The anonymisation decision-making framework

My name is Mark Elliott from Manchester University, National Centre for Research Methods (NCRM) and lead of the UK anonymisation network.

In the previous talk I described to you the sort of basic concepts of an anonymisation and this one I'm going to go through a framework that we call the anonymization decision-making framework. And that's a framework which helps you to make decisions about your data situation and to anonymize effectively. Ok so what's the anonymisation decision-making framework or the ADF? It's a system for developing an anonymization policy and it's a practical tool for understanding your data situation. The important points to emphasize at this stage are that it isn't a checklist. You can't go through the ADF ticking boxes and pop out with a nicely anonymized data set at the end. There's a context into play between different considerations that you will need to take account of and sometimes you're going around in circles between different steps in the process until you get to a resolved data situation.

What's your responsibility? You're holding some data and you want to make sure that it's safe and all you particularly want to share it or release it or disseminate it in some way. What you need to do is understand how a privacy breach might occur with your data in your data situation. Understand the possible consequences of that breach and then reduce the risk of that breach occurring to a negligible level. And the concept of negligibility is quite important. We're not asking for risk to be reduced to zero and we are asking for risk to be reduced to a very low level. Now fortunately getting risk down from moderate to negligible is a lot easier than getting risk from negligible to 0. So negligible is a slightly flexible concept and it means really that a risk that a reasonable person would ignore.

Ok so in fact there is a 10-point process to the anonymization decision-making framework. I list them off here now we're going to go through each of these in terms. So I'm not going to go through them in on this particular slide but what I will say is there's three different activities. The first is the data situation which is made up of the first five points. The second is disclosure risk assessment and control and that's the technical process the technical part of the process. And the third is impact management and that is the soft processes that you have to go through to make sure that the any impact of any breach is reduced should things go wrong. Ok so what do you actually have to do? The first thing is to understand and describe your date situation. So is the situation static or dynamic? By static we mean simply we're looking at an existing dataset in a particular environment and say actually is that safe? So you might do that as part of a risk review. The second is thinking about data moving around so dynamic data situation is one where data is being shared or released into new environments. So what are the environments that the data is moving around it and how does the data you have relate to that environment. So here's an example so here we have a data situation where we're moving data in fact into 3 environments so we might imagine an organization that's collecting data passing data to a local authority as part of some legal process and the local authority wants to release some aggregates from that data that's collected from the third party. And releasing aggregates means essentially publishing open data. Now each of those data processes will have its own risks associated with it and each of those environments will have their own risks associated with it as the data moves through them. So you need to understand in your particular situation and this sort of flow diagram often helps that.

Second point is you need to understand the legal and governance issues that surround your particular data sets. What legislation is relevant often it's the data protection act but it may

well be that other legislation governs the particular uses that your data may be put to. How did the different legislation interact in terms of what it you are allowed to do and not to? Related to this is what happened to the data before it reached you? Are you the primary collector of the data or are you a secondary user? And what are the governance processes that were affecting the original owners of the data before they handed it over to you and how does that interact with what you're allowed to do with the data now this can be quite a complex stage.

Now know your data actually may be this particular idea well of course I need to know my data but actually just getting a map of your data and in the framework we have a template for doing this, actually having a map of all the different property of data and then thinking how these relate to, notions of risk, is actually quite important and doing that as an exercise in early stage is very important and can then frame your risk management process. So where have the data come from, how were they collected, who are the data controllers for the data, are there any other parties involved as data processes, what are their responsibilities, is the data about people and is the based data personal data. The data could be about people and not be personal if it's been anonymized and similarly sometimes data can appear not to be about people but actually still be personal data. A good example of this was the data collected by the Department for Communities local government on fires by the Fire Rescue Service and that data ostensibly was about fires but fires often have very close association with individual people because they happen in locations often people's houses and therefore those data was still personal data even though ostensibly they were not about people.

Data subjects who are they are they're vulnerable as or a sensitive group, what is the relationship between the data subjects and the data, have they given any sort of consent to its reuse and so on. What type of data you have, is it quantitative qualitative, is it in the form of micro data or aggregates or some other form. What type of variables do you have do you have any variables that will be regarded as a sensitive either in law or just generally in terms of how they are understood. Do you have any standard identifier? Properties of the dataset that might be relevant the quality of the data actually this is slightly paradoxical but lower quality data is actually lower risk. It's less easy to find somebody if the data is of lower quality. Is the data time linked i.e. is it a hierarchical or flat, is it drawn from multiple sources, is it a population or is it a sample of the population and all of these properties can affect the risk.

Now understand the use case and again you may not be entirely clear about why you need to do this. What were the data to be used for? There may be that there's a specific request to share the data as a specific organization who wants to use it for a specific purpose. Actually understanding that purpose in detail will allow you to arrive at what is effectively a minimum specification for the data that are needed. So what variables are needed, is all the data needed or what a sample suffice, who will hold the shared data who will access it and how. Essentially these are definitions of the data and the data environment in that that new situation. If you've got well understood use case and then you go back and start thinking about what sort of data you're able to release in terms of its risk then you can have a dialogue between yourself and the potential user. Now in more general use cases where perhaps disseminating a data set for research purposes or as open data you can still usefully think about how users would like to use this data and to sort of think about what actually is of the most value.

Understanding your ethical obligations beyond the legal constraints is also important where

the loci of consent with these data and this can be quite complicated. So the data subjects may not have been involved in any direct consent process. So this happens when often when there are multiple levels of date subjects. So for example GP data consent is often given by GPs to access those data but not by the patients. And the data actually are data about GPs and they're also data about the patients as well. So a complex mixes of different types of data subjects. So who's consented to what is actually quite important in terms of understanding our ethical obligations and who is aware of what not just to do with the notion of consent but also awareness and other reasonable expectations that the data subject might have as to what is going to happen to their data about them once they've had it over for one purpose. Is it reasonable for them to expect that it won't be used for another purpose or would it be in the normal expectation of a data subject that actually their data would be reused and that's a very fuzzy area but if something is important to understand thinking about in terms of your data flow? Is the data situation sensitive so the topic of the data sensitive, Is it about a particularly stigmatizing disease for example, is the population that the data is about a vulnerable population perhaps it's a data about children and are there any sensitive variables again either legally or in terms of what's generally understood to be sensitive.

Ok now we move on to the technical disclosure control part of the framework. Identify the process you will need to use to assess the disclosure risk. Now there's a separate talk on that on disclosure risk assessment and control. Here we're just going through very briefly the main points. The first of these is scenario analysis that's very important. You are answering the question you set yourself up earlier on which was how my privacy breaches occur. Until you know that you can't possibly go about measuring the risk. It's not an abstract notion of risk assessment but a very located one this is the thing that I'm imagining happening and this is how it might happen and then you can measure the risk of that particular event. Now usually there is some form of re- identification but what resources is this imaginary person who's going to do this re-identification going to be using. And that's where we start thinking about the data environments. So these data are going to be in this environment so I potentially will have access to these resources in order to do the re-identification and you can think about the mapping between those and that's the function of scenario analysis. Systems disclosure risk assessment is a formal process of measuring the risk. Once you've defined the set of key variables and again I'll talk about that in more detail in the particular talk on disclosure risk assessment. Penetration tests are a simulation of an actual attack. So they're here we say to an individual ok here's the data set you see if you can find somebody in there and there are processes formalizing how you go about doing those penetration test either in-house or indeed as a crowdsourced hacking challenge.

Comparative data situation analysis takes the idea that you consider your own situation at the moment has to be safe and secure and therefore operate it as a gold standard. So if you're going to do a one-to-one share with another organization this can be particularly relevant because you can think about your data as you currently have it in your data environments and then what's the comparative risk in a different data environment and if that is less than the risk in your current data environment it's probably sufficient to say that it is sufficiently safe.

Final point here is, just consider using a thermostat approach to this this is a really good strategy if you're thinking about releasing open data. Go for a really cautious level of risk to start with release that and hopefully nothing will happen and the environment will become used to your data being existent the population will become used to they'll be inquiries made about it you'll get an understanding about demand for different types of data so you can enrich your understanding of the use case and then you can go for a slightly more liberal

approach just tweaking up the thermostat little. And this technique has been used for example by the German statistical agency determining the data that will be released into their research data centres. Once you've done your risk analysis then you know where it is you've got to bring the risk down, down to that negligible level and what controls you're going to apply and essentially there's two types of controls you can restrict access in some way or other who how what and where and to do what or you can place controls on the data so you might for example only release a sample of the data rather than the full dataset. You might decide to aggregate or suppress variables or you might deter by adding some noise to the data in some way. I am not going to go into the details of those and again just there will be some more on that in the next talk.

Output disclosure control can also be applied so if you are allowing access in a restricted environment there is still the question of what people take out those environment we don't just do analysis for the fun of it we usually want to publish we usually want to use our outputs for other purposes. So what outputs you're going to let out if your means of controlling risk is to restrict access so that only a particular environment is used such as a data centre.

Ok now we're moving on to the impact sides of the anonymization decision-making framework. First point is identifying who your stakeholders are and plan how you will communicate with them. Who needs to know about your share, which needs to know about the release of data, is the data subjects' the wider public the users these are questions you need to address. Once you've identified then what is the engagement you're going to going to carry out. Will they be involved in the design of the data and that might be particularly relevant to users but could involve data subjects as well. Is there going to be a consultation and for some way you might be changing policy around particular use of particular data. Consultation with the wider public might be really important in terms of thinking how that's going to play out when it gets publicized not by you. And transparency might well be important so being transparent , your own public announcements about the data, about the data processes, who are sharing with and why, are all important in terms of building understanding about why you're doing what you're doing and this will reduce the impact of a breach should it happen. What do they need to know? Well as well as about the data are you going to publish details of your anonymisation process? There's an advantage of that it might reassure people about the security however if you do publish details about some anonymization processes, this actually increases the risk of a breach.

Plan what happens next once you've shared and release the data. You should be monitoring use and you should be considering continuing to consider risk. Risk will change over time some elements of time will make the risk go down the data will be themselves getting older but actually some elements will increase the risk. New data are entering the data environment new technologies are available to do linkage and so on. I need to constantly keep looking at that. Also another issue is why you are not generally just considering releasing a single dataset this will be part of a series of releases as new data comes in on the same subject. So actually you're setting up a precedent and a policy which will then have to shift back from if you now consider the risk to be too high and that was itself need to be managed.

Finally you need to plan what you're going to do if things do in fact go wrong. Now we remember we have accepted that we're not working at zero risk therefore by definition there is a residual risk therefore there is a possibility that they'll be a breach and therefore you need to plan for that as part of your anonymisation strategy. Avoid the lure of catastrophisation.

These aren't the same sort of problems as a nuclear power station blowing up and it's very easy when thinking about data privacy to cast them in that light. Disclosure event mapping is a key thing right what is it that happens and if you've done your breach scenarios well you should have a map for that already. What is it that's going to happen and what's going to happen next. So it's not simply that there is an event and that's the end of it well actually something will happen next there'll be a player in the media there'll be your own communications whether you talk to anniversary if it's an actual attack that's led to the breach. The idea is to be active and be planned.

Ok to conclude the anonymisation decision-making framework is a tool which allows you to think constructively about your data situation. It moves closer to a harmonized idea of anonymisation which ties together the technical and the legal aspects of that process. Now we have an open-access book forthcoming and if you look on ukanon.net you'll see more information about the likely release date.

Thank you.