

In Conversation With...Kaveh G. Shojania, MD

September 1, 2011

In Conversation With...Kaveh G. Shojania, MD. PSNet [internet]. 2011.

<https://psnet.ahrq.gov/perspective/conversation-withkaveh-g-shojania-md>

Editor's note: *Kaveh G. Shojania, MD, is the Canada Research Chair in Patient Safety and Quality Improvement and the Director of the University of Toronto Centre for Patient Safety. Dr. Shojania is a leading expert on evidence-based patient safety strategies and translating research into practice. He led the team that produced the 2001 AHRQ evidence report, [Making Health Care Safer](#), and was awarded a 2004 John M. Eisenberg award for his contributions to safety research. This year, he became the editor of BMJ Quality and Safety (formerly Quality and Safety in Health Care), one of the world's leading safety journals. He has written extensively about many issues in safety, including the [role of incident reporting](#) and other [strategies for measuring errors](#) and harm.*

Dr. Robert Wachter, Editor, AHRQ WebM&M: What were the hopes for incident reporting systems when they first became an important part of building safety systems?

Dr. Kaveh G. Shojania: The most plausible hope was that by analyzing certain types of incidents, that you could learn things that would mitigate risks in the future. You're supposed to try to identify teamwork, communication, or patient identification problems and use the incident to get a handle on some problems in your institution. The plausible part of this hope is that some institutions would have idiosyncratic problems, which is to say problems other than those highlighted in the literature. An incident reporting system would be a great way of getting at these local, institution-specific problems, or ones not easily studied because they're rare. There were some other hopes, but the most plausible one is that you would get a handle on system problems usually through studying relatively uncommon, but potentially catastrophic, problems, just like in other high risk industries, where major errors are uncommon, making incident reporting the preferred approach.

RW: What problems have emerged over the last decade as we've looked at the recent history of incident reporting systems?

KS: The first problem is that they tend to be underused, and physicians almost never fill out incident reports—it's mostly nurses and sometimes pharmacists. There's a certain lack of engagement in the process. Another problem is that the wrong types of events tend to get reported. In most of the institutions I've ever worked, both in the United States and Canada, probably the single most commonly reported

incident is falls. Falls are a classic example of something that there's almost no point in reporting. It is almost an epidemiologic problem. It happens often enough that we should probably approach falls almost the way we approach infection control. What we really need is very targeted data collection about some of the known risk factors and mitigating circumstances and so on. You wouldn't want to do a root cause analysis about a fall every time it happens. Incident reports are really designed for things that are the equivalent of aviation catastrophes—uncommon but catastrophic events. So, I would say reporting the wrong type of events is another problem, meaning that incident reports are underused and, when they are used, often capture problems better suited to other strategies.

RW: Can you make that distinction a little clearer? What is it about a fall that's different than a wrong-site surgery that—assuming it was painless and costless—you wouldn't want detailed and nuanced individual information about the event to help you understand what to do?

KS: I've thought about this a fair amount. Falls are a common enough problem at every hospital that you should stop doing incident reports. Maybe another way to think about it is, we're ready to move on from incident reports. Another thing is that it's a common problem for which there isn't a single great solution, so it's not clear what an institution is supposed to do after analyzing one or more individual falls. But the more general point is that, once you decide this thing is prevalent, you stop doing isolated case reports and case analyses, and you start having more targeted reports. What we really need to know is almost like a checklist from a menu of risk factors and interventions. Were they on sedating medications? Was it dark? Had they had a history of falls? That's a very structured type of data collection. It's ironic that the single most commonly reported incident at most hospitals is not the type of event that incident reporting was designed to capture. You're really hoping to get things like wrong-site surgery, where it's so rare that it's precious information—every-defect-is-a-treasure kind of thing. So I guess it's a combination of it being a well-known, relatively frequent problem, and there's already a fair amount known about the risk factors for the event, so you can almost have a structured data collection. I'm just saying that you wouldn't want to do the whole aviation style, let's get the black box and find out exactly what happened when Mrs. Jones specifically fell. The identity of that person and the providers caring for her are not that important. You don't need to do a detailed interview of everybody to find out exactly what was going on at the hospital at that time the way you do when someone's wrong leg gets operated on. The same issue applies to many medication safety problems, which is another common category of incident reports. Many such reports, for instance ones involving opiates or anticoagulants, capture problems that occur frequently enough that analyses of individual incidents serve little purpose.

RW: What other concerns do you have?

KS: A frequent problem is that, even for serious stuff, despite all the problems with incident reporting, there still are a handful of really important problems that probably get reported, at least at academic hospitals. Even when you do have an incident that is worth investigating, by the time it comes across the right person's desk and the right things are set in motion, weeks to months have gone by. Then it's really difficult to do a root cause analysis. There's been a lot of misdirected activity around incident reporting systems in the last 5 to 10 years—let's get Web-based incident reporting and let's produce fancy pie charts. It really misses the point: None of the data from incident reporting as an aggregate is useful. A lot of institutions

make that mistake. Incident reporting was designed to identify important incidents that need to be investigated ASAP, so that people who knew what happened can be interviewed before they forget. Sometimes you can't even figure out who all the relevant people are that were caring for the patient. You almost need a hotline, which would be in many ways the ideal incident reporting system—leave an anonymous message about the name of the patient and why you're calling and let a SWAT team just descend on the case within 48 hours. Because otherwise, you'll never really find out what happened and any resulting root cause analysis will suffer as a result.

RW: There are lots of issues you're bringing up; is one of them just the volume? I can tell you at my hospital we get 15,000 incident reports a year now. And we do a lot more root cause analyses than we used to.

KS: There are two problems there, you're right. I guess I was talking simultaneously about two of them. As you know from some of the things we [published together](#) where we were helping institutions do a root cause analysis, even though we anonymize the institution, obviously behind the scenes we're doing a root cause analysis or helping them do one and then publishing it. We know that it's very difficult to track down what's happened if weeks to months have gone by. Even just for this first step of establishing the chronology of events, you need to be able to talk to people. It's sometimes surprisingly difficult to even know whom to talk to, and if you can identify them 6 months later, they're probably not going to remember a lot of what happened. So some of the value of a rapid response to an incident is actually just for the sake of doing a high quality investigation, but then there is this new volume overload issue. Interestingly, the volume overload issue happens elsewhere, not just in places like California, because at other hospitals they don't even have the staff that you have at UCSF. So if there's one person in the quality department, like at many hospitals, then even without a big push to have incident reporting and root cause analysis, they still end up overwhelmed. Incident reporting only captures probably 5% of the target events, but even with that a lot of hospitals experience information overload.

I have to say some of it is motivated by political considerations—every time someone is harmed, it will look bad if we don't do an investigation. But, from a true patient safety perspective, harm isn't necessarily the issue. Harm doesn't actually play a role in the value of an incident. I recognize, of course, that an institution can't be seen as not investigating something where someone got harmed. Of course you have to do some kind of investigation on some of these incidents just because the patient was harmed. That's an ethically good thing to do and a wise thing to do. But there does need to be some more safety-oriented framework for how you would prioritize the incidents that get the full guns approach. You could imagine what some of them would be, any time there's a wrong patient involved or a wrong medication, or any of the really serious or never events. In fact, the original definition of a "critical incident" was basically an incident involving the potential for system learning. And, this potential has nothing to do with harm; it has to do with the nature of the errors and system defenses against errors in a given situation. No one wants to look at a Swiss cheese model with just one slice of cheese and a couple of holes.

But the biggest limitation for incident reporting is not the incident reporting itself; it's the actual fixing part. It's relatively easy to increase the reporting rates if you really want to; you can go out and make a full court press and increase the incident reporting. The problem is that then you have to actually fix stuff, and that is

the hard part. It's very easy to always come up with another protocol, or have the nurses go to another educational in-service, or send out a memo to all the doctors reminding them to do such and such. You can always look like you're doing something. But if you actually want to fix these systems problems that come up over and over again, the deficiencies in teamwork and communication or concrete problems like patient identification issues, then you really do need to get the whole hospital to embrace taking this on. They're not going to do that 100 times a year, they're probably not even going to do it 40 times a year. So, it's sort of pointless to do a full court press on multiple investigations if you're not going to come up with solutions for them.

RW: The defenders would say that sure, you can only do the full court press only so often, but if you distribute the information about the incidents, that can help. For example, all of the medication errors go to somebody in the pharmacy whose role is to look at them and look for patterns, and when a pattern emerges take action, as opposed to necessarily reflexively jumping after each one. So that the institution is smarter and more likely to fix things than it would be if it had no idea that these things were happening.

KS: I'm speaking more about the root cause analysis marriage with incident reporting. I think that there is a role for managers in different departments to do something with that data. In general you want to look for patterns across incident reports. Otherwise, how are you going to see a latent problem that's come up over and over again? I've been surprised at how often institutions, even ones with their hearts in the right place, put a lot of effort into these things. There's really no formal system for looking across the incident investigations that they've done over the past 5 or 10 years. It often relies on individual memory. If the person who did the incident investigation is gone or no longer the head of the department of safety or quality, it's just random whether someone notices that this came up 2 years ago and there was this pattern emerging. In fact, we've been conducting a national study, interviewing people at about 30 or 40 different hospitals across Canada on their experiences with this type of thing. What we're learning is that they're mostly conducting isolated incident investigations and there is no great way to learn across them. For some of them like medications and falls, where if there can be more structured data gathering, it's actually a win-win. On one hand it would be inappropriate to investigate each of those in detail, so you might as well just capture the basic information, and on the other hand it allows you to review these incidents and look for patterns or any alarming trends, increases, and decreases.

RW: You mentioned this epidemiologic issue and the fact that these are not true rates. When you hear of a hospital and they tell you their incident reports have gone up 30% in the last year, is your bias that they're getting safer, less safe, or I don't know?

KS: Well, the received view is that they're getting safer, or the culture is improving. I certainly don't believe they're less safe. The naïve view that the safety field has tried to debunk is that these incident reports are not measures of true harm at an epidemiologic level, and, if anything, an increase in incidents is a better thing because it shows engagement. I would say that this is probably true, but the increased engagement won't last. If you increase your reporting rate in the short term, you better be careful because in the long term, if you don't actually fix things, people can get quite demoralized and quickly become disengaged.

Sometimes institutions make the mistake of not closing the loop with the incident reporter. The VA has done a very good job with getting back to the reporter. In fact, not only do they get back about

acknowledging and personally thanking them for doing the report, when they do their root cause analysis they actually run it by the person who reported it to get their input. It shows them that not only have they followed up on what they took the time to report, but they're even asking for their input based on their experience with the actual incident. But many institutions don't even get to that part, so people get demoralized after a year or two: "I keep filling out these reports, I never hear back from anybody, and now I'm not even going to bother filling out those reports anymore."

The initial spike is usually a good thing and shows more engagement and more attention to it. But you have to actually follow through and fix things. We see the same phenomenon with executive walk rounds. Some people go into the executive walk rounds with just the PR thing—like let's show that management cares. And you can get away with that for a year or two. But, after a while people notice that they keep complaining about these things at these walk rounds but no one is actually fixing things. And that ends up being demoralizing. It ends up almost being counter-productive. So the bottom line is that I think that the higher incident reporting rate is probably a good thing, at least as a surrogate marker for engagement and culture, but if you do increase things you have to actually follow through on fixing things.

RW: So just to summarize, if a hospital came to you and asked for a consult about what makes a good incident reporting system, it sounds like obviously a very strong follow-up procedure, not only in terms of listening and doing, creating a meaningful action plan to fix things but getting back to the reporter, and not using incident reporting as the only lens into your safety hazards.

KS: Sometimes it is worth the investment up front of somebody who's knowledgeable reviewing the incident, even just the actual briefest description of it within 24 to 48 hours and deciding what is best to be done with that incident. In other words, is this something that needs a really timely root cause analysis done right now—something that because of its severity and/or the likelihood that we're not going to be able to figure out what happened in 3 months that we need to get on the ball with this right now—or is this something that can go through some kind of more succinct root cause analysis? Some kind of triage of these, probably by the safety and quality department—they do this at some hospitals and people get good at doing this. The problem with that is that then it really does depend on the institutional culture, because then it can become some perverse incentive, a nurse manager or a physician chief may not want to say that some terrible thing happened on a unit and may just catalogue things. But however you decide to skin the cat, there does need to be a set of eyes looking at these reports sooner rather than later. At a place where you have thousands of them, the solution is to have them categorized electronically by the reporter as wrong site surgery, patient misidentification, or a major medication overdose. Once your volumes get high, you can't have somebody screening each of these, but I do think some combination of an automated screening and a human being up front could really enhance the process. That way, the full resources of your root cause analysis or quality department and the clinicians involved descend upon a relatively small number of the highest yield incidents.

RW: Is there anything particularly important about the technology that either helps support a good incident reporting system or can get in the way of it?

KS: I think people are often distracted by the technology. A lot of Web-based programs cost several hundred thousand dollars a year, and with just investing in that, hospitals sometimes feel like they did their

job. But the real work is almost independent of what you do with technology. With anything else, you could have the same thing achieved other ways. In Australia, they've invested in call centers—which would probably never work in the United States—where they have, for a whole region, a small cadre of people trained to receive these calls. Other hospitals in North America would just have a hotline, and other hospitals have made a paper-based system work well. I don't think the technology is so important. Although in the future there will be more hospitals where they have a borderline overwhelming number of incident reports, and there might be some neat stuff that can be done using a combination of data-mining software to enhance a human being's ability to see patterns across multiple incidents or identify incidents that need a quick or timely investigation. You can only do so much with the menu driven option. You may almost need a natural language search engine to be able to be trolling around in your 15,000 incidents and look for certain key words.

RW: I imagine that as technology gets better, so will the technology that captures the process of clinical care, potentially making a system in which the providers initiate all the reports not necessary, because they might get captured by the system.

KS: I've seen some neat things at my own institution where, to try to involve clinicians more, they're trying to embed these things within our electronic signout system. Whenever you do anything that involves physicians, the key is to make it really painless. But if there's a very simple way of initiating an incident report from within the sign out system, then the sign out system actually does have a lot of rich clinical information too. So sometimes someone on the safety/quality department side can actually do something with that information. Again, I think a lot of physicians wouldn't mind being contacted by someone a day or two later to say "we got your message and we want to hear a little more about this incident." What they don't want to do is fill out 20 different questions on exactly what was going on and do you think the patient was harmed and all this kind of stuff. They just want to let somebody know that something important happened, and if somebody in a position to fix it comes in contact with them later, that's okay.

RW: As you know, there's a movement now to create [Patient Safety Organizations](#). One of the rationales behind them is to create larger aggregations of organizations, multiple hospitals, or regions where in some ways the incidents are shared. If we have these problems learning from what's going on under one roof, does expanding the number of roofs involved in the system make this even more daunting, or are there opportunities that we'll see through the creation of these PSOs?

KS: That was one of the interesting themes that emerged when we interviewed safety experts about their views about monitoring and addressing internal safety problems—these experts had very interesting differences in opinion. There were a number of things that they thought the same thing, basically there was a commonality, but to what extent should incident reporting systems and root cause analyses be unit-specific and to what extent should they be very corporate or central to the organization? The arguments for the centrality and the commonality type of approach are obvious. Then you can pool across things, you can get a bird's eye view. From the senior management perspective obviously that's what they would like. But I'd say about half of the experts thought that a lot of stuff in medicine is very different in some clinical contexts than others. An obstetrician is not going to want to fill out the same type of form that a general internist might or an emergency doctor might and so on and so forth. I think the truth is somewhere in between. Some of the people I spoke with at the Brigham, for instance, have managed to come up with that

balance where there's a certain amount of stuff that has to be standardized in terms of how it's collected. But then there is an ability to allow people to respond to things that are more unit-specific. It's tempting to want to pool across these incidents and then across institutions. The reality is that there is probably a mixture of things that are quite idiosyncratic to an institution or even to a unit within an institution. And then some things are more common, human factors type stuff, especially equipment design. That's the type of thing where you really would want to pool across institutions because no institution is going to solve that problem on their own anyway. It would be really important to know if the same type of smart pump is causing the same type of problems at multiple institutions. They're going to need a whole group to go to the company and say, "Redesign your pump." Whereas a lot of other incidents and serious things like wrong-site surgery where there will be so many local factors that I'm not sure it's worth the effort of forcing standardization when a lot of the incidents will still not really be good learning opportunities for other institutions.