Privacy Commissioner
Te Mana Matapono Matatapu

PRIVACY COMMISSIONER JOHN EDWARD'S SPEECH TO THE PAEDIATRICS SOCIETY OF NZ CHILD PROTECTION SPECIAL INTEREST GROUP

NAPIER WAR MEMORIAL AND CONFERENCE CENTRE, MARINE PARADE, NAPIER
TUESDAY 18 NOVEMBER 2014

INTRODUCTION

"First, do no harm". I've always admired the deceptive simplicity of that most basic of injunctions to the medical practitioner.

It's deceptive, because it can be interpreted as 'don't stuff it up'. Don't do anything wrong. Don't do anything bad.

But that can lead to problems of its own. Because what happens when you only get part way through that sentence? 'Don't do anything' is *one* way of not doing anything wrong. But sometimes we have to act, we have to *do*. And that entails risk.

"Risk" is a word that you will all be familiar with, as in recent years it has become a byword for an industry of consultants and policies and management models that is impossible to avoid. But you may not be aware that the word "risk" apparently comes from the Ancient Greek word for 'something you *really* want to avoid when you're sailing' – appropriate for the title of this meeting – "Navigating the journey through childhood".

It's a useful way of thinking about it; because once you're on the water, sailing along in your trireme, you need to get where you're going to.

Turning back is not an option unless you want to be roundly mocked at the *taverna*.

Risk is an inevitable part of the delivery of health and social services and is as much a part of keeping children safe as it was of navigating the seas between Scylla and Charybdis.

Much of what I am going to talk about today is an approach to assessing risk, assessing privacy risks and information risks, by answering the easy questions first. And then letting that assessment of risk guide your actions.

What I aim to unfold for you is a way of charting a path around the things you really want to avoid and getting to where you need to be, which is safer and happier children and young people.

I have some familiarity with some of the risks, tradeoffs and difficult ethical and legal decisions faced by the caring professions on a weekly basis. For many years I was a District Inspector for Mental Health. This meant I had an independent oversight role in the provision of compulsory mental health services. I inspected hospitals, accommodation providers, documentation, and I investigated complaints and undertook enquiries.

One young woman was an inventive and determined self harmer. She took every opportunity to cut herself with any fleck of glass or discarded paperclip that she could get her hands on. If she had leave from the ward she would conceal the tiniest of weapons, and would later be found in a pool of blood in a toilet. It was extremely distressing for her, for her family, and for the staff who had to care for her.

The staff decided that the only way to properly manage the risk she presented was to increase supervision, and security. She was cared for mostly in the locked, intensive care part of the ward, and increasingly, in seclusion. She asked me to investigate her care. Now I'm not a doctor, or any kind of health professional, and it was not my place to second guess the judgement of the team whose job it was to make clinical judgements about her. So I enlisted the help of a senior child and adolescent psychiatrist from another region, and also engaged an expert from personality psychotherapy services. These two people told me that in an effort to contain the risk this patient presented, the clinical team were in fact causing greater harm, and exacerbating the risk. This was a very difficult message for her treatment team to hear. All they wanted to do was to ensure her safety. They were acting with her best interests at heart, and in utmost good faith. But they were getting it wrong, and causing iatrogenic harm, that is, harm caused by the clinical intervention.

My advisers told me that clinicians had a limited ability to control the patient's behaviours in a hospital environment, and that things would only get worse. If someone wants to kill themselves, they will. She presented obvious risks that everyone understood, but those risks had to be managed in a different way, with cognitive behavioural therapy, and a management plan in which responsibility was shared among all those who would be involved with her care. A meeting was convened between the Police, the community mental health team, the family, the ambulance services, the emergency department, the GP and others. Everyone was told what the plan was, and what the risks were. She would be released, she would have programmes, and support, and structure. And she would cut.

And when she cut, she would be taken by ambulance to the emergency room, and stitched up, assessed, and returned to her accommodation as soon as possible. Everyone agreed with that plan. Everyone accepted the risk that she might die. And everyone knew that she might anyway.

With everyone in agreement they could move her out of her isolation, and back to the community without thinking they were acting recklessly and endangering her life. They would not overreact when she cut herself, and she would not be readmitted, just to begin the cycle again.

That anecdote or vignette illustrates two or three things. I hope it demonstrates that I am not talking to you here from an ivory tower of theory and legal niceties. Second, that risk comes in many forms, and can be a double edged sword. Taking action, or not taking action to avoid a legal risk can exacerbate a clinical risk and vice versa. Third, very often a team of people need to come together to provide all the services a person needs, and that means sharing information, across disciplines, agencies, and family members.

Guidance Material

The need for agencies to share information has been highlighted in many reports and recommendations of inquiries, Health and Disability Commissioner investigations, district inspector investigations, and sadly, in coroners findings.

There have been a number of impediments to good information sharing. Some of these are process, some are institutional, some are about concerns of legal ability and liability. This last often coalesces around the Privacy Act, for which I am responsible. I'm not going to stand here and say that "you're getting it all wrong, its easy", because it is not always easy. You deal with some very complicated human interactions, and difficult medical and social problems. Why would we expect the law to be simple?

While there will be difficult complex ethical and legal questions in a few cases, often there are simple answers. Since taking up my role as Privacy Commissioner, I've adopted the mission of "making privacy easy".

Not long after I took up my appointment we contacted Russell Wills, the Children's Commissioner, and suggested there could be some benefit in putting our heads together to provide some practical guidance for people working in multi-agency, multi disciplinary teams providing services for vulnerable children and their families.

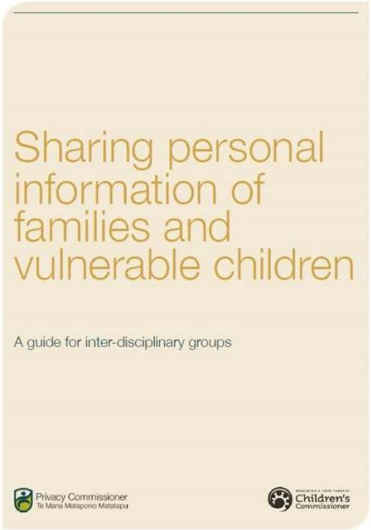

Our discussions eventually coalesced in this document, *Sharing Personal Information of Families and Vulnerable Children* which can be found on our website. It's aimed at multi-agency teams, and is intended to give those teams a clear and straightforward shared understanding of how they can use information to work together. Today I am launching this guidance document together with a practical decision making tool, and a few ancillary matters like predictive risk modelling and the mysteries of the AISA, before we throw the floor open for questions.

Escalation Pathway

The escalation pathway is a tool that Sebastian Morgan-Lynch of my office developed with people from (Children's Commissioner) Russell Will's office, and it's pretty simple.

The guidance is about "personal information". I often hear conversations on this topic get off on the wrong foot. When you hear someone say "who owns the information", you're heading down an unhelpful path. Individuals and agencies have rights and responsibilites in relation to personal information, and

those rights and responsibilities do not depend on any concept of “ownership”. Rather, they are defined by concepts such as who holds the information, and for what purpose, and by what authority.

Let’s say you want to disclose a bit of information. First question; can you keep it anonymous? Sometimes you just want a second opinion, or maybe even to know whether your vague feelings of unease have a basis. If so, no problem. Anonymous disclosure is almost always ok. That's the first step on the pathway, 'not naming any names'.

But anonymity is often not that helpful, particularly when you're trying to leverage the collective wisdom of a group. If you have six professionals around a table, each of whom might have a piece of the puzzle, you can't afford to be too cagey about your own pieces. So anonymity isn't going to work - move along the pathway to the next step. Asking someone for their permission to share their information is a surprisingly powerful tool. I say surprisingly, but it's really not that surprising because it gives people power over their own information and it gives you the opportunity to explain the process.

Children’s actions teams recognise this by only bringing in clients who want to be there, or are willing to be there. Earlier in the year I and my staff spent some time talking to the people who are setting up the trial children’s action teams in Whangarei and Rotorua and a strong message that came back was the importance of consent, of authorisation. It's not an absolute, we'll see that as we keep going along the escalation pathway, but consent is the best starting point. They’ve also been set up that way of necessity, because the legal instrument that would provide with authority to share with everyone around the table whether consent has been granted or not, the Approved Information Sharing Agreement is not yet in place.

But let's say consent hasn't worked, because it won’t always. You're dealing with people with difficult and complicated lives, and they may not trust the use you're going to make of their information or they may not trust the people who they think are going to see it.

You will always have a reason for wanting to disclose the information, because otherwise why are you doing it? In a children's action team you are wanting to improve someone's life. You'll also always have a reason why you collected that information, it can be a bit trickier to be specific about that, but it's hopefully the same reason. Now if those two reasons match up, and they should, then step three of the escalation path says you should be able to disclose to an appropriate agency for that purpose. You can always disclose information when the disclosure is to help accomplish the purpose for which you obtained it. This has some wider application when you've told someone that you're going to be sharing their information. At that point and from then on, sharing of the information is one of the purposes of collection. Purpose disclosure and openness will get you a long way in the information privacy world; know your purpose and then be open about it, and that's step three of the escalation path.

But sometimes it doesn't work. You've found out that an estranged husband is hitting his former partner, or vice versa, and you need to do something now. You can't keep it anonymous because that won't do anyone any good, though it might be handy to get a second opinion before you irrevocably damage your clinical relationship with one or both. Neither of the two have any interest in you telling anyone about it, so that's step two out of the way. You've never suggested passing on this kind of information, because they seemed to be fine up until now, so you've never told them you might

consider a disclosure and that's step three out of the question. So you move to step four - is there a serious threat here?

I don't want to give you lengthy legal lecture, but I do need to tell you that the law around serious threats was changed recently, in response to the blurred lines that disclosures of threats can involve. In the past the baseline for disclosure of information about a threat was that the threat had to be both serious and imminent. The trouble with that was 'imminence' could be a very hard bar, and that didn't match how we actually assess threats very well. If you are absolutely certain one of your patients is going to kill their partner within six months, that's a threat that is serious but not imminent. If you have concerns that a child is unsafe in their current home environment but can't put your finger on any threat that is going to endanger their safety in the immediate future, then that would also not meet the old test.

So the law changed, and now the bar for disclosure of threats is that the threat must be "serious". You can disclose information about a serious threat to anyone who can act to prevent or lessen that threat. The important bit is that serious is now defined, and you can look at the threat's seriousness, its imminence and its likelihood. An underlying theme that runs through all these decision making aids is "proportionality". It is seldom necessary to open your filing cabinets or information system to a team member to have full access to in all circumstances. At each of these steps you need to make an assessment of how much information you need to disclose in order to achieve the purpose you have in mind, or to mitigate the risk that you have identified and disclose only that much, only to those who need it.

That's step 4 of the escalation pathway. Is there a serious threat?

You sometimes hear folks say "safety trumps privacy". Others have said "I'd rather explain myself to the Privacy Commissioner than to the Coroner. 'Safety trumps privacy' can be a useful rule of thumb placeholder for the legal provisions that allow disclosure of information to prevent or lessen threats or to address risks to maintenance of the law or public safety, like the one I just described. But it also sets up a dichotomy between holding onto information and being open about it, as if if you want safety you can't have privacy and vice versa. I don't accept that dichotomy. Privacy is the bargain between the agencies with the power and the individuals with the information, and that bargain consists of a set of values. There is seldom a stark choice of "do you want privacy or do you want safe children". We want, and are entitled, to both.

Sometimes, though, there's information which you want to disclose to someone that doesn't reveal a serious threat, but you think another person, professional or family member needs to know. You can't do it anonymously, you can't get their agreement for whatever reason, you haven't told them the disclosure is going to happen. That's step five of the pathway and the stakes are getting higher. This is where the rest of the law comes in, because step five asks 'is there another legal provision that might apply?' It's a big question, because it covers any of the multifarious laws that apply to information. For instance: child safety disclosures can be made to any member of the Police or CYF Social Worker by anyone who has a legitimate concern for the child's safety or wellbeing. You can disclose health information to a child's parent or guardian on request. If you're a health practitioner you can disclose to family members or contact people, in line with recognised professional practice. And so on; I don't want to just reel them off, but there are a lot of laws that allow you to disclose information where you need

to. Step five is where you have the more involved discussions, where the compliance stakes get higher. And also where you call our office helpline: 0800 803 909, please do, that's what it's there for.

The advantage of applying the escalation tool, though, is that it is guiding you to the simplest, least risky way of doing what you have to do, without having to understand all the different legal relationships between provisions like ss. 15 & 16 of the CYPF Act, or s.66 of that Act or s.22C or 22F of the Health Act or rule 11 of the Health information Privacy Code etc etc. If you can keep things simple then do. Use anonymous information, get peoples agreement to the disclosure, tell them what you're doing. Get the easy stuff out of the way first, is the message of the escalation tool, and both I and Russell hope it's a useful one.

If you get to the end of the pathway and still don't have your answer, you're likely to be in one of those rare complex situations, probably where agencies are at odds, and some strong personalities are involved. We haven't been able to completely eliminate the need for lawyers I'm afraid to say!

AISAs

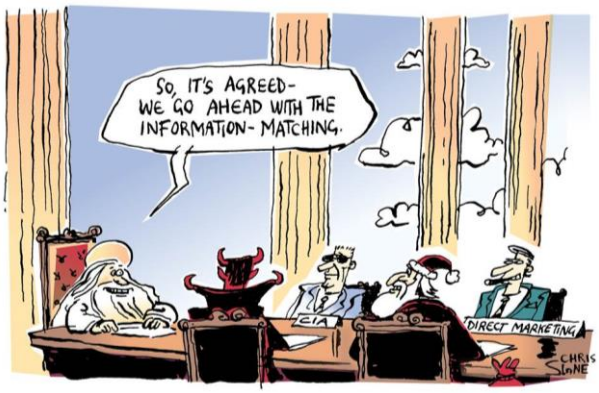

One of the many things that sits at that last stage of the escalation path is the AISA. An AISA is an approved information sharing agreement. That acronym begs a few legitimate questions you might want to ask - who agrees? who approves? what information is shared? and with whom?

Maybe a better way of thinking about AISAs is as miniature laws, designed for a specific identified set of activities and relationships, rather than the general all encompassing principles of the Privacy Act. The Law Commission recommended them back in 2008, and they're important because there's a big one coming out soon that regulates information sharing for child safety. The Law Commission's reasons can be read in its report, which is available online, but one of those reasons was to acknowledge the truth that the big simple laws around information sharing didn't generally accomplish their goals. Inevitably they get burdened with exceptions, and complexities, and uncertainties, and you end up in a worse place than you were when you started.

The reason for having AISAs is to let agencies be comfortable about what they can and can't share. To make information sharing simple. So you have a list of agencies that sign up, a shared purpose for the sharing - keeping children safe - and some agreed constraints on how the information is treated. In the child safety and well-being space there is an AISA being developed by the Ministry of Social Development, and we're helping out with some advice. When it's done we'll tell the Minister whether we think it is a sensible balance between getting the job done and keeping the information of children, their families and whanau safe.

Once the Minister approves the AISA, the law has changed, openly, safely, and transparently in a way that tells the people sitting round the table and discussing what they need to do to help the child or young person whose file, and whose life, is in front of them.

Often the AISA might not be an essential legal enabler. As I demonstrated with the escalation pathway, there's often a simple answer to the question 'can I share', and you can find it by the path. But what the AISA does is sit at the end of the path and provide comfort to agencies that have a legitimate concern for how the information they hold about vulnerable children and other individuals is going to be used.

Predictive Risk Modelling

Predictive Risk Modelling, which is one of the topics the MSD AISA may end up covering, is right at the edge in terms of public policy in the delivery of social services internationally. The concept has been fictionalised in movies like “Minority Report” in which pre-crims are identified before they have done their offending. We are not approaching that level of dilemma, but the research that has been undertaken into the role and reliability of predictive risk modelling as a tool to assist in identifying and designing interventions for vulnerable children is still interesting.

The theory of PRM is that all children and adults present a certain risk of suffering, or inflicting, harm. And so if you look at the characteristics of those children and adults you can notice common elements among at-risk individuals and act to prevent the harm.

So far that’s no different to any process of assessing risk; where the difference comes is in the scale. Big data. Perhaps administrative, rather than clinical data. Because it is one thing to look at your own patients or clients and make a judgment about the risk they are under or the risk that they pose to others on a bespoke sort of basis, but it is another to run a clever algorithm of the data of everyone in the country and have it pop out a list of the top ten, or twenty or a hundred of the people who need help right now.

But there are dangers. One danger is that the algorithm becomes the single source of the truth. Futurist Marshall McLuhan suggested that the devices we use shape whatever we do, and eventually become prosthetics. So a car is a prosthetic for our legs, a tv is a prosthetic for our imagination, clothes are prosthetics for our skin. So the danger is that we're liable to make a prosthetic for our judgment. Are we ready for that? If we're not, should we be?

One of the main reasons for controlling how personal information is used and disclosed is to keep control of the information with the people concerned. We do that because it's important that we are not prisoners of our past. We need to be able to change, and if we are always held to our previous actions.

There's a line to walk, between personal autonomy and personal responsibility, of course, and we can see that line being drawn in statutes like the Clean Slate Act.

Another risk is that we've got this great technology but we don't really know how to use it. If you have a little popup on your screen telling you that kid with the estranged parents and the bad truancy and the police report is 13.784% more likely to need government assistance five years from now, how do you use that information? If you get a list of 4,259 potentially at-risk children the PRM algorithm thinks need some attention and your caseload is already packed to the brim, how do you use that information?

What happens to the family who your judgement tells you is in danger but doesn't "ping" on the algorithim? How do you justify allocating resource to them over someone who "rates highly"?

These questions will have answers, of course, and I'm reasonably comfortable that we'll eventually find answers to the other questions that PRM might pose in the future, but I have an feeling we need less excitement in the potential of the technology and more confidence in its capabilities. Right now predictive risk modelling feels like an MRI machine that has been air dropped into a village. It could save lives, but … where do you even plug it in? It's great technology, but there's no one there who knows how to use it, and even if they do manage to plug it in and get a scan, there's no one able to read it or do anything with the results.

Before I conclude and open the floor to questions, I'll finish the story I started telling before:

I should tell you the outcome of the vignette. The woman at the centre of that investigation came out of hospital, and moved towns. I lost touch with her. Two or three years later she wrote to me. She had recovered. She had finished a tertiary qualification and was looking forward to beginning her career. She was reflecting on that earlier time in her life and wrote to me to ask me for a copy of the papers I had about her, and to express her gratitude to the people who had come together to give her a future.

Risk is an inescapable part of life. You take risks every day, on behalf of us all, to reduce the risks faced by some of the most vulnerable people in our community. We are grateful for that. I hope the material I have launched today will help you to apply your professional judgement appropriately to manage and reduce those risks.