

PREDICTIVE RISK MODELLING AND CHILD MALTREATMENT AN ETHICAL REVIEW

25 SEPTEMBER 2013

Associate Professor Tim Dare

Department Of Philosophy



**THE UNIVERSITY
OF AUCKLAND**

NEW ZEALAND

Te Whare Wānanga o Tamaki Makaurau

Table of Contents

Summary and Recommendations	1
Introduction	7
The Vulnerable Children Study	10
Preliminary Points	13
1. Technical arguments for the predictive power of the model	13
2. The significance of implementation decisions	13
3. Legal Issues	14
4. PRM and alternative approaches to child maltreatment	14
Ethical Issues	16
1. Overview	16
2. Specific Ethical Issues	18
a) Universal vs Targeted Responses	18
b) Over and Under Identification	22
c) Stigmatisation and the Costs Associated with Identification as At-Risk	25
d) Mandatory vs. Voluntary Engagement with Services	33
e) General Preconditions for Ethical Screening	36
f) Resource Allocation	48
g) Privacy	51
h) Effect on Social Service Professionals and other Frontline Staff	61
i) Rights and the Vulnerable Children PRM	62
Conclusion	66
Bibliography	68

Summary and Recommendations

The Vulnerable Children Report¹ recognizes that the application of PRM to child maltreatment raises significant ethical issues. The report contains a brief summary of these issues, but includes a recommendation that “[a] full ethical evaluation of PRM is necessary before implementation” and that “[a]dditionally, an ethical framework should be developed to guide agencies in their responses to the use of automated child risk scores”. This report contains that evaluation and begins to develop an ethical framework for the implementation of the Vulnerable Children PRM.

***Summary Evaluation:** The application of predictive risk modelling to child maltreatment does raise significant ethical concerns. Many of these concerns can be significantly mitigated or ameliorated. Remaining concerns may plausibly be regarded as outweighed by the very considerable potential benefits of the Vulnerable Children PRM. In sum, the application of predictive risk modelling to child maltreatment is ethically justified provided the recommendations below are addressed.*

This report does not comment on the accuracy of the technical arguments for the predictive power of the model made in the Vulnerable Children Report and does not provide legal assessment.

It is important to bear in mind that the Vulnerable Children PRM addresses issues that have been and are being managed by way of a variety of alternative methods and approaches. Consequently ethical questions about the model are often *comparative*, asking how costs and benefits associated with the model compare with those of plausible alternatives.

¹ Vaithianathan et al., 2012.

Universal vs. Targeted Responses

The Vulnerable Children PRM offers clear potential benefits. It appears to allow social services to identify children with an elevated risk of substantiated maltreatment, and to offer targeted services to reduce that risk. Accurately identifying such children would also allow the targeted, and so it is hoped more effective, allocation of child protection resources.

***Recommendation 1:** That targeted intensive preventive intervention is offered to children identified as at high risk of maltreatment.*

Specific areas of ethical concern discussed and my recommendations in relation to them are as follows:

Over and Under Identification

As with any risk prediction tool, the Vulnerable Children PRM will inevitably make some errors at any threshold for referral, identifying as low risk some children who go on to experience abuse or neglect, and identifying as high risk some children who do not.

***Recommendation 2:** That the databases upon which a child maltreatment PRM draws are expanded to include as many New Zealand children as possible .*

***Recommendation 3:** That current early identification referral routes including those initiated by health professionals and other front-line social service professionals, should be maintained alongside the Vulnerable Children PRM.*

***Recommendation 4:** That ways of reducing the consequences of mistaken identification as at high-risk are explored, including:*

- *Providing opportunity for experienced social services professionals to exercise judgement about appropriate responses to a family's identification as at risk;*

- *Ensuring that such professionals understand the potential of the Vulnerable Children PRM to miscategorise families;*
- *Providing training to guard, in so far as possible, against confirmation bias in the professional engagement with families identified as high-risk.*
- ***See also recommendations 5 and 6***

Stigmatisation

Particular ethical issues may arise when significant burdens, most notably the burdens of stigmatisation, are placed upon individuals or families as a consequence of their being identified as falling within the Vulnerable Children PRM's high-risk deciles.

Recommendation 5: *That interventions be at the minimum level necessary to achieve the benefits offered by the Vulnerable Children PRM.*

Recommendation 6: *That information produced by the Vulnerable Children PRM is disseminated as narrowly as possible, consistently with achieving the benefits of the programme. I recommend that only senior and experienced staff have access to such information and that they be carefully trained as to how to manage the information they possess.*

Recommendation 7: *That consideration be given to what level of detail is required to be disseminated to make effective use of the model's predictions.*

Recommendation 8: *That training and implementation emphasise that those identified as at risk have committed no wrong and that most of them will not go on to do so. Interventions must be preventive and supportive and not punitive.*

Recommendation 9: *That interaction with high-risk families is as similar as possible to that with other families, at least in the external presentation of those interactions.*

Recommendation 10: *That ways of engaging with the media over child maltreatment and the Vulnerable Children PRM are explored with an eye to*

minimising stigmatisation and promoting as accurate an account of the PRM as possible.

Mandatory vs. Voluntary Engagement

There are a number of risks associated with mandatory engagement policies, including: effects on the efficacy of engagement; concerns about the feasibility of mandating engagement; threats to the default rights of individuals who have not been found to have done anything wrong, many of whom will not do so; the potential harms of mandatory engagement; and the possibility that such engagement will disproportionately burden minority groups likely to be over-represented in those identified as at high risk.

***Recommendation 11:** That engagement with high-risk families is on a voluntary basis.*

Ethical Constraints on Screening

The Vulnerable Children PRM should meet the requirements of a set of general preconditions for ethical screening. Subject to issues raised elsewhere in this report I believe that the model is largely compliant with those requirements. Issues of concern under the screening principles include: concerns about the efficacy of interventions in response to identification of high risk, the capacity of social services to respond appropriately to identified need, and the public acceptability of the Vulnerable Children PRM. These issues are addressed by recommendations made under other headings.

Resource Allocation Issues

I identify and discuss two resource allocation issues:

A concern that the Vulnerable Children PRM will lead to neglect of the needs of lower risk children and hence that a universal allocation of child protection resources should be preferred.

A concern that the Vulnerable Children PRM will identify a degree of need that will overwhelm child protection services to the overall detriment of vulnerable children.

Recommendation 12: *That the Vulnerable Children PRM is used as an opportunity to deliver additional intensive intervention to high risk families and that existing universal services remain in place.*

Recommendation 13: *That child protection resources and workload are managed to ensure response to identified risks.*

Privacy and Confidentiality

The Vulnerable Children PRM poses some threat both to general moral rights to privacy and to legal and moral rights to confidentiality.

Recommendation 14: *That invasions of privacy which could appear discriminatory are monitored and minimised, consistently with delivering the benefits of the Vulnerable Children PRM.*

Recommendation 15: *That staff having access to the information provided by the Vulnerable Children PRM are made subject to a specific duty of confidentiality.*

See also Recommendation 6 and 7 above

Effects on Social Services Staff

I address two possible effects of the Vulnerable Children PRM on social service professionals and other frontline staff, the concern that it might reduce engagement between such staff and their clients, and concerns about who is to have responsibility for monitoring the Vulnerable Children PRM and engaging with families.

Recommendation 16: *That the Vulnerable Children PRM is not seen as a replacement for the judgement and engagement of experienced social service professionals. (See also Recommendation 12.)*

***Recommendation 17:** That implementation decisions around the Vulnerable Children PRM identify staff or services who will have responsibility for monitoring the Vulnerable Children PRM and engaging with families, taking into account the ethical issues relevant to those decisions raised in this report.*

The Vulnerable Children PRM and Rights

A number of moral and legal rights are relevant to the Vulnerable Children PRM, including rights under the United Nations Convention on the Rights of the Child, New Zealand's Bill of Rights Act and Human Rights Act, and moral and legal rights around privacy and confidentiality issues. These rights are important to the ethical status of the Vulnerable Children PRM but do not provide fixed or determinative solutions. The issues they raise are addressed by recommendations made under other headings.

Introduction

Child maltreatment is both intrinsically wrong, and has serious, pervasive, and long-lasting adverse consequences.² The Vulnerable Children Study³ was commissioned by the New Zealand Ministry of Social Development (MSD) to find out whether it was possible to use administrative data to identify children at risk of maltreatment. The study was undertaken by a cross-university team of researchers based at the University of Auckland's Centre for Applied Research in Economics who were provided with a data set linking administrative records from the income support benefit system and the Child, Youth and Family child protection system. Both agencies are under the auspices of the MSD and both systems hold information collected nationwide. The researchers developed a predictive risk-modelling (PRM) tool using an algorithm that appears to have significant capacity to ascertain and stratify children's risk of experiencing maltreatment in the future, generating a risk score that could be sent to frontline staff or agencies for response.

The Vulnerable Children Study recognises that the application of PRM to child maltreatment raises significant ethical issues. The study report contains a brief summary of these issues, but includes a recommendation that “[a] full ethical evaluation of PRM is necessary before implementation” and that “[a]dditionally, an ethical framework should be developed to guide agencies in their responses to the use of automated child risk scores”.⁴ This report contains that evaluation together with a number of recommendations.

It will be useful to indicate in broad terms what such an evaluation might involve. At the broadest level, consideration of the moral status of efforts to reduce child maltreatment requires us to address and balance a number of different moral perspectives and interests, including:⁵

² Norman et al., 2012.

³ Vaithianathan et al., 2012.

⁴ Ibid., p.3. The brief ethical analysis in the report is to be found at pp.32-35.

⁵ My aim in these introductory remarks is to give a sense of the evaluative task, not to provide a comprehensive list of relevant moral perspectives or interests.

- The rights of the child, both institutional under instruments such as the United Nations Convention on the Rights of the Child (UNCRC) and relevant domestic law, and moral or human rights which exist, if they do, independently of legal or conventional systems.
- The rights of parents and other adults, including institutional and moral rights forbidding discrimination, protecting privacy and autonomy, and the like.
- Moral interests grounded in appeals to the idea that morality requires us to maximise some value such as happiness or preference satisfaction, and which would no doubt regard child maltreatment as morally forbidden while also attaching value to respecting autonomy, privacy, and the like.
- Moral interests generated by various sources of moral value, which might, for instance, treat as morally desirable conditions and decisions conducive to human flourishing while condemning those that were not (e.g., moral values grounded in aretaic theories).
- Various obligations, permissions and duties specific to the roles of those social service professionals, such as, for instance, the obligations to gather information about vulnerable children, to preserve specific confidences, and to give priority to the interests of children.
- The moral interests of the larger community, interests that might be grounded in any of the moral perspectives or theories gestured at so far, and which might encompass not only a concern for children but also for the efficiency and fairness of their community's institutions, and so on.

There are potential conflicts between many of these interests and perspectives, both within and across these groupings. An adult's right to privacy, for instance, might conflict with a child's right to be protected and with a professional's role-duty to determine what is happening in a home. At a broader level, there is also potential conflict between moral or legal rights on the one hand, and those moral

interests founded upon consequences on the other: the high utility value of protecting children from harm may call for conduct that is inconsistent with the rights of those affected.

I have said that consideration of the moral status of efforts to reduce child maltreatment requires us to *balance* competing moral perspectives and interests. To be clear, and to explain this reference to balancing,⁵ I do not think it is appropriate in the context of the current evaluation to grant lexical priority to moral considerations *because they are moral considerations of a certain kind*. It is often thought, for instance, that recognition of an interest *as a right* entails that it operate as ‘trump’⁶ over competing non-rights based moral interests (such as those defended on, for instance, consequentialist or aretaic grounds). There is, however, a fairly straightforward reason to think that one cannot avoid a more pluralist approach in the current context. It is clear that there are important rights-interests at stake in child protection – interests which will not be captured by any simple cost-benefit or consequentialist approach – but it is also clear that there are potential and significant conflicts *within* the clusters of interests and perspectives that bear upon child maltreatment policy decisions. One cannot avoid balancing competing interests by granting lexical priority to rights-based interests, then, since there are unavoidable conflicts within the cluster of interests plausibly protected by rights. One is driven, in my view, to a more pluralist approach. Moral evaluation of the Vulnerable Children PRM requires us to identify and balance the relevant moral perspectives and interests, aiming to give appropriate weight to each, in order to arrive at a view of the ethical status of the model. That is what I attempt to do in this document.

I begin with a brief sketch of the Vulnerable Children PRM, intended to inform subsequent discussion.

⁶ “Rights are best understood as trumps over some background justification for political decisions that states a goal for the community as a whole” (Dworkin, 1984, p. 153).

The Vulnerable Children Study

The core of the predictive risk model presented in Vaithianathan et al., 2012 is an automated algorithm that predicts the probability that a child will be maltreated by the age of 5. As noted, the algorithm was developed and validated using an anonymised dataset linking administrative records from New Zealand's welfare benefits and care and protection systems for children who were born between January 2003 and June 2006 and had a benefit spell before the age of 2, a sample of 57,986 children comprising about 33% of all children born in New Zealand during that period. The Vulnerable Children Study randomly divided this sample into two groups, 70% in one and 30% in the other. The study's algorithm was developed by identifying variables which are associated with child maltreatment⁷ in the 70% sample. 132 variables – including demographic and historical features of a child, their family, household and community – were found to make a statistically significant contribution to the model and were therefore retained in a 'core algorithm' which was then tested on the 30% validation sample. The algorithm generated a risk score at the start of each new benefit spell for each child in the sample. Those risk scores were then segmented into deciles, with 10 indicating a child as being within the top 10% of risk, down to 1 as being in the bottom 10%.

The algorithm produces a risk score at the start of any new benefit spell which occurred before the child's second birthday, and so recognises dynamic risk factors. A new spell starts whenever the benefit system records alterations in a family's circumstances indicated by such changes as a change of benefit type, the arrival or departure of a partner, or movement of a child from one caregiver to another. A higher PRM score than previously generated signals that risk is escalating.

⁷ A maltreatment finding is defined as a substantiated finding of neglect or emotional, physical, or sexual abuse by age 5. Substantiation requires clear and sufficient evidence that abuse has occurred. Normally, substantiating evidence will be gathered by social service professionals.

As a partial indication of the model's accuracy, when the children in the validation sample were 'followed' through until their fifth birthday it was found that 48% of children with risk score of 10 (the top 10% risk category), and 29% with a risk score of 9 had maltreatment findings by age 5, compared with around 2% of children in the lowest risk score. The Receiver Operating Characteristic (ROC) curve is commonly used as a more thorough indication of a predictive model's accuracy.⁸ The Area Under the Curve (AUC) indicates the accuracy of a model. For comparison's sake: a perfect predictive model, which gives all and only true positives and true negatives, would receive a 100% AUC, while a model that can discern a true positive or negative no better or worse than chance would receive a 50% AUC. The AUC for this model is approximately 76%, which is technically regarded as fair, approaching good, and is similar to the accuracy of a mammogram conducted without prior risk indication of cancer (as is done in New Zealand).

Using data supplied by the MSD and birth records, the researchers calculated that 5.4% of all New Zealand children have a substantiated maltreatment finding by age five. 83% of those children would have appeared in the project sample (that is, started a benefit spell before age 2), and so received a risk score under the proposed algorithm. Given that the sample comprised approximately 33% of all children born in New Zealand during the study period, offering intensive prevention services to all of the children who received a risk score would be a massive undertaking and one that would involve a very high false positive rate –

⁸ Because one can adjust the threshold set for a positive (i.e. by changing the minimum risk score), a statistic like the AUC that shows accuracy over a range of thresholds is required. The ROC curve compares two figures across this range. The first figure is the True Positive Rate (TPR), or sensitivity. In the present case, it equals the percentage of households that received a score over a given threshold and ended up with a maltreatment finding among all those that received a maltreatment finding. The second figure is the False Positive Rate (FPR). In the present case, it equals the percentage of households that received a score over the threshold but did not have a maltreatment finding. Setting the threshold for 'maltreatment' low (e.g., risk score 1) gives a very high TPR, since it makes it unlikely any cases of maltreatment will be missed, but it also generates a very high FPR since in making sure it captures every true positive it includes many false positives (at some level the TPR and FPR would both be 100%). Setting the threshold very high, to the point of stringency where few households met the criteria positive, will deliver TPRs and FPRs of or approaching zero, since it excludes all positives – true and false). At any threshold where the TPR and FPR are the same, the model is only as good as chance.

prompting unnecessary interventions – since most of those children appear in the lower risk deciles. The Vulnerable Children Report therefore suggested focusing intensive intervention⁹ on the top two decile groups – children with risk scores of 9 and 10 – which contained 37% of total national maltreatment findings, while making up 5% of the total population. If social services were to refer all cases in the top two deciles they would refer 5% of the total population to the intervention, with a false positive rate of 63%, and capture rate of 37% of national maltreatment cases.¹⁰ If we assume that assistance would be offered in all cases of a 9 or 10 risk score and assume moreover that 50% of subjects took up the offer, with a nominal prevention-per-intervention rate of 46%,¹¹ then we could expect 280 cases of maltreatment to be prevented up to age 5 for every two years of conducting the programme.

The report also offers a 'business case' for the model based on these assumptions, concluding that on the assumption that each intervention cost \$8,210, the cost per maltreatment avoided would be \$48,000 if the programme were offered to all children in the top two risk deciles (i.e. deciles 9 and 10), and \$32,000 if it were offered to children in the top decile only.

⁹ Services provided to children and their families generally fall into three categories: universal, targeted and specialist. Universal services are offered to all children and families (primary prevention). New Zealand's Well Child/Tamariki Ora services are examples (see footnote 14 below for a little more detail). Targeted services (secondary prevention) are aimed at families with identified needs. The Nurse Family Partnership Programme developed by Olds and his colleagues in the US, and referred to in Vaithianathan et al., 2012 where low income young mothers with no previous live births are offered intensive home visiting (Olds et al, 1997). Specialist services (tertiary prevention) are offered to cases where maltreatment has already occurred. In New Zealand such services are offered where various statutory care and protection responses are triggered.

¹⁰ We will see that there is reason not to adopt this recommendation to target the Vulnerable Children PRM's top two risk deciles, but instead to target children identified as being at high risk: See section 2.a below. I retain Vaithianathan et al's references to these risk deciles for the purposes of describing their report.

¹¹ This figure is derived from the US Nurse Family Partnership Programme, used in Vaithianathan et al., 2012 as an illustration of possible intervention models. See Olds et al., 1997. I say more about this programme and alternatives in section 2.e.i. below.

Preliminary Points

1. Technical arguments for the predictive power of the model

The power of the predictive risk model outlined in the Vulnerable Children Study is of central importance to an ethical evaluation of its application to the assessment of child maltreatment risk. The technical grounds for confidence in the power of the model are set out in Vaithianathan et al., 2012. This review does not revisit or evaluate those technical arguments. Instead it evaluates the model as presented in the report, proceeding on the assumption that the claims made in the report as to the power of the model are accurate.

It is important to note, however, that accepting the claims made as to the power of the Vulnerable Children PRM in Vaithianathan et al., 2012 does not require us to avoid the significant ethical issues raised by the accuracy of the model. At its best predictive risk modeling of child maltreatment will not identify every at risk family, and maltreatment findings will not be made against all of those it does identify as at risk – no predictive model has perfect specificity and sensitivity. I return to this issue below. (See in particular the discussion under section 2.b, ‘Over and Under Identification’, but the issue arises throughout our discussion).

2. The significance of implementation decisions

Important implementation decisions remain to be made about PRM in the child maltreatment area. This review makes some suggestions as to implementation, but cannot provide a full assessment of possible implementation protocols. The details of such protocols will have considerable significance for the ethical issues generated by the Vulnerable Children PRM. Implementation decisions to be taken, some addressed in more detail than others in this report, seem likely to include:

- What information is given to social service professionals through the PRM tool. Whether such professionals are provided, for instance, with risk-scores or with lists of families with the expectation that they will follow up and offer services where they judge it appropriate to do so;

- Whether follow-up on the part of agencies to high risk scores is mandatory or discretionary;
- Whether assistance given to high-risk families is mandatory or voluntary;
- Which existing agency will have responsibility for monitoring and responding to the risk scores produced by the Vulnerable Children PRM, or whether a new agency or division will be created to take on that role;
- What particular assistance programmes are offered, and their effectiveness;
- What amount of funding is attached to the programme;
- How enacting the programme affects overall workloads of agencies;
- How the programme is presented to and understood by the public and professionals;
- What data sets and which people end up being included.

3. Legal Issues

Many of the ethical issues discussed in this paper raise legal issues. Although this report identifies and touches upon those legal issues where necessary, it does not purport to offer a comprehensive coverage of relevant legal issues.

4. PRM and alternative approaches to child maltreatment

It is important to bear in mind throughout this review that while the application of automated predictive risk modelling to child maltreatment is new – the authors of the Vulnerable Children Study identify only one researcher studying the use of information and communication technology modelling to predict future child maltreatment¹² – institutional or state sponsored social service agencies have been responding to concerns about child maltreatment since at least the 1870s.¹³

¹² Schwartz et al., 2008

¹³ The world's first organization devoted entirely to child protection seems to have been the New York Society for the Prevention of Cruelty to Children formed in 1875 (Myers 2008). Britain's first act of parliament for the prevention of cruelty to children, commonly known as the "children's charter" was passed in 1889, providing powers to enter homes if children were thought to be in danger and arrest those found to be ill-treating children. It is clear, however, that there was both common law and statutory

A review of the history of child protection is beyond the scope of this paper. A fuller account of the range of responses is available in the literature review that accompanied the Vulnerable Children Study report.¹⁴ For current purposes, it will suffice to emphasise that every attempt to prevent child maltreatment has proven to carry its own significant risks: of failing to identify at-risk children in time; of initiating unnecessary and disruptive intervention in families who would in fact have never maltreated their children; of erecting barriers between social service professionals and families; of exposing families and agencies to sometimes venomous media attention; of demoralising child-protection professionals; of stigmatising families; of fostering confirmation bias; and so on.¹⁵

The point of this bleak litany is not to condemn all that has gone before, or to suggest that PRM promises any kind of miraculous solution to the ‘wicked problem’ of child maltreatment.¹⁶ Rather, it is to emphasize that whatever the actual costs and benefits of using PRM in this area, the approach must be considered in light of alternatives which carry very real costs of their own. The ethical questions are essentially and unavoidably comparative; questions not simply about the costs and benefits of PRM, but about how those costs and benefits compare, from an ethical perspective, with the costs and benefits of plausible alternatives.

protection for children before these general initiatives. Myers cites a 1642 Massachusetts law that gave magistrates the authority to remove children from parents who did not "train up" their children properly (Myers 2008, 450). (Although, as the quoted reference to training suggests, the Massachusetts' law was not primarily concerned with child welfare issues. The full title of the Act – the 'Massachusetts Bay law for the training of children and servants in the home' clearly signaled its primary goal of ensuring that householders trained children and servants "in learning, and labor and other employments which may be profitable to the commonwealth". I thank a referee for drawing my attention to the focus of the Massachusetts' law.)

¹⁴ Vaithianathan et al., 2012.

¹⁵ Munro 1996; Mansell et al., 2011.

¹⁶ The term ‘wicked problem’ seems bleakly apt from a normative perspective. In this context, however, it is a technical term describing the complex composition of the child maltreatment problem and its consequent resistance to linear or step-wise problem-solving, such that responses to one of its strands often reveal others, even more complex and challenging. (Devaney 2009)

Ethical Issues

1. Overview

I remarked at the outset that providing a moral evaluation of the Vulnerable Children PRM requires us to identify and balance the relevant moral perspectives and interests in order to arrive at a view of the ethical status of the model.

The most significant consideration on the *positive* side of this ethical balance is the possibility that the Vulnerable Children predictive risk model will allow social services to identify children with an elevated risk of future maltreatment, and having done so, to offer services that will remove or reduce that risk. The Vulnerable Children PRM might be able to make a difference to the lives of some of the most vulnerable members of our community. Accurately identifying such children might also allow the targeted, and so it is hoped more effective, allocation of child protection resources. *If* PRM did allow the early identification of vulnerable children, and *if* effective interventions were available, and *if* such identification allowed for the more effective allocation of preventative resources, then the potential benefits of the Vulnerable Children PRM would be considerable and of obvious moral value.

However the application of predictive risk modelling to child maltreatment also has very clear ethical risks and costs.

- The tool has a fair to good degree of accuracy: 48% of those placed in the highest risk decile went on to have a maltreatment finding by the age of 5. It follows, however, that 52% of the families placed in that highest risk decile did not have such a finding (although when these children are followed to age 7 slightly more than 50% have a substantiated maltreatment). The tool then, would call for intervention (of a sort yet to be determined) with families who would not go on to maltreat their children. Some of these unnecessary interventions would impose burdens on families.

- Notwithstanding its accuracy, the tool will fail to identify some at risk children. One obvious reason the Vulnerable Children PRM will fail to do so is its reliance on the particular databases from which it draws information. Families who do not have contact with any of the agencies under the MSD umbrella will not appear in the databases available to the tool. Moreover, the figures used to validate the Vulnerable Children PRM are based on cases of *substantiated* maltreatment. Vaithianathan, et al., 2012 relies upon evidence that 5.4% of all children born in New Zealand during the study period had findings of maltreatment by age 5, but acknowledges that those children are “only a subset of those who actually have some maltreatment occur to them”.¹⁷ Not all maltreatment will be substantiated. More generally, no predictive tool has perfect sensitivity.
- One potential burden associated with the Vulnerable Children PRM is the possible stigmatisation of risk scored individuals and families. A very high proportion of substantiated child maltreatment in New Zealand occurs in the homes of benefit recipients. This is not surprising. It has long been appreciated that there is a correlation between the sorts of background social conditions that are likely to lead a family to require social support and the conditions that contribute to child maltreatment. However, beneficiaries already face considerable social stigmatisation – a recent poll identifies them as the group that suffers the most discrimination in New Zealand¹⁸ – and it is possible that the predictive risk model would capture and reinforce that stigmatisation. If it did so, it might actually contribute to child maltreatment by increasing the pressure and social isolation of targeted families.¹⁹
- The promise of the tool rests upon it being possible to design effective interventions, but this may not be straightforward. Given the complexity of

¹⁷ Vaithianathan et al., 2012, p.17

¹⁸ Dickison, 2013

¹⁹ The issue connects with the second bullet point above, arising in part because the Vulnerable Children PRM is focused entirely upon social welfare benefit recipients. The MSD is currently exploring the feasibility of broadening out the contributing databases to include all births.

the child maltreatment problem, we cannot assume that every identified case will be preventable. Some of the highest risk families may not be receptive to assistance, mandatory or otherwise, and, even if they are, available interventions may not succeed in preventing maltreatment.

- Effective intervention strategies require engagement with at risk families. Many such families will have no reason to have a current and active relationship with a social service professional, and it is not easy to see how such a relationship will be initiated and by whom. Furthermore, even where there are existing relationships, perhaps with professionals providing services under the Wellchild/Tamariki Ora programme, those relationships may not easily be broadened to include interventions recommended by the Vulnerable Children PRM. Doing so might threaten the effectiveness of the existing relationship.

The application of predictive risk modelling tools to child maltreatment, then, raises a considerable range of significant ethical concerns. Can those ethical costs be ameliorated or completely addressed? Are those that cannot be addressed outweighed by the very considerable ethical benefits that might be delivered by the tool? These are the broad questions to bear in mind as we proceed through the detail of the evaluation.

2. Specific Ethical Issues

a) Universal vs Targeted Responses

The Vulnerable Children PRM is proposed as part of a targeted response to child maltreatment. It seeks to generate information that will allow social service professionals to make early identifications of at risk children, to focus intensive preventative efforts on those children and their families, and to inform targeted allocations of scarce child protection resources. At this level of description, the approach can be contrasted with various forms of universal assistance that seek to reduce recognized maltreatment risk factors (such as poverty) across entire populations. At least some early criticisms of the Vulnerable Children PRM have

been motivated by the belief that universal approaches should be preferred over targeted alternatives.

However there are significant limitations to any plausible universal response to child maltreatment.

New Zealand already offers some forms of universal assistance to families through programmes such as the Well Child/Tamariki Ora service,²⁰ and Vaithianathan et al., 2012 assume that delivery of intensive services to high-risk families would not lead to changes to these existing programmes. However, these programmes involve relatively 'light-touch' engagement by contrast with the intensive interventions which evidence suggests may be effective in addressing child maltreatment.²¹ If these programmes were maintained at their current level of intensity, there is little reason to think they would address New Zealand's child maltreatment problem: after all, they are already in place and available to all New Zealand families.

It is likely to be difficult and risky to attempt to transform the content and delivery of existing programmes in ways would have them deliver the sort of interventions that there is some reason to believe might be effective in reducing maltreatment risk. It would be difficult, for instance, to maintain the important existing relationships between families and the nurses and community health workers who deliver the Well Child/Tamariki Ora programme while changing the focus and intensity of that programme. Doing so may well place that programme at risk if families felt less willing to engage.²² These difficulties are unlikely to be unique

²⁰ "Well Child/Tamariki Ora services are a screening, education and support service offered to all New Zealand children and their families or whānau from birth to five years. The services assist and support families and whānau to improve and protect their children's health. Well Child services are provided by registered nurses and community health workers/kaiāwhina with specific training in child health. The current Well Child Schedule includes 12 health checks (contacts) with eight of these offered to children aged from six weeks to five years. At present not all children receive all 12 contacts. First time parents and families requiring extra support are offered additional contacts. Some contacts are carried out in the home and others in clinics or community facilities" (Well Child/Tamariki Ora National Schedule).

²¹ Though see discussion below, section 2.e.i.

²² "Many nursing professionals also worry that the role in screening for child maltreatment can make building a working relationship between them and expectant parents more difficult, as parents may perceive health visitors in a policing rather than supporting capacity and therefore antagonise further intervention and advice offered by health visitors." Browne et al, undated, 12.

to the Well Child/Tamariki Ora programme. New universal programmes would face similar difficulties.

The differences between the relatively light-touch universal programmes it would be feasible to offer population wide and those that have been effective²³ in reducing child maltreatment have significant resource implications. Since they are population-wide, universal programs have significant costs. Given the relatively low incidence of child maltreatment (5.4% is an appalling rate of maltreatment, but still it is a low absolute percentage) much of the cost of such programmes would be directed to families whose children were not at any elevated risk of maltreatment. For these families, there would be no gain (in terms of reduced maltreatment outcomes) from resources directed to them. However such an allocation would mean that fewer resources would be available to treat families whose children were at elevated risk. Focus on universal service provision to the exclusion of targeted programmes, then, could lead to higher rates of maltreatment as high-risk families did not receive the sort of intense intervention that might have reduced their risk-status.

Even if it were possible to provide the sort of intervention that might reduce the risk of maltreatment on a universal basis, it is not clear that it would be appropriate to do so. We will see below that the 'false-positive' problem suggests that even in the case of high-risk families, intervention should be moderated to reduce the burdens it would impose on families incorrectly identified as at risk. That same argument would apply in the case of a universal programme. It is not just that one could not feasibly offer an effective programme universally: arguably one should not do so since a universal model would involve intensive intervention with many families who were not fact at elevated risk of child maltreatment.

Of course similar resource allocation concerns can be raised against any preferred distribution. Vaithianathan et al., 2012 propose a focus on the two highest risk deciles. Adopting that proposal would lead to less intensive focus on lower-risk

²³ Again, discussed further in section 2.e.i, below.

families, perhaps at the cost of interventions which would prevent lower risk families escalating into higher-risk cases and perhaps eventual maltreatment.

Given such considerations, from an ethical perspective, resource distributions should be justifiable in terms of their ability to protect important moral interests while minimising moral burdens posed by, for instance, unnecessary intervention.

In my view the proposal to focus intensive interventions on children identified as at high risk, while retaining the current universal assistance programmes, seems to be plausibly justified.

There is reason to be wary of the specific proposal in Vaithianathan et al., 2012., 2012 to focus intervention on their two highest risk deciles. Doing so may impose arbitrary restrictions on interventions that may prove to be more effectively offered to a wider or narrower group. The advantages of targeted intervention are gained by focusing on children identified as at high risk, without committing to the inevitably somewhat abstract groupings into risk deciles. We do not need the apparent precision of Vaithianathan et al's deciles to make the general point that targeted intervention is ethically justifiable, and that predictive risk modelling may be an appropriate way of allocating child protection resources. Nonetheless, those deciles serve to spell out the benefits of targeted intervention. Targeting Vaithianathan et al's top two deciles would see intensive intervention offered to 5% of the total population, with a false positive rate of 63%, and capture rate of 37% of national maltreatment cases. If 50% of families in those deciles took up the offer, with a nominal prevention-per-intervention rate of 46% derived from the US Nurse Family Partnership Programme, we could expect 280 cases of maltreatment to be prevented up to age 5 for every two years of conducting the programme. The 'false positive' rate here is not trivial, and we will return below to ways in which its ethical significance might be addressed. In sum, however the balance of moral benefits and burdens appears plausibly to support a targeted programme broadly along the lines proposed in Vaithianathan et al., 2012.

b) Over and Under Identification

Notwithstanding the predictive power of the Vulnerable Children PRM, it cannot avoid the fundamental difficulty faced by social service professionals attempting to address child maltreatment: decision making in child protection is intrinsically difficult, uncertain, and subject to the practical limits to our ability to predict future outcomes.²⁴ As with any risk prediction tool, the Vulnerable Children PRM will inevitably make some errors at any threshold for referral, identifying as low risk some children who go on to experience abuse or neglect, and identifying as high risk some children who do not. It will not identify all children who go to experience maltreatment because information that should inform the assessment of risk may not be available in the administrative data, or may not be able to be taken into account because of difficulties with data linkage. Children whose families have little or no prior contact with social agencies will be under-represented among those identified as high risk simply because less will be known about them.

Nonetheless, there is evidence that the Vulnerable Children PRM is on the right track, that standardised ‘actuarial’ assessment tools based on empirical evidence are more accurate than consensus-based models or professional judgment in the assessment of risk of future harm to children,²⁵ and that computerised predictive risk models have clear advantages over existing risk assessment tools. In addition to the evidence supplied in Vaithianathan et al., 2012 predictive risk modeling has a substantial history in health care that provides good data on the power of such modelling.²⁶ Predictive risk models utilise actual and current data for the specific population at risk and outcome of concern, and allow population-wide screening that is reasonably cost-effective. Standardised assessment tools tend not to be validated for the population to which they are applied,²⁷ are often under-utilised

²⁴ See for instance Munro, 2011 and Leventhal, 1988.

²⁵ Gambrill et al., 2000; Barlow et al., 2012; D’Andrade et al., 2008.

²⁶ Panattoni et al., 2011

²⁷ Schwartz et al., 2008

where they are resource intensive to apply, and are vulnerable to a range of operator errors at the data-entry point.²⁸

What, from an ethical perspective, should we make of the ineliminable predictive shortcomings of the Vulnerable Children PRM? I will address a number of aspects of this question below (those concerning stigmatisation, for instance). For now it is important to note that alternative methods of identifying at risk families are no better and quite probably worse. The reasons canvassed in the previous section for preferring targeted responses to child maltreatment over universal alternatives may simply amount to reasons to prefer predictive risk modelling if it is the best risk identification tool available.

If that is right, the ethical onus should be on seeking ways to mitigate potential harms that result from the Vulnerable Children PRM's predictive limitations.

One obvious response to some of those limitations is to broaden the databases upon which the Vulnerable Children PRM draws. Currently, families who do not have contact with any of the agencies under the MSD umbrella will not appear in the databases available to the tool. They may appear in other databases, however, such as those generated by the health system and New Zealand births, deaths and marriages registry. The MSD is currently assessing the feasibility of a child maltreatment PRM which includes all children whose births are registered, a cohort which would capture 94% of all New Zealand live-born children, and yield records for in excess of 60,000 children in each annual birth cohort.²⁹

One should also ensure that other routes of referral are created or maintained. The Vulnerable Children PRM should be seen as an additional method for predicting risk, not a replacement. Current early identification referral routes include those initiated by health professionals and other front-line social service

²⁸ Vaithianathan et al., 2012

²⁹ New Zealand Government, 2012, Volume II, p.81, para.60.

professionals, and these should be maintained alongside the Vulnerable Children PRM. Again, the MSD is proposing to maintain other referral routes.³⁰

These responses – broadening the data bases available to the Vulnerable Children PRM and maintaining alternative referral routes - are essentially efforts to provide as broad and accurate a referral process as possible. Other limitations on the accuracy of the Vulnerable Children PRM, in particular those which lead it to incorrectly classify families as high-risk – the false-positive problem – call for different responses. What might one do to reduce the ethical significance of false positives under the Vulnerable Children PRM?

Most obviously, one should reduce the false-positive rate as far as possible. Doing so is not simply a matter of improving the predictive accuracy of the tool. Most obviously, as Vaithianathan et al., 2012 recognise, one can also do so by choosing higher thresholds for intervention: limiting intervention to children identified as at high-risk means that fewer families would be mistakenly identified as falling above an intervention threshold. The desire to reduce false-positive rates is one important element in the balance to be struck between reducing child maltreatment as far as possible while minimising the imposition of unwarranted burdens.

Beyond this, the significance of the false-positive problem turns in large part upon the *consequences* of being mistakenly identified as at high-risk. We will see below, for instance, that false-positives are an extremely serious issue for the PSA test for prostate cancer because a positive finding prompts an invasive and risky biopsy and, potentially, unwarranted stress.³¹ One might reduce the ethical significance of false-positives in the child maltreatment context, then, by reducing the

³⁰ New Zealand Government, 2012, Volume II, p.80, para.58: "... Predicting maltreatment is not easy, and risk scoring will be just one component of a wider system to target support to vulnerable children:

- ...
- Children who are not picked up by the model will still be able to be identified as at risk and prioritised for assessment and services through identification by frontline professionals, including care and protection social workers.

³¹ Section 2.e.

significance of the consequences of misidentification. Some ways to do that appear to include:

- Providing opportunity for experienced social service professionals to exercise judgement about appropriate responses to a family's identification as at risk;
- Ensuring that such professionals understand the potential of the Vulnerable Children PRM to miscategorise families;
- Providing training to guard, in so far as possible, against confirmation bias in the professional engagement with families identified as high-risk;
- Offering rather than requiring engagement as a consequence of identification as at risk;
- Ensuring that intervention triggered by identification as at risk is as non-intrusive as possible consistent with the overall aims of reducing child maltreatment risk.
- Ensuring that intervention triggered by identification as at risk is positive and supportive rather than punitive;
- Identifying and minimising the adverse effects of identification as at risk, such as, for instance, possible stigmatisation.

c) Stigmatisation and the Costs Associated with Identification as At-Risk

Predictive risk tools seek to identify the probability that some state of affairs will come into existence in the future. Particular ethical issues arise when significant burdens are likely to be borne by individuals or groups *as a consequence* of being identified as at risk of realizing that state of affairs. Such costs may range from those that are fairly straightforward and transparent, such as increased difficulty in obtaining insurance, to the more complex and diverse burdens of social stigmatisation.³² There are a number of distinct issues here:

³² Note at the outset that it is important to guard against a possible misinterpretation of the 'product' of the predictive risk model. Some of the possible burdens associated with identification as at risk – such as

- While we might accept that it is proper for an individual to bear costs associated with actual realised wrong doing, it is difficult to justify the imposition of such costs in anticipation of conduct that might never come to pass, merely on the basis of that individual's membership of a high-risk group.
- In many cases, the burdens that follow from being identified as a member of a group arise from false beliefs about what that membership means. For example, social stigmatisation of people with HIV/AIDS involves mistaken views about the infectiousness of the disease and a lack of awareness of the effectiveness of modern treatment options. The allocation of burdens in this case is illegitimate because HIV/AIDS sufferers pose little or no risk to those with whom they have casual contact.³³ It seems quite likely that identification as an at risk individual or family under the Vulnerable Children PRM would be open to similar misinterpretation. More generally, some of the data and predictor variables used by the Vulnerable Children PRM are very likely to be misinterpreted by at least some audiences, in part because of the existing stigma borne by groups likely to be over-represented in those identified as at high-risk.

There is an obvious point worth emphasizing here. While the allocation of burdens in such cases is ill founded, they are no less onerous for those who bear them. Such burdens cannot be discounted in an ethical assessment simply because they rest on prejudice and misunderstanding. Policy makers who trigger such costs must take them into account even though they may think them completely without foundation and even though, in some sense, the policy makers are not responsible for their imposition on the bearers of illegitimate stigma.

being given an opportunity to engage with social services – are properly seen as responses to a status which *will* have been realised at the time of any such offer, namely that of being identified as an at-risk family. It should not be supposed that those are burdens that attach to having had a maltreatment finding, the future state of affairs predicted by the model.

³³ Stigmatisation will not always be driven by 'beliefs' and this may matter. Racism, for instance, is likely to be significantly grounded in sub-doxastic (i.e. non-belief-driven) states or unreflective emotions that cannot be easily corrected through education.

- In some cases the burdens associated with identification as an at risk individual or group may actually increase risk of the adverse outcome. The stigmatisation of HIV/AIDS sufferers act as a barrier to disclosure, for instance, making sufferers reluctant to take measures to reduce harm to their partners or to seek effective treatment and hence increasing risks both to themselves and to others.³⁴ In other cases the connection between identification as at risk and increased risk of realisation may be less direct, perhaps, for instance, leading health, social service, and enforcement professionals to engage differently with stigmatised individuals.³⁵

The sorts of burdens associated with identification as an at risk group may also have more general implications for health and well-being. Jonathan Mann, Director of the World Health Organization's Global Program on AIDS, suggests that the health effects of stigmatisation in the HIV/AIDS case were illustrative of a broader correlation, showing "a consistent pattern through which discrimination, marginalization, stigmatization and more generally a lack of respect for the human rights and dignity of individuals and groups heightens their vulnerability.... In this regard, HIV/AIDS may be illustrative of a more general phenomenon in which individual and population vulnerability to disease, disability and premature death is linked to the status of respect for human rights and dignity."³⁶

- The burdens of stigmatisation following identification as an at risk group or individual often fall upon those who are already the subject of social disapproval or demarcation, 'appropriating and reinforcing pre-existing stigma'.³⁷ It is easy to see that this may be an issue in the current context. Inevitably, beneficiaries will be over-represented in those identified as at

³⁴ "The widespread expectation of stigma combined with actual experiences with prejudice and discrimination exerts a considerable impact on [people with HIV], their loved ones and care-givers. It affects many of the choices [they] make about being tested and seeking assistance for their physical, psychological and social needs. . . . Fearing rejection and mistreatment many . . . keep their serostatus a secret" (Herek, 1999).

³⁵ See the large literature on confirmation bias, and Gambrill et al., 2000.

³⁶ Mann, et al., 1994.

³⁷ Parker and Aggleton, 2002.

high-risk by the Vulnerable Children PRM even if the data available to the model's algorithm are broadened beyond that in the welfare and CYFS databases. A recent New Zealand Human Rights Commission report showed high levels of discrimination against beneficiaries³⁸, echoing the results of a similar study in the United Kingdom.³⁹ These pre-existing vulnerabilities and stigmas clearly raise a particular ethical duty of care when considering action that may exacerbate the social disapproval and isolation already experienced by groups and individuals.

i. Possible Responses to concerns about stigmatisation

- I remarked at the outset that the costs and benefits of the Vulnerable Children PRM must be considered in light of alternatives that carry very real costs of their own. That point is germane here. It is appropriate to be concerned about the costs the model might impose upon those identified as at risk. However those costs are not unique to the Vulnerable Children PRM. It would be naive to suppose, for instance, that negative conclusions were not already drawn from correlations between child maltreatment and socio-economic position, that existing approaches to child protection did not carry risks of confirmation bias, of unwarranted intrusion on families who were not at risk, of appropriating and reinforcing existing stigma.

The point again is not to suggest that these costs can be disregarded. Rather it is to emphasise the importance of weighing the costs and benefits of implementing PRM against those of alternatives. Arguably, for instance, the Vulnerable Children PRM may reduce some of these potential burdens, allowing social service professionals to avoid confirmation bias more effectively, and allowing more effective targeting of services which, while not eliminating unwarranted intrusion, may reduce it, and so on.

³⁸ Dickison, 2013.

³⁹ Baumberg, et al., 2012.

- I began my attempt to spell out the possible costs of identification as an at risk family or individual with the obvious concern that *predictive* risk models seem likely to allocate burdens prior to the realization of predicted states of affairs and when it remains possible that they will not be realised. This concern about the ‘anticipatory allocation’ of burdens raises much broader issues about when it is legitimate to interfere with individual liberty. The standard liberal position is straightforward. There is, John Stuart Mill argued, “one very simple principle, as entitled to govern absolutely the dealings of society with the individual in the way of compulsion and control. ... The only purpose for which power can be rightfully exercised over any member of a civilized community, against his will is to prevent harm to others”.⁴⁰ Although it is clear even in this initial statement of the harm principle, Mill goes on to make explicit that the principle may warrant preventative interference. Speaking of crime, though in a way that makes clear his comments encompass future harms in general, he writes, “it is one of the undisputed functions of government to take precautions against crime before it is committed, as well as to detect and punish it afterwards”.⁴¹

This does not, of course, give states *carte blanche* to interfere wherever they perceive a threat of harm. With characteristic suspicion of intrusion into individual liberty, Mill recognizes that preventative intervention is particularly threatening: “The preventive function of government”, he writes “is far more liable to be abused, to the prejudice of liberty, than the punitory [sic] function; for there is hardly any part of legitimate freedom of action of a human being which would not admit of being represented ... as increasing the facilities for some form or other of delinquency”.⁴² On Mill’s account, we must preserve as much liberty as possible, even when we are justified in interfering, choosing the least intrusive means necessary to prevent harm: if voluntary

⁴⁰ Mill, 1859, Chapter 1, para. 9.

⁴¹ *Ibid.*, Chapter 5, para 5.

⁴² *Ibid.*

engagement is an effective way of reducing the risk of child maltreatment, it is to be preferred over more liberty-limiting alternatives.

More generally, we should prefer less rather than more intrusive interventions *provided that* they are sufficient to prevent the targeted harm and do not threaten more harm than that we seek to prevent. It is important not to understate the significance of the first leg of this proviso. Mill's starting point is that we are entitled to interfere to prevent harm to others. Mill would think – and so should we – that more rather than less intrusive interference would be justified only to the extent that it was reasonable to believe that it was necessary to address a threat of serious harm.

Much of this discussion assumes that liberal communities must make judgements about probability of harm. Even in the absence of PRM, social service professionals are constantly and unavoidably doing so. We might hope that PRM gives us better grounds for these judgements than some of the alternatives. There is a very general point here. It should not be supposed that it is unethical to make social policy, even coercive policy, where outcomes are less than absolutely certain. Policy makers will often have to act under conditions of uncertainty since these are the inescapably the conditions in which we find ourselves much of the time. Indeed, it may be an unethical abrogation of responsibility to refuse to act other than in conditions of certainty.⁴³

The mere fact that predictive risk modelling will place burdens upon families and individuals who have done no wrong and who may never do so, then, seems not of itself sufficient to rule against such modelling.

- The most obvious response to at least some of these concerns about stigmatisation is to maintain careful control over the dissemination of the 'product' of the predictive risk model. I suggest that such information should be disseminated as narrowly as possible, consistently with achieving the

⁴³ See Dare, 1998 proposing vaccination as an example.

benefits of the programme. This may mean, for instance, that only senior social service professionals should have access to such information, that they be carefully trained as to how to manage the information they possess, so as to limit opportunity for stigmatisation of individuals or groups, and that consideration be given to what level of detail is required to make effective use of the model's predictions.

- Social service professionals are at a general risk of confirmation bias whenever evaluating the risk that a child may be harmed. They may be more inclined to accept information that confirms the positions they originally held, and ignore information contradicting them. As Eileen Munro puts it: “Whether suspicious or optimistic about a family, social workers tend to be biased in their attitude towards new information.”⁴⁴ One can imagine how a PRM programme might aggravate this issue: If a social service professional begins interaction with a family knowing that they are in the top risk decile, they may more readily identify evidence that reinforces this risk attribution. However, we can also envisage that the programme may reduce confirmation bias. Because the system is more rigorous and transparent than any of our internalised heuristics about 'what sort of people' are more likely to be at risk, the system may act as a yardstick against which social service professionals can objectively evaluate their preconceptions.⁴⁵
- Given that the Vulnerable Children PRM predicts risk, rather than wrongdoing, it is crucial that any resulting engagement is supportive and preventative, rather than punitive. Not only will a significant portion of families not go on to record maltreatment findings, none of the families need have done so at the time at which they appear in the Vulnerable Children PRM's risk rankings. Responses must not prejudice individuals or families or appear to do so.

⁴⁴ Munro, 1996, p. 799. See also Browne, undated, 12: “ ... a child injury are [sic] more likely to be regarded as *non-accidental* if professionals are aware of them being classed as 'high-risk' and this is dangerous.”

⁴⁵ There is an interesting issue here. It will difficult to attribute the absence of maltreatment findings to interventions prompted by the Vulnerable Children PRM (since for any individual case it is possible no maltreatment finding would have occurred). Social service professionals could very easily reinforce confirmation bias by assuming that no maltreatment occurred only because they had intervened.

Carefully ensuring that services are supportive and preventive may also increase the acceptability and so effectiveness of engagement.

For the same reasons, it is important that classification is expressed in ways that do not suggest that wrongdoing has occurred. The authors of a European study of early prediction and risk detection models in Child Protection follow this line of reasoning to conclude that while labeling families as ‘high-risk’ would be unethical and stigmatising given the modest power of the tools they were considering (all considerably less accurate than the Vulnerable Children PRM), nonetheless screening to identify children and families in need for more support was justified, provided the classification of families was expressed positively, as, for instance, ‘high priority for services’ rather than “high risk”.⁴⁶

- An essential element of stigma is the strong delineation and distancing of the stigmatised group. One general strategy for mitigating stigma, therefore, would be for the state to treat high risk families as similarly as possible to other families, at least in the external presentation of their interactions. I have suggested that this cannot amount to a universal as opposed to a targeted programme.⁴⁷ I do suggest, however, that agencies implementing the Vulnerable Children PRM seek interventions that do not themselves mark families or individuals as high-risk.
- The New Zealand media, like their international counterparts, have shown considerable interest in child maltreatment cases. In at least some instances, media coverage has clearly contributed to the stigmatisation of struggling families, generalising from tragic cases in ways that attribute undesirable characteristics to socially identifiable groups. Agencies with responsibility for implementing the Vulnerable Children PRM should explore ways to engage more productively with the media over child maltreatment, and, more generally, the presentation of social service agency clients. The media are also likely to be interested in the Vulnerable Children PRM itself. Given that the

⁴⁶ Browne, undated, 41.

⁴⁷ Section 2.b.

scientific basis of the system is not entirely straightforward, and given that misinterpretation of the system could come at the cost of additional stigmatization of people identified as being at high-risk, it is important to ensure that the Vulnerable Children PRM is presented as accurately as possible. Done poorly, media presentation of the Vulnerable Children PRM could create a false association of a high risk score with some form of present moral wrongdoing, overlooking the fact that many high risk families will not record a maltreatment finding, and a false impression that risk scores are individual, rather than population based, and thus that any child with a high risk score is at actual risk from their parents. Done well, one might hope, these public responses might be largely avoided. Hence there is an ethical responsibility to ensure that the media receive a sober and accurate impression of the system to convey to the public.

d) Mandatory vs. Voluntary Engagement with Services.

Vaithianathan et al., remark that: “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 were used to engage high-risk families in voluntary services.⁴⁸ I can now spell out the considerations that support this preliminary conclusion in more detail. There are a number of reasons to prefer voluntary engagement.

It is widely accepted in the social service professions that mandating engagement compromises efficacy. In a recent survey article Snyder and Anderson write that “[t]he literature is relatively uniform in suggesting that mandated clients are more resistant to treatment than voluntary clients”.⁴⁹ They go on to note, however, that several studies highlight the “undeniable treatment successes in ... many areas of mandated therapy”, including that directed at addressing child abuse. The studies supporting mandatory engagement, however, are few and focus on court ordered

⁴⁸ Vaithianathan et al., 2012, p.2.

⁴⁹ Snyder and Anderson, 2009, p. 279.

engagement – a feature, we will see in a moment, that distinguishes them from possible mandatory engagement processes that might be considered in response to a predictive risk score. On balance, it seems reasonable to assume for current purposes that the common opposition to mandatory engagement is well founded, especially in light of the other reasons to reject such an approach that follow.

One reason to favour voluntary engagement is that it is not easy to see what sort of engagement and compulsion might be possible in the current context. There is a set of related issues around this uncertainty.

As I have already mentioned studies supporting mandating focused on court ordered engagement. The position of families or individuals who have come before the courts is importantly different from that of those identified as at high risk by a PRM tool. First, courts can significantly curtail the liberty of individuals only *after* it has been established by due process that these individuals have breached some legal standard. In the course of that process they will have the opportunity to hear the charges or claims made against them, to offer a defence with appropriate assistance, to be judged in a public and neutral manner, and so on. The position of the individual who is merely predicted to pose a risk is fundamentally different. It is perhaps worth casting this point in terms that relate directly to apparently analogous tests found in law. We have seen that 48% of children with decile 10 scores had maltreatment findings by the age of 5. That risk assessment figure may seem to approach the ‘balance of probabilities’ test that applies in civil proceedings, sometimes expressed in percentage terms as being met at probabilities over 50%, where it warrants quite significant intrusions on individual liberty. However, the civil law’s balance of probabilities test applies to a determination that some wrong *has* occurred, and it is proof of that to the appropriate standard that warrants interference with individual liberties.

There is also an important pragmatic difference between the positions of those before the courts and those identified as at high-risk by a predictive risk modelling tool. The former will have an existing and active relationship with institutional or state agencies. Families who receive a high risk score might not have an existing

relationships if, for instance, their child is say too young for school and too old for a Plunket nurse. In the absence of pre-existing relationships it is harder to see how discussion about engagement might begin, or why families or individuals would cooperate.

One might also fear that mandatory engagement would actually do more harm than good. We have seen in other contexts that attempting to compel engagement occasionally drives families further from possible sources of support, either literally as in the case of 3-year-old Liam Williams-Holloway, whose parents took him into hiding to prevent the boy's cancer from being treated,⁵⁰ or by alienating them from such services.

The limited literature on mandated engagement – focusing, we have seen, mainly on court-ordered clients – identifies a further difficulty, namely the over-representation of minority groups among those compelled to engage with services. This outcome is an aspect of a broader problem that must be addressed in a thorough analysis of the Vulnerable Children PRM.⁵¹ Given the contingent overrepresentation of minority groups among those requiring support from New Zealand's benefit systems, it is inevitable that the same over-representation of disadvantaged and minority groups will be found among the individuals and families identified as at high risk by the Vulnerable Children PRM. The over-representation of minorities in processes in connection with social problems such as child maltreatment is likely to contribute to the stigmatisation of such groups. In addition, in the context of mandatory engagement, it has been suggested that the over-representation of such groups threatens the efficacy of mandatory engagement since 'Cultural factors inherent in this clientele are often not well

⁵⁰ In an illustration of just how difficult these issues may be, Liam's case had a tragic influence on a related case, deterring a DHB from seeking court order to compel the parents of 13-year-old Tovia Lafau to bring him for treatment. It seems likely that Tovia's parents would have cooperated with such an order, but the hospital cited Liam's case as a reason for not seeking one. Tovia died of his cancer and his parents were convicted of failing to provide the necessities of life. No charges were brought against Liam's parents (Dare, 2009).

⁵¹ The MSD is sensitive to the cultural ramifications of these issues. It was initially proposed that they be addressed in this report, but the author advised that he did not have the necessary expertise or standing among the appropriate communities. Consequently the MSD has commissioned a separate report from more appropriately positioned and qualified reviewers.

understood by ... therapists, and this may cause strain in the therapist-client relationship”⁵² and to “carry social power imbalances into the helping relationship”⁵³ which may serve to undermine such relationships.

e) General Preconditions for Ethical Screening

The standard treatment of health screening defines it as:

[T]he presumptive identification of unrecognized disease or defect by the application of tests, examinations, or other procedures which can be applied rapidly. Screening tests sort out apparently well persons who probably have a disease from those who probably do not. A screening test is not intended to be diagnostic. Persons with positive or suspicious findings must be referred to their physicians for diagnosis and necessary treatment.⁵⁴

It is widely recognized that despite its capacity to deliver considerable benefits, allowing efficient allocation of resources and early diagnosis and intervention, health screening also carries some costs, including those associated with the burdens of compliance, over-diagnosis, misdiagnosis, and the creation of anxiety or unwarranted confidence. Such concerns led the World Health Organization to publish ten prerequisites, proposed by Wilson and Jungner, which should be met by any ethical screening program. Similarities between screening and PRM have recently led Lewis et al to suggest that equivalent caveats should apply to a predictive risk modelling approach to stratifying a population according to their risk.⁵⁵ The WHO Screening Principles and suggested PRM adaptation are set out in in **Table 1** below.

⁵² Synder and Anderson, 2009, p. 279.

⁵³ O'Hare, 1996.

⁵⁴ Wilson et al., 1968.

⁵⁵ Lewis et al., 2013.

Table 1

	The WHO Principles	The PRM Adaptation
1	The condition should be an important health problem.	The adverse event should be an important problem. ⁵⁶
2	There should be a treatment for the condition	There should be an intervention that can mitigate the risk of the event.
3	Facilities for diagnosis and treatment should be available.	Resources and systems should be available for timely risk stratification and for preventive interventions.
4	There should be a latent stage of the disease.	There should be sufficient time for intervention between stratification and the occurrence of the adverse event.
5	There should be a test or examination for the condition.	There should be a sufficiently accurate predictive risk model for the adverse event.
6	The test should be acceptable to the population.	The predictive risk model should be acceptable to the population.
7	The natural history of the disease should be adequately understood.	The natural history of maltreatment should be adequately understood by the organisation offering the preventive intervention.
8	There should be an agreed policy on whom to treat	There should be an agreed policy on whom to offer the preventive intervention.
9	The total cost of finding a case should be economically balanced in relation to medical expenditure as a whole.	The total cost of risk stratification should be economically balanced in relation to child protection expenditure as a whole.
10	Case-finding should be a continuous process, not just a "once and for all" project.	Stratification should be a continuous process, not just a "once and for all" project.

The WHO Principles and the proposed PRM adaptation provide a useful starting point for the ethical evaluation of PRM considered as a screening programme, indicating areas in which PRM ‘performs well’, areas in which further work is required, and areas which may pose persisting ethical challenges. Some of the prerequisites in table 1 appear to be comfortably met:

- Child maltreatment is an important problem (Principle 1);
- Early identification of elevated risks of maltreatment provides opportunity for intervention (Principle 4);

⁵⁶ Lewis et al., retain the term 'health problem', rather than simply 'problem' in their adaptation, but the adaptation seems intended to extend PRM beyond health contexts.

- The Vulnerable Children PRM has considerable predictive power (Principle 5);
- There is a good economic case for the application of PRM to child maltreatment (Principle 9);
- The model provides a dynamic assessment of individuals, updating with changes to information held on the supplying databases, and is proposed as an continuous process at the population level (Principle 10).
- Although the vulnerable children PRM does not purport to identify or rely upon the causes – or the natural history – of child maltreatment, the model does show the effects of a large range of variables on the probability of maltreatment. Further, the list of variables considered for inclusion in the model was selected on the basis of evidence in the literature about risk factors for maltreatment (Principle 7).
- Although a decision has not yet been made concerning to whom intervention would be offered, the Vulnerable Children Project recommends that intensive intervention be offered to those identified as at high risk and hence gives grounds for an informed agreement (Principle 8).
- With respect to Principle 2, there is some evidence for the effectiveness of existing interventions,⁵⁷ and further trials are being considered for two sites in New Zealand. There are likely to be outstanding concerns about the efficacy of interventions prompted by the model, given factors such as the difficulty of engaging with some high-risk families and the recognized complexity and difficulty of the child maltreatment problem. I discuss these concerns in more detail in section 2.e.i. below.
- The Vulnerable Children PRM may provide a positive contribution to concerns about the availability of resources and systems under Principle 3, insofar as it provides information that could allow services to be targeted to children at high risk. Here again, though, there are some outstanding

⁵⁷ Vaithianathan, et al., 2012.

concerns about the ability of social services to respond to identified risks. Again, see section 2.e.i. below.

- As indicated elsewhere in the report, particularly in the discussion of stigmatisation, I think there will be some persisting public concern about the application of PRM to child maltreatment (Principle 6). While it is unlikely that these concerns will be completely removed, I hope to have suggested adequate responses to them.

The WHO principles and hence the PRM adaptation of them should be regarded as necessary rather than sufficient conditions for ethically acceptable screening programmes: ethically acceptable programmes must meet those conditions, but meeting them will not suffice to show that a programme is ethical. There are other ethical constraints. Consider the controversy over prostate cancer screening. Prostate cancer is the most frequent malignant tumour in men in the western world. In 2009, there were 3369 new cases and 562 deaths in New Zealand. The five-year relative survival among men with cancer confined to the prostate or with only regional spread is 100% compared with 31.9% among those diagnosed with widespread metastases.⁵⁸ Advanced stage disease is generally not curable. A screening programme that could identify asymptomatic men with aggressive localized tumours, then, could substantially reduce prostate cancer morbidity, including urinary obstruction and painful metastases, and mortality. Prostate-specific antigen (PSA) testing seemed to be just such a programme and it became widely adopted for cancer screening by the early 1990s. A director of the most important study of prostate cancer screening, the European Randomised Study of Screening for Prostate Cancer (ERSPC),⁵⁹ has recently written that the programme

⁵⁸ Hoffman, 2013

⁵⁹ The European Randomised Study of Screening for Prostate Cancer (ERSPC) is the world's largest prostate cancer screening study and involves eight countries – Belgium, Finland, France, Italy, Netherlands, Spain, Sweden and Switzerland. Participants in the randomised study totalled 182,000, of which 162,000 men contribute to the core age group 55–69. Men randomized to the group being offered screening were tested using the prostate specific antigen (PSA) marker, every two or four years with an average follow-up of 11 years.

took the WHO Principles as its 'benchmark'. "PSA based screening followed by prostate biopsies", wrote Chris Bangma,

... was the most appropriate method to find early prostate cancer at a curable stage. Curative treatments like radiotherapy and radical surgery were available, and later on brachytherapy was added. Technically, radical surgery had improved, though side effects like stress incontinence and erectile dysfunction remained considerable. Also, radiotherapy had its acute and chronic side effects. PSA screening appeared to be cost effective, non-invasive, and judged a good method for mass screening.⁶⁰

However it is now widely accepted that population wide PSA screening raises serious ethical concerns. The ERSPC has issued positive reports – the most recent concludes that "a man who undergoes PSA testing will have his risk of dying from prostate cancer reduced by 29%"⁶¹ – but the authors acknowledge that PSA screening leads to very significant over-diagnosis, the diagnosis of cancers that do not pose any threat to the patient because they are slow growing or 'indolent'. Approximately 30% of detected cancers are unlikely to progress and cause a patient's death. The United States Preventive Services Task Force analysed the ERSPC data and estimated that, for every 1,000 men aged 55 to 69 years who were screened every 1 to 4 years for a decade: 100 to 120 men would have a false positive test leading to a biopsy; about one-third of the men who had a biopsy would experience at least moderately unpleasant symptoms from the biopsy; 110 men would be diagnosed with prostate cancer; about 50 of these 110 men would have a complication from treatment, ranging from erectile dysfunction and urinary incontinence, to serious cardiovascular events and death; and 0 to 1 deaths from prostate cancer would be avoided.⁶² On balance the costs of PSA screening are now widely thought to outweigh its benefits and many organizations, including the Urological Society of Australia and New Zealand, advise against routine population screening.⁶³

⁶⁰ Bangma, 2012.

⁶¹ Schröder et al. 2012

⁶² U.S. Preventive Services Task Force. Screening for Prostate Cancer.

⁶³ Urological Society of Australia and New Zealand, PSA Testing Policy (2009).

The PSA case appears to identify additional prerequisites for ethical screening programmes, at least some of which are relevant to the Vulnerable Children PRM. We have seen that the Vulnerable Children PRM will generate false positives. Referring all cases in the top two deciles would refer 5% of the total population to the intervention, with a false positive rate of 63%, and capture rate of 37% of national maltreatment cases. However the PRM false positive rate is considerably lower than that for the PSA test, which is in the order of 75%: only 25% of men who meet the standard cut off for biopsy (4.0 ng/mL) are found to have prostate cancer.⁶⁴ We can express these rates in terms of the respective programmes' areas under the ROC curve. The Vulnerable Children PRM, we have seen, has an area of 76%, considered 'fair' approaching 'good'. The PSA test, by contrast has an area of 67% (considered 'poor' approaching 'fair') for discriminating between cancer and no cancer.⁶⁵

Equally significantly, we have seen that the interventions triggered by the positive PSA results were extremely burdensome: approximately one third of men who received a false positive result in the ERPSC study experienced moderately unpleasant symptoms from the biopsy. There is no reason to suppose that the likely interventions in the PRM case threaten harms of comparable magnitude. We have discussed ways in which the harms of such an intervention might be ameliorated elsewhere in this report (by making it voluntary, and part of a generalized offering of service, for instance), and that discussion will be relevant here. Indeed it is possible that appropriate interventions may benefit even families who would not have gone on to produce maltreatment findings.⁶⁶ Not only will some percentage of false positives 'by the age of 5', go on to have maltreatment findings by the age of 7, but identification as at high-risk indicates that families are

⁶⁴ Thompson, et al. 2005.

⁶⁵ Ibid.

⁶⁶ "Note also that a 'false positive' does not indicate that no needs are present. For example, a family with employment and alcohol issues may rightly be considered not appropriate for a statutory child protection response since the children are safe and loved. The decision to undertake a statutory investigation was 'false' in relation to the need for statutory response but not false in the sense that there are no needs or that no other form of support would be welcome to achieve a better outcome for this family" (Mansell, et al., 2011, p.2080).

in need. That is, if the Vulnerable Children PRM had a wider range of outcomes included as adverse events that the sorts of interventions we contemplate were expected to reduce (such as teen pregnancy, truancy or criminal activity) then the prognostic strength of the PRM against this wider set of events could potentially be much higher. Intensive interventions designed to address those needs may both reduce child maltreatment among true positives and benefit false positives.

Finally, the application of principles developed for screening to PRM serves as an important reminder of the distinction between identification of risk status, confirmation or diagnosis, and the response: “A screening test is not intended to be diagnostic”. The diagnostic step in the case of PRM will look rather different from its analogues in many, but not all, other screening contexts. In the PSA case for instance, diagnosis consists of a biopsy that confirms or disconfirms whether a patient has cancer. There will be no such discrete diagnostic step in the PRM case, since the ‘finding’ will always be a future event. However, screening for cardiovascular risk has a similar profile (screening picks up indicators of a risk of a future event, not an existing condition), and social service professionals or other frontline staff will be responsible for fine-tuning the assessment of risk and subsequent response to children identified by the PRM tool, steps which could be seen as analogous to the further diagnostic step of a biopsy.

i. Principle 2: The obligation to provide effective intervention

The availability of an effective intervention or response to risks identified by screening programmes is widely seen as crucial to the ethical status of such programmes. The requirement appears as Principle 2 of the WHO Screening Principles (described by Wilson and Jungner as “perhaps the most important” of the principles⁶⁷) and its PRM analogue; the idea that, in order for screening to be ethical there should ‘a treatment for the condition’, or ‘an intervention that can mitigate the risk of the event’.

⁶⁷ Wilson et al., 1968, p.27.

Although the point has not as far as I know been made elsewhere, it is important to see that Principle 2 may be motivated by two rather different concerns, concerns which carry significantly different ethical implications.

It might seem that without an effective intervention screening may be unethical for the same reason that it would be unethical to know that someone needed help and to refuse to come to their aid. Some such intuition seems to lie behind that idea that “[o]nce it has been determined that a child is at risk of maltreatment, there is an ethical imperative to respond with an appropriate intervention that is likely to reduce that risk.”⁶⁸ The idea here seems to be that a duty to render assistance is triggered or crystallised by *cognisance* of an existing need.

It is not easy to see how this justification of Principle 2 gives a reason against screening. Reading it as doing so seems to assume that it might be better, from an ethical perspective, not to know about genuine but unmeetable needs. It seems likely that such a principle could only be justified by a concern for the moral well-being of those who would intervene if only they could do so: people in need are not themselves made worse off by knowing others know of their plight but cannot help.⁶⁹ One might make the abstract moral point to impotent would-be rescuers that they cannot be morally obliged to do something they cannot do: ‘ought implies can’ as Kant famously argued⁷⁰ – and hence that their *moral* position is not changed by discovering an unmeetable need, though of course such a discovery might be psychologically troubling. More importantly and practically, it seems unlikely we are better off, as individuals or communities, not knowing about unmeetable needs. We may be able to meet them in the future, and knowing of and about them may provide an incentive to improve our capacity to help and information about how to do so.

⁶⁸ Mackay, Ross, 2013.

⁶⁹ The situation is made more complicated if, as in the child maltreatment case, the person in need is a very young child: it is harder to determine what they know and what significance they give to that knowledge.

⁷⁰ “The action to which the ‘*ought*’ applies must indeed be possible under natural conditions” Kant, 1781, p. 473.

More plausibly, the availability of an effective intervention is a prerequisite to an ethical screening programme because screened populations and individuals should not be made worse off by screening. Screening programmes might make people worse off because the programme itself is burdensome, because the only available interventions are burdensome, or because it delivers information that is burdensome. Here the screening process or some limited set of responses to it, are themselves seen as potentially harmful, and hence as requiring some countervailing benefit. It would be unethical, the idea goes, to impose the potential harms of screening unless doing so enabled the delivery of compensating benefits; unless, in other words, screening was worthwhile given the burdens it imposed. Most of the concerns about the implications of Principle 2 for prostate screening programmes appear to flow from this understanding of the basis of the Principle.

My tentative conclusion is that Principle 2 is best seen as resting on ethical justifications which recognise that screening programmes which might themselves generate harms must be capable of delivering countervailing benefits. This conclusion has practical implications for the ethical evaluation of predictive risk modelling of child maltreatment, directing inquiry, on the one hand, to the potential harms of such modelling – harms such as stigmatisation, of unavoidable false positives and false negatives, of intrusion upon the privacy of ‘screened’ families and individuals, and the like, identified and discussed elsewhere in this report – and, on the other, to potential countervailing benefits.

This conclusion supports the rejection of a reading of Principle 2 according to which it poses a simple ‘yes/no’ test: it is not that screening can never be justified unless there is ‘cure’ for the condition screened for, or if the risk of the event screened for can be eliminated. The question is rather, ‘is there an intervention capable of mitigating the risk of the event screened for and delivering benefits to screened populations that outweigh the burdens of screening?’ Such an approach seems consistent with Wilson & Jungner’s original formulation of Principle 2, which specified that “[f]or declared disease there is, of course, the ethical obligation to provide an *accepted treatment whether or not this is of scientifically*

proved value”, and their statement of the conclusion of the discussion of the principle, in which they wrote “[i]t is axiomatic, therefore, that case-finding should only be undertaken when the prospects for treating the condition *are at least reasonable*”.⁷¹ Wilson and Jungner seem not to have seen Principle 2 as requiring absolute certainty of completely effective interventions.

Reading Principle 2 in light of the justifications of this sort also suggests, I propose, that we should take what we might call an encompassing view of the benefits and burdens of screening programmes. While one should not justify screening programmes by reference to benefits that are not related to the stated aim of the programme, we should not take too narrow a view of those benefits. I have remarked, for instance, that it seems likely that appropriate interventions may benefit even families who would not have gone on to produce maltreatment findings. Evidence of the efficacy of available child maltreatment interventions is *arguably* sufficient to satisfy Principle 2 so read. A handful of programmes have shown some evidence of capacity to reduce child maltreatment, but there are grounds for caution in each case:

- I have mentioned the *Nurse Family Partnership*, a U.S. home visiting programme for first time, low-income mothers, which has been shown to have reduced substantiated reports of abuse and neglect, but while the programme was evaluated in three pilot demonstration sites and has replicated many positive effects on children’s development, it was possible to test its impact on substantiated maltreatment reports in one site only.⁷²
- *Early Start*, a New Zealand home visiting programme, reduced parents’ reports of harsh or abusive punishment and parents’ reports of contact with child protection services operated by Child Youth and Family (CYF),

⁷¹ Wilson et al., 1968, pp28 and 29, emphasis added.

⁷² Olds et al., 1997. Despite these concerns about evidence for its effectiveness, note that Mikton and Butchart, 2009 identify the Nurse Family Partnership in the USA “as the only home visiting programme whose effectiveness has been unambiguously demonstrated. A randomized controlled trial showed a 48% reduction in actual child abuse at 15-year follow-up” (Mikton and Butchart, 2009, 354). The Nurse Family Partnership programme is the programme used to calculate possible effects of the Vulnerable Children PRM.

but there were no differences between the treatment and control groups on CYF-reported contacts.⁷³

- *Triple P*, a multi-tiered family intervention system, trialled at a population-level in a selection of US counties, found increased levels of maltreatment in control counties, compared with no change in counties where the treatment was implemented,⁷⁴ but further evaluation of *Triple P* is needed on higher-risk populations and some aspects of the analysis remain unclear.⁷⁵
- A coordinated, hospital-based, parent education programme in New York state targeting parents of new-born infants and aimed at reducing violent infant shaking, did show reductions in the incidence of abusive head injuries among infants and children aged under 3 years, but did not use a randomised controlled trial design leaving open the possibility that confounding variables had an effect on the outcome.⁷⁶
- The *Safe Environment for Every Kid* (SEEK) model of paediatric primary care, implemented in Baltimore MD, reduced rates of care and protection services reports of abuse or neglect and parent-reported harsh punishment, but was implemented in a single clinical practice on a relatively small sample, with high rates of refusal to participate and non-completion of the laboratory protocol.⁷⁷
- The *Chicago Child-Parent Centers* (CPCs) programme reduced child abuse and neglect between ages 4 and 17, based on the children's subsequent reports at ages 20-21, but the evaluation of the CPCs did not use a randomised controlled trial design, so may have been subject to confounding variables.⁷⁸

⁷³ Fergusson, Horwood & Boden, 2012.

⁷⁴ Prinz *et al.*, 2009.

⁷⁵ MacMillan *et al.*, 2009.

⁷⁶ Dias *et al.*, 2005.

⁷⁷ Dubowitz *et al.*, 2009.

⁷⁸ Reynolds *et al.*, 2002.

While meta-analyses emphasise the ‘sparse and mixed nature of the evidence’⁷⁹ that such interventions reduce the “direct measures”⁸⁰ of child maltreatment, there is greater confidence that they do deliver benefits which contribute indirectly to that goal. Mikton & Butchart conclude that meta-reviews of intervention programmes “suggest that early home visitation programmes are effective in reducing risk factors for child maltreatment, but whether they reduce direct measures is less clear-cut”,⁸¹ and Howard & Brooks-Gunn (2009) that “the evidence ... is stronger with respect to parenting and the quality of the home environment [than with respect to maltreatment]” and that “home visits impart positive benefits to families by way of influencing maternal parenting practices, the quality of the child’s environment and children’s development”. Even if intervention with high-risk families cannot be proven to *directly* reduce the risk of maltreatment, then, there does seem reason to think that appropriate interventions may be effective against a broader range of conditions plausibly related to maltreatment risk. The World Health Organisation notes, for instance, that there is:

... some strong evidence to show that programmes that promote safe, stable and nurturing relationships between parents (or caregivers) and children reduce child maltreatment and its life-long negative consequences for mental and physical health, social and occupational functioning, human capital and security and, ultimately, for economic development.⁸²

If this is right, there seems reason to think that the Vulnerable Children PRM meets the requirements of WHO’s ethical screening Principle 2 and its PRM analogue.

⁷⁹ I take this description and the summaries of interventions and concerns about the evidence for their effectiveness from Mackay, 2013.

⁸⁰ Mikton and Butchart, 2009, 354

⁸¹ Ibid.

⁸² World Health Organisation, 2009, p.10

f) Resource Allocation

In early discussion of the Vulnerable Children PRM it was suggested that the tool could have ethically significant resource allocation implications. There seem to be at least two rather different threads to these concerns.

i. Targeted vs Universal Services

Some ethical concerns about the Vulnerable Children PRM are motivated by the worry that the model will give priority to targeted interventions to high risk families at the expense of lower risk families, and that such targeting will, over all, produce worse child maltreatment outcomes. Such a concern might be grounded in the belief that a universal allocation of child maltreatment resources would reduce the incidence of maltreatment at all levels (perhaps because intervening with lower risk cases prevents them becoming higher risk), or in the belief that while very serious cases of maltreatment – most obviously those involving child deaths – are tragic, they are relatively rare and hence that it is a mistake to respond to the undoubted public pressure such cases generate in ways which reduce capacity to respond to less serious but still significant harms suffered by children at lower risk.

We have seen above, however, that universal programmes carry risks of their own, including that of failing to identify and respond to the specific needs of high-risk children. Beyond this, these concerns about the possibility that the Vulnerable Children PRM may lead to an harmful focus on some children to the detriment of others, would constitute a significant challenge to the Vulnerable Children PRM and the commitment to giving it an important place in child maltreatment initiatives if and to the extent that doing so really did lead to a significant reduction in response to lower risk children, and if doing so did lead to missed opportunities to intervene at lower risk levels in ways which reduced transition to higher risk. It is important that such concerns be taken into account in the implementation of the Vulnerable Children PRM, which should be considered in the context of overall responses to vulnerable children. It is important, for

instance, that the implementation of the tool does not entail a reduction in existing universal efforts to promote child wellbeing. The increased information provided by the Vulnerable Children PRM may allow more effective responses to higher risk cases, perhaps freeing resources for allocation to other child maltreatment efforts. The concerns which motivate these objections to the Vulnerable Children PRM can be read as good reasons to think there is an ethical obligation to ensure in so far as possible that the programme does deliver improved overall outcomes for child maltreatment.

ii. Workload Capacity

Many writers have commented on the way in which child protection services in Western countries have swung between a range of policy responses, sometimes emphasising prevention, sometimes detection, sometimes risk mitigation.⁸³ Often, these writers suggest, focusing on identifying or encouraging reports of maltreatment actually reduces capacity to respond to vulnerable children. Discussing the side effects of a UK initiative requiring domestic violence incidents to be reported to child protection services, Humphreys and Stanley write that

... a blanket referral policy in which most domestic violence cases involving children are referred to the statutory agency may in fact increase the danger to children, as the response is flattened out to deal with volume rather than seriousness. Children at real risk of significant harm may then be lost among the 'debris of referrals'.⁸⁴

On one reading, this appears to be a version of a problem we saw in our discussion of screening. Ethical screening programmes, we saw, required interventions capable of treating or mitigating the risk of identified states of affairs. There seem at least two plausible motivations for this prerequisite in the screening case. First, we might think that people are made worse off by being told they have an untreatable condition or high risk of an adverse event, than they were when they simply had the condition or risk. Second, screening programmes that identify

⁸³ Mansell, et al., (2011).

⁸⁴ Humphreys, et al., 2006, p.43. Cited Peckworth, 2013. And see Mansell, 2011, p.2076: "Child protection services have become so 'forensic' that capacity to provide adequate services is diminished and demand then overwhelms caring services".

conditions or risks that cannot be treated or mitigated might be an unethical waste of resources. The significance of the first motivation in the child maltreatment context is unclear, in part because it is unclear what information would be passed to families and in part because it is harder to accept that risk of maltreatment is, like prostate cancer perhaps, a condition that it is occasionally better to leave unaddressed. The second motivation does seem relevant, and makes clear that social service authorities have an obligation, having identified risk, to respond appropriately, to have adequate triaging systems, to provide adequate resourcing to meet predictable need.

We also need to bear in mind, again, that problems caused by increasing detection of maltreatment is not specifically an issue for the Vulnerable Children PRM. As I said when I made this general point at the outset, that is not to say that it is not a genuine ethical problem. It does show, however, that the ethical issue is essentially a comparative one. The question is not simply 'can social service authorities respond appropriately to identified risk', but 'would the Vulnerable Children PRM improve or worsen their capacity to do so relative to existing alternatives?', and arguably, insofar as it allows more precise and efficient targeting and use of resources, the Vulnerable Children PRM might help rather than hinder here. While efforts to identify or encourage reports of maltreatment (either through blanket referral policies such as those noted in the passage quoted immediately above, or through mandatory reporting requirements) generate high volumes of notifications to child protection agencies, many of these reports concern children who are not in fact at risk of maltreatment. There is reason to think that the Vulnerable Children PRM would not generate a similar mix of referrals, since it refers only where there is an actuarially determined connection to the likelihood of future maltreatment.

g) Privacy

i. The Distinction Between Privacy And Confidentiality

The Vulnerable Children PRM raises a number of issues concerned with the privacy of risk scored individuals and families and with the confidentiality of their information. Although privacy and confidentiality are interrelated, it is useful to maintain a distinction between them since they pick out two rather different sets of concerns relevant to the current discussion. Confidentiality is concerned with data or information and its security. My doctor is required to keep my health information confidential: she may not pass it on to others without my consent. Privacy need not involve information or the threat that it might be misused. The person who peers into my bedroom invades my privacy, and I need take no comfort from the fact that they will never speak about what they have seen. Unfortunately, for terminological neatness at least, confidentiality is sometimes called ‘informational privacy’. New Zealand’s Privacy Act 1993, on this account, is primarily concerned with ‘confidentiality’ or ‘information privacy’, rather than with the moral interest in privacy as a concern with the protection of a sphere of private conduct. Nonetheless, the distinction is important because measures to protect and determine the limits of confidentiality may not address invasions of privacy. I discuss the two types of concern in the following subsections.

ii. The Moral Right to Privacy and the Vulnerable Children PRM

The moral right to privacy – understood here as the value that is compromised by trespasses into the private or intimate spheres of a person’s life – has been defended on a number of grounds. Defenders have argued that it is essential to human dignity,⁸⁵ crucial to intimacy⁸⁶ or the development of meaningful

⁸⁵ Bloustein, 1964

⁸⁶ Gerstein 1978; Inness, 1992

interpersonal relationships,⁸⁷ as a right giving us the capacity to control the access others have to us,⁸⁸ or as necessary to personal expression and choice.⁸⁹

Grant for the moment that people do have a moral right to privacy in this sense; that they may be wronged by trespasses into their private affairs or the intimate spheres of their lives, even though trespassers receive no ‘information’ and threaten no breach of confidentiality. If this is so, the Vulnerable Children PRM may pose ethical risks not remedied by, for instance, controlling the dissemination of information or ensuring that interventions pose as little burden as possible. The breach of such rights consists not in the collection or misuse of information or the imposition of cost but in the trespass itself.

It is possible that such rights might also ground a perception that risk scored individuals or groups have been discriminated against or treated unequally. Insofar as the Vulnerable Children PRM enables a government department to generate a risk score, based on a large set of demographic and historical features of a child, their family, household and community, risk scored families might plausibly feel that their privacy rights have been given less respect or weight than those of others. If those who defend privacy rights by appeal to their significance to human dignity are correct, such invasions might be of very considerable significance.

Nonetheless it is important to note that there is reason to be wary of the potential of moral rights to privacy to shield (often domestic) domination, degradation, and abuse. The concern is raised most clearly in a feminist critique of privacy, which recognises that drawing a strong distinction between public and private spheres is not without risk. “What puzzles feminists”, writes one commentator:

“...is how to make sense of an important and valuable notion of privacy that provides them a realm free from scrutiny and intervention by the state, without reverting to the traditional public/private dichotomy that has in the past relegated women to the private and domestic sphere where they are

⁸⁷ Fried, 1970; Rachels 1975

⁸⁸ Gavison, 1980; Allen, 1988; Moore 2003.

⁸⁹ Schoeman, 1992.

victims of abuse and subjection. The challenge is to find a way for the state to take very seriously the domestic abuse that used to be allowed in the name of privacy, while also preventing the state from insinuating itself into all the most intimate parts of women's lives."⁹⁰

The feminist challenge is particularly notable in the context of a discussion of child maltreatment since the Vulnerable Children PRM is intended to contribute to a reduction in domestic maltreatment. The challenge draws attention to the need, noted early in this review, to balance the importance of moral rights to privacy against other important interest, some of them rights (such as a right to freedom from abuse or neglect), some of them consequentialist (such as a concern for the long term impact of maltreatment, and the efficient use of scarce resources). Acknowledging that there are important privacy rights in this moral sense, does not settle whether predictive risk modelling – even if it is conceded to involve some intrusion into privacy rights – is or is not ethical.

Some feminists argued for the rejection of privacy rights and the distinction between public and private spheres altogether,⁹¹ but we need not go so far. Legitimate *public* interest in the protection of the vulnerable provides a principled way of specifying limits to the moral right to privacy. Monitoring and even interference may be justified when necessary to protect the vulnerable, harm to whom *should* be a matter of public concern. We should have proper regard to the interests that ground the moral importance of privacy, but those interests must be balanced against important countervailing concerns. Ruth Gavison makes the point nicely. “The interdependence of public and private”, she writes:

... provides another argument against the presumptive entitlement of noninterference in the family realm. Because private family arrangements may well be among the most consequential factors in individuals' development and may largely determine their options in the public world, public attempts to optimize these arrangements may be desirable. As usual, these reasons may be balanced against other constraints, including the wish to protect, to some extent, the privacy of families, and the limited effectiveness of public regulation in this area. But these are very different from "jurisdictional" objections to interference with private lives.⁹²

⁹⁰ DeCew, 2012

⁹¹ MacKinnon, C., 1989

⁹² Gavison, R., 1992-1993, p.24

Privacy interests do not, in my view, exclude a carefully managed predictive risk modeling programme, though they do provide further reason to identify and minimize the possible harms and intrusion posed by such a programme. Concerns about the moral right to privacy I think, can be met by the steps proposed elsewhere in this review to minimize the moral risks of using predictive risk modelling in the child maltreatment area.

iii. Confidentiality or Information Privacy and the Vulnerable Children PRM

I have suggested that whereas privacy is concerned with interests in protecting a sphere of private conduct, confidentiality is concerned with the use of information. Breaches of confidentiality (as noted, sometimes called information privacy) always involve the disclosure of information. As such they can be addressed by controlling such transmission. The Vulnerable Children PRM raises clear confidentiality issues insofar as it generates predictions of risk of child maltreatment by drawing from databases which hold information (some of which is gathered for purposes other than child protection) and which will be accessed without the consent or knowledge of those who supplied that information or the families to which it pertains. Under what circumstances might we think it legitimate to disclose that information in the manner contemplated by the Vulnerable Children PRM?

The first question is just what sort of disclosure *is* being contemplated? One might think that the more detail which is disclosed and the wider the disclosure the greater the confidentiality issues. I have already suggested that disclosure of information generated by the Vulnerable Children PRM should be as narrow as possible, consistently with achieving the important moral goals of the tool.

A similar approach might be taken toward the *content* of the information disclosed. What information about a family might senior social service officials receive from the tool? At one end of a range of possibilities, they might be provided with full details of a family's circumstances, including the risk ranking, and if the tool is able to provide it, details of changes in circumstances that have

triggered any changes in a family's risk rating. Somewhere toward the middle of the range of possibilities, we might imagine officials being provided with a list of names and risk scores. Toward the other end, we might imagine social service officials receiving *only* a list of names of those who appear as at the highest risk, without further information from the tool. Under this latter approach, social service professionals would be called upon to initiate an assessment of the circumstances of the family, in order to determine what action should be taken. While such an assessment might be triggered by a PRM risk score, the assessment and response could be identical to those carried out within current social work practice.

I remarked at the outset that consideration of the moral status of efforts to reduce child maltreatment required us to address and balance a number of different moral perspectives and interests, and this seems, again, to be a point where that balancing comes to the fore. We should, in my view, prefer the narrowest disclosure of content necessary to protect important interests of vulnerable children: if the final option sketched above achieves that goal, while intruding as little as possible upon the legitimate confidentiality interests of those whose information is accessed by the Vulnerable Children PRM, then that option should be preferred. Of course this is a contingent recommendation: it may be that without further information, about, for instance, a family's risk ranking, front-line professionals cannot offer effective services to families. The onus is on social service authorities to show why they require more rather than less information, and they should do so by showing why more information is essential to the provision of effective services.

Perhaps it will seem that this approach gives too little value to confidentiality: one might think that confidentiality interests are absolute in a way which counts against the sort of trade-offs I have in mind. In fact, however, almost all of the plausible accounts of confidentiality (or information privacy) allow for exceptions and limitations. It is possible to defend the moral value of confidentiality from a number of philosophical perspectives. Three leading alternatives are:

- A deontological, or rights-based, perspective that argues that respecting the conditions under which I make disclosures to you is a way of respecting my autonomy or moral agency. If I give you information about myself (my health, my preferences, my history, my domestic situation, etc) on the understanding (perhaps an understanding engendered by an undertaking on your part) that you will use that information only for certain purposes, then respect for my agency requires that you limit your use of that information accordingly.
- A consequentialist perspective that argues that a policy of respecting confidentiality has good consequences. The standard defence of confidentiality in professional contexts has a consequentialist structure. Physicians and lawyers, according to the standard argument, cannot provide high quality professional services to their clients or patients without full and frank disclosure from those clients or patients, who, it is supposed, would often be reluctant to be entirely candid if they thought their professionals might pass information on to others. The common professional obligations of confidentiality are justified, on this account, because having them in place makes it more likely that professional relationships will produce the benefits of competent and informed professional representation.
- A ‘mixed’ model, according to which professional rules requiring confidentiality function as (more or less) absolute constraints on professionals – so placing them under a duty to maintain confidentiality and granting to clients a correlative right that they do so – *because* placing professionals under such obligations has good consequences.⁹³

On any but the narrowest readings of the deontological approach (a reading which, arguably, even its most famous advocate, Immanuel Kant, did not endorse⁹⁴), there is room for exceptions to the moral duty of confidentiality.

⁹³ Luban 1989, Dare, 2009

⁹⁴ Mahon, 2009

A deontological approach, for instance, still allows one to ask just what the conditions of disclosure were in order to work out what is required by a duty to respect the autonomy of those who disclosed information: we have to know what an agent chose in order to respect their choice.

From here, in the present case, there may appear to be a straightforward justification for limited disclosure of at least a good deal of the information that might be provided to social service professionals under the Vulnerable Children PRM. Standard welfare benefit application forms include advice that information is being collected for the purpose of providing for the 'care and protection needs of children' and for 'providing support and services for you and your family'.⁹⁵ Arguably, then, using information collected in such forms in order to reduce or remove the risk of child maltreatment *is* consistent with the autonomous choices of those providing the information. There are, however, at least two problems with accepting that those who have completed benefit application forms containing such clauses have given consent for the information to be used for risk profiling purposes.

First, it must be recognised that those completing such forms may feel they have little real choice: they are seeking essential support for themselves and their children and are hardly in a position to haggle over the terms imposed by social service agencies. Given this, we might hesitate to regard their apparent acceptance of the terms on the forms as clearly autonomous. This concern can be countered, however, with the response that it is objectionably paternalist to disregard the apparent choice of those who supply information having been informed of such terms. And of course, many of the choices we all make are more or less compelled, in the sense that we choose from a range of options we would prefer to have been broader. Paradigmatically autonomous agents often choose reluctantly from among constrained alternatives, and their choices are no less autonomous for that.

⁹⁵ See for instance the Domestic Purposes Benefit Application:
<http://www.workandincome.govt.nz/documents/dpb-sole-parents-application.pdf>

Second, however, even if we accept that those filling out such forms have consented to have their information used to provide “support and services” and for the ‘care and protection needs of children’, it would not necessarily follow that they had consented to have their information used for the purposes of a predictive risk model. A referee of this report drew attention to this aspect of the issue when she wrote that

I would ask "do you think that the individuals ... understand that their information could be used to classify them as high risk for child maltreatment?" You would need to be able to answer "yes" to this question for their consent to the use of their personal information to be considered informed.

It is a good idea to frame the issue in terms of such a hypothetical question, though, with respect to the referee, I do not think her version of the question is quite right. It is correct to think that the clauses on the standard application forms would indicate consent to inclusion in the predictive risk model only if those filling out the forms knew what that meant, but putting them in that position would involve a little more than telling them their information ‘could be used to classify them as high risk for child maltreatment’. Rather it would require a description of the tool, of the possibility that it might show the applicant to be at high-risk, and an explanation of what such a classification might lead to. The referee thought we would be unlikely to answer ‘yes’ to her question, that it was unlikely, that is, that families completing the current forms would appreciate that the information they provided could be used to classify them as high risk for child maltreatment. I think she is probably right. It is less clear to me, however, we could not answer ‘yes’ to my fuller version, treating it not merely as a matter of whether families *understood* that their information could be used to generate a child maltreatment risk score, but that they consented to that use on the understanding that risk scores would be distributed narrowly, that interventions prompted by a high-risk classification would be optional, supportive, and non-punitive.

Nonetheless, I am not convinced that the ‘Privacy Statements’ in current social welfare and similar documentation can safely be relied upon to evidence informed consent. In addition to concerns about whether the apparent acceptance of the

terms contained in such clauses, evidenced when a welfare support applicant completes and signs the forms, can safely be regarded as informed and voluntary, there is the further difficulty that the Vulnerable Children PRM is likely to draw on further sources of information – such as birth records – in contexts which are likely to generate equally or more difficult concerns about the quality of consent. New Zealand Parents are required to register births, for instance, and it is difficult to see how one could regard such compulsory registration as evidence of the voluntary acceptance of ‘use of data’ conditions. Further, it might seem that the existing clauses in the standard benefit forms allow a broader use of the information collected there than I have earlier suggested would be appropriate for child maltreatment predictive risk modelling. I suggested that we should, for instance, prefer the narrowest dissemination and disclosure of content necessary to protect important interests of vulnerable children. It might seem that clauses which contemplate the use of information for the ‘care and protection needs of children’ and for ‘providing support and services for you and your family’ license a more expansive use.

My tentative view, then, is that it is not appropriate to place too much weight on the ‘Privacy Statements’ in current social welfare and similar documentation. This is not to say that one cannot mount a deontological defence of the use of information gathered through such forms to inform predictive risk modelling to address child maltreatment. As noted at the outset, there are conflicts between significant rights at issue in the context of child maltreatment. The New Zealand Privacy Act itself recognises such conflicts, and the possibility that privacy rights might give way to other important interests, when it provides exceptions to allow the Privacy Commissioner to authorise an agency to collect, use, or disclose personal information in ways would otherwise be in breach of the privacy protection principles of the Act, if satisfied that “the public interest ... outweighs, to a substantial degree, any interference with the privacy of the individual that could result from that collection or, as the case requires, that use or that

disclosure".⁹⁶ Of course one cannot simply assume that the Commissioner would approve the use of information gathered in the various data bases for the purposes of the predictive risk modelling, but the existence of such exceptions shows an appreciation that there may be significant conflict between privacy rights and other important interests and the possibility that privacy rights might give way in the face of such conflicts. One might argue that the rights of vulnerable children should take priority over privacy rights, particularly if the use of information is appropriately constrained in ways suggested elsewhere in this review.

The availability of exceptions under consequentialist and mixed justifications is more straightforward and familiar. The confidentiality clauses of the Australian and New Zealand Association of Social Workers, for instance, provides that '[c]lient integrity is preserved by maintaining client confidentiality ...' before allowing disclosure in emergencies "in which it is in the client's best interests", and "where the client or someone else (such as a child) may be endangered or harmed by non-disclosure".⁹⁷ The analogous clause in the Code of Ethics of the NZ Medical Association provides that:

Doctors should keep in confidence information derived from a patient, or from a colleague regarding a patient, and divulge it only with the permission of the patient except when the law requires otherwise, or in those unusual circumstances when it is clearly in the patient's best interest or there is an overriding public good.⁹⁸

See as well the New Zealand Medical Association's Position Statement on Patient Confidentiality and HIV/AIDS, which remarks that "[e]thically, decisions around [rare situations in which a patient with HIV poses risks to others] demand that the doctor maintains an extremely delicate balance between his or her clear and primary ethical duty to his/her patient to maintain confidentiality and the 'public good' to which a doctor also owes a duty of care" and which, having quoted the confidentiality clause above, goes on "[a]gain the issues are around what

⁹⁶ Privacy Act 1993, s.54

⁹⁷ *Code of Ethics of the Australian and New Zealand Association of Social Workers*. Sections 3.12 and 3.13.

⁹⁸ NZ Medical Association, Code of Ethics (2008).

constitutes an overriding public good. It can also be inferred from our code that each circumstance has to be taken individually and in its context. There can be no complete and rigid guideline".⁹⁹

The point here is not simply to draw attention to the institutional fact that these clauses provide exceptions to the duties of confidentiality they establish. It is rather to draw attention to the way in which they illustrate how broadly consequentialist reasoning provides exceptions to such clauses. The duties of confidentiality in these cases hold *subject to* emergencies "in which it is in the client's best interests", or "where the client or someone else (such as a child) may be endangered or harmed by non-disclosure", or "when it is clearly in the patient's best interest or there is an overriding public good", or to the 'delicate balance' between a clear and primary ethical duty to patient confidentiality and the 'public good'. In hard cases, at least, one simply cannot apply such clauses without weighing the duty of confidentiality against the patient's interests or the public good. They are intended to be read, that is to say, in ways that allow the sort of limited balancing approach to confidentiality recommended above in respect of the Vulnerable Children PRM.¹⁰⁰

h) Effect on Social Service Professionals and other Frontline Staff

Concerns have been raised that the Vulnerable Children PRM might reduce engagement between social service professionals and their clients, leading such professionals to trust computers rather than their own judgment and marginalising professional judgement. However, the Vulnerable Children PRM may be used to supplement rather than replace frontline staff. It does not entail the replacement of social service professionals or shifting them off the frontline. It may have the opposite effect. Many risk assessment tools rely on social service

⁹⁹ NZ Medical Association Position Statement on Patient Confidentiality and HIV/AIDS (2008)

¹⁰⁰ A referee of this report suggested that the "evaluation of the value of confidentiality by considering the plausibly justified trade-off between confidentiality and the public good in emergency cases might seem a little forced. Presumably such gains as might be made by breaching confidentiality are those of child health and not the avoidance of emergency". While I take the point, I am inclined to think credible risks to identifiable children might provide sufficient warrant to override confidentiality.

professionals correctly identifying and reporting risk factors. Such "operator driven" tools are resource intensive and significantly threatened by what we might crudely call data entry problems. Better, one might think, to leave the complex and time consuming data collection and validation tasks to an automated system and its designers (properly informed by social service professionals), freeing front line staff to exercise experience and judgment in decisions about the proper response to information delivered by the tool. The point reinforces the significance of implementation decisions to the ultimate ethical justification of PRM in the child maltreatment area. Those decisions should ensure a place for social service professional judgement and experience.

There is at least one other issue under this broad heading that we have touched upon elsewhere, namely the question of who takes on the responsibility of monitoring the Vulnerable Children PRM and engaging with families. As we have said, some families will have no current active relationship with a social work professional. Others may be receiving services, perhaps through the Wellchild/Tamariki Ora programme, their GP or other local health or services provider. In the first case, effective engagement must be initiated and maintained. It may seem in the second case that those who are already engaged with the family should manage any initiative promoted by the Vulnerable Children PRM. On occasion, however, the functions which intensive intervention calls for in response to a high-risk rating may sit uneasily beside the existing or primary professional relationship between the professional and the family or individual. Staff may feel uneasy, and with good cause, about taking on dual roles. The detailed implementation issues this problem generates are beyond the scope of this paper. Here as elsewhere, however, decisions on these matters will be crucial to the overall ethical justification of the application of the Vulnerable Children PRM.

i) Rights and the Vulnerable Children PRM

As a conceptual matter, rights may be regarded as especially powerful normative claims, claims that entail that some other person or agency is under a duty to act

consistently with the asserted right. Rights so understood are extraordinarily important: they protect the important interests of individuals (and perhaps groups¹⁰¹) from disputed views as to what would be fair or just and from calculations of social utility.¹⁰²

I have for the most part avoided explicit appeal to moral theory in this document, but it will be obvious that many of the issues we have discussed have raised broadly 'consequentialist' considerations, considerations about what would have the best outcome for children, families and social service professionals. Such arguments will often carry the day. However it is tempting to think that we might avoid some of the uncertainty around those issues by simply pointing to a right, playing a normative trump. It might seem, this is to say, that appeal to rights would allow us to resolve the apparent conflicts between the different moral interests and perspectives we sketched in the introduction to this paper. There are some obvious rights to which one might look. New Zealand ratified the United Nations Convention on the Rights of the Child (UNCRC) in 1993, three years after the UN General Assembly adopted it. Article 19 of the Convention reads as follows:

States Parties shall take all appropriate legislative, administrative, social and educational measures to protect the child from all forms of physical or mental violence, injury or abuse, neglect or negligent treatment, maltreatment or exploitation, including sexual abuse, while in the care of parent(s), legal guardian(s) or any other person who has the care of the child.

Some contributors to the child maltreatment literature have argued that greater emphasis should be placed upon such rights. Richard Reading et al describe the dominant approach to child maltreatment as "similar to the conventional medical model of diagnosis and treatment", that emphasizes "identification, assessment, and intervention to treat and prevent further harm",¹⁰³ an approach which would include the Vulnerable Children PRM. They argue that equal focus should be given to maltreatment conceived of as a violation of children's rights, as they are spelled

¹⁰¹ Dare, 2001

¹⁰² "Rights are best understood as trumps over some background justification for political decisions that states a goal for the community as a whole." Dworkin, 1984), p. 153.

¹⁰³ Reading, et al. 2009.

out in the UNCRC. Reading et al do not seek to displace the view of child maltreatment as a public health issue, but instead to add the unique rights-based approach to our understanding of the problem. “Rights-based and public-health approaches to child maltreatment”, they conclude:

... are complementary, and when harnessed in concert they can act as a highly effective instrument of change in policy, professional activity, and public values. The unique strength of a rights-based approach is the legal status of rights conventions, and thus the accountability and transparency this facilitates.¹⁰⁴

In practice, however, the UNCRC demands less accountability than this description might suggest. States' obligations under the Convention have been described as those of “progressive implementation”. Article 4 requires ratifying countries to “undertake all appropriate legislative, administrative and other measures for the implementation of the rights recognised in the ...Convention.” Even rights such as the right to life are qualified with terms such “to the maximum extent possible” and “subject to progressive realisation and resource availability”.

None of this is to deny that the UNCRC specifies important moral considerations for New Zealand or that the required reports on New Zealand’s performance under the Convention are valuable. The point is rather that the moral wrongs done to children who are maltreated, or who are perhaps not protected as well as they ought to be, are wrong independently of the UNCRC. The most recent UNCRC report on New Zealand was very critical of our “staggering rates of child abuse and poverty”. The White Paper for Vulnerable Children,¹⁰⁵ in which the Vulnerable Children PRM plays a significant role, can be seen as a response to some of these concerns, addressing a number of the recommendations in the UN Committee on the Rights of the Child’s response to New Zealand’s report to the Committee, which included the establishment of “mechanisms for monitoring the number of cases and the extent of violence, sexual abuse, neglect, maltreatment or exploitation...”.¹⁰⁶

¹⁰⁴ Ibid, 340

¹⁰⁵ Ministry of Social Development, 2012

¹⁰⁶ UNCRC, Fifty-sixth session.

We have also encountered a number of other rights claims as we have sought to provide an ethical evaluation of the Vulnerable Children PRM.

- Stigmatisation, I suggested, may be interpreted as implicitly raising issues of discrimination and so as raising issues under the Bill of Rights and Human Rights Act which together prohibit discrimination on grounds that include 'being a recipient of a benefit'.¹⁰⁷
- I suggested that one important reason to be wary of mandatory engagement was the implication that it might appear to treat those identified as at risk as though it had already been established that they had committed some wrong, leaving them vulnerable to compulsion and the removal of significant social liberties and rights.
- I suggested that the cluster of concerns about privacy and confidentiality raised a number of issues usefully understood to rest upon moral rights to privacy and legal rights under privacy legislation and legally enforceable professional duties of confidentiality.

Rights have important roles in the ethical environment around the Vulnerable Children PRM and child maltreatment more generally, though they should not, in my view, be given normative priority. Instead, they compete in the moral balance of reasons with a range of other competing interests and perspectives. As remarked at the outset, the relevant rights are often in tension not only with competing moral interests grounded in alternative normative theories and perspectives, such as those grounded in broadly consequentialist theories, but also with other rights; the moral right to privacy, for instance is in direct tension with children's rights to protection. One cannot avoid balancing competing interests by granting lexical priority to rights-based interests, then, since there are unavoidable conflicts within the cluster of interests plausibly protected by rights. One is driven, in my view, to a more pluralist approach, one that takes the rights at issue in child protection policy and practice seriously, without giving them

¹⁰⁷ Bill of Rights Act 1993, s.19(1) and Human Rights Act 1993, s.21(1)(k)(ii).

determinative normative status. Moral evaluation of the Vulnerable Children PRM requires us to identify and balance the relevant moral perspectives and interests, aiming to give appropriate weight to each, in order to arrive at a view of the ethical status of the model.

Conclusion

The Vulnerable Children Report recognizes that the application of PRM to child maltreatment raises significant ethical issues. The report contains a brief summary of those issues, but includes a recommendation that “[a] full ethical evaluation of PRM is necessary before implementation” and that “[a]dditionally, an ethical framework should be developed to guide agencies in their responses to the use of automated child risk scores”. This report contains that evaluation and makes some recommendations that might form an ethical framework for implementation of the Vulnerable Children PRM.

As with any risk prediction tool, the Vulnerable Children PRM will inevitably make some errors at any threshold for referral, identifying as low risk some children who go on to experience abuse or neglect, and identifying as high risk some children who do not. That acknowledged, quantifiable, limitation requires the retention and support of current early identification referral routes such as those initiated health professionals and front-line social service professionals. The consequences of mistaken identification as at high-risk must be reduced as much as possible by, for instance, providing opportunity for experienced professionals to exercise judgement about appropriate responses, ensuring that such professionals understand the potential of PRM to miscategorise families, and providing training against confirmation bias.

The Vulnerable Children PRM raises particularly pressing concerns about the possible stigmatization of individuals or families as a consequence of their being identified as at high-risk, and consequently there is an ethical obligation to address the causes and effects of such stigmatization. We might do so by ensuring

that interventions are at the minimum level necessary, and that information produced by PRM is disseminated as narrowly as possible, with as little detail as possible, consistently with achieving the benefits of the intervention. Furthermore, since those identified as at risk will have committed no wrong and that most of them will not go on to do so, interventions must be preventive and supportive, not punitive. Interaction with high-risk families should be as similar as possible to that with other families needing support, although not at risk of maltreating their children, agencies with responsibility for implementing the Vulnerable Children PRM should explore ways of engaging with the media to minimise stigmatisation, and engagement with high-risk families should be on a voluntary basis.

The Vulnerable Children PRM is largely compliant with requirements for ethical screening programmes, though such principles emphasise the importance of developing and maintaining effective interventions and the capacity of child protection services to respond appropriately to identified need.

The Vulnerable Children PRM must not lead to the neglect of the needs of lower risk children or overwhelm social services to the overall detriment of vulnerable children. It should be used as an opportunity to deliver additional intensive intervention to high-risk families, not to reduce or neglect existing universal services. In a similar vein, the Vulnerable Children PRM must not be seen as a replacement for the judgement and engagement of experienced social service professionals.

The Vulnerable Children PRM poses some threat both to general moral rights to privacy and to legal and moral rights to confidentiality. Such threats must be monitored and minimised to ensure that they are justified and consistent with relevant moral and legal constraints.

I remarked at the outset that consideration of the moral status of efforts to reduce child maltreatment requires us to address and balance a number of competing moral perspectives and interests. That balance is not easy to identify, and it will

not be easy to maintain. Nevertheless, the issues at stake in child protection are of overwhelming importance. There is a powerful ethical imperative to find a way to protect some of the most vulnerable members of our community. While the application of predictive risk modelling to child maltreatment does raise significant ethical concerns, in my view those concerns can either be significantly mitigated by appropriate implementation strategies or are plausibly outweighed by the potential benefits of such modelling.

Bibliography

- Allen, A. *Uneasy Access: Privacy for Women in a Free Society*, Totowa, N.J.: Rowman and Littlefield, 1988.
- Allen, A. "Privacy and Medicine", *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/privacy-medicine/>
- Australian and New Zealand Association of Social Workers *Code of Ethics*
http://anzasw.org.nz/social_work_practice/topics/show/158-summary-of-the-code-of-ethics
- Bangma, Chris 'To Screen or Not to Screen', *ERSPC News Update*, 15 March 2012.
<http://www.erspc-media.org/>
- Barlow, J., Fisher, J.D. and Jones, D. (2012). *Systematic Review of Models of Analysing Significant Harm*. (Research Report DFE-RR199). London: DfE [online]. Available from:
<https://www.education.gov.uk/publications/standard/publicationDetail/Page1/DFE-RR199>
- Baumberg, Ben , Kate Bell, Declan Gaffney, Rachel Deacon, Clancy Hood, and Daniel Sage. *Benefits Stigma in Britain* (Elizabeth Finn Care, University of Kent, 2012)
- Bayer, Ronald. "Stigma and the Ethics of Public Health: Not Can We but Should We." *Social science & medicine* 67, no. 3 (2008): 463-72.
- Bayer, Ronald, and Jennifer Stuber. "Tobacco Control, Stigma, and Public Health: Rethinking the Relations." *Journal Information* 96, no. 1 (2006).
- Bell, Kirsten, Amy Salmon, Michele Bowers, Jennifer Bell, and Lucy McCullough. "Smoking, Stigma and Tobacco 'Denormalization': Further Reflections on the Use of Stigma as a Public Health Tool. A Commentary on Social Science &

- Medicine's Stigma, Prejudice, Discrimination and Health Special Issue (67: 3)." *Social science & medicine* 70, no. 6 (2010): 795-99.
- Bloustein, E., 'Privacy as an Aspect of Human Dignity: An Answer to Dean Prosser' *New York University Law Review* 1964: 39: 962–1007
- Browne, K. and Chou, S. (undated) *A Literature Review on Systems for Early Prediction and Risk Detection in Child Protection in Europe*. Deutsches Jugendinstitut. Available from: http://www.dji.de/bibs/Expertise_Browne.pdf
- Bruce, G, and Jo C Phelan. "Conceptualizing Stigma." *Annual Rev Sociol* 27, no. 1 (2001): 363-85.
- Burris, Scott. "Stigma, Ethics and Policy: A Response to Bayer." *Social science & medicine* 67 (2008): 473.
- Dare, Tim. "Mass Immunisation Programmes: Some Philosophical Issues." *Bioethics* 12, no. 2 (1998): 125-49.
- "Group Rights and Constitutional Rights" in Grant Huscroft and Paul Rishworth (eds.) *Litigating Rights* Oxford, Hart Publishing, 2001
- . "Parental Rights and Medical Decisions" *Pediatric Anesthesia* 19, no. 10 (2009): 947–52.
- , *The Counsel of Rogues? A Defence of the Standard Conception of the Lawyer's Role* (Ashgate, 2009)
- DeCew, Judith, "Privacy", *The Stanford Encyclopedia of Philosophy* (Fall 2012 Edition), Edward N. Zalta (ed.), <http://plato.stanford.edu/archives/fall2012/entries/privacy/>
- Dias, Mark S., Smith, Kim, deGuehery, Kathy, Mazur, Paula, Li, Veetai & Shaffer, Michele L. (2005). "Preventing head trauma among infants and young children: A hospital-based parent education program". *Pediatrics*, 115(4): e470-e477.
- Dickison, Michael 'Beneficiaries 'attacked on all sides'', *New Zealand Herald* February 6, 2013. http://www.nzherald.co.nz/social-welfare/news/article.cfm?c_id=322&objectid=10863719&ref=rss
- Dubowitz, Howard, Feigelman, Susan, Lane, Wendy & Kim, Jeongeun. (2009). "Pediatric primary care to help prevent child maltreatment: The Safe Environment for Every Kind (SEEK) model". *Pediatrics*, 123(3): 858-864.
- California Department of Health Services, *A Model for Change: The California Experience in Tobacco Control*, Sacramento, Calif: Tobacco Control Section, California Department of Health Services, 1998.
- Devaney, J. & Spratt, T. 'Child Abuse as a Complex and Wicked Problem: Reflecting on Policy Developments in the United Kingdom in Working with Children

- and Families with Multiple Problems', *Children and Youth Services Review* 31, no. 6 (2009): 635-41.
- Dickison, M. (2013, February 6). Beneficiaries 'attacked on all sides'. *The New Zealand Herald*, p.A3.
- Dworkin, Ronald, *Taking Rights Seriously*, (London: Duckworth, 1978)
- Dworkin, Ronald 'Rights as Trumps,' in *Theories of Rights*, ed. by Jeremy Waldron (Oxford: Oxford University Press, 1984), pp. 153-167 New York: Oxford University Press.
- Fried, C., *An Anatomy of Values*, Cambridge: Harvard University Press, 1970
- Fergusson, David M., Horwood, L., John & Boden, Joseph. (2012). *Early Start evaluation report: Nine-year follow-up*. Wellington: Ministry of Social Development.
- Gambrill, Eileen, and Aron Shlonsky. 'Risk Assessment in Context', *Children and Youth Services Review* 22, no. 11 (2000): 813-37.
- Gavison, R., 'Privacy and the Limits of Law' (1980) *Yale Law Journal*, 89: 421- 71
- Gavison, R., 'Feminism and the Public/Private Distinction' (1992-1993) *Stanford Law Review* 45.1: pp.1-46,
- Gerstein, R., 'Intimacy and Privacy' *Ethics*, (1978) 89: 76-81
- Gilbert, R., Fluke, J., O'Donnell, M., Gonzalez-Izquierdo, A., Brownell, M., Gulliver, P., Janson, S. and Sidebotham, 'Child maltreatment: variation in trends and policies in six developed countries', *The Lancet* December 2011 (online).
- Goffman, Erving. *Stigma: Notes on the Management of Spoiled Identity* New York J. Aronson, 1963.
- Herek, Gregory M., "Aids and Stigma." *American Behavioral Scientist* 42, no. 7 (1999): 1106-16.
- Herek, Gregory M, and John P Capitanio. "Aids Stigma and Sexual Prejudice." *American Behavioral Scientist* 42, no. 7 (1999): 1130-47.
- Herek, Gregory M, and Eric K Glunt. "An Epidemic of Stigma: Public Reactions to Aids." *American Psychologist* 43, no. 11 (1988): 886.
- Humphreys, S. and Stanley, N. (eds) (2006) *Domestic Violence and Child Protection: Directions for Good Practice*, London, Jessica Kingsley. 2006,
- Hoffman, Richard M. 'Screening for prostate cancer' UpToDate, (Wolters Kluwer, 2013) <http://www.uptodate.com/contents/screening-for-prostate-cancer> (accessed 12 March, 2013)
- Howard, Kimberly S. & Brooks-Gunn, Jeanne. (2009). "The role of home-visiting programs in preventing child abuse and neglect". *The Future of Children*, 19(2): 119-146.
- Inness, J., *Privacy, Intimacy and Isolation* (Oxford: Oxford University Press, 1992)

- Irueste-Montes, Ana Maria, and Montes Francisco. "Court-Ordered Vs. Voluntary Treatment of Abusive and Neglectful Parents." *Child Abuse & Neglect* 12, no. 1 (1988): 33-39.
- Kant, Immanuel, *Critique of Pure Reason*, (1781/1787), trans. Norman Kemp Smith (London; New York: Macmillan, 1929)
- Leventhal, J. 'Can child maltreatment be predicted during the perinatal period: Evidence from longitudinal cohort studies?' *Reproductive and Infant Psychology* (1988) Volume 6, Issue 3, 139-161.
- Lewis, G. Kirkham, H. Duncan, Ian, Vaithianathan, R, 'How Health Systems Could Predict and Prevent 'Triple Fail' Events That Are Costly And Yield Poor Outcomes And Patient Experience' Forthcoming, 2013
- Link, Bruce G., and Jo C. Phelan. "Conceptualizing Stigma." *Annual Review of Sociology* 27, (2001): 363-85.
- Luban, David Lawyers and Justice (Princeton University Press, 1988)
- Mackay, Ross 'Predictive risk modelling of child maltreatment: The ethical obligation to respond to children identified as being at risk of future maltreatment' (Unpublished, Draft Working Paper, 2013)
- MacKinnon, C., 1989, *Toward a Feminist Theory of the State*, Cambridge: Harvard University Press
- Mahon, James "The Truth About Kant On Lies," in *The Philosophy of Deception*, ed. Clancy Martin (Oxford University Press, 2009), 201-224
- Mann, J. M, L Gostin, S Gruskin, T Brennan, Z Lazzarini, and H Fineberg. "Health and Human Rights." *Health and Human Rights: An Inter-national Journal* 1 (1994): 7-23.
- Mansell, James, Rissa Ota, Ricus Erasmus, and Kip Marks. "Reframing Child Protection: A Response to a Constant Crisis of Confidence in Child Protection." *Children and Youth Services Review* 33 (2011): 2076-86.
- Mill, John Stuart. *On Liberty*, London, GBR: Electric Book Company, 1859.
- Ministry of Social Development, 'Discussion Predictive risk modelling of child maltreatment: The ethical obligation to respond to children identified as being at risk of future maltreatment', Unpublished, 2013.
- Moore, A., 'Privacy: Its Meaning and Value' *American Philosophical Quarterly* (2003) 40: 215-227.
- Munro, Eileen. "Avoidable and Unavoidable Mistakes in Child Protection Work." *British Journal of Social Work* 26, no. 6 (1996): 793-808.
- . "The Dangers of Information Sharing." *Social Policy Journal of New Zealand* 31 (2007): 41-55.
- . *The Munro Review of Child Protection: Final Report, a Child-centred System*. Department for Education (The Stationery Office, London, UK., 2011).

- Mikton, Christopher and Alexander Butchart 'Child maltreatment prevention: a systematic review of reviews' *Bulletin of World Health Organ* 2009; 87: 353–361
- Myers, John E.B. 'A Short History of Child Protection in America' *Family Law Quarterly* 42, no. 3 (2008): 449
- New Zealand Government (2012) *White Paper for Vulnerable Children: Volumes I and II*. Available from: <http://www.childrensactionplan.govt.nz/the-white-paper>
- New Zealand Medical Association *Position Statement: Patient Confidentiality and AIDS/HIV* (2008)
<http://www.nzma.org.nz/policies/advocacy/positionstatements>
- . *Code of Ethics* (2008).
http://www.nzma.org.nz/sites/all/files/Code_of_Ethics.pdf
- Norman, Rosana E, Munkhtsetseg Byambaa, Rumna De, Alexander Butchart, James Scott, and Theo Vos. 'The Long-Term Health Consequences of Child Physical Abuse, Emotional Abuse, and Neglect: A Systematic Review and Meta-Analysis', *PLoS medicine* 9, no. 11 (2012): e1001349.
- O'Hare, Thomas "Court-Ordered Versus Voluntary Clients: Problem Differences and Readiness for Change." *Social Work* 41, no. 4 (1996).
- Olds, D. L., Eckenrode, J., Henderson, C. R., Jr., Kitzman, H., Powers, J., Cole, R. et al. (1997). Long-term effects of home visitation on maternal life course and child abuse and neglect. Fifteen- year follow-up of a randomized trial. *JAMA*, 278, 637-643.
- Panattoni, L.E., Vaithianathan, R., Ashton, T. and Lewis, G.H. Predictive risk modelling in Health: Options for New Zealand and Australia. *Australian Health Review* (2011) 35, 41-51.
- Parker, R. Aggleton, P. "Hiv / Aids-Related Stigma and Discrimination : A Conceptual Framework and an Agenda for Action." *Social Science & Medicine* 57, no. 1 (2002): 1-28.
- Peckworth, Sue, 'Domestic Abuse, Safeguarding Children and Public Health: Towards an Analysis of Discursive Forms and Surveillant Techniques in Contemporary UK Policy and Practice' *Journal of Social Work* (2013) 1–18 doi:10.1093/bjsw/bct042
- Prinz, R. J., Sanders, M.R., Shapiro, C.J., Whitaker, D.J. & Lutzker, J.R. (2009). "Population-Based Prevention of Child Maltreatment: The U.S. Triple P System Population Trial". *Prevention Science*, 10(1): 1-12
- Rachels, J., 'Why Privacy is Important' (1975) *Philosophy and Public Affairs*, 4: 323–33

- Reading R, Bissell S, Goldhagen J, et al. Promotion of children's rights and prevention of child maltreatment. *Lancet* (2009) 373: 9660, 24–30 January, Pages 332–343
- Reidpath, Daniel D, and Kit Yee Chan. "Hiv, Stigma, and Rates of Infection: A Rumour without Evidence." *PLoS medicine* 3, no. 10 (2006): e435.
- Runyan, D., Wattam, C., Ikeda, R., Hassan, F., and Ramiro, L. (2002) Child abuse and neglect by parents and caregivers. In: Krug E, Dahlberg LL, Mercy JA, Zwi AB, Lozano R, editors. *World Report on Violence and Health*. Geneva, Switzerland: World Health Organization.
- Scambler, Graham. "Health Related Stigma." *Sociology of health & illness* 31, no. 3 (2009): 441-55.
- Schoeman, F., *Privacy and Social Freedom* Cambridge: Cambridge University Press, 1992
- Schröder FH, Hugosson J, Roobol MJ, et al. Prostate-cancer mortality at 11 years of follow-up. *New England Journal of Medicine* 2012;366(11):981-990
- Schwartz I., Jones, P., Schwartz, D. & Obradovic, Z. 'Improving Social Work through the Use of Technology and Advanced Research Methods', in D. Lindsey & A. Shlonsky (eds) *Child Welfare Research: Advances for Practice and Policy*, (New York: Oxford, 2008) pp.214-30.
- Sei-Hill, Kim, and James Shanahan. "Stigmatizing Smokers: Public Sentiment toward Cigarette Smoking and Its Relationship to Smoking Behaviors." *Journal of health communication* 8, no. 4 (2003): 343-67.
- Snyder, Christine M J, and Stephen A Anderson. "An Examination of Mandated Versus Voluntary Referral as a Determinant of Clinical Outcome." *Journal of Marital and Family Therapy* 35, no. 3 (2009): 278-92.
- Thompson, Ian M., Donna Paula Ankerst, Chen Chi, M. Scott Lucia, et al. "Operating Characteristics of Prostate-Specific Antigen in Men with an Initial Psa Level of 3.0 Ng/ml or Lower." *The Journal of the American Association of Medicine* 294, no. 1 (July 2005): 66-70.
- U.S. Preventive Services Task Force. Screening for Prostate Cancer: Final Recommendation Statement. AHRQ Publication No. 12-05160-EF-2. <http://www.uspreventiveservicestaskforce.org/prostatecancerscreening/prostatefinalrs.htm>
- United Nations Committee on the Rights of the Child, Fifty-sixth session. 'Consideration of Reports Submitted by States Parties under Article 44 of the Convention: Concluding observations: New Zealand' <http://www.unhcr.org/refworld/category,COI,CRC,STATEPARTIESREP,NZL,4dcb8beb2,0.html>

- Urological Society of Australia and New Zealand, PSA Testing Policy (2009).
<http://www.usanz.org.au/uploads/29168/ufiles/newsfiles/factsheets/psatesting.pdf>
- Vaithianathan, R., Maloney, T., Jiang, N., Dare, T. de Haan, I., Dale, C. & Putnam-Hornstein, E. "Vulnerable Children: Can Administrative Data Be Used to Identify Children at Risk of Adverse Outcomes?": Ministry of Social Development, New Zealand 2012.
<http://www.msd.govt.nz/documents/about-msd-and-our-work/publications-resources/research/vulnerable-children/auckland-university-can-administrative-data-be-used-to-identify-children-at-risk-of-adverse-outcome.pdf>
- Well Child/Tamariki Ora National Schedule (<http://www.wellchild.org.nz/?t=34>).
- Wilson, J.M.G., Jungner G. *Principles And Practice Of Screening For Disease* (World Health Organization; Geneva, 1968)
http://whqlibdoc.who.int/php/WHO_PHP_34.pdf
- World Health Organisation. *Violence prevention the evidence: Preventing violence through the development of safe, stable and nurturing relationships between children and their parents and caregivers* (World Health Organization; Geneva, 2009)
http://www.who.int/violence_injury_prevention/violence/4th_milestones_meeting/evidence_briefings_all.pdf