Vulnerable Children

Can administrative data be used to identify children at risk of adverse outcomes?

Centre for Applied Research in Economics (CARE)
Department of Economics
University of Auckland

September 2012

Project Team

Rhema Vaithianathan, PhD (lead investigator)
r.vaithianathan@auckland.ac.nz
Associate Professor and Director (CARE)
Department of Economics
University of Auckland

Economists:

Tim Maloney, PhD
Professor and Head
Department of Economics
AUT University

Nan Jiang, PhD
Lecturer
Department of Economics
AUT University

Child Protection:

Irene De Haan, PhD
Lecturer
Department of Social Work
University of Auckland

Claire Dale, PhD
Research Fellow
Department of Economics
University of Auckland

Emily Putnam-Hornstein, PhD
Assistant Professor
School of Social Work
University of Southern California

Ethics:

Tim Dare, PhD
Associate Professor
Department of Philosophy
University of Auckland

SUMMARY OF FINDINGS

- Predictive Risk Modelling (PRM) is an automated algorithm which harvests data from a variety of sources. In this analysis, we use PRM to generate a risk score for the probability of a maltreatment finding for each child at the start of any main welfare benefit spell involving the child.
- A maltreatment finding is defined as a substantiated finding of emotional, physical or sexual abuse or neglect by age 5.
- We consider a “core algorithm”, applying it to children under age 2 and predicting the risk of a maltreatment finding. Although 5.4% of all New Zealand children have a finding of maltreatment by age 5, the rate is substantially higher for children seen on a main welfare benefit (13%) than for children never seen on a benefit by age 2 (1.4%).
- Of all children having a finding of maltreatment by age 5, 83% are seen on a benefit before age 2, translating into a very high “capture” rate.
- The national prevalence rate of maltreatment finding for under 5 year olds in New Zealand is more than 20 times the risk of breast cancer in women aged 50 to 60 for whom routine screening is offered.
- Performance of Predictive Risk Models is usually summarised by the percentage area under the Receiver Operator Characteristic (ROC) curve. A model with 100% area under the ROC curve is said to have perfect fit. The core algorithm applied to children under age 2 has fair, approaching good, strength in predicting maltreatment by age 5 with an area under the ROC curve of 76%. This is similar to the predictive strength of mammograms for detecting breast cancer in the general population.
- The most at risk children identified by the model represent 37% of all children in New Zealand who will have a maltreatment finding by age 5. This group comprises 5% of children overall.
- The highest risk group of children identified are almost 30 times more likely to have a maltreatment finding than the lowest risk group identified.
- If children with the 20% most risky benefit spells are offered a programme which can reduce maltreatment by 10%, we would need 27 families to take up the programme in order to avoid 1 child having a maltreatment finding. This is called the Numbers Needed to Treat (NNT).
- If the 20% most risky benefit spells are offered a programme such as the Nurse Family Partnership which has been shown in overseas studies to reduce maltreatment rates by 46%, we estimate that it would cost $48,000 per maltreatment finding avoided.
- If programme participation begins as soon as children start a spell in the top 20% most risky spells, only a small fraction will have a maltreatment finding within the first year. Therefore, the PRM identifies risk early enough for prevention services to be effective.
- A full ethical evaluation of PRM is necessary before implementation. Additionally, an ethical framework should be developed to guide agencies in their responses to the use of automated child risk scores. Preliminary ethical analysis suggests that mandatory policies for high risk families need to be treated extremely cautiously; we anticipate far fewer ethical concerns if scores are used to engage high risk families in voluntary services.
- Operationalising the PRM requires strong engagement with providers and frontline staff. Therefore, a careful, deliberate and phased implementation (following the ethical evaluation) is necessary.

CONTENTS

BACKGROUND

The Purpose of Scoring Children for Risk of Maltreatment Findings

Correctly assessing the likelihood that a child will have a maltreatment finding at some future time enables scarce child protection and early intervention preventive resources to be strategically targeted. It means that a suite of appropriate programmes of varying intensity can be offered to children and caregivers at all levels of risk. An illusrative example of a possible intervention would be the Nurse Family Partnership programme described in more detail in Box 1 below. We are not advocating for this particular programme, but merely using it as an example of the kind of service that could be offered to high-risk families. We will build a business case around this programme later in this report to demonstrate the possible advantages from risk stratification for the strategic targeting of resources.

While a plethora of “operator driven” risk assessment tools exist (sometimes referred to as “actuarial” risk tools in the literature), these are inadequate for a number of reasons. One concern is that operator driven tools rely on the social worker or frontline agency correctly applying the model. Compliance is dependent upon an agent who is sufficiently trained and motivated to apply the model, and to then respond to the estimated risk. A second concern is that operator driven risk assessment tools are infrequently validated for the population being risk rated.

The objective of this project is to test whether an automatic risk scoring tool – or *Predictive Risk Model* (PRM) – can be developed and validated for New Zealand children. This tool would automatically generate a risk score for children either (1) when they arrive on the benefit system or (2) when their circumstances change once supported by a benefit. The risk score would be generated by a computerised algorithm and, if implemented, could be automatically sent to frontline staff, an external provider, or a central agency for response.

The use of Predictive Risk Modelling is most advanced in healthcare and it has not, to our knowledge, been used anywhere in stratifying children based on their risk of maltreatment. However, there is no reason why the same principles that have been successfully applied to the healthcare arena cannot be applied to other areas like child maltreatment.

The principal requirements for the utilisation of a PRM include: (1) a sufficiently wide net of the target population captured in the systems from which data are harvested; (2) comprehensive and timely data on risk factors; (3) risk scores that can be generated immediately; and (4) outcomes that can be predicted with sufficient accuracy. In the case of child maltreatment, it is particularly important that the protocols followed once the risk score is generated are ethical.

> The Nurse Family Partnership Programme is one of the best researched early intervention programmes in the field. It is a voluntary home-visiting programme that targets low-income, first-time mothers and their babies. Home visits begin during pregnancy and continue through a child's second birthday. The programme aims to improve prenatal, birth, and early childhood outcomes. This programme has five times the overall benefit to society when offered to higher risk mothers compared to low risk mothers. Risk stratification is therefore crucial to maximising its impact. (M R Kilburn and L A Karoly, 2008).

BOX 1: NURSE FAMILY PARTNERSHIP PROGRAMME

The Theory of Risk Stratification

Risk stratification of a population enables programmes to be designed around the needs of each sub-population (Figure 1).

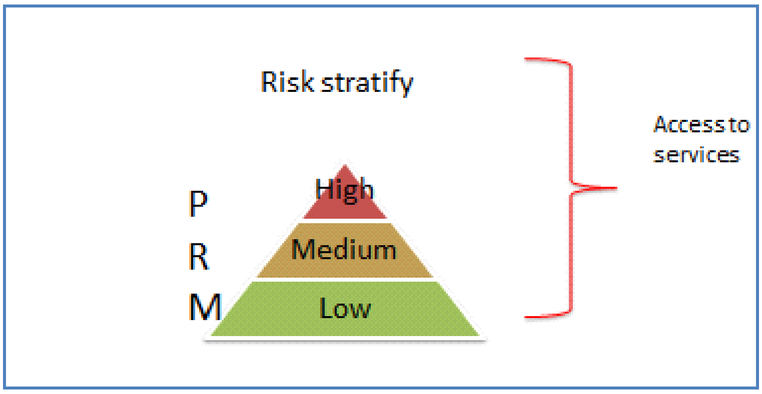

FIGURE 1: STRATIFICATION OF A POPULATION BY RISK/NEED

High Risk/Need: The objective of programmes offered to these families is to ensure that they do not have an adverse outcome.

Medium Risk/Need: The objective of programmes offered to these families is to enhance protective factors and reduce impact of cumulative risk.

Low Risk/Need: Programmes offered to these families are cost-effective broad community-based programmes.

The Algorithm

At this stage in the project, the prototype risk scoring tool we are developing is notionally to be used whenever a child starts a period of inclusion in a main[1] welfare benefit[2].

We initially build a *core algorithm* which is applied to children before they reach the age of 2 years old and predicts each child's risk of having a substantiated finding of maltreatment by age 5 (see Figure 2).

[1] Children included in the prototype risk scoring analysis are those who are included in main, working-age benefits such as the Unemployment Benefit and Domestic Purposes Benefit. Children included in supplementary payments or family tax credits but not included in a main benefit are excluded.
[2] The tool could also be run on a daily basis on all children currently in the benefit system, although the value of this would need to be determined.

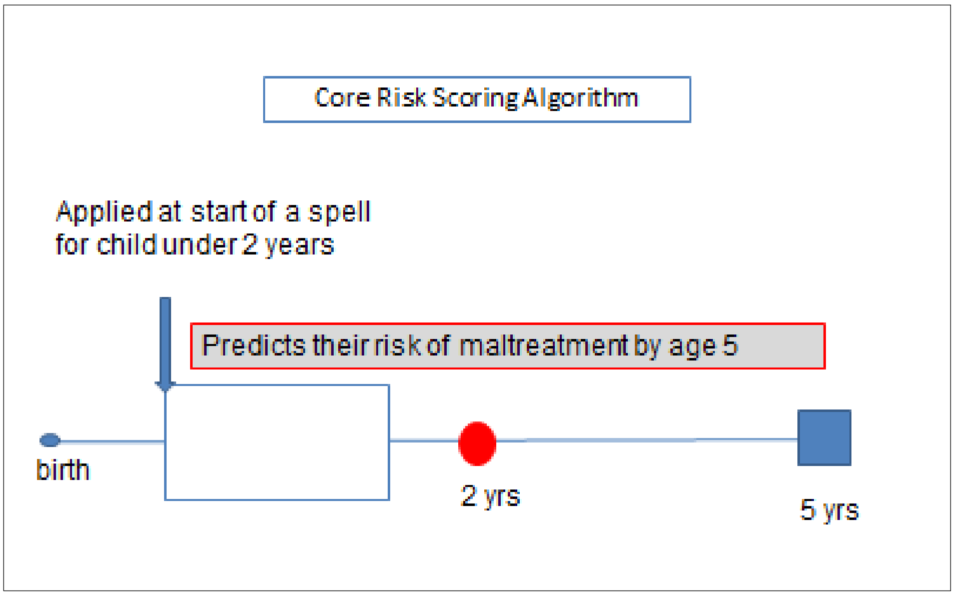

FIGURE 2: RISK SCORING FOR A MALTREATMENT FINDING BY AGE 5

A **maltreatment finding** is defined as a *substantiated*[3] *finding* by age 5 of

- ✓ Emotional abuse
- ✓ Physical abuse
- ✓ Sexual abuse, or
- ✓ Neglect

We use the term *maltreatment finding* to remind the reader of the fact that our outcome of interest is an administrative definition which relates to there being a finding of some emotional, physical and/or sexual abuse, or neglect.

Along with the core model, which predicts the risk of a maltreatment finding, we also develop separate **outcome specific algorithms** that can be run separately. These models predict each child's risk of having a substantiated finding of *neglect*, *emotional abuse* and *physical/sexual abuse* by age 5 and *behavioural problems* by age 7. Note that the behavioural problems are not included in the generic maltreatment findings outcome. We are interested in this for the outcome-specific algorithms because of a concern that for many children, evidence of behavioural problems emerge at the start of schooling.

An algorithm is a list of "predictor variables" and estimated weights that accompany each predictor variable. These weighted variables are used to generate a probability score for each spell. This probability score is the

[3] The term "substantiate" means "support with proof or evidence". In the local context, it is the social worker's responsibility to substantiate abuse (i.e., gather clear and sufficient evidence to determine that abuse has actually occurred). Substantiated maltreatment refers to maltreatment where there has been a finding of physical abuse, sexual abuse, emotional/psychological abuse or neglect. If substantiated, these are entered into the record system under these categories as "findings".

chance that the child who has started the spell will have an adverse outcome by some given age. In this case, the outcome variable is a maltreatment finding by age 5.

To develop this algorithm we extract historical data and, using statistical methods, look for patterns in the data.

We choose a set of predictor variables which closely correlate to a substantiated maltreatment event in the child's future. We then test the accuracy of the predictions on a subset of data (called a validation sample). The accuracy of the predictions can guide users to consider adopting the algorithm, to look for better predictor variables, or to abandon a Predictive Risk Modelling approach altogether.

Risk Scoring of "Spells"

The algorithm repeats the risk scoring exercise every time the child enters or re-enters the benefit system, changes caregivers while included in benefit, or a new spell is started for some other reason (we later define exactly what triggers these spells).

That means that a child's risk is dynamic and can increase and decrease over time. For example, a child who starts a new spell because their caregiver has a new partner who has a history of neglect might have a risk score that is initially low and then increases dramatically.

It is important to note that it is spell-starts that we are risk scoring. As discussed below, most children seen on benefit by age 2 have just one spell start by that age, and the majority start their first spell at birth.

DEVELOPING THE RISK SCORING ALGORITHM

In this section we describe how we develop our core algorithm.

Data

We use a linked dataset provided to the researchers by MSD under confidentiality agreements with the University of Auckland. All data are de-identified by MSD so that researchers could not identify individuals. This dataset links administrative records from the benefit system (Work and Income) and the Child, Youth and Family system, providing a wide range of data on the history and outcomes of children who have had contact with these services.[4]

Inclusion criteria are based on the child:

Children are included in the analysis where:

1. The child is identified with a family that has had a benefit spell begin between the child's birth and 2nd birthday (pre-birth pregnancy-related spells are also included where these are able to be identified because the caregiver received Sickness Benefit for reason of pregnancy);[5]
2. The child is born between January 2003 and June 2006 (this allows all children to have reached the age of 5 by the time the sample period ends).

We were provided with an 80% sample on which to base our research.

We develop an algorithm suitable for use at the start or change of every benefit spell, in particular where:

1. A claim for Sickness Benefit related to pregnancy is made; or
2. A child is newly included in benefit or re-included in benefit after a period not supported by the benefit system; or
3. Within a spell of continuous inclusion, there is a change in (i) benefit type or (ii) caregiver(s) (e.g., partner newly included in the benefit of a formerly single client, or the child moves from one benefit caregiver to another).

Potential predictor variables which are in the data include the following:

1. Care and protection and benefit history of the subject child (e.g., indicators of findings of abuse or neglect, child protection notifications, investigations, Family Group Conferences, Child, Youth and Family Assessments, court orders; proportion of time on a benefit).
2. Care and protection histories of other children in the family at the start of spell and other children included in benefit with the subject child's caregivers in the past.
3. Characteristics of caregivers at start of spell (e.g., gender, age, school or post-school qualifications).
4. Characteristics of family at start of spell (e.g., single vs. dual caregivers, number of younger and older children, birth intervals to next youngest and oldest child, multiple-birth children, age of caregivers when oldest child and subject child were born).

[4] CSRE (2012) provides more information on exactly how this linked database was constructed.
[5] These pre-birth trigger events are included only for children who are included in a benefit upon birth.

5. Care and protection and benefit histories of the subject child's caregivers before age 16, and caregivers' benefit histories in adulthood.

Probit Regression

We randomly split the sample provided into 70% and 30%. We then use the 70% sample to develop the core and outcome specific algorithms and the 30% sample to test how well the algorithm predicts a finding of maltreatment.[6]

We estimate the coefficients using a stepwise probit regression implemented in Stata. We construct 224 variables that could be used to predict maltreatment by age 5. We then drop variables that are not statistically significant (using a p-value of 0.02) or are perfectly correlated. This is done automatically in Stata stepwise regression.

The estimated model is then used to calculate the performance of the algorithm in the 30% "validation sample".

The process selected 132 variables for inclusion in the core algorithm. Given the choice of probit regression, what makes it to the model and what does not, and what comes out as a positive or negative predictor, might be sensitive to a number of issues such as choice of variable selection technique (forward, backward, stepwise), or slight changes in entry and exit criteria.

Interpreting these coefficients requires some care. They are not necessarily "causal" as our aim in devising the algorithm is not to ask "what contributes to maltreatment". Rather we ask "what variables can help us best discriminate between spells that are high risk and ones that are low risk?"

Baseline Level of Risk

Figure 3 shows the proportion of children who are on a benefit prior to age 2 who have a maltreatment finding by age 5 of neglect, emotional, physical or sexual abuse, and behavioural problems by age 7. Of the specific outcomes, findings of emotional abuse is the highest (9.4%) and the lowest is physical or sexual abuse (1.7%).

Using data on care and protection outcomes for all children in the birth cohort, we estimate that 5.4% of all children in New Zealand are maltreated by age 5.[7] We also estimate that the average rate of maltreatment amongst children who were never on a benefit is 1.4%.

To get some sense of whether it is appropriate to screen for such rare events, the prevalence of breast cancer in females aged 50 to 60 (who are routinely screened) is 0.2 to 0.3% (Ann Richardson et al., 2005) and Government policy is to provide mammograms for all women in this age group.

[6] This is done to prevent an over-fitting of the model.
[7] The approach taken to estimating the total number of children ever present in the cohort by age 5 follows that set out in CSRE (2012).

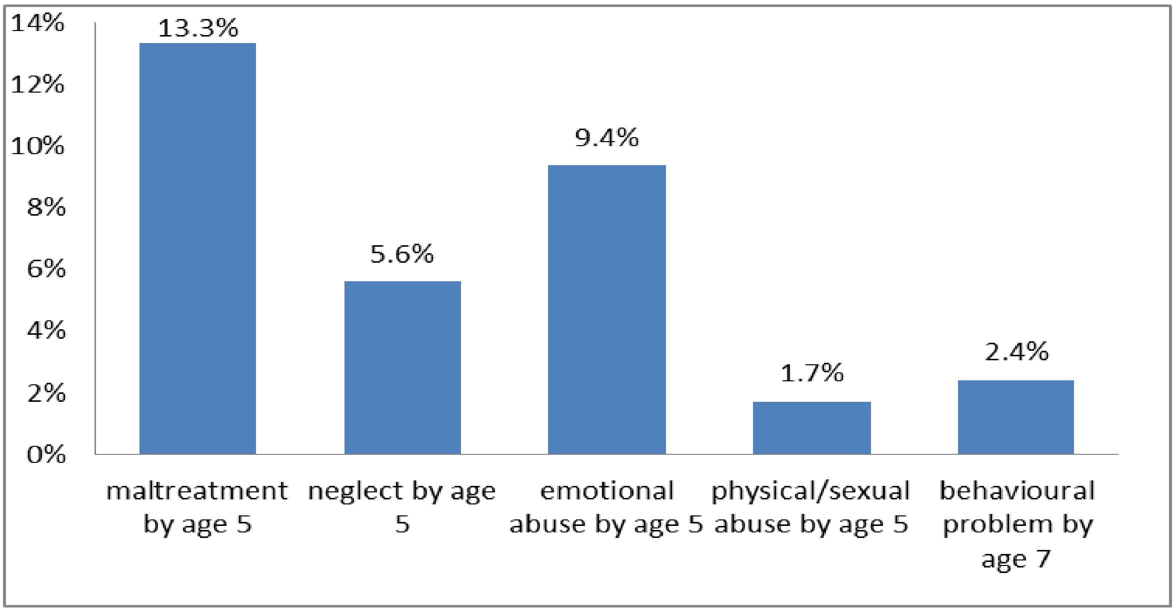

FIGURE 3: PROPORTION OF CHILDREN APPEARING ON THE BENEFIT BEFORE AGE 2 WITH MALTREATMENT FINDINGS

Capture Rate

An important question we need to verify before going down the path of risk scoring children on a benefit is whether the benefit net is wide enough to capture a sufficiently large percentage of children who are found to be maltreated.

It ought to be remembered throughout this paper that the percentage of maltreatment that is found and substantiated is lower than the percentage actually maltreated since there is a considerable proportion of hidden maltreatment that might never come to the notice of child protection services. Our model cannot pick this up.

For the cohort in our sample (i.e., children born between January 2003 and June 2006) there were 11,900 children who had a finding of maltreatment by age 5. Our data indicates that 9,800 of these cases had been seen on a benefit prior to turning 2 and therefore would have received a risk score through our algorithm. Therefore, an algorithm that risk scores children while on a benefit in the first 2 years of their life has the ability to capture 83% of all children with maltreatment findings by age 5.

There are 10,300 children maltreated by age 5 and seen on a benefit by age 5, comprising 86% of all children with maltreatment findings by that age – suggesting that the majority of children with maltreatment findings are on the benefit relatively early in their lives since we only lose 3% when we ignore children who arrive on the benefit after age 2.

Over the course of the first 2 years of life, we estimate that 33% of all children in the cohort would have come in contact with the benefit system.

Understanding Spells

There are 103,397 spells in the under 2 sample, representing 57,986 unique children. Many of these children have multiple spells, although the majority (some 56%) have only a single spell (see Table 1).

Number of spells	Percent of children on a benefit by age 2
1	56%
2	22%
3	14%
4	5%
5 or more	4%

TABLE 1: NUMBER OF SPELLS (BEFORE AGE 2)

The majority (some 61%) of children who had a spell starting before age 2 started their first spell at birth. This suggests that for most children, services can be offered at birth. This is particularly important given that children who start their spell at (or before) birth have relatively higher rates of maltreatment (Figure 4).

Age at start of first spell	Percent of children on a benefit by age 2
Pre-Birth	17%
Birth	43%
0-6 months	18%
6-12 months	10%
12-18 months	6%
18-24 months	5%
Total	100%

TABLE 2: AGE AT START OF EARLIEST SPELL (BY AGE 2)

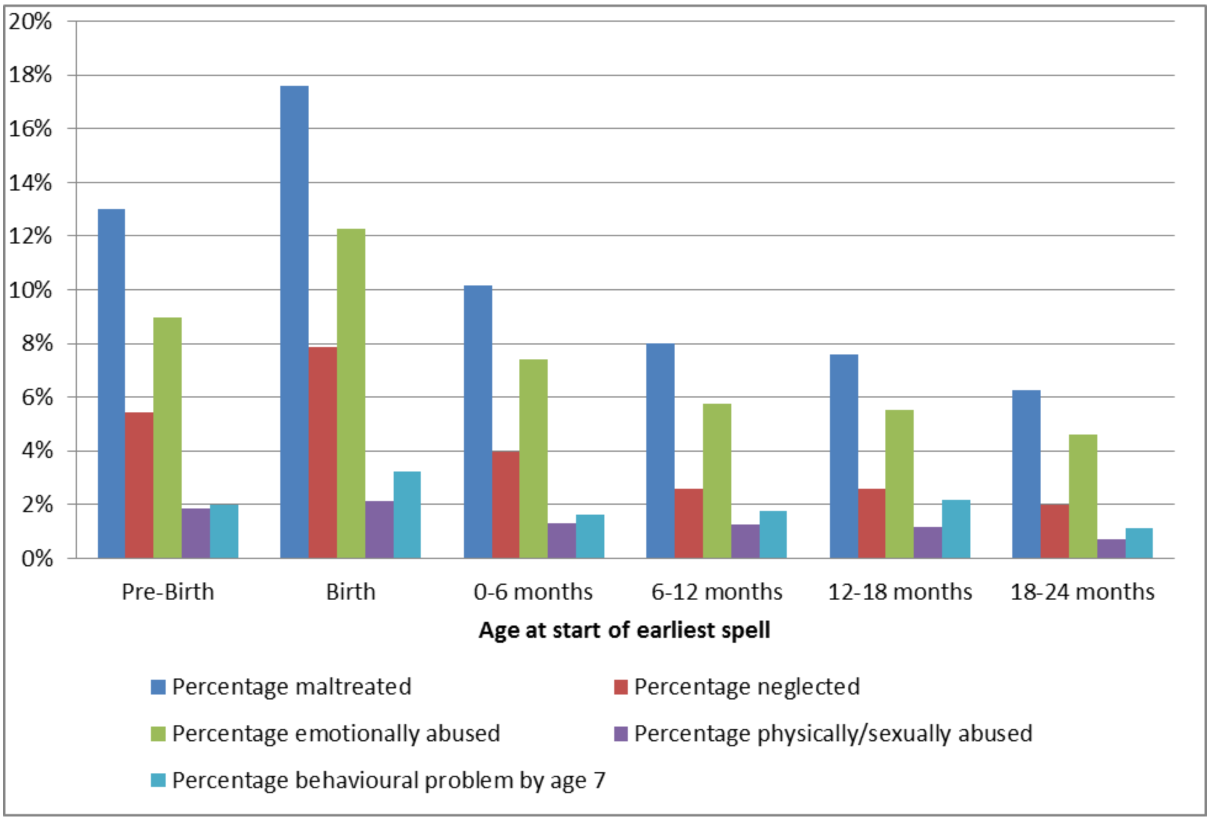

FIGURE 4: AGE AT START OF FIRST SPELL AND PERCENT WITH SUBSEQUENT MALTREATMENT FINDINGS BY AGE 5

PERFORMANCE OF THE CORE ALGORITHM

As noted above, performance of Predictive Risk Models is usually summarised by the area under the Receiver Operator Characteristic curve (ROC). A model with 100% area under the ROC is said to have perfect fit, whilst a model with 50% area under the ROC is no better than tossing a coin when predicting who is more likely to have a finding of maltreatment.

The core algorithm has an area under the ROC of 76% with a 95% confidence interval of [76%, 77%]. This is considered "fair, approaching good" in terms of predictive power. The accuracy is similar to digital or film mammography to identify risk of breast cancer in women who show no symptoms of the disease (Pisano, Gatsonis, Hendrick, Yaffe, Baum, Acharyya, et al, 2005).

Traditionally, there are two ways in which the results of a predictive risk algorithm can be presented. One is as a "probability". For example, a probability of 20% is interpreted to mean that if we take 100 spells in which children have a maltreatment probability of 20%, then on average, 20 of the spells will involve children who will actually have a substantiated maltreatment finding by age 5.

Another way of presenting the score is as a risk decile. In this context, a decile score of 10 means that the spell is in the top 10% of riskiest spells. The way to interpret this decile score is to say that if we took 100 spells at random, the spells in the 10th decile, would be the 10 spells most likely to be predicted to have a finding of maltreatment by age 5.

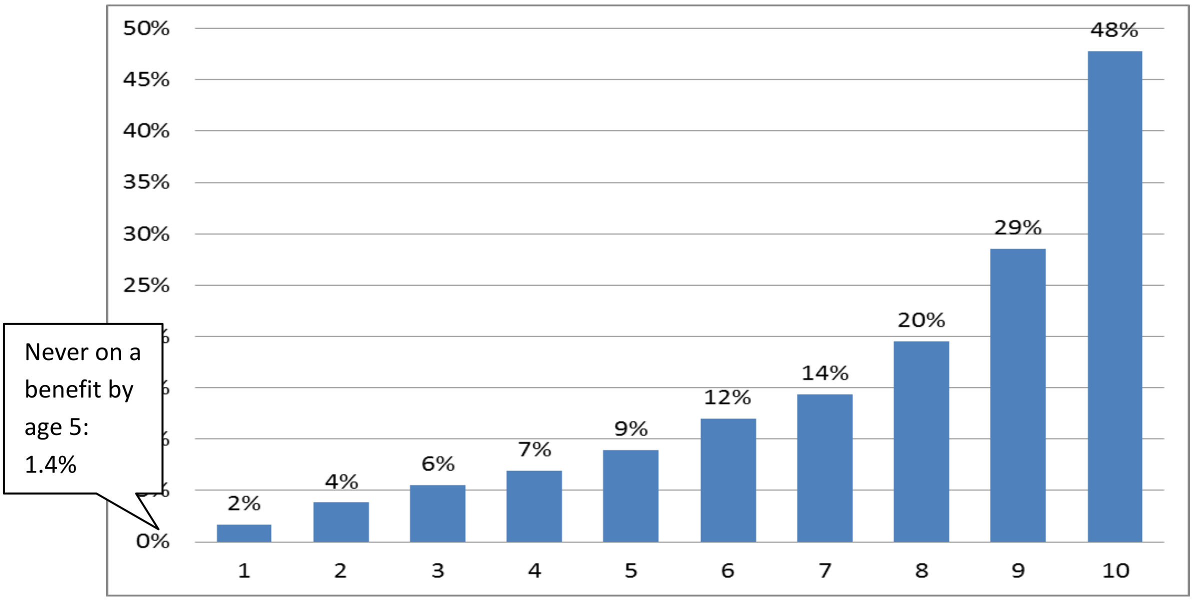

FIGURE 5: MALTREATMENT RISK DECILE AT START OF FIRST SPELL BEFORE AGE 2 AND PROPORTION OF CHILDREN WITH MALTREATMENT FINDINGS BY AGE 5

Figure 5 illustrates the percentage of children who experience a maltreatment finding by age 5 based on the child's accompanying risk decile at the start of their first spell. Approximately 48% of children in the highest decile had a maltreatment finding. The children in this tenth decile have over three times the risk of having a maltreatment finding by age 5 as the child having an average first spell. Conversely, around 2% of children with a first spell in the lowest decile had a maltreatment finding. The children with a first spell in this first decile are seven times *less* likely than average to have a finding of maltreatment by age 5.

For comparison, we calculated the maltreatment finding rate of children who never appeared on a benefit by age 5, using the number of children who were maltreated in the child protection data but not observed on the benefit prior to maltreatment. We divide this number by the estimated number of children in New Zealand in the cohort who were never seen on the benefit by age 5.[8] The average risk for a non-beneficiary child is 1.4%.

Deciles of the first spell equal to or greater than…	Sensitivity Maltreatment by age 5 amongst children on benefit by age 2	Cumulative proportion Percentage of all children on benefit by age 2
1	100%	100%
2	99%	88%
3	95%	76%
4	91%	66%
5	86%	56%
6	79%	47%
7	70%	37%
8	59%	26%
9	44%	16%
10	25%	7%

TABLE 3: SENSITIVITY AND CUMULATIVE PROPORTION OF CHILDREN AT EACH MALTREATMENT RISK DECILE

Another way of assessing the predictive power of the algorithm is to consider the sensitivity of the risk score at various deciles (see Table 3). For example, suppose we considered all children whose first spell was equal to or greater than the 9th decile of risk. This means targeting the highest risk 20% for some service. According to Table 3, we would have included 16% of the children, who represent 44% of all children seen on benefit by age 2 with a maltreatment finding by age 5, in the service group. Of course, we are only discussing the proportion captured compared to the children who are seen on a benefit. In Table 4 we reproduce the sensitivity in terms of all children in the cohort.

[8] The approach taken to estimating the total number of children ever present in the cohort by age 5 follows that set out in CSRE (2012).

Since only 33% of children are on a benefit by age 2, if we target first spells greater than or equal to 9, we will only target 5% of all children in that cohort. These children will account for 37% of all children with a maltreatment finding by age 5.

Deciles of the first spell equal to or greater than…	Sensitivity Maltreatment finding by age 5 amongst all New Zealand children	Cumulative proportion Percentage of all children in New Zealand aged under 2
1	83%	33%
2	82%	29%
3	79%	25%
4	75%	22%
5	71%	18%
6	66%	15%
7	58%	12%
8	49%	9%
9	37%	5%
10	21%	2%

TABLE 4: SENSITIVITY AND CUMULATIVE PROPORTION OF CHILDREN AT EACH MALTREATMENT RISK DECILE AS A SHARE OF ALL NEW ZEALAND CHILDREN IN THE COHORT

Of course, those children who have findings of maltreatment are only a subset of those who actually have some maltreatment occur to them, but this is beyond the scope of the present study.

Maltreatment Risk Scores and Specific Outcomes

An interesting question is what the predictive power of the maltreatment finding risk scores might be for other specific outcomes such as findings of neglect, emotional abuse, physical and sexual abuse by age 5 or behavioural problems by age 7.

In Table 5 we show, for example, that if we offer services to the top 20% highest risk group for findings of maltreatment overall (i.e., deciles 9 and 10), we will be offering services to 39% of those children on a benefit who have a physical or sexual abuse finding by age 5. These percentages are even higher for neglect (54%) and emotional abuse (43%) by age 5, and behavioural problems (41%) by age 7. Thus, our maltreatment risk score correlates well with specific forms of maltreatment findings by age 5, as well as the additional adverse outcome of findings of behavioural problems by age 7.

Maltreatment risk decile of the first spell equal to or greater than…	Maltreatment by age 5	Neglect by age 5	Emotional abuse by age 5	Physical/sexual abuse by age 5	Behavioural problem by age 7
1	100%	100%	100%	100%	100%
2	99%	99%	99%	97%	97%
3	95%	97%	95%	92%	93%
4	91%	94%	91%	88%	87%
5	86%	90%	86%	82%	82%
6	79%	86%	79%	73%	75%
7	70%	78%	70%	64%	66%
8	59%	68%	58%	52%	58%
9	44%	54%	43%	39%	41%
10	25%	32%	23%	23%	21%

TABLE 5: FINDINGS OF MALTREATMENT, NEGLECT, EMOTIONAL AND PHYSICAL/SEXUAL ABUSE AND BEHAVIOURAL PROBLEMS BY MALTREATMENT RISK DECILES AT FIRST SPELL

Numbers Needed to Treat (NNT)

Another approach to determining the strength of a model (and one that is familiar to health sector analysts) is the concept of "numbers needed to treat" or NNT.

This figure indicates the number of people who will have to take up a service offer to avoid 1 child having any maltreatment finding. Of course, this depends on both the likelihood of maltreatment in the population taking up the service and the effectiveness of the service in reducing maltreatment events. Approximately 20,000 children aged under 2 start a first spell every year. In Table 6 we calculate the number of children who will be offered a service if the cut-off for the service is a first spell in the 8th, 9th or 10th decile.

The true positives are the number of children who, if offered the service, will have a maltreatment finding by age 5. For example, if the 9th and 10th deciles are offered the service, we estimate that 3,284 children will be offered the service in a year and 1,211 will have a maltreatment finding before age 5 and 2,073 would not.

Decile of the first spell equal to or greater than…	Number of children offered the service	True Positives	False Positives
8	5,398	1,624	3,773
9	3,284	1,211	2,073
10	1,425	681	744

TABLE 6: NUMBER WHO WILL BE OFFERED A SERVICE EVERY YEAR BASED ON MALTREATMENT RISK SCORE AT FIRST SPELL AND TRUE AND FALSE POSITIVES (FOR MALTREATMENT FINDINGS BY AGE 5)

	Efficacy of the Service		
Decile of the first spell equal to or greater than…	**10%**	**25%**	**50%**
8	33	13	7
9	27	11	5
10	21	8	4

TABLE 7: NUMBERS WHO HAVE TO TAKE UP THE SERVICE IN ORDER TO AVOID ONE CHILD HAVING ANY MALTREATMENT FINDING BY 5

Table 7 shows the NNT for various levels of efficacy of the service and decile cut-offs. A service with a 25% efficacy that is offered to the 9^{th} and 10^{th} decile will be required to be taken up by 11 people in order to avoid one child having any finding of maltreatment by age 5.

Note that not all families who are offered a service will take it up. However, this does not affect the NNT since this relies on the numbers who take up the service (rather than the number who are offered the service).

Return on Investment

Another way of establishing whether the PRM offers a way of cost-effectively addressing maltreatment is to consider a “business case” for an intervention based on offering the service to those children whose first spell is in the 9^{th} or 10^{th} decile.

In this section we undertake a business case example of the return on investment from the Nurse Family Partnership Programme (Olds, 2010). This is a programme that has been shown to be effective, reducing the risk of substantiated maltreatment by 25 percentage points (from 54% to 29%). This suggests a 46% reduction in the baseline rate of maltreatment.

We apply the same reduction in baseline rate to a programme population chosen by our risk score. We consider two risk score cut-offs. The first is a cut-off in which the programme would be offered to all children with a first spell in the top 20% of risk (i.e., 9^{th} and 10^{th} decile); the second is to choose the children with a first spell in the top 10% of risk.

The average maltreatment risk for children with a first spell in the top 10% of risk is 48%, and the average for the children with a first spell in the top 20% of risk is 37%. We then assume that this risk will be reduced by a factor of 46% if the children are offered the programme.

The cost of the programme is calculated using the time spent by the nurse who visits each child. We assume that a nurse visits each child for 1 hour per week over a 2 year period. With a salary of $75,000 and assuming that a nurse is able to undertake 5 visits per day, with 10 weeks for continuing education per annum and 4 weeks’ leave, the cost of the programme in New Zealand will be $8,210 per child.

We calculate the costs and benefit for 1 year of the programme – in other words, the cost to provide the intervention for children under age 2 arriving on the benefit in a 1 year period. The programme would be offered to the cohorts for 2 years following the first time that they arrive on the benefit. We assume 50% will take up the programme (although the cost per maltreatment avoided does not depend on the take-up rate). We outline the cost-effectiveness of the intervention for the two groups based on this simple, but illustrative "back of the envelope" calculation. The purpose is not to make a serious attempt at developing a business case for an intervention. Rather, it is to demonstrate how a risk-stratified population can help an agency develop the appropriate business case for a service and can identify an appropriate target population.

Offered to the top 20% at first spell and a 50% take-up	
Average maltreatment	37%
Number of children in a one-year birth cohort who would have received the programme	1,642
Cost to Government for cohort over the 2 years of the service	$13,481,000
Costs p.a. to Government	$6,741,000
Expected maltreatment if no programme	606
Expected maltreatment with programme	325
Number of maltreated children avoided	280
Cost per maltreatment avoided	$48,000

TABLE 8: ILLUSTRATIVE BUSINESS CASE FOR SERVICES OFFERED TO RISK DECILES 9 AND 10

Offered to the top 10% at first spell and a 50% take-up	
Average maltreatment	48%
Number of children in a one-year birth cohort who would have received the programme	712
Cost to Government for cohort over the 2 years of the service	$5,850,000
Costs p.a. to Government	$ 2,925,000
Expected maltreatment if no programme	341
Expected maltreatment with programme	157
Number of maltreated children avoided	184
Cost per maltreatment avoided	$32,000

TABLE 9: ILLUSTRATIVE BUSINESS CASE FOR SERVICES OFFERED TO RISK DECILE 10

If the programme was offered to children with a first spell in the top 20% of risk, and 50% of the children took up the programme, the cost to the government would be approximately $13.5 million over two years. We estimate that this would reduce the number of children with maltreatment findings by age 5 by 280, at a cost per maltreatment avoided of $48,000. If the same service was offered to children in just the top 10% of risk, the cost to the government for the programme would be reduced to approximately $5.8 million over two years. We estimate that this would reduce the number of children with maltreatment findings by age 5 by 184, at a cost per maltreatment avoided of $32,000.

OUTCOME SPECIFIC ALGORITHMS

Using the same methodology as in the case of the core model, we also developed a set of *outcome specific algorithms* which are designed to predict the probability that a child has a finding of a specific outcome: namely, neglect, emotional abuse or physical/sexual abuse by age 5, or behavioural problems by age 7. This yields four additional algorithms to the core one.

Model	Area under the ROC of the Model	95% C.I.	
Core model (maltreatment findings by age 5)	76%	76%	77%
Findings of neglect by age 5	79%	78%	80%
Findings of emotional abuse by age 5	75%	74%	76%
Findings of physical/sexual abuse by age 5	68%	66%	70%
Behavioural findings by age 7	71%	69%	74%

TABLE 10: PERFORMANCE OF OUTCOME SPECIFIC ALGORITHMS

The outcome specific algorithms (OS) have similar performance to the maltreatment algorithms (Table 10). The best fitting model is the one which predicts findings of neglect by age 5 (with area under the ROC of 79%) while the worst performing algorithm is that for findings of physical or sexual abuse by age 5 (area under ROC of 68%).

An important question is whether there is anything to be gained from operationalizing outcome specific algorithms in addition to a core model. To test this, we look at the percentage of children by outcome specific risk score who went on to have a substantiated finding of that outcome. We compare this percentage to findings for the same outcome, but where risk is classified using the more general core model.

Figure 6 to Figure 8 illustrates the comparison between the core and OS models for findings of physical or sexual abuse, neglect and emotional abuse by age 5 and behavioural findings by age 7. The figure suggests that the more general maltreatment risk decile is extremely good at identifying children who are also at highest risk of each of the specific outcomes. We conclude that there is little immediate value in operationalizing outcome specific algorithms.

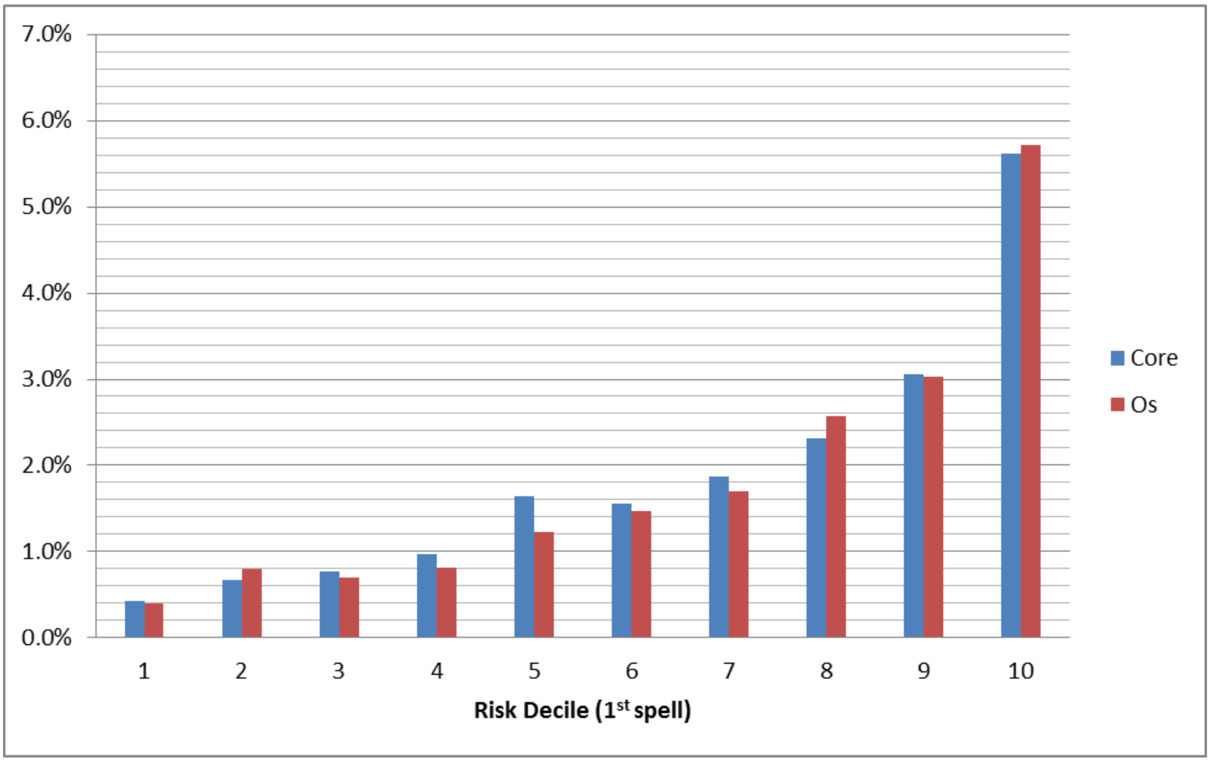

FIGURE 6: OUTCOME SPECIFIC VS. CORE MODEL DECILES AND PERCENTAGE WITH PHYSICAL OR SEXUAL ABUSE FINDINGS BY AGE 5

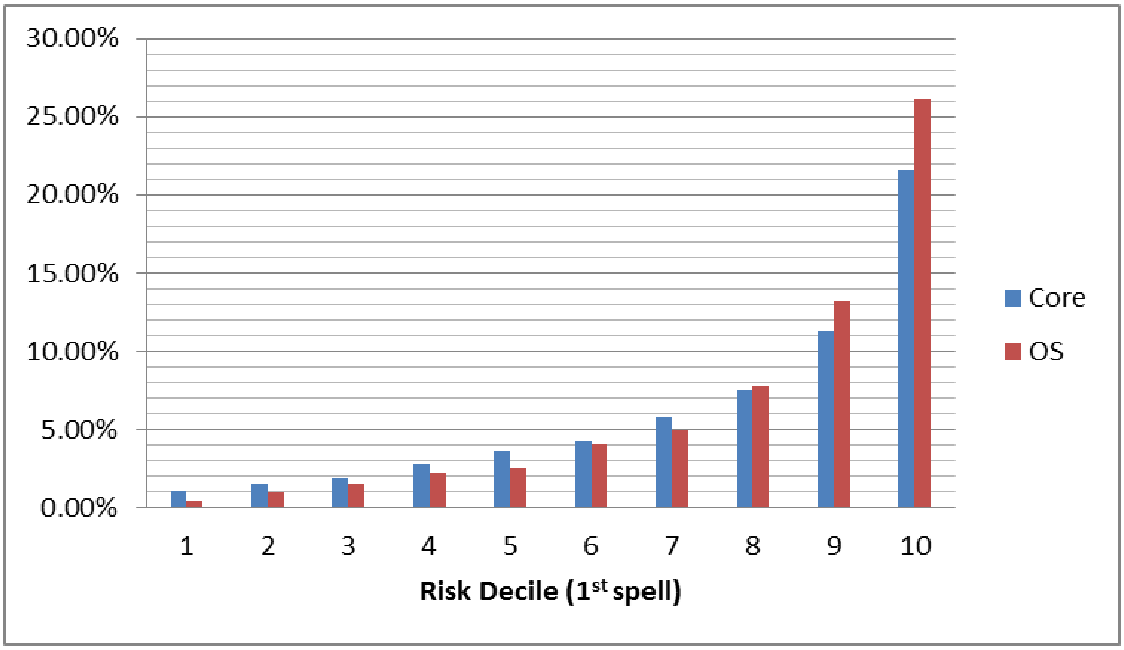

FIGURE 7: OUTCOME SPECIFIC VS. CORE MODEL DECILES AND PERCENTAGE WITH NEGLECT FINDINGS BY AGE 5

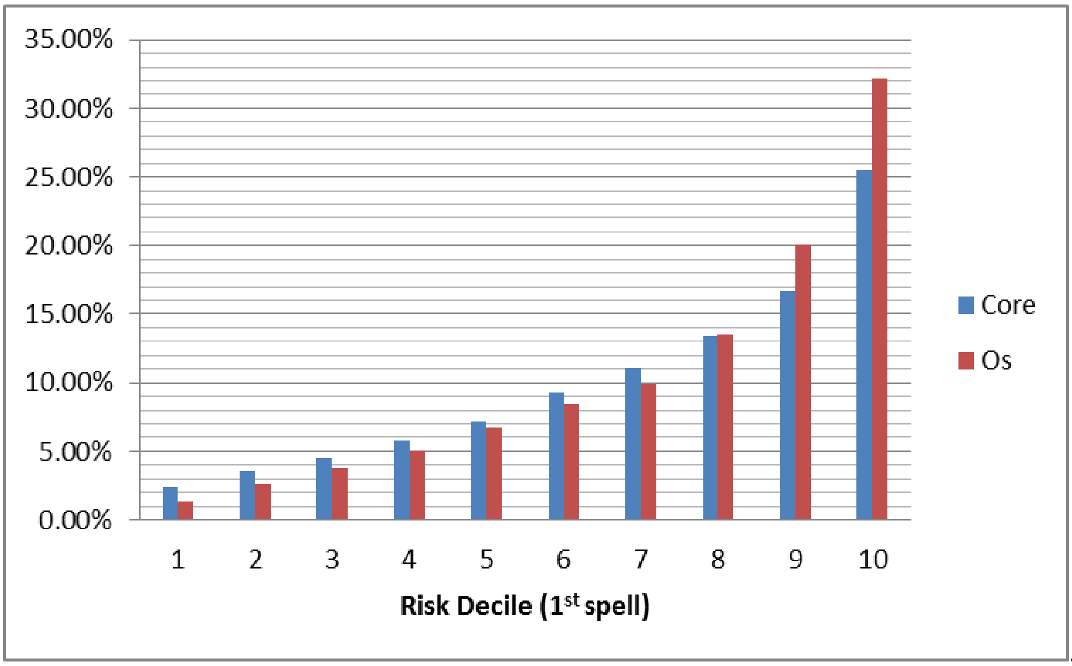

FIGURE 8: OUTCOME SPECIFIC VS. CORE MODEL DECILES AND PERCENTAGE WITH EMOTIONAL ABUSE FINDINGS BY AGE 5

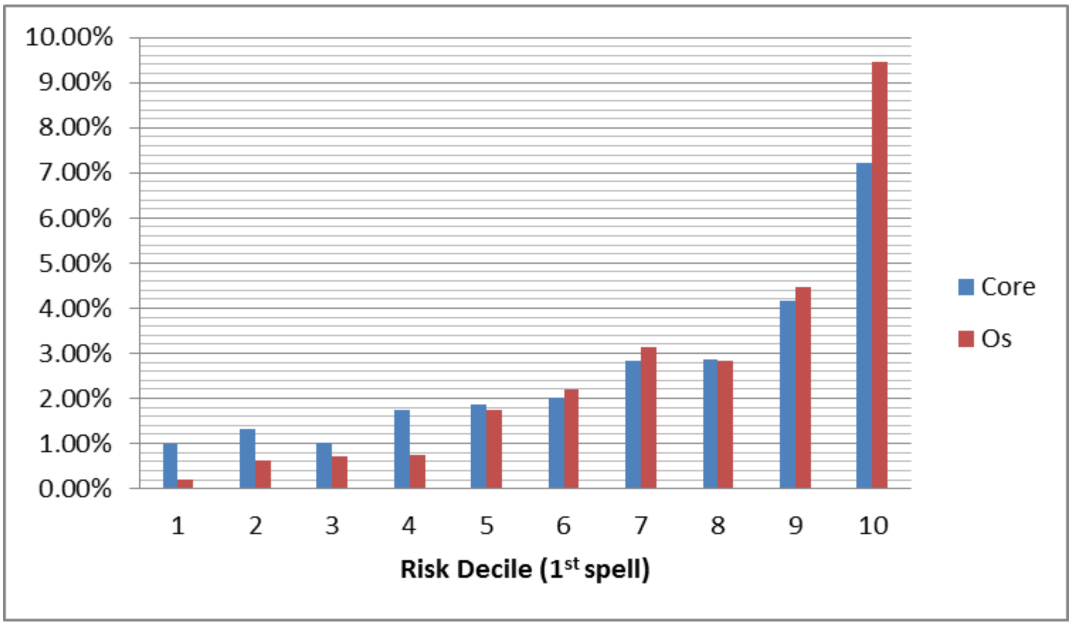

FIGURE 9: OUTCOME SPECIFIC VS. CORE MODEL DECILES AND PERCENTAGE WITH BEHAVIOURAL FINDINGS BY AGE 7

Targeting High Risk Groups for Services

In this section we consider what we could expect if we recruited the highest risk groups into a programme. Suppose a programme is offered to a child when they hit the top 10% or 20% target. We first consider the age of the subject child when the computed risk score would place them for the first time in the high risk group. Figure 10 shows that nearly three-quarters will appear in the top 20% of risk at birth. Slightly fewer (approximately 65%) will appear in the top 10% of risk at birth. This suggests that factors that push children into the top 10% of risk are more likely to be recorded after the birth of the child (e.g., care and protection notifications involving the subject child and other children). It is important to note that, in practice, there might be some weeks of lag between the birth of the child and his or her inclusion in a benefit.

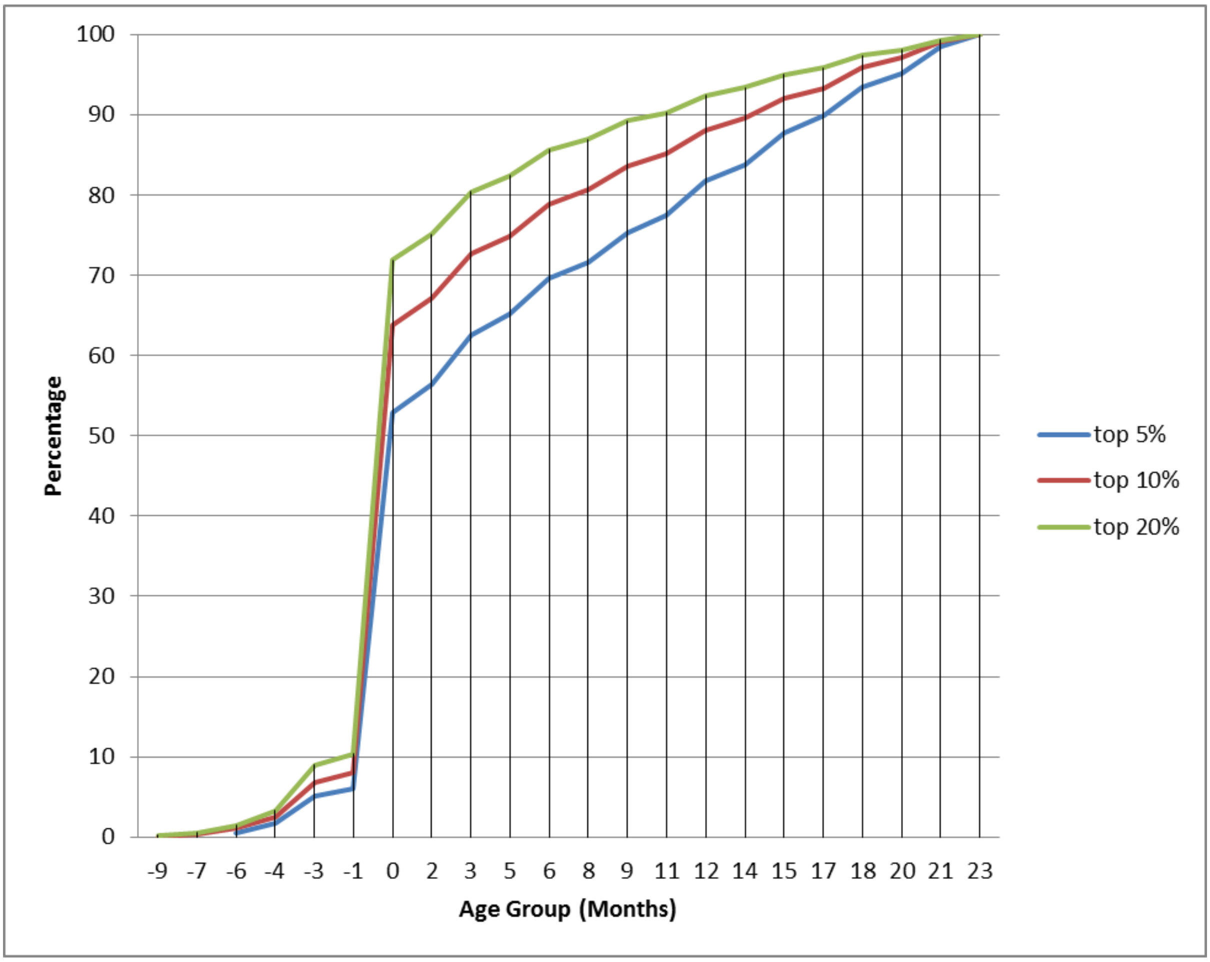

FIGURE 10: AGE AT WHICH CHILDREN WILL BE OFFERED A SERVICE WITH A RECRUITMENT TARGET OF TOP 20% AND TOP 10% OF RISK

We next consider the amount of time available to provide services to families once the child has been identified as high risk. If PRM is to be used to target and offer services where risk is high, then it is important to determine the amount of time before the maltreatment finding occurs. If the majority of children are identified as high risk very late, we would have insufficient time to prevent the occurrence of maltreatment. In Figure 11 we illustrate the percentage of children who have a maltreatment finding with increasing time from the start a benefit spell that places the child in the top 20% of risk.

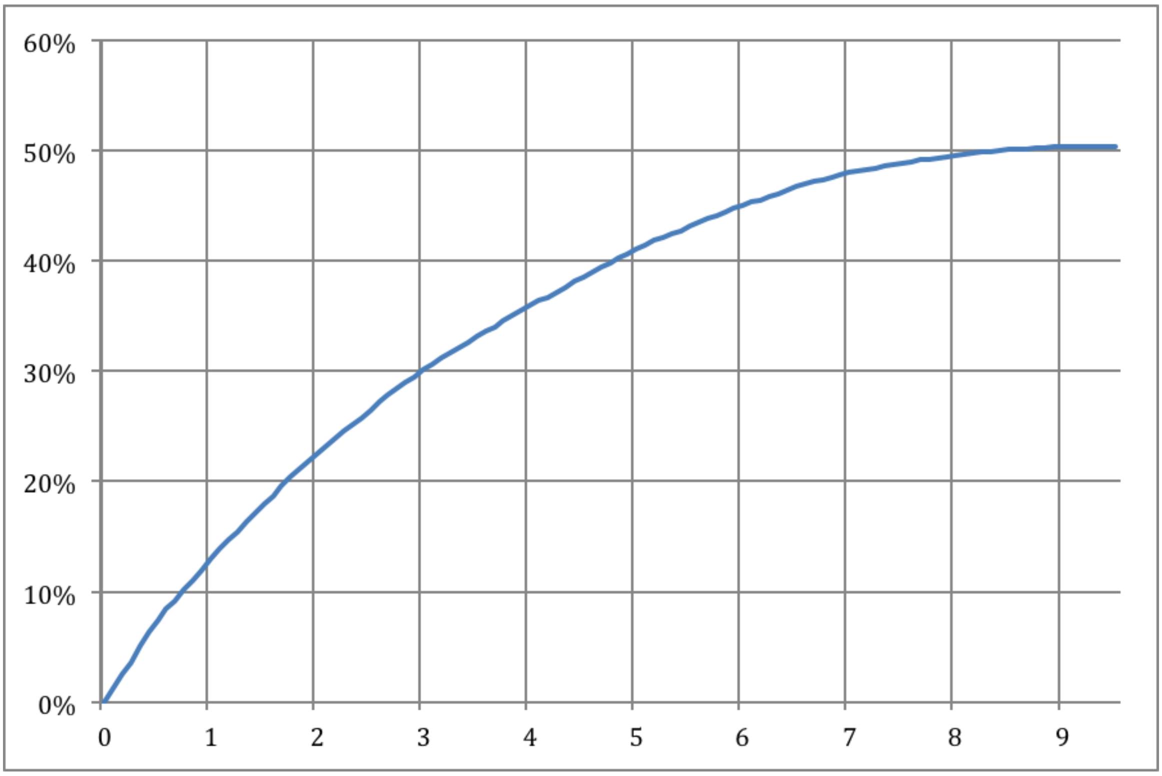

FIGURE 11: ELAPSED YEARS AFTER HITTING THE TOP 20% RISK TO THE FIRST SUBSTANTIATED MALTREATMENT FINDING

After hitting the top 20% of risk, slightly over one-fifth of these children who eventually are maltreated will have experienced this maltreatment within the first year. Less than half of these children will have been maltreated within the first two years. After five years following their inclusion in the highest 20% of risk, slightly over four-fifths of these children will have experienced a maltreatment event. With a median time to maltreatment after being found hitting the top 20% of risk of approximately two and one-half years, there appears to be sufficient time to enable a prevention strategy to be offered that will have some impact on maltreatment.

LITERATURE REVIEW

A separate stand-alone literature review has been conducted (R. Vaithianathan et al., Forthcoming). In this section we summarise the key issues from that review.

The following sections describe preventive interventions, then international use of Actuarial Risk Assessment (ARA) tools in child protection, and use of PRM in the health sector. This provides the foundation for a discussion of risk-modelling in child protection, including the New Zealand model. The final section's brief discussion of risk factors for the child, the caregiver, the family and the community reinforces the difficulty frontline staff have in compiling an accurate risk analysis in an ARA. The limited literature in the area of PRMs suggests this is a potentially useful way to improve child protection.

Generally 'primary prevention' is defined as any intervention designed for the purpose of preventing child abuse before it occurs" (L. Bethea, 1999). The terms 'primary', 'secondary' and 'tertiary' demarcate 'pure prevention' at one end of a spectrum, to 'response to substantiated incidents' at the other. 'Primary' prevention refers to large scale, broad brush strategies aiming at true prevention of maltreatment such as the 'Never shake a baby' and 'It's not OK' campaigns and some parenting education. 'Secondary' prevention aims to work with particular families when there is some indication of risk; including the cumulative stress (often due to poverty) that can tumble a family into what has been referred to as a 'cascade of risk'. Intervention at an early stage in a family's difficulties may prevent things getting worse – and ideally help things get better (e.g. NGO family support services; some targeted parenting courses). 'Tertiary' prevention is the kind of work Child Youth and Family does to investigate reports of concern and then, if necessary, develop a plan for intervention that will keep a child safe from future maltreatment.

Consensus Approach

One approach to risk assessment is the use of professional assessment, or a consensus approach. This approach is favoured by child protection authorities in New Zealand, in contrast to the US and other jurisdictions. The Risk Estimation System (RES), introduced into New Zealand practice in 1996, was 'designed to have social workers build an analysis about risk, not locate it as a concrete state' (T. Stanley, 2007). To come to an understanding of the level of risk to which an individual child was exposed, social workers using RES were expected to encompass the findings of a range of areas of inquiry and reflection: a social work assessment with families; analysis of case files; professional social work theory and models; and consultation with a supervisor.

RES focussed on risk of future harm at the expense of considering a child's needs and the family's potential to care for the child safely. RES has now been replaced by a suite of tools (Family Strengths and Risks Assessment; Assessment Framework; Child and Family Consult; Three Houses) designed to enhance professional decision-making by focussing on strengths and needs as well as risks; and thus promoting comprehensive understanding of the unique situations of children and families.

One concern with a consensus approach is that the large number of factors that frontline staff face in determining risk and their own cognitive biases may induce staff to focus on only a few key factors and ignore the complex multiplicity of factors that influence risk (A Shlonsky and C Friend, 2007)

Actuarial Risk Assessment (ARA) Tools

Wald and Woolverton (1990) identify three main characteristics of valid risk assessment instruments: they are able to differentiate cases in which subsequent maltreatment is likely from those in which it is not; they accurately and reliably measure predictor variables; and they are able to determine the risk of subsequent maltreatment in the context of specific interventions.

ARA tools:

1. Are "Operator driven" and require frontline staff (using checklists) to enter the variables that are used to predict risk;
2. Provide a coarse classification of risk;
3. Are validated on other populations, often in other jurisdictions.

PRM tools:

1. Use routinely collected administrative data to exploit historical correlations and patterns;
2. Assign a precise risk score, enabling early detection of high risk.

There is a plethora of ARA tools in use internationally in child protection (K Broadhurst et al., 2010). For example, in the US context, the Child Abuse Potential Inventory (CAPI) is well established (A.M Begle et al., 2010). CAPI (Form VI) is a 160-item, self-report questionnaire that is answered by the frontline worker in an agree/disagree, forced choice format (J.S. Milner et al., 1984). While a review of all ARAs is beyond the scope of this project, many arguments for and against the increased use of actuarial models can also apply to the use of PRM. To this extent, we review the literature.

A complaint about ARAs that applies to some extent to PRMs is that historic or static factors are given weight at the expense of dynamic factors. Historic factors can trap caregivers as prisoners of their past, while dynamic factors alter spontaneously and they are more susceptible to intervention and manipulation (I.M Schwartz et al., 2008). Schwartz, Jones, Schwartz and Obradovic (2008) report widespread belief that future improvements in the predictive validity of risk assessments resides in continued reassessment of dynamic risk factors, thus increasing opportunities for effective prevention.

A different set of criticisms of ARAs revolve around their reliance on the inputs of frontline staff. While some authors argue that the use of ARAs undermines the role of the frontline social worker, others have found that risk tools are not utilised by frontline practitioners in the manner in which policymakers or designers envisage (K Broadhurst, C Hall, D Wastell, S White and A Pithouse, 2010); (P Gillingham and C Humphreys, 2010). For example, an Illinois study found caseworkers were artificially inflating initial risk scores in order to ensure children's acceptance for ongoing child protection services (C.G. Lyle and E. Graham, 2000). Conversely, in Queensland, where ARA use is mandatory, a study found that when tools consistently scored children as being at high risk but there were not enough services to be offered to these clients, the frontline staff became discouraged by the risk tool (B Lonne, 2006).

While one concern with a consensus approach is the large number of factors that frontline staff face in determining risk and their own cognitive biases that may induce them to focus on only a few key factors and ignore the complex multiplicity of factors that influence risk (A Shlonsky and C Friend, 2007), at the same time, ARAs might prevent social workers from using their professional judgement, leading to missing some important risk factors that are not in the tool (Philip Gillingham, 2006). Gillingham also notes other concerns: practitioners might feel less accountable because they can rely on the risk tool as an "excuse" for not exercising their own judgement; and risk tools might also be used by managers and organisations to protect themselves from blame when there is an adverse outcome.

Main Findings

- Actuarial risk models are common in frontline social work and have been increasingly popular because it is believed that they reduce the “cognitive biases” of frontline social workers.
- Critics claim these tools undermine professionalism and could be used to reduce the accountability of frontline staff.
- Interviews with frontline staff reveal that frontline workers often do not adhere to the tools.
- Ensuring fidelity to risk tools requires a consultative approach with frontline staff so that the tools are seen to be complementary to professional judgement and helpful to the work of staff.

Predictive Risk Modelling in Health Care

In health care, *case-finding* is the process of identifying patients to recruit for an intervention or service (J. Billings et al., 2006) such as case management, intended to reduce the risk of an expensive adverse event like a fall or acute exacerbation requiring hospitalisation. There are two important goals of case-finding:

1. To identify patients at sufficiently high risk of an event; and
2. To identify high-risk patients early enough so that there is time to intervene.

The effectiveness of health care services is found to be drastically reduced by poor case-finding. An example is the UK’s Evercare trial, designed to provide primary care management for frail elderly patients in order to reduce emergency hospitalisation (Hugh Gravelle et al., 2007). Evaluation of the pilot concluded that one reason for the failure of Evercare could be the lack of appropriate case-finding and targeting of the intervention.

Evidence of poor case-finding emerged more recently among US hospital clinicians. When asked to rate whether patient A or B is more likely to return to hospital, frontline staff at a teaching hospital were found to be no more accurate than a coin toss (Nazima Allaudeen et al., 2011).

Evidence of the inadequacy of case-finding methods, and of frontline clinicians’ identification of high risk patients, led to an increased interest in an automated system of risk stratification such as PRMs.

PRMs exploit historical correlations and patterns in routinely collected administrative data to assign a risk score for an adverse event, such as readmission to hospital. The English National Health Service (NHS) has been at the forefront of these, now having implemented the third generation of PRM models for hospitalisation risk. Box 1, reproduced from Panattoni et al (2011) describes more fully the English predictive risk models.

PARR MODEL

The Patients At Risk of Re-hospitalisation (PARR) tool can be downloaded free of charge by NHS organisations in England, and runs off hospital episode data, data from the census and a geographical indicator of deprivation. These are routinely collected in the English NHS. PARR generates a risk score between 0 and 100 for each patient that reflects their risk of readmission in the subsequent 12 months. For high-risk patients (risk score of >50) the tool has a sensitivity of 54.3% and a positive predictive value of 65.4%. For very high-risk patients (risk score of >80) the sensitivity is lower at 8.1% but the positive predictive value rises to 84.3%.

COMBINED PREDICTIVE MODEL ("COMBINED MODEL")

The Combined Predictive Model is an example of a model designed to produce predictions for the entire population, not just those who have had a recent hospitalisation. In addition to the datasets used in PARR, the Combined Model uses variables from the primary care electronic medical record (EMR). EMR data are collected and collated differently across England, so the Combined Model has to be adapted to suit local circumstances. Using the Combined Model, people in the 0.5% of the population with the highest predicted risk are 18.6 times more likely than the average patient to have an emergency admission in the year following prediction.

TABLE 11: SUMMARY OF PREDICTIVE RISK MODELS BEING USED IN HEALTH CARE IN ENGLAND

The Auckland hospitals (Waitemata, Auckland City and Counties Manukau) are currently running a PRM pilot developed as a joint initiative between researchers at the University of Auckland (led by Rhema Vaithianathan) and the Great Auckland Integrated Health Initiative. This project generates a stratified list of patients where each patient, upon discharge, is risk scored for the probability that they will be readmitted within 365 days. This list is then sent to General Practitioners (GPs) on a monthly or weekly cycle. GPs are encouraged to review the high risk patients, and to determine whether they need any further treatment or services.

Main Findings

- A number of international studies on Case Management programmes in health care designed to prevent patients from entering hospital have found to be poor at identifying patients who are at risk.
- One international study found that frontline clinicians were no better than a random coin toss in identifying patients at risk of future readmissions to hospital.
- PRMs which risk score patients for their probability of being readmitted to hospital in the next 365 days have been used in the English NHS for 4 years.
- The University of Auckland, working with Auckland Regional Hospitals, have implemented a PRM pilot in the Auckland Region which generates risk scores for all patients who are discharged from a regional hospital.

Risk Modelling in Child Protection: a Place for PRMs

The same challenges which confront health care in identifying high risk cases to offer more intensive intervention are also present in child protection. The general desire for a tool that accomplishes goals similar to PRM is supported by a number of authors. For example, Schwartz et al. (2008) emphasise:

> *potential lies in using robust technologies to mine and generate valuable data and information for professionals that will enhance their ability to deliver services more efficiently and effectively. Also, the use of sophisticated and advanced research protocols, which are more commonly used in such fields as engineering, epidemiology, and medicine, have the potential to lead to more-promising service interventions and programs, particularly when such protocols are integrated with the technology infrastructure that is now available to us*. (page 214)

However, despite this desire for improved information and communication technologies (ICT) tools, and the recognition that a vast array of data is now available that can be exploited to improve targeting, we did not find any jurisdiction that is using, or has used, PRM in child protection.

An extensive grey and published literature search was conducted, networking by telephone, targeted international List Servs (electronic mailing lists for subscribers) and person-to-person emails. This search revealed no indication that PRM systems are, or have been, used anywhere in practice to predict child abuse; nor is there evidence of automated PRM being run on large administrative databases. Rather, risk assessment in the child protection field is ARA-based.

The published literature (also searched in French, Spanish and German) also showed no reports of PRM being used in this way. The nearest discussions concerned enhanced use of ICT including databases. These are responsive, operator-driven tools for social workers to record children onto; and they remain actuarial, albeit with longer checklists. In key ways, these databases differ from PRMs, a tool that surveys a much wider range of children, with details passed to social workers or other service delivery agents when an electronic "flag" is raised.

In the UK, lack of information sharing had been identified as a failing in numerous high profile child death cases. This led to attention being given to "e-practice" in child protection. Garrett (2005) highlights the envisaged databases for children featured in the Children Act 2004 (UK); the proposed utilisation by local authorities of the Risk of Offending Generic Solution (RYOGENS); and the pending introduction of the Integrated Children's System (ICS). The RYOGENS 'practitioner' uses the software to select one or more options from a predefined list of '40 reasons for concern' (with 120 sub-categories) derived from the Youth Justice Board's relevant form and the Department of Health's Framework for Assessment of Children in Need and their Families. The ICS was a ContactPoint database aiming to improve information sharing between professionals and across agencies.

Critics of the scheme claimed it was evidence of a "big brother state" which, according to the House of Lords and House of Commons Joint Committee on Human Rights, threatened rights to privacy under the European Convention on Human Rights. While the database was halted by the newly elected coalition government in May 2010 (W. Shafiq, 2011), the ICS was introduced as an IT system that implemented record-keeping and reporting functions. However, Ince and Griffiths (2011) call the ICS "a case study in how not to develop such systems":

> *After a short time in operation, the system is now regarded as deficient and, disturbingly, there is a body of evidence to suggest that the impact of the implementation has all too often been antithetical to core social work values and ambitions.*"

A project aiming to use neural networks to take risk assessment beyond these limited models has been recently undertaken and reported by Schwartz and colleagues (D.R Schwartz et al., 2004, I.M Schwartz, P.R Jones, D.R Schwartz and Z Obradovic, 2008). For this project an artificial neural network (a computational intelligence technique) was utilised with a sample of 1,767 cases from the US's Third National Incidence Study of Child Abuse and Neglect and “challenged to predict those children meeting the Harm Standard for abuse – the most severe classification of demonstrable harm from the case file data” ((D.R Schwartz, 2011), page 1081). Results indicated that the trained network was approximately 90% accurate, with relatively few false positives and false negatives.

Building on this work, Schwartz (2011) conducted a data mining and analytics project for the New York State Office of Children and Family Services (OCFS). This project tested several predictive models with a sample consisting of 55,934 cases of child maltreatment reported to the NYS hotline between 2000 and 2010. Results were promising and warranted the trial project being extended to utilise state-wide OCFS data. This work is still under development and has not yet been formally reported in the literature.

Main Findings

- Currently no jurisdictions actively use or have used PRMs in child protection.
- Some experts have called for better utilisation of data and modern ICT tools.
- New York State appears to be the most advanced in terms of implementing a PRM for screening calls to its maltreatment hotline.

Determining Risk Factors

The current University of Auckland project draws on over 200 variables to predict risk. The literature on risk factors is very large and mainly beyond the scope of this project. The main focus in the majority of the relevant published literature is the debate regarding the most critical risk factors (V Carter and M.R. Myers, 2007). A recent literature review conducted on risk factors for subsequent child maltreatment for MSD (Anne Kerslake Hendricks and Katie Stevens, 2012) also provides a list of potential risk factors adapted from a World Health Organisation (WHO) report (A Butchart and A.P. Harvey, 2003).

Identified individual risk factors of caregivers number 17, including ‘difficulty bonding with newborn child’ and poverty; the 9 child-specific risk factors include high needs and abnormalities; and community and family risk factors number 21, including family breakdown and family violence.

One useful lesson from the risk factor analysis is to consider whether there are other sources of data available to MSD that might be utilised to improve the predictive power of the PRM. We also consider those risk factors that might be added to the currently available factors in the risk modelling work and ones that are most likely to be observed after the initial PRM screening.

The analysis suggests that there is some potential for health data and community level data. The latter would be derived by using the location of the household and then mapping that on to location-specific variables such as police data and census data on housing, poverty and other community level effects.

Main Findings

- Including the place of residence of the household with neighbourhood variables derived from census data would be very helpful and relatively easy.
- Data from health records, police records and school records are also potentially useful sources of risk factors.

ETHICAL EVALUATION

We believe that if this approach is going to be implemented, then a strong ethical framework is important.

The application of PRM to child maltreatment raises a number of significant ethical issues. We do not believe these issues pose an insurmountable barrier to the application of the model. This summary of ethical issues is offered, rather, as a guide to matters which we believe need to be considered by those applying the model.

Ethical issues raised by the application of PRM to child maltreatment include:

- **Obligations Where Risk Is Identified.** Once risk is identified, we need to think carefully *what ethical obligation is there to offer intervention to children who may be identified as at risk by the approach, and how will inevitable cut-offs be identified and defended?*

  Lesson: Consider an ethical framework for how agencies will and will not respond to predictive risk.

- **How to Treat the "Risk" vs. Actual Maltreatment.** PRM contemplates interventions or offers of intervention to children identified at risk of *future* findings of maltreatment. However, prospective intervention raises very different ethical issues (around certainty, the status of default liberties, for instance) than intervention following established maltreatment. The risk of transgression is not the same as the act of transgression. Risk of maltreatment finding (however high and however accurate the model) therefore should be carefully distinguished from the act of maltreatment.

  Lesson: Consider how agencies can educate stakeholders and frontline staff about the difference between a risk assessment and substantiation. Ensure monitoring regime and governance arrangements that can ensure that risk assessment is used appropriately.

- **Concerns about PRM as a Policy Approach and Principles of Natural Justice.** Families or individuals identified as at risk will not normally have had notice that they are the subject of such assessment. They will not know (or understand clearly) the way the model works. Some critics have written as though risk modeling should always be considered as research and hence subject to familiar ethical requirements of informed consent (Baldwin, 1993; Taylor, Baldwin, et al., 2008). While this might be a step too far, a concern remains that agencies making PRM assessments are acting quasi-judicially, making decisions which affect individual and family access to significant social benefits, and hence are bound by principles of natural justice including rights to a fair and transparent hearing, to be heard, and to be accorded a process free from bias.

  Lesson: Ensure stringent confidentiality, transparency and governance is maintained. Ensure those subject to risk assessment are given appropriate support to understand and appeal the process.

- **Concerns Relating to the Stigma of Child Maltreatment.** Although PRM has been applied in healthcare utilisation, its application to child maltreatment findings raises distinct ethical concerns. In most health contexts identification as "at risk" carries relatively little social stigma. The assessment that a child is at risk of maltreatment, by contrast, will stigmatise the family and is therefore not "costless" to the individuals or families concerned (Baldwin, 1993, p 358). This suggests a need for confidentiality, and the need for sensitivity in approaches to families and individuals identified as at risk.

Lesson: Ensure that the benefits are large enough to warrant the stigmatisation and the false-positives that might be inevitable from a risk assessment. Similar issues are dealt with in, for instance, HIV tests or genetic tests for inherited diseases, and these could provide some guidance for the appropriate administrative structures. Ensure steps are in place to mitigate the harm of stigmatisation as much as possible, including clear communication and stringent confidentiality.

- **PRM and Mandatory Interventions.** A number of the ethical concerns noted above highlight the importance of clearly distinguishing between intervention or offers of intervention made on the basis of Predictive Risk Modelling, and mandatory intervention imposed following established breaches of laws against child maltreatment.

Lesson: In general, mandatory or imposed interventions should not be considered in response to predictive risk.

- **Predictable Misinterpretation of Data.** It must be acknowledged that some of the data and predictor variables used by the proposed model are highly likely to be misinterpreted by at least some audiences. The decision not to report coefficients in this report, for instance, was based in part upon the belief that the insignificant contribution those factors make to the power of the tool was outweighed by the likelihood of crude and misleading interpretations of that information given existing social prejudices and stereotypes.

Lesson: Develop an ethical evaluation of the predictor variables that ought not to be used in the algorithm. Develop procedures to reduce the risk of misuse of data insofar as possible. Clear communication and appropriate confidentiality are again likely to be central features in such procedures.

- **Non-Beneficiaries are Not Risk-Assessed**. Due to the current limitations of data, we are unable to implement a Predictive Risk Model that captures all children who are at risk – in particular those children who never appear on a benefit. This means that beneficiary status and child maltreatment assessment become linked.

Lesson: MSD could commit to a programme of work which would extend risk assessment to all children and evaluate the benefits of this.

- **Privacy Issues.** A more challenging question is who ought to be able to access the information. For example, a child might receive high risk scores because their caregiver's partner was themself a victim of substantiated maltreatment when young, or have a history of previous children being maltreated while in their care. To the extent that the caregiver is not aware of this, ought the public agency alert the partner to this? Would they be culpable if they did not alert them to the risk posed by the partner, and the child is eventually maltreated? An equally challenging issue is the extent to which child protection agencies have an obligation to intervene once a risk score is identified. For example, if a high risk family is offered services which they refuse, then would the agency have any additional obligation to the child? Should they increase their surveillance of the family? In this case, there could be major concerns that the risk score is pre-empting maltreatment findings (rather than simply risk stratifying the population).

Lesson: Careful consideration needs to be given on the rights and responsibilities of the agencies, the children and the families.

NEXT STEPS

Refinement of Algorithm and Expanding to Non-Beneficiary Children

The algorithm discussed in this report could be operationalised, but we would recommend that more work could be done on both the data and the model to maximise the predictive power. In particular, exploring how all children (not just those on benefits) can be included would be a useful step in refinement process.

A Full Ethical Evaluation

There are a number of very difficult scenarios that could occur which need to be carefully thought through. For example, if a caregiver's risk is increased due to the history of her partner, would MSD be obliged to tell her about this history? If a very high-risk family is offered services which they turn down, and MSD evaluates the child as having an extremely high risk, and the child is subsequently abused, what is the culpability of MSD?

There are a large range of such scenarios that need careful consideration, and both legal and ethical evaluation before a PRM can be operationalized.

We would recommend a full ethical evaluation of the programme and a set of guiding principles designed to steer implementation and the interventions offered.

Designing the Interventions to be Offered and Non-Compliance

A range of evidence-based effective interventions are available to reduce child maltreatment. However, one of the issues that we need to confront is non-compliance and non-take up. The "traditional" public policy approach to non-compliance and non-take up is to subsidise take-up, levy fines or taxes on non-take up, or regulate compliance with the threat of benefit withdrawal.

These measures were based on a "rational" economic framework which implicitly assumed that by changing incentives and relative prices, we could steer agents to choose the "right" outcomes. Quite apart from any ethical objections, as outlined in the previous section, these approaches have come into disrepute because they often do not work. In particular, people who are going through a stressful period in their life may not respond particularly "rationally" to incentives presented by policy-makers.

With the advent of behavioural economics and social marketing, as well as improved communication technology, the favoured approach to compliance in service design is improved **programme design**. The idea is that compliance rates can be increased through better design (e.g., opting out of a programme vs. opting in or overcoming non-cost barriers which create initial inertia). Customising programme design is also favoured. For example, classifying non-compliers into a typology, and designing for each type the best way to offer the intervention, the script to use when the intervention is offered, who should be providing the service etc, has been shown to be effective. It is also relatively costless as there is minimal administrative burden.

Mandatory approaches, on the other hand, are expensive for the State to administer and run the risk of reducing the effectiveness of the programme.

References

Allaudeen, Nazima; Jeffrey Schnipper; E. Orav; Robert Wachter and Arpana Vidyarthi. 2011. "Inability of Providers to Predict Unplanned Readmissions." *Journal of General Internal Medicine*, 26(7), 771-76.

Baldwin N; Spencer N.J. 1993. "Deprivation and child abuse: implications for strategic planning in children's services." *Children and Society* 4:357–375.

Begle, A.M; J.E Dumas and R.F Hanson. 2010. "Predicting Child Abuse Potential: An Empirical Investigation of Two Theoretical Frameworks." *Journal of Clinical Child and Adolescent Psychology 2010;39(2):*, 39(2), 208 - 19.

Bethea, L. 1999. "Primary Prevention of Child Abuse." *American Family Physician*, 59(6).

Billings, J.; J. Dixon; T. Mijanovich and D. Wennberg. 2006. "Case Finding for Patients at Risk of Readmission to Hospital: Development of Algorithm to Identify High Risk Patients." *BMJ*, 333(7563), 327.

Broadhurst, K; C Hall; D Wastell; S White and A Pithouse. 2010. "Risk, Instrumentalism and the Humane Project in Social Work: Identifying the Informal Logics of Risk Management in Children's Statutory Services." *British Journal of Social Work*, 40, 1046 – 64.

Butchart, A and A.P. Harvey. 2003. "Preventing Child Maltreatment: A Guide to Taking Action," World Health Organization and International Society for Prevention of Child Abuse and Neglect.

Carter, V and M.R. Myers. 2007. "Exploring the Risks of Substantiated Physical Neglect Related to Poverty and Parental Characteristics: A National Sample." *Children and Youth Services Review*, 29(1), 110 - 21.

Garrett, P.M. 2005. "Social Work's 'Electronic Turn': Notes on the Deployment of Information and Communication Technologies in Social Work with Children and Families." *Critical Social Policy*, 25(4), 529 - 53.

Gillingham, P and C Humphreys. 2010. "Child Protection Practitioners and Decision-Making Tools: Observations and Reflections from the Front Line." *British Journal of Social Work*, 40(8), 2598 - 616.

Gillingham, Philip. 2006. "Risk Assessment in Child Protection: Problem Rather Than Solution?" *Australian Social Work*, 59(1), 86-98.

Gravelle, Hugh; Mark Dusheiko; Rod Sheaff; Penny Sargent; Ruth Boaden; Susan Pickard; Stuart Parker and Martin Roland. 2007. "Impact of Case Management (Evercare) on Frail Elderly Patients: Controlled before and after Analysis of Quantitative Outcome Data." *BMJ*, 334(7583), 31.

Ince, Darrel and Aled Griffiths. 2011. "A Chronicling System for Children's Social Work: Learning from the Ics Failure." *British Journal of Social Work*.

Kerslake Hendricks, Anne and Katie Stevens. 2012. "Safety of Subsequent Children International Literature Review," Wellington: New Zealand Families Commission

Kilburn, M R and L A Karoly. 2008. "The Economics of Early Childhood Policy: What the Dismal Science Has to Say About Investing in Children," *Occasional papers.* RAND Corporation.

Lonne, B. 2006. "Rethinking and Reforming Child Protection Systems," *Conference: Strength Based Strategies.* University of Queensland, Brisbane, Australia.

Lyle, C.G. and E. Graham. 2000. "Looks Can Be Deceiving: Using a Risk Assessment Instrument to Evaluate the out-Comes of Child Protection Services." *Children and Youth Services Review*, 22(11-12), 935-49.

Milner, J.S.; R.G. Gold; C. Ayoub and M.M. Jacewitz. 1984. "Predictive Validity of the Child Abuse Potential Inventory." *Journal of Consulting and Clinical Psychology; Journal of Consulting and Clinical Psychology*, 52(5), 879.

Panattoni, L.E; R Vaithianathan; T Ashton and G.H Lewis. 2011. "Predictive Risk Modelling in Health: Options for New Zealand and Australia." *Australian Health Review*, 35, 45 - 51.

Richardson, Ann; Brian Cox; Thelma Brown and Paul Smale. 2005. "The Impact of Breast Cancer Screening on Breast Cancer Registrations in New Zealand." *NZMJ*, 118(1209), 6.

Schwartz, D.R. 2011. "Child Protective Services Modeling: New York State Office of Children and Family Services," New York: Child Protective Services, 11.

Schwartz, D.R; A.B Kaufman and I.M Schwartz. 2004. "Computational Intelligence Techniques for Risk Assessment and Decision Support." *Children and Youth Services Review*, 26(11), 1081 – 95.

Schwartz, I.M; P.R Jones; D.R Schwartz and Z Obradovic. 2008. "Draft Improving Social Work through the Use of Technology and Advanced Research Methods," D. Lindsey and A. Shlonsky, *Child Welfare Research.* Oxford Scholarship Online,

Schwartz, I.M.; P.R. Jones; D.R. Schwartz and Z. Obradovic. 2008. "Improving Social Work Practice through the Use of Technology and Advanced Research Methods." *Child welfare research: Advances for practice and policy*, 214.

Shafiq, W. 2011. "Can Citizen Analytics Transform the Public Sector&Quest." *Journal of Direct, Data and Digital Marketing Practice*, 13(2), 148-55.

Shlonsky, A and C Friend. 2007. "Risk Assessment in the Context of Child Maltreatment and Domestic Violence: The Challenge of Prediction." *Medscape Today*, 7(4), 253 - 74.

Stanley, T. 2007. "Risky Work: Child Protection Practice." *Social Policy Journal of New Zealand*, 30, 163.

Taylor, J., N. Baldwin, et al. (2008). "Neonatal and child: Predicting child abuse and neglect: ethical, theoretical and methodological challenges." *Journal of Clinical Nursing*, 17.

Vaithianathan, R.; Maloney, T.; Dale, C.; Jiang, N.; Putnam-Hornstein, E. and Dare, T. Forthcoming. "Predictive Risk Modelling for Maltreatment in Children: A Literature Review," Auckland: University of Auckland.

Wald, M.S. and Woolverton, M. 1990. "Risk Assessment: The Emperor's New Clothes?" *Child Welfare: Journal of Policy, Practice, and Program*.