

Episode 60: Detecting Harm Early - Drug Early Warning Systems Explained

This transcript was created by AI from audio and lightly edited for readability. Treat it with caution, and assume occasional transcription errors may remain.

Transcript

Welcome back to the Tox Lab. I'm Rebecca. I'm Rob. This week we're gonna be taking a look at drug early warning systems. And I think when most people think of early warning systems, they think of klaxons and sirens.

And actually we're gonna be looking at the systems behind those and how they apply to drug harms. And really the big question we're going to be looking at this week is who do we listen to first when it comes to drug early warning systems? So before we dive in properly to the science, I wanna run a little thought experiment, Rebecca. Okay.

You're in a quiet Somerset village, doesn't have to be Somerset somewhere. If you don't take Somerset, I like cheese. The milkman says there's something going on in the village. People aren't opening their doors like they used to.

People's routines are disrupted. Okay. The vicar says church attendance is down, but we see this in winter, so it's nothing to worry about. Okay.

Who do you trust? The vicar? The upstanding member of the community. He had an explanation for what was going on.

He had an explanation exactly. And so therefore more weight goes on the vicar. Well, it turns out this village is actually riddled with flu and people were in fact staying in their house because they were unwell. The milkman noticed, however, people trusted the vicar that there was nothing going on because he provided an explanation.

But we're not saying the vicar was wrong. He just had a different view that was treated with more weight. And we're going to be looking at that idea a little bit later on about who we listen to and not just about who is right or who knows all the information, but who is best placed to identify something is up early. So before we jump into talking about early warning systems today, if you want to stay up to date with our content, you can find us on Instagram at @the_tox_lab.

We also have a page on LinkedIn and we're fairly active on there. And for those of you who don't use social media but still want to get in touch with episode suggestions or comments about episodes, you can email us at thetoxlab@gmail.com. So thinking very generally then to begin with, what is an early warning system? It's not alarms and sirens that tell you that a tsunami's coming.

They're part of it, but they're the final part. What we're going to look at is the systems that occur before that point. And crucially, because this is a toxicology podcast, how that applies to drugs, drug markets and drug harms. So the UN Office for Disaster Risk Reduction defines an early warning system as integrated systems of hazard monitoring, forecasting, risk assessment, communication and preparedness that enable timely action before harm occurs.

And the key idea of an early warning system is to act early to reduce harm rather than reacting after damage has been done. There are a few core principles of effective early warning systems. An early warning system should be integrated and use multiple data sources. They should be timely, action

orientated, people-centered.

And there are four kind of key pillars that we're going to be talking about today that underpin early warning systems as a whole. So this includes risk knowledge, detection, monitoring and forecasting, response capability and warning communication. So lots of nice management terms there. Really, in principle, what we're talking about then, aren't we?

Systems that detect signals, risk assess those signals, translate those signals into actions, all with the aim of reducing harm before harm has occurred. So we'll start by talking about pillar one, which is risk knowledge. And it's all about understanding what harms exist, who is affected and why. So in environmental early warning systems, which is what a lot of early warning systems are based on, this involves looking at past disasters, climate change, urbanization patterns.

And when looking at drugs, this information is informed by monitoring drug trends and horizon scanning to anticipate future risks. So by monitoring trends, we can look at what's already happening and help to establish a baseline of expected behaviors, harms and market features. And there are a variety of data sources that you can look at for this, which includes repeated cross-sectional surveys, drug related death data, hospitalizations and ambulance attendances, drug and alcohol treatment service data, wastewater analysis, syringe residue analysis, drug seizure and arrest data and digital drug markets and online community discussions.

A limitation with this is that trends do take time to emerge, this data is often very retrospective. Yeah, I mean, using drug related death data, for example, by the time it's been compiled, you're looking at data that's two or three years out of date sometimes. And although this is old data, it does still matter because it can help to inform policy development and service commissioning. It helps to provide context when those acute signals emerge, when we need to act, and it can help tailor population specific harm reduction messaging.

So as I mentioned a bit earlier, we can use horizon scanning to look ahead to identify emerging risks. And this kind of focuses on mega trends. So long-term global forces that might shape the drug market. So this can include climate change and environmental degradation, rapid technological advancement, migration and population displacement.

And these forces all affect drug supply chains, market volatility, patterns of drug use and access to treatment and harm reduction services. So let's put it in the current context of where we are now, sat here in the Southeast of the UK. It's about weighing up what we already know to try and predict what we think might happen in the future. So what have we seen over these last few years?

Well, we've seen an increase in novel opioids. We know we're seeing significant changes to heroin production coming out of Afghanistan. Considering facts like that to help inform where we focus our next pillar, is that what I've taken that to mean? Yeah.

This is one area where we as professionals working in the field probably have views that might actually be useful, mightn't they? We see the end result sometimes when we're looking at things like drug-related harms. But as a result, we probably do formulate opinions on where we need to be focusing our thoughts because we see at the cold face. So our next pillar is looking at detection, monitoring and forecasting.

So in an environmental early warning system, this pillar is aimed at identifying and tracking acute threats in near real time and predicting likelihood, severity and timing of disasters. So you can kind of think about like earthquake monitors, where they're looking at that data and predicting how strong an earthquake is gonna be and who might be impacted so you can kind of plan ahead as to who might need to be evacuated and things like that. So how does this relate to a drug early warning system? So

in a drug early warning system, we're basically looking at signals.

So a signal is data suggesting a new or changed relationship between a drug and a serious acute health risk. So signals may involve the appearance of new substances, unusual or severe clinical presentations, sudden spikes in overdoses or deaths, high potency or high dose products, drug substitution or mislabeling, contamination or adulteration, new routes or methods of drug use. So in other words, does a very short term effect precede a bigger harm? And that's what we're calling a signal.

Yeah. I think these signals are the crucial bit, aren't they? What constitutes a signal? Is a small cluster of overdoses something that needs to be concerned about?

Is a run of hospital admissions? Is a change in appearance of a drug enough to warrant early warning? This is very similar to trend monitoring, isn't it? But much shorter term data?

The key difference between looking at signals and trend monitoring is signals that aim to identify probable cause and effect relationships really quickly to kind of inform where there might be a change that is having acute harms that then needs to be actioned. Why don't you give us an example? So we can look at lots of different real time data sources for the signal detection. So we can look at analytically confirmed sources.

This includes drug checking services, wastewater analysis, syringe residue, emergency departmental psychology data, and ambulance and poison-centered data. So very similar to the sort of data we were looking at for the trend monitoring, but the most up-to-date versions of that that we can get. There are other non-analytically confirmed but rapid sources of data, and this can include anecdotal reports from people who use drugs and reports from frontline services. And these anecdotal systems are used in places like Scotland and New Zealand.

The problem with these is that there aren't any analytically confirmed results with this reporting, but it does often offer the earliest indication of an emerging harm. This is the postman versus the vicar.

Yeah. So in other words, then, all of this real-time data is monitored as part of these systems, and then decisions are made as to if what is being seen could constitute a signal worthy of early warning, which takes us to networks.

So pillar three looks at response capability and networks and readiness. So it refers to the ability to prepare for, respond to, and recover from hazardous events. So in environmental systems, these response plans are pre-designed and extensively tested. I'm sure a lot of us have had those days where they send out a national alert on our phones to kind of test the early warning system.

So in the event of a massive natural disaster, we all get the message and we can all take the appropriate action. In drug early warning systems though, the emphasis is on collaboration and coordination between different teams. And we want rapid data sharing, collective interpretation of these signals, so then we can decide if and how to respond either like issuing a drug alert, for example. And this is where things start to get messy, don't they?

Information sharing can sometimes be quite tricky depending on who holds what information and what that information is. Sometimes people, for confidentiality reasons, don't want to share it. There can be lots of good reasons for why people don't want to share information, in fact. So for a drug early warning system to be really effective, which is what we all want, it requires diverse expertise.

So you want to include in this community organizations, people who use drugs, toxicologists, addiction medicine specialists, epidemiologists, emergency clinicians, public health agencies, and law enforcement. Drug early warning systems can also operate at a variety of levels. So you can have them

at local levels, national, regional, and international levels.

And local systems tend to be more flexible and are more likely to include anecdotal data in the early warning system, whereas national or international systems are more formalized and often have quite strict evidence thresholds for what they consider to be a signal that's reliable to act on. So many sources of analytically confirmed data and things like that. And I think this is where we often find ourselves as toxicologists now working in a way that we never were before. Over the last 10 years, I've seen my role become more and more public health orientated.

That's clearly not the primary purpose of my job, but historically we had no input in public health systems and now we do. So our fourth and final pillar looks at warning communication. So in environmental disasters, this pillar converts threat data into clear, actionable messages. So like I was talking about with the alerts on your phone.

In drug early warning systems, this typically takes the form of public drug alerts. So what makes a really good drug alert? It needs to be clear, timely, actionable, relevant to real world drug use and including people who use drugs in the design of these alerts can help to increase trust and improve uptake of harm reduction strategies. And we can get different types of warnings and responses as well.

So we can have public facing alerts for communities and then those who use drugs, we can have technical alerts for clinicians and services, which is more appropriate for those groups. Regulatory responses can also include scheduling or rescheduling substances, putting in place temporary bans or developing new psychoactive substances legislation, which we've seen. And then there are also practical harm reduction responses as well. So for example, distributing naloxone or distributing drug checking tools or testing strips to communities.

So really then in summary, what is an early warning system? Systems that monitor what has occurred before us, try and look in real time data as much as possible for signals, communicate and assess those signals. And then if relevant, take action. And actually as toxicologists, we probably sit at all of those levels to a degree, don't we?

Possibly not number four. Number one, we've probably got an opinion on what has come before us and what may be ahead of us. Yeah, and we're definitely helping to provide the data that will be looked at for that pillar. Yeah, for sure.

Number two, possibly, depending on the work we're doing, we might be at the analytical front where we are detecting something new that could constitute a signal. Three, we need to be working with our colleagues to facilitate these within the confines of what we are allowed to share of course, because as I say, sometimes there are good reasons for not sharing information. Four, action. Well, obviously we're not gonna be putting drug alerts out or talking to individuals.

We might need to take action in terms of updating libraries. Yeah, that's definitely a really good point, especially for a lot of the new novel drugs. If we don't know to look for it, we're not gonna spot it. So making sure our libraries are up to date based on drug alerts that's gone out or information we have preceding that is always really useful.

So that's the basis of an early warning system and where as toxicologists, I think we fit into this. I'd love to hear from people if you've got other thoughts on how we fit into these models. Why don't we look then at a little bit more detail on the milkman versus vicar side of things, incorporation of people who use drugs in these systems? Yes, this paper was conceptualized by the AIVL, which is the Australian Injecting and Illicit Drug Use League.

And it's a national peer-led organization representing a network of peer-based harm reduction organizations for people who use drugs. And it also involved the PRN, which is Promote Response Network. This is an interdisciplinary and inter-sectoral national network of practitioners who contribute to early warning systems and drug monitoring. And the aim was to further explore the potential development of a community drug reporting system in Australia.

And also the aim to explore early warning system practitioners and community representative perspectives. And they wanted to kind of look at, one, the potential value of a community drug reporting system targeting people who use drugs. Two, the risks associated with that system. And three, the key principles and approaches to design and implementation.

So for this study, they held semi-structured interviews with individuals involved in early warning system activities. An advisory group with members of community-based and peer-led organizations for people who use drugs was also formed. And they advised on the approaches for developing a community drug reporting system. They then basically looked at these interviews and undertook a thematic analysis of the data.

So they had 20 interviewees and 14 advisors who predominantly supported the development of such a reporting system. And there was a general belief that, including community reports in drug-related early warning systems is a matter of principle and a broader movement towards recognizing experiential knowledge as a legitimate source of data. So many benefits were highlighted in these interviews, including the fact that community reports can accelerate the identification of drug-related concerns to inform more effective early warning system communications.

Creating a system for previously ad hoc reporting mechanisms may help mitigate the fragility of reporting arrangements based on personal relationships. And it could also help to detect concerns that may have previously been missed. So the general consensus was incorporation of people who use drugs into these systems actually was net benefit. So not just relying on scientific data as a signal, but actually incorporating experiential data.

A couple of examples that were brought up in these interviews, one of which was a cluster of nitricine overdoses in 2024 where the government was informed and an alert was put out and they end up getting confirmation of the nitricines through syringe testing. But this was all advised by a peer support worker who was working with people who use drugs who suggested that these things get done to analytically confirm what was being seen so an alert could be put out. So somebody right at the front line was working with people who use drugs who sensed something weird was off. They were the milkmen in the scenario.

And eventually, I feel like we've stretched this analogy a little bit too far. I think we might have. The vicar caught up and confirmed it in the lab. But crucially here, that signal came from people right at the front line.

And again, community reports have the ability to capture trends before they're captured by formal data systems. One example that was talked about was an increase in GHB use in those who had no prior experience of the drug and an increase in overdoses which were noted by community observations before the trend was observed by other sources. So again, some clear examples there of very early signals being picked up before more traditional systems have actually caught up and provided what we would call, in inverted commas, solid data.

And another example related to the purity of methamphetamine which was being questioned by the community of people who use drugs. And they felt that the purity was increasing but this took quite a while to be verified. There was lots of data saying that the purity hadn't changed. But eventually, that

was analytically confirmed.

And again, there was a delay in public health responses and putting out an alert and a lot of the community felt like their concerns were being dismissed. Which is a problem because they're the ones who are at risk and are trying to do something about it by letting people know and then being told that actually what they think they're seeing isn't correct. Yeah, that's a really good point actually. I hadn't considered the impact on ignoring people being that they are also the people impacted.

If I get ignored, it feels bad but it doesn't actually directly affect me. Yeah, no, that's a completely fair point. However, it was highlighted that community reports can lack some accuracy or specificity to directly inform a response and government public health departments were advising that an alert wouldn't be put out based on a community report alone. They didn't feel like that was enough of a solid set of information to be able to act on that.

That's interesting. So they provided helpful signals but there was still a consensus that these signals were not in and of themselves sufficiently strong to warrant an intervention. And I do kind of see their point where they lack a bit of accuracy and people aren't really sure what they're seeing. They know they're seeing something but they don't know what it is possibly putting out a vague alert could potentially do more harm than good.

That was going to be my next point actually. What is the downside of triggering the system and having an intervention? So there is a risk if you power a drug alert that turns out to not be anything quote unquote worthy of an alert that you could potentially damage the community's trust in the organisation and could potentially diminish its credibility.

Uncoordinated or overly reactive responses to community reports could also undermine the effectiveness of early warning system communications because if you're putting out tons and tons of alerts on things that don't need to be alerts then it just becomes background noise and people don't pay attention to the things you want them to be paying attention to. It's completely fair, isn't it? You sort of drown people with alerts. The effectiveness of the alert becomes less effective.

Yeah, for definite. And there could also be multiple reports related to the same single event. For example, like one particular overdose and if you've got multiple people reporting on it and it seems like it's a bigger cluster than it is then you are making more noise than you might necessarily need to. Yeah, that's fair.

So going back to a couple of final benefits that were highlighted from these interviews. Community reports can provide context to assisting complex assessments surrounding issuing an alert. Collecting community reports can also be used to help tackle misinformation through harm reduction education and community reports help to bridge the knowledge gap of novel psychoactive substances with an undocumented FOMCO kinetic and FOMCO dynamic profile. So by getting experiences from users you can learn how to assess a little bit of like how strong this drug might be and what its sort of effects are.

Whereas otherwise we'd have absolutely no idea what we were dealing with. So I did highlight in these interviews that there were implementation challenges and risks. So there were concerns over safety and wellbeing of individual reporters and how they might be affected coming forward and reporting. One, the fact they use drugs and two, the effects of those drugs.

Two, the effectiveness and sustainability of the reporting system itself. If you're getting a lot of reports coming in that you've then got to sift through all of that data. How sustainable is that going to be long term? And they were also looking at reputational risks for the organization responsible for the system.

And it seems to be a recurring theme with a lot of these groups dealing with different organizations about maintaining a good reputation so that when you're putting out that information it's considered credible and reliable. Then again they did point out there might be a risk to a wider drug-related early warning system if you are putting out a lot of alerts based on community reports that aren't credible or backed up. Importantly they did bring up some barriers to reporting, reach and equitable engagement. So the success of such a system may be contingent on its ability to engage a large and diverse cross-section of people who use drugs.

Drug criminalization is a major barrier to engagement in these systems. If individuals' reports are accessible to law enforcement and for law enforcement purposes that could lead to legal ramifications of people who use drugs and this risk would need to be mitigated to allow for participation. They also highlighted the importance of equity of access and they highlighted different groups who might be affected. So women individuals with different levels of technological literacy or access.

Those in remote communities, different cultures and language barriers and these can all pose a challenge to people being able to report to these systems. So you don't just need one way of reporting. You're going to have to have people who can report online via like phone call or text, email, post. You have to make sure it's as accessible as possible so you're capturing all the different people who are affected by drugs.

This is a fair point, isn't it? That there's sometimes people who use drugs is not one community. There could be lots of different groups of people under that broad category. And it's entirely possible for a system to focus on one group and not actually necessarily pick up a different group, for example.

It was also brought up by those being interviewed and the authors of this paper that the system would also need to be a system that can record neutral or beneficial slash pleasurable drug experiences rather than concerns. And this would be critical for engagement and would help to keep community interest. As probably rightly so, people who use drugs would be interested in this sort of information rather than just the negatives and the downsides of the drugs that are out there. This feels like a much more controversial suggestion.

Yes. Up to that point, I think there was no doubt. I don't know. What do you think?

I can see where they're coming from because you want to keep that community engaged in reporting to the system. And if they're only ever reporting when something is negative, they might, I don't know, they might be less inclined to report, but it does open up a massive issue with stakeholders who don't want to be associated with that side of drug use and kind of want to focus on like harm reduction and the negative surrounding it. So it does become quite a messy melting pot. But I do sort of see where they're coming from and that being able to capture those experiences might help keep communities engaged.

Yeah. Okay. Fair point. Because actually, I had another thought relating to drug alert risks actually that you put out a drug alert that says heroin in circulation is stronger than normal or is contaminated with nitazenes and therefore take harm reduction practices.

There is a risk that some users will seek out this deliberately stronger product. Oh, that's a fair point. I hadn't considered that. That it's seen as a quality signal, not as a warning signal.

That's likely going to be a minority of users, but there is a risk of that. So it's a really interesting idea, isn't it? The idea of using the thoughts and feelings of people who use drugs to act as that signal for early warning systems. And I think it's crucial that all people involved in the process do provide some input to that signal because I think it is really important.

Clearly, people who use drugs are the ones who are going to notice something is different in the drug supply long before the lab has tested it and confirmed it. I think all in all, it's a really, really hard system to try and run, isn't it? I occasionally get asked my opinion on signals. There'll be a cluster of overdoses and somebody will email me and say, what do you make of this, Rob?

And I'll provide my thoughts. I'm quite glad I don't have to be in that room to make the decision. Yeah. Because that I think is really hard to be able to weigh up the evidence and decide whether or not a drug alert is warranted, given the risks of what we've said about either the risks of diluting alerts or the risks of putting alerts out as a quality signal and things like that.

I'm quite glad I don't have to make that call because it must be really hard to weigh up. Yeah, for definite. And then there's always the risk that if you don't take action, that more people are negatively affected. And then there's a chance to put alerts out down the line, but you're always going to think the back of your mind should I have done something sooner.

But you are working with the information you have at the time that you're making the decision. So it is a really tough position to be in. And that was going to be my other thought, actually. So I've been involved in a signal.

I've spotted a new drug. I have informed the relevant people and said, this is a pattern I have noticed. Should we take action? No action was taken.

Subsequently, there were sadly more deaths related to this signal. And I carry that with me a little bit. You know, that always sits with me. Now, I did all I could within the systems.

People assessed it. And this isn't a criticism of the people. They assessed it and decided action shouldn't have been taken. And yet there was a negative outcome.

I do think you only see the negative outcomes because if you put a drug alert out, you will never know how many people that drug alert has saved versus if you make the decision not to, you will always see the cases like this where you think, oh, should I have put an alert out? So I do think you get a biased response in that sense. Yeah, you're definitely looking at it through a very skewed lens. But here is the question then, and we've touched on this a little bit.

Is it not okay to have a system that accepts you're going to have a few false positives? Or is the risk considered too great on credibility? Because this seems to be what it boils down to, doesn't it? Getting it right versus getting it fast, being fast versus being credible.

Yeah, and it always seems like you can't have the best of both worlds. I suppose you can't if you've got to try and analytically confirm everything that always takes time. I can take weeks sometimes. You get something new and you think, well, I think it's this, but I need a reference material.

You could be looking at weeks for a solid confirmation. Yes, I'm happy that this peak that I thought was this new drug is this new drug. I think it boils down to the specific information and the context you have at the time and all that information together, whether putting that information out there quickly versus waiting for it to be credible is more important. And I do think it's really context dependent.

A hundred percent. Yeah, no, it is definitely. So what's the UK picture like? We've got a little bit of fragmented systems, local systems, sometimes intersecting with broader organisations.

It's a bit of a mess, isn't it? Yeah. We are, I think, probably closer than we've ever been to a fully functioning early warning system. And certainly I'm seeing relationships with external stakeholders in a way that I never was ten years ago, but we're not quite there yet.

So I think as toxicologists working in these systems, it's an interesting time, isn't it? There's a lot going on out there. We see increasing reports from external service people about changes in drug patterns, changes in feeling. We probably need to put a fairly balanced hat on, don't we?

We need to listen to these signals, but also perhaps respond in a way that helps solidify that data. If we work in a hospital, for example, it's trying to roll out toxicology testing where appropriate and testing for new compounds and things like that. But that's easier said than done, obviously. Well, a lot to take home, I think, a lot to unpack.

And if you've got any thoughts on where toxicologists sit in early warning systems, do get in touch. Do let us know your thoughts on this or share this episode with people you think might be interested, because we'd really like to get the reach out there and hear people's thoughts and feelings on this emerging area. Hope you all have an amazing week, and we'll see you next time. Bye.