maternal death events have been included and therefore, mortality rates reported here may be lower than the true population prevalence. The sample size of the universal tier is 52,520. The sample size for Medicaid Eligible is 14,562.

Table 4 shows the graded relationship between the various measures of adversities and the Hello Baby eligibility groups. A full one-third (31%) of children in the Priority group have a child welfare case opened, compared with 7.7% in the Family Support group and 0.36% in the Universal group. Homelessness service use, maternal jail booking, maternal mortality, child mortality, post-neonatal infant mortality and violent and accidental child deaths all show similarly graded patterns. Even when we restrict attention to Medicaid-eligible families – defined as families where the birth was funded by Medicaid, and therefore a naturally high need population – we find that the more than 4 ER visits and behavioral health crises have a considerably higher prevalence in those in the Priority group (12.8% and 40.8% respectively) compared with Medicaid-eligible families who are in the Hello Baby support group and those in the Hello Baby program.

Table 4: Rates of adversities by eligibility group (PPV)

Validation outcome	Prevalence of adversity in eligible population (per 1,000)		
	Hello Baby Priority	Hello Baby Family Support	Hello Baby Universal
Universally measured adversity indicators			
Abuse and neglect (case opening)	313.0 per 1,000	77.2 per 1,000	3.6 per 1,000
Homelessness services needs assessment	65.8 per 1,000	14.6 per 1,000	0.8 per 1,000
Maternal jail booking	119.9 per 1,000	13.9 per 1,000	0.35 per 1,000
Maternal mortality	5.0 per 1,000	0.7 per 1,000	0.0 per 1,000
Any adversity	391.8 per 1,000	94.1 per 1,000	4.5 per 1,000
Child mortality			
Child mortality (any cause)	19.0 per 1,000	11.7 per 1,000	3.2 per 1,000

Child post-neonatal mortality	6.1 per 1,000	3.7 per 1,000	0.5 per 1,000
Violent, accidental and maltreatment related child mortality/near mortality	8.0 per 1,000	3.8 per 1,000	0.6 per 1,000
Medicaid-eligible population only			
Maternal behavioral health crisis	128.3 per 1,000	53.8 per 1,000	16.4 per 1,000
Maternal ER visits for medical care	408.3 per 1,000	241.1 per 1,000	64.6 per 1,000

Notes: The sample count for the Universal Indicators comprises a total of 52,520 births distributed across Priority (2,626), Family Support (13,130) and Universal (36,764). For the Medicaid eligible population, the sample size is 14,562 births distributed across Priority (2,322), Family Support (8,817) and Universal (3,423).

The validation analysis suggests that even though we trained the model for a home removal by age three, the children and families identified for Priority services are likely to experience a broad range of other adversities.

Section 6: Alternative Approaches to Eligibility

To further address the issue of whether the PRM approach is valuable, we compare alternative ways of determining eligibility. Table 5 provides the relative risk of each of the outcomes based on a variety of eligibility rules.

These rules include a benchmark of a randomly selected 5% of families as well as traditional approaches to eligibility. We tested two “counter-factual” eligibility rules. The first was a randomly selected subgroup from the Medicaid-eligible population, and the second was teen mothers. They are all designed to choose roughly similar number of families (the teen mother group is slightly smaller because only 4% of births were to teen mothers).

Table 5: Alternative approaches to Eligibility

	Families selected using PRM eligibility for Hello Baby Priority (vs. other families)	Randomly selected 5% of families (vs. other families)	Randomly selected Medicaid-eligible families (vs. other families)	Teen mothers (vs. other families)
Number of eligible children	2,626	2,626	2,719	2,285
Training Outcome				
<i>Relative Risk and 95% confidence interval</i>				

Abuse and Neglect (removals)*	22.24 [17.50, 28.25] ^[1]	0.99 [0.74, 1.32]	3.05 [2.56, 3.64]	3.57 [2.99, 4.27]
Universally Measured Adversity Indicators				
<i>Relative Risk and 95% confidence interval</i>				
Abuse and Neglect (case opening) *	13.6 [12.6, 14.8]	1.00 [0.83, 1.22]	3.08 [2.73, 3.47]	3.16. [2.78, 3.58]
Homelessness services assessment*	14.81 [12.18 17.99]	1.31 [0.88, 1.95]	3.08 [2.34, 4.08]	2.27 [1.62 3.18]
Maternal jail booking*	30.54 [25.66, 36.34]	0.84 [0.55, 1.28]	2.95 [2.31, 3.79]	2.38 [1.79, 3.18]
Maternal mortality	27.44 [11.74, 64.15]	2.94 [0.87, 9.92]	5.39 [1.99, 14.59]	2.2 [0.51, 9.4]
Any adversity*	13.97 [13.02, 14.98]	1.00 [0.85, 1.20]	3.09 [2.78, 3.44]	3.11 [2.78, 3.48]
Child mortality				
Child mortality*	3.48 [2.58, 4.69]	0.65 [0.36, 1.19]	1.27 [0.82, 1.98]	1.37 [0.87, 2.18]
Child post-neonatal mortality	4.47 [2.6, 7.7]	0.93 [0.34, 2.54]	2.20 [1.10, 4.38]	2.31 [1.12 4.79]
Violent, accidental and maltreatment related child mortality/near-mortality	5.54 [3.4, 9.0]	0.6 [0.20, 2.01]	2.83 [1.54, 5.20]	3.26 [1.78, 5.96]
Medicaid-eligible population only				
Maternal behavioral health crisis *	2.96 [2.59, 3.39]	0.97 [0.72, 1.30]	1.03 [0.87, 1.21]	1.37 [1.13, 1.66]
Maternal ER visits for medical care *	2.13 [2.0, 2.26]	1.02 [0.90, 1.17]	1.06 [0.98, 1.14]	1.47 [1.36, 1.6]

[1] The relative risk for the trained outcome for the Priority group is restricted to the test data.

* implies that the results are statistically significantly different between the PRM method of selection and the other methods.

Table 5 shows that the counterfactual methods of selecting roughly 5% of the birth cohort result in a group of eligible families that have lower relative risks of adversities both in the training outcome and in the range of adversities that we have available in the administrative data. Even for those outcomes that are restricted to Medicaid-eligible families, the relative risk of the families chosen by the PRM model is higher.

Section 7: Eligibility Fairness & Bias

Due to segregation and racism (institutional and individual), Black and other racial/ethnic minority families have been deprived of equitable opportunities in the labor markets, and chronically over-surveilled by criminal justice and child protection systems. As such, the rates of real and system-defined adversities are notably higher for Black families. Table 6 shows the base rates of foster care placement (the training outcome) and experiencing any of the universally measured adversity indicators presented in Table 3 (i.e. case opening, homelessness, maternal jail booking, maternal mortality).

The ability of a program like Hello Baby to address and reverse racial disparities depends, ultimately, on its efficacy in serving the needs of Black families. The PRM tool is only a small part of the system – a much larger part is played by the providers who will have to engage and serve families with culturally appropriate services. The County is fortunate in having engaged Healthy Start, Inc. as the provider for the Priority group. They have a strong track record in reaching Black families and a commitment to reducing racial disparities.

With regard to the PRM tool, we need to take steps to ensure that the eligibility criteria are fair and establish that we are not failing to serve children from a specific racial subgroup with Priority services. This includes 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.

There are multiple ways of defining “fairness” across race/ethnicity – and other demographic groups. The definition depends on whether the action taken in response to the PRM tool is punitive (e.g., refusing parole) or, as in this case, beneficial.

A way of conceptualizing fair eligibility rules when the objective is to offer a beneficial program is that White and racial/ethnic minority families within each eligibility strata should have similar levels of associated adversities and needs. To be clear, this does not mean that the overall service needs across different racial and ethnic families with newborns is the same.

Table 6. Base Rates of placement in foster care and adversities (any), across racial subgroups.

Subgroup	Base rate of placement in foster care	Base rate of any adversity (case opening, homelessness, maternal jail booking, maternal mortality)
White	1.4%	3.0%
Non-White	3.0%	9.1%

Notes: The n for White is 38,265 and for non-White is 14,255.

Given the higher levels of removals and adversities experienced by non-White children, we would expect a larger share of non-White children to be eligible for the Priority and Family Support tiers.

Table 7 shows the racial breakdown of the three eligibility groups. Of the overall birth cohort, 27.1% are non-White. Within the Hello Baby Priority group, 56.2% of the eligible group are non-White and within the Family Support group, 50.3% are non-White. The implication is that around 10% of the non-White cohort are eligible for Hello Baby Priority compared with 3% of the White cohort. To put this into context, a common eligibility rule used for programs similar to Hello Baby Priority is to restrict it to children who are Medicaid eligible. Amongst Medicaid-eligible families in the County, 54% are non-White. As we show in the next section, Medicaid-eligible children experience far fewer adversities.

Table 7: Distribution of eligible groups by race

	Hello Baby Priority	Hello Baby Family Support	Hello Baby Universal	Total
Non-White	1,476 (56.2%)	6,602(50.3%)	6,177 (16.80%)	14,255 (27,1%)
White	1,150 (43.8%)	6,528 (49.7%)	30,587 (83.2%)	38,265 (72.9%)
Total	2,626 (100%)	13,130 (100%)	36,764 (100%)	52,520 (100%)

When being eligible for the Priority group means accessing a higher level of service, ideally, we want those who are eligible to all have similar risk of harm irrespective of race. When evaluating the fairness of predictive risk models, across racial sub-groups, it is standard to look at a range of different measures.

In the context of using PRM to match services to families, *predictive parity* is the most suitable measure. This requires that White and non-White Priority children within each eligibility category are at equal risk of the harms even if we capture different proportions of each racial sub-group within the categories.

As shown in Table 8, the risk of harm is associated with the Priority group rather than with race. The exception is that the rate of removals amongst the White population in the Priority group is higher than for the non-White population (248.3 per 1,000 versus 160.3 per 1,000). However, the measures of adversities and mortality are similar between White and non-White children in the Priority group. This suggests that in general, the needs amongst White and non-White populations (within each strata) are fairly similar.

Table 8: Prevalence of removals, any adversities and child mortality by Race and Priority Group

Validation outcome	Prevalence of selected adversity in eligible population by race (per 1,000 and 95% C.I.)					
	Hello Baby Priority		Hello Baby Family Support		Hello Baby Universal	
	White	non-White	White	non-White	White	non-White
Removal by Age 3 ⁵	248.3 [198.6, 298.0]	160.3 [122.7, 198.0]	33.0 [24.2, 41.8]	31.9 [23.4, 40.5]	0.7 [0.1, 1.2]	1.8 [0.0, 3.9]
Any universally measured adversity	404.3 [375.9, 432.8]	382.1 [357.3, 406.9]	87.5 [80.6, 94.3]	100.7 [93.5, 108.0]	3.3 [2.7, 4.0]	10.0 [7.6, 12.5]
Child mortality (any cause)	20.0 [11.9, 28.1]	18.3 [11.4, 25.1]	10.0 [7.5, 12.4]	13.5 [10.7, 16.3]	2.8 [2.2, 3.4]	5.3 [3.5, 7.2]

In the literature on fairness of predictive risk models, other concepts of parity that are used include error rate balances of false positive rates and error rate balances of false negative rates. When the PRM is used to identify people for more punitive interventions, such as in the case of recidivism tools that are used to refuse parole to offenders, these error rate balances are more important. The problem as outlined by Chouldechova (2017) and Kleinberg, Mullainathan and Raghavan (2016) is that when the base rates are different, it is not mathematically possible to meet predictive parity along with equal error rate balances. In forthcoming technical reports we provide a more detailed analysis of the PRM tool against all these different definitions of equality.

However, in the case of the Hello Baby tool, our analysis has shown that White and non-White children who are prioritized are at similar risk of harm. Whether the program will be as effective for both subgroups is the question and an impact evaluation is needed to answer this question.

While a different context, the County does have some experience with the impact of PRM tools on racial disparities. In 2016 the County deployed the Allegheny Family Screening Tool which was evaluated by Professor Goldhaber-Fiebert (Stanford University). His interim findings showed that an impact of the tool was to reduce racial disparities in case openings (see the interim impact evaluation [here](#)).

Some critics might argue that programs like Healthy Start are not unalloyed benefits because they might result in more referrals to child welfare. All providers of prevention programs - including Healthy Start workers - are mandated reporters. This means that if they observe abuse or neglect in the families they serve, they are required to report this to Child Welfare. This phenomenon is referred to as “surveillance

⁵ Calculated on test set only.

bias” and is sometimes raised in opposition to prevention programs especially those that involve home visiting. However, the evidence suggests that if there are these effects, they are very small, and certainly not of a size to outweigh the real benefits of prevention programs⁶. Additionally, the County has undertaken extensive training of Healthy Start staff to ensure that they understand the signs of abuse so that they are not unnecessarily referring families to Child Welfare.

Section 8: Implementation

In this section we provide more details on how the Hello Baby program works.

The Hello Baby process begins when a pregnant woman gives birth to her child(ren) at any one of the four birthing hospitals in Allegheny County. After delivery, hospital staff provide the new parent(s) with information about the Hello Baby website and programmatic supports, such as a Baby Specialist at 2-1-1 and the opportunity for a referral for targeted supportive services. During this conversation, the parent(s) are told how they can opt-out from their information being used to determine eligibility for additional services. They are assured that even if they opt-out of the PRM tool, they can still receive services through traditional referral pathways.

The PRM tool will be run weekly basis. Once the County receives birth records, these will be loaded into a custom-built application called the Hello Baby Admin Application (Admin App), which allows DHS staff to download mailing labels with the mother’s name and address, as listed on the birth certificate. These labels are affixed to postcards which explain the Hello Baby program and opt-out process in further detail. The postcard also has an area for an opt-out deadline date

Opting Out of the PRM

Parent(s) have 20 days from the time the postcard is sent to opt-out from their information being used in the PRM tool. Parent(s) can access the Hello Baby website and click the opt-out link at the footer of all pages. They are taken to an electronic form to complete and submit with mother and child’s date of birth. There is additional information on the opt-out page about the DAL, Family Centers, 2-1-1 and the voluntary nature of any services.

By contacting the DAL, a parent can speak directly to a representative and provide the mother and child’s name and date of birth to opt out. No additional information is necessary to complete the process. The DAL representative searches and marks the birth records as opted out in the Admin App in real-time with the parent on the phone, as long as the birth record has been loaded into the Admin App.

If a parent has not opted out or opted out late, their information will be used in the PRM tool after the opt-out deadline has passed. The family details are used by the County’s internal case management application, Synergy. Synergy is used by all the Family Centers within Allegheny County and all their existing cases are housed within this application. This case management tool will only contain data for

⁶ 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.

clients who are eligible for services. Since only eligible clients are loaded, there is no need to redact this data.

Family Centers and Healthy Start will receive electronic referrals through Synergy for parent(s) with a. Those families who are eligible for Priority group will be electronically referred to Healthy Start through Synergy and those families who are eligible for the Family Support group will be electronically referred to a Family Center.

When a referral is sent electronically to Healthy Start or the Family Centers, they will receive basic demographic information for the mother, child(ren) and the father, within the case in Synergy. Demographic information includes first and last name, date of birth, gender, and Social Security Number for each individual and the mother's address. All of this information is directly from the child(ren)'s birth record.

The Hello Baby referrals within Synergy are flagged as such so the supervisor and worker who receive the referral know the case was generated because of the PRM tool and can follow outreach protocols accordingly. Details of the family's exact risk priority will not be shared with staff or shown in Synergy.

Storing of Data

The predictive model is run in the analytics cluster area of the Data Warehouse. The Data Warehouse combines data from many sources and is used specifically for analytics. The PRM tool risk classification will not be shared or retained outside this system. These databases will be maintained permanently for all clients for research and validation purposes as noted. They are not accessible to end users outside the analyst teams who perform these activities.

As stated above, the system that will be used to provide and track services is the Synergy 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 Priority or Family Support services. Referrals will not have risk levels attached, and there will be no documentation for families falling below the service eligibility threshold. Clarification of these standards will be included in the County's forthcoming practice guidelines and policies. Since only eligible clients are loaded, there is no need to redact this data.

Training Staff

Comprehensive training has been provided to all Hello Baby Family Center and Priority staff prior to each receiving their first list of clients. The training provides a brief overview of PRM and the application of it within Allegheny County to give participants an understanding of what risk modeling is, how the PRM tool was built, and the predictive power of the model.

The training also outlines the business process by which the providers will receive referrals. Much of the training is dedicated to building worker understanding of the policy and practice for using the PRM tool to determine eligibility. Some of the key points emphasized in trainings include:

- The Allegheny County Office of Community Services will use integrated data and a PRM tool to screen Allegheny County families with new babies (except those who choose not to participate) and to tier the families based on need.
- All services are voluntary
- The DHS data team will have access to the raw data while community partners will only have eligibility levels.
- The information generated by the tool will never be used for purposes other than referring families who are eligible for services. The data will not be used in the child welfare intake or investigative decision-making processes or for any other child welfare purpose. Child welfare staff will never have access to that information. It will be used only to determine the family's service eligibility level after birth (assuming the family has not opted out of the program).

Next Steps

DHS is committed to ongoing improvements in both the PRM tool and the Hello Baby program. The Allegheny Family Screening Tool was modified several times as part of the County's 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.

Appendix: Features used in the Hello Baby PRM tool

Below is the list of features that have a non-zero weight in the PRM tool. Readers should be aware that as systems change, there will be tweaks to the model that could result in additions and subtractions to the list. The list of features published below is current as of November 2020. It represents the first update to the features list, replacing the original list that was current in September 2020.

Feature (ordered alphabetically)	Feature description
PRI_CHILD_BIRTH_WT	Birthweight of child in grams
PRI_CHILD_BIRTH_WT_LOW	Low birthweight (<2500g) (indicator)
PRI_CHILD_BREASTFED_NO	No breastfeeding at birth (self-reported) (indicator)
PRI_CHILD_PRETERM_BIRTH	Obstetric estimate of gestation 32 and 37 weeks (indicator)
PRI_CHILD_VERY_PRETERM_BIRTH_MISSING	Obstetric estimate of gestation missing (indicator)
PRI_DAD_CRT_ALL_COUNT_2	Months father was active in courts records in the last 2 years (count)
PRI_DAD_CRT_ALL_COUNT_3	Months father was active in courts records in the last 3 years (count)
PRI_DAD_CRT_PROB_COUNT_2	Months father was active in courts (probation) in the last 2 years (count)
PRI_DAD_CYF_REF_OTHE_SCI_2_COUNT	Prior child welfare referrals for father in the last 2 years with a role of "other" that were screened in (count)
PRI_DAD_CYF_REF_OTHE_SCI_3_COUNT	Prior child welfare referrals for father in the last 3 years with a role of "other" that were screened in (count)
PRI_DAD_CYF_REF_OTHE_SCI_EVER_COUNT	All prior child welfare referrals for father with a role of "other" that were screened in (count)
PRI_DAD_CYF_REF_PERP_SCI_3_COUNT	Child welfare referrals for father in the last 3 years with a role of "perpetrator" that were screened in (count)
PRI_DAD_CYF_REF_PERP_SCI_EVER_COUNT	All prior child welfare referrals for father with a role of "perpetrator" that were screened in (count)
PRI_DAD_CYF_REF_PERP_SCO_EVER_COUNT	All prior child welfare referrals for father with a role of "perpetrator" that were screened out (count)

PRI_DAD_CYF_REF_VICT_SCI_EVER_COUNT	All prior child welfare referrals for father where his role was “childhood victim” that were screened in (count)
PRI_DAD_CYF_REF_VICT_SCO_EVER_COUNT	All prior child welfare referrals for father where his role was “childhood victim” that were screened out (count)
PRI_DAD_EDUC_BA_OR_HIGHER	Father's education is Associate Degree, BA, MA or Doctorate (indicator)
PRI_DAD_EDUC_MISS	Father's education is missing or unknown (indicator)
PRI_DAD_HH_SHELTER_DAYS_DUMMY_EVER	Days father was ever in homeless shelter program (count)
PRI_DAD_HISTORY_MISS	No ID for father that allows association with data from County Data Warehouse (indicator)
PRI_MED_DEL_MTHD_ROUTE_CESAREAN	Final method of delivery is cesarean (indicator)
PRI_MED_PREG_INFECTON_YES	Infections present/treated at delivery (indicator)
PRI_MOM_AGE	Age of mother in years
PRI_MOM_AGE_25_29	Mother's age is 25-29 years (indicator)
PRI_MOM_CRT_ALL_COUNT_1	Months mother was active in courts records in the last year (count)
PRI_MOM_CRT_ALL_DUMMY_EVER ⁷	Months mother was ever active in courts records (count)
PRI_MOM_CRT_CM_PLS_DUMMY_EVER ⁷	Months mother was active in Common Pleas Court in the last 3 years (count)
PRI_MOM_CRT_CM_PLS_DUMMY_NOW	Mother was active in Common Pleas Court at/or just prior to the time of the birth (indicator)
PRI_MOM_CRT_MAJ_DIST_CRM_COUNT_1	Months mother was active in Magisterial District Court in the last year (count)
PRI_MOM_CRT_MAJ_DIST_CRM_COUNT_3	Months mother was active in Magisterial District Court in the last 3 years (count)
PRI_MOM_CRT_MAJ_DIST_CRM_DUMMY_NOW	Mother was active in Magisterial District Court at/or just prior to the time of the birth (indicator)

⁷ Feature mislabeled as indicator. It is coded as a count.

PRI_MOM_CRT_MAJ_DIST_NTR_COUNT_1	Months mother was active in Magisterial District Court Non-Traffic cases in the last year (count)
PRI_MOM_CYF_REF_ACCEPT_FOR_SERVICE_3_COUNT	Prior child welfare referrals for mother that were accepted for service in the last 3 years (count)
PRI_MOM_CYF_REF_ACCEPT_FOR_SERVICE_EVER_COUNT	All prior child welfare referrals for mother that were accepted for service (count)
PRI_MOM_CYF_REF_PARE_SCI_1_COUNT	Prior child welfare referrals for mother in the last 1 year with a role of “parent” that were screened in (count)
PRI_MOM_CYF_REF_PARE_SCI_3_COUNT	Prior child welfare referrals for mother in the last 3 years with a role of “parent” that were screened in (count)
PRI_MOM_CYF_REF_PARE_SCO_3_COUNT	Prior child welfare referrals for mother in the last 3 years with a role of “parent” that were screened out (count)
PRI_MOM_CYF_REF_PERP_SCA_3_COUNT	Prior child welfare referrals for mother in the last 3 years with a role of “perpetrator” where the family already had an active child welfare case (count)
PRI_MOM_CYF_REF_PERP_SCA_EVER_COUNT	All prior child welfare referrals for mother with a role of “perpetrator” where the family already had an active child welfare case (count)
PRI_MOM_CYF_REF_PERP_SCO_3_COUNT	Prior child welfare referrals for mother in the last 3 years with a role of “perpetrator” that were screened out (count)
PRI_MOM_CYF_REF_VICT_SCO_EVER_COUNT	All prior child welfare referrals for mother where her role was “childhood victim” that were screened out (count)
PRI_MOM_EDUC_BA_OR_HIGHER	Mother’s education is Associate Degree, BA, MA or Doctorate (indicator)
PRI_MOM_EDUC_LESS_HS	Mother’s education is 12th grade or less, no Diploma (indicator)
PRI_MOM_FIRSTBORN_MISSING	Firstborn child is unknown or missing (indicator)
PRI_MOM_HA_RES_ACHA_DUMMY_EVER ⁷	Months mother was ever in public housing support from Allegheny County Housing Authority (ACHA) (count)
PRI_MOM_HH_PSH_EPISODES_DUMMY_EVER ⁷	Episodes mother was ever in Permanent Supportive Housing program (count)

PRI_MOM_HH_SHELTER_DAYS_2	Days mother was in a homeless shelter program in the last 2 years (count)
PRI_MOM_HH_SHELTER_DAYS_DUMMY_EVER ⁷	Days mother was ever in a homeless shelter program (count)
PRI_MOM_HH_SHELTER_EPISODES_2	Episodes mother was in a homeless shelter program in the last 2 years (count)
PRI_MOM_HH_SHELTER_EPISODES_DUMMY_EVER ⁷	Episodes mother was ever in a homeless shelter program (count)
PRI_MOM_HH_TRANSITIONAL_EPISODES_3	Episodes mother was in homelessness Transitional program during the last 3 years (count)
PRI_MOM_MARRIED	Mother is married (indicator)
PRI_MOM_PAY_PRIVATE	Source of payment for delivery is private insurance (indicator)
PRI_MOM_PR_LIVE_BIRTHS_COUNT	Previous live births (count)
PRI_MOM_PRENATAL_VISIT_YES	At least one prenatal care visit (indicator)
PRI_MOM_SMKD_3MTH_PRIOR	Cigarette smoking in the three months prior to pregnancy (indicator)
PRI_MOM_SMKD_3MTH_PRIOR_MISS	Cigarette smoking before pregnancy is missing (indicator)
PRI_MOM_SMKD_DURING_PREG	Cigarette smoking during pregnancy (indicator)
PRI_MOM_WIC_YES	Mother received Special Supplemental Nutrition Program for Women, Infants, and Children (WIC) food prenatally

Challenge with interpretability of LASSO

Regression models are in general considered interpretable. This is because typically, a person can contemplate the entire model at once and understand a prediction by performing simple calculations in a reasonable time (Lipton, 2018). LASSO (Tibshirani, 1996-A) is a type of regression that aims to construct sparse linear models by shrinking some coefficients and making other coefficients zero, while optimizing for prediction accuracy. Although LASSO produces a smaller set of non-zero weights than some alternative approaches, individually interpreting each of these weights can be misleading - especially when the features are highly correlated. 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.

To help explain this fact, let's consider two popular solvers for the quadratic optimization problem posed by LASSO: 1) the least angle regression algorithm (Efron et al., 2004) and 2) the coordinate descent algorithm (Friedman et al., 2010). The least angle regression works as follows (Tibshirani, 1996-B): with all coefficients being initially equal to zero, the coefficient of the predictor variable which is most correlated with the outcome is iteratively increased in the direction of its correlation with y , until another predictor variable is found that has as much correlation with the residuals of $y - \hat{y}$ calculated along the path of the first predictor variable. The coefficients of these two variables are then iteratively increased in their joint least squares direction, until a third predictor is found that has as much correlation with a residual in this direction. This process is continued until all predictors are in the model. It is clear that each coefficient can only be interpreted along with all the other coefficients from previously selected variables, as these coefficients are changed in their least square direction all together. The coordinate descent algorithm, on the other hand, optimizes by taking partial derivatives of the objective function with respect to one predictor variable at a time. As a partial derivative, this computation relies on having all the other coefficients at a fixed value - a fact that is only true if there is minimal correlation between the features. In both cases, the constraint on the LASSO objective function will force some of these coefficients to be exactly equal to zero.

For example, in the case of the Hello Baby PRM tool, paternity missing is a feature as is father's education. It is impossible for us to understand the impact of paternity missing (PRI_DAD_HISTORY_MISS) on its own, because by definition, if paternity is missing then all the "father" features are set to zero. To calculate the impact of paternity missing, while holding all other features constant - which is how one traditionally interprets marginal effects in a regression context - one would first have to aggregate the impact on a range of features which were implications of "father missing". Readers more familiar with step-wise regression methods will understand to some extent why feature weights are not valuable for understanding the model.

Notwithstanding the fact that feature weights are not a useful guide to feature contributions within the PRM tool, for the sake of full disclosure the County is willing to provide the list of weights on request.

References

- Boller, Kimberly, Debra A. Strong, and Deborah Daro. "Home visiting: looking back and moving forward." *Zero to Three* 30.6 (2010): 4-9.
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2), 153-163.
- Schwalbe, Craig. "Re-visioning risk assessment for human service decision making." *Children and Youth Services Review* 26, no. 6 (2004): 561-576.
- Daro, D., & Dodge, K. A. (2010). Strengthening home-visiting intervention policy: Expanding reach, building knowledge. *New directions for America's preschool policies*, 79-86.
- Foster-Fishman, P. G., Nowell, B., & Yang, H. (2007). Putting the system back into systems change: A framework for understanding and changing organizational and community systems. *American journal of community psychology*, 39(3-4), 197-215.
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1), 1
- Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*.
- Mattox, Teryn, and M. Rebecca Kilburn. "Evidence-based Practices in Child and Family Services." (2014). RAND Corporation.
- Tibshirani, R. (1996-A). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267-288.
- Tibshirani, R. (1996-B). A simple explanation of the Lasso and Least Angle Regression. <http://statweb.stanford.edu/~tibs/lasso/simple.html>
- Vaithianathan, R., Rouland, B., & Putnam-Hornstein, E. (2018). Injury and mortality among children identified as at high risk of maltreatment. *Pediatrics*, 141(2).
- Wulczyn, F., & Lery, B. (2013). Family Group Decision Making in Allegheny County.