

## BroadcastMed | Grand Rounds: Ethics, Health Informatics and Translational Science

RICHARD SHARP: Good afternoon, everybody. I'm Richard Sharp, director of the Biomedical Ethics Research Program, and it's my privilege to introduce our speaker this afternoon, Dr. Ken Goodman. Dr. Goodman is the founder and director of the University of Miami Bioethics Program, and its Pan-American American Bioethics Initiative. He's also Professor of Medicine at the University of Miami, and holds joint appointments in the departments of philosophy, health informatics, public health sciences, and the Department of Electrical and Computer Engineering, as well as the School of Nursing and Health Studies, and the Department of Anesthesiology.

You can tell from that introduction alone that his interests are very broad and multi-disciplinary. In addition, I want to share with you, he is a remarkably accomplished figure in the field of biomedical ethics. His prior service activities, for example, include having chaired the ethics committee of AMIA, the American Medical Informatics Association. He's also a fellow of the College of Medical Informatics. In fact, to the best of my knowledge, he's still the only ethicist who's been elected as a fellow to the college.

He's the past chair of the American College of Epidemiology's Ethics Committee. More generally he writes widely on topics in biomedical ethics, and has published multiple volumes, or co-edited collections. He wrote one of the definitive books on a prominent case study, involving a young woman named Terry Schiavo, a very provocative and national case, in which Dr. Goodman was one of the leading voices from the ethics community throughout that saga.

His current work includes work examining ethical components of evidence-based medicine, public health ethics, and medical computing. Again, I can't stress to you enough the impressiveness of his resume and CV, and the influence that Dr. Goodman has had on a number of us in the field of medical ethics, as one of those leading voices. And it really is a privilege to welcome him here to Mayo Clinic, and I look forward very much to talk on health informatics, and translational science. Please join me in welcoming him.

[APPLAUSE]

**DR. KENNETH GOODMAN:** Thank you very much Dr. Sharp, you were very kind and it is a privilege to be here. I am from Miami. In fact, yesterday as I was driving here, I had the uncanny feeling that as I was making my way down the road, snow was melting. I think it probably was melting, my delusion was I was causing it.

Raising great issues, of course, about how we make inferences about the world, how we identify causal connections, as opposed to correlations. How we make observations, and how we try and systematically link them together in a way that's useful-- useful for patients, and useful for public health. When we do it systematically, and we do it right, we use machines. And these are intelligent machines, these are computers of one sort or another.

I used the phrase, "intelligent machine" intentionally, and not entirely to be provocative, although it shouldn't be that. The computers we have today, the data mining algorithms we use, the inferences they're able to draw are like nothing in the history of science. And in fact, we couldn't do half-- well, we couldn't do much-- far more than half of what we do now. Precious little work is done on three by five cards anymore, even though some of us really like to carry them around, they correspond to individual patients.

What we try and do is we try and figure out how to metabolize all this information. And, as you know, it raises many ethical issues. Our focus today is on translational science though, which is an opportunity, I believe, to try and do something really quite creative, and progressive, and exciting in the world of information technology, and the growth of knowledge.

Let me begin with disclosure slides. I have nothing to disclose that I'm gonna tell you about. We have these learning objectives. It's a very interesting question about whether or not an ethics talk should have measurable goals and objectives. We have an annual ethics conference, and every year I want to make sure there's CMEs available for it. And they say, well, the requirement is having measurable goals and objectives, and you have to measure behavior change. And I said at one point, if a physician were to change her behavior based on listening to me for an hour, she's probably made a mistake.

Saint Isidore, thought by some to be the last philosopher of antiquity-- you see, he lived in the sixth century-- took it upon himself, along with many others, to note, one that Western civilization seemed to be collapsing. And they were really quite concerned about how the artifacts of civilization would be preserved. And so they hired all of these amanuenses, people that copy manuscripts. And they copied paintings, they try to copy sculptures.

But when it came to music, he said, famously-- well, famously to me anyway, and wrongly, because others, in fact, had found a way to try and transcribe music-- that melodies cannot be written. It seemed too complicated, or there was no formalism to do it. That I believe-- someone will know better-- is from a Beethoven string quartet. In fact, we can do it. In fact, now we have we have machines that, if you have one of these, you plug it into your electronic keyboard, if you type a note, it will actually display the note, and you can write music without writing anything.

Against that background, bear in mind that I-- as all people who give talks about information technology-- will use metaphors. And the metaphors become so real to us, we forget that they are metaphors. So when someone says there's a firewall, or when someone says they're using a cloud, or when someone says they've written something. In fact, they may have done-- and those are no such thing, we're using the word in a different way. And it carries with it, by the way, the meaning that it's metaphoric on, which might or might not lead us to some mischief.

Here's what I hope to cover. I want to talk about the role of standards as ethical issues. The very important role of standards, and for our purposes, in biomedical informatics. I want to talk a little bit about what I think is a wildly neglected area at the intersection of science and ethics. Namely, the work that our colleagues in laboratories do when they write code.

If anybody here is responsible, or works with the RCR program, the Responsible Conduct of Research program, which we all have come to know and love-- A, because we really came to know and love it, B, because we thought it was intellectually compelling, C, because we really wanted this for our graduates students, and/or D, because National Institutes of Health requires it. In any case, I think an under-addressed, if not completely neglected area of RCR curriculum is software engineering.

A learning health care system is another metaphor. And in fact, we are trying now to figure out how to take our systems, our hospitals, our clinics, our public health data, and to merge it in ways that are useful. As you know, the eMERGE network that the Mayo participates in is doing some very important work in trying to figure out how best to make use of genetic information, genomic information, in electronic health records.

And then I'm not sure whether ethics is better when it's provocative or not, but I hope to be a little provocative about the duties of our citizens. We worry a great deal about our duties, we care about ethics, we want to make sure that our colleagues who do science and practice are mindful of the ethical issues that their work raises. But I also want to suggest that the people we do this for-- the populations and patients that we care for-- have obligations as well. Obligations which they can very easily discharge.

And then to conclude with a discussion of system interoperability. That's a geeky term from informatics, but it means a lot in our context. Right now, if a patient from Miami were to show up here, and someone said, well, why don't you send her record over, someone is going to have to print it, and scan it, and email it, or fax it. I find that state of affairs in 2015 to be striking. And we'll talk a little bit more about that. All of these I want to be understood as ethical issues, to broaden the discourse, to make sure that in translational science, we're mindful of the very important roles of health information technology, which, in fact, themselves raise very interesting ethical issues. OK.

So standards as ethics. For instance, here is a group, HL7, Health Level 7, is trying to write standards, along with everybody else, for functional electronic health record systems. They involve decision support, they involve privacy, they involve the representation of information, that sort of thing. There's an old joke in informatics, it says that the problem is not having standards, we've got lots of standards.

In fact, trying to get our heads and hands around what, in fact, should be standardized is a great debate. We were talking last night with the famous example of why it is standards matter. And it has to do with the Baltimore fire of 1904, not that I remember, thank you very much. But the idea was this-- somebody is walking down the street in Baltimore, flips a cigar on the sidewalk, it falls into a grating on the sidewalk, and six hours later, Baltimore is on fire.

They send out a call-- dit, dit, dit-- for fire companies, and they came from Washington, they came from New York, they came from all over the middle Eastern seaboard. And they arrived in their fire trucks, and they got out with their hoses, and they walked over to the fire hydrants, and their hoses didn't fit the hydrants. And Baltimore burned to the ground.

So do we want to say, well, you wouldn't want to regulate fire hose standards, because that would have a chilling effect on creativity, and fire hose development. I don't mean to be snide, but we hear this all the time in other domains. Would you want to say the creative genius of fire hose thread manufacturers and designers was being impeded by an attempt to regulate their product, and so forth?

No, we wouldn't say such a thing. We now regulate fire hose couplings. And when you think of all the things that work, you see that most industries are far ahead of the informatics industry in developing standards which are either voluntarily adopted, or adopted as a condition of certain other benefits.

This is not a convenience. The example from Baltimore is that cities burned and people died. That when it comes to information technology, we are talking about a number of important values-- trust of the people whose information is there, their safety, the quality of our work and our research, whether it works or not, and so forth.

I want to cast patients for you as not mere stakeholder-- stakeholder, by the way, [INAUDIBLE]-- would be a metaphor. That they're not mere stakeholders, that they are drivers of the engine. If we mean what we say, when we say that our best efforts, our excellence, is devoted to patients, and to improving the care of populations, then I think we have a collective opportunity, an obligation, to make them fuller partners in this, and to listen to them.

That if we want to be able to move toward evidence-based, patient-centered systems-- that sounds like a slogan, forgive me for it, but no one is going to argue against it-- you get the idea, we want to make sure that the best available evidence is available at the bedside. And moreover, that we do this not for the sake of-- I don't mind systems that improve our ability to bill, or to schedule, or to code-- but at the end of the day, the reason we should be doing it, is because we believe, on empirical grounds, it improves patient care.

And for the philosophers who have the taste, or the stomach for this sort of thing, I'm of the view that there are uncontroversial both utilitarian and rights-based foundations for this. Normally when Dr. Sharp and I-- if we were inclined to disagree-- one of the things that we'd say is, if we disagree with each other, at least one of us is wrong. We can't be two right answers. And moreover, though, that when, according to the largest and most powerful tools of doing moral philosophy are in alignment, then you probably got the right answer. There's more to be said about that at the ethics workshop, that's elsewhere.

Here are practical examples. At an institution with which I have a relationship-- that's vague, because in Miami there are several of them, and so you shouldn't infer which one-- they have a fetal monitor. And the fetal monitor is manufactured by Corporation A. It dawns on me that we have a lot of CME rules for disclosure in mentioning of drugs and devices. And in some very interesting sense, there's a question about whether the Food and Drug Administration should regulate electronic health records as medical devices.

If the answer to that ends up being yes, then the companies that manufacture them then become like the companies that manufacture drugs. Which some people think should happen. But it also means for the sake of a presentation like this, I don't want to take a chance, so I'm not going to mention any of their names.

So Company A makes a fetal monitor, Company B manufacturers the electronic health record. Same hospital, same patients. How do we get the output from the fetal monitor into the electronic health record? It should, of course, be seamless, and electronic. It flows from the fetal mon-- someone might want to click it to edit it first, or somehow make sure that something adds some information, some added data perhaps. And there you go, it automatically flows into the electronic health record. It does not.

What happens is someone prints the output from the fetal monitor, and hands it to someone else, who scans it, makes a PDF of it, and the image of the fetal monitor output-- we've become so used to doing that, that it actually seems like a reasonable thing to do. It was, like the story I told you, everybody comes to realize their vision is failing in different sorts of ways. Sometimes it's the back of a classroom.

In my case, it was when I pulled over, walked out of the car, walked up to the street sign to read it, and then I said to myself, you know something, this is probably aberrant. I had gotten used to compensating for-- we're doing the same thing now. And we are not insisting it be any other way. We're soldiering on, like the good soldiers that we are, and tolerating a situation where, well, fetal monitor data is very clumsily included in the record that we actually use to keep track for our patients.

So there you go, that's A and B. I was looking to try-- so the relationship between A and B is, if you go into windings, you can find lots of arrows that are not exactly-- images of arrows that don't get you from here to there. I had a great deal of fun looking those up.

Here's another example. I have spent a great deal of time at the hospitals with which I have a relationship, trying to say, we need to do a better job representing code status in the electronic health record. If, for example, it says full code, or no code, that's useful to an extent. But maybe I'd like to see an advance directive. Maybe I'd like to see some of the supporting documentation.

If you use a post or a most form, or something like that, then surely I should be able to find that, in the same way that I find every thing else that seems to be related to what I'm looking at. And I have spent no small amount of time in two major electronic health record systems, trying to arrange for it to be easy to find supporting documentation, to guide a clinician in deciding what to do if a patient's heart stops. This should not be difficult.

And yet, we have no standards for it, it's completely idiosyncratic. As we transfer patients increasingly around, or share their information, that strikes me as an interesting bit of electronic primitivism in our environment. So there are-- pardon me, one second. It's much more humid in Miami, I've noticed.

Like Oscar Wilde, I can resist everything but temptation. Here are examples of failure to plan, failure to note where the parameters were actually located, failure to be attentive to risks that are forthcoming, and failure to use tools appropriately.

[LAUGHING]

There are two ways of thinking about ethics. One of them is prohibitive. We very often love to warn our colleagues-- stop, slow down, don't do that. And that, I think there are times and that's absolutely appropriate. We sometimes share the concern that our colleagues want the ethics community to be in that role. And we work with our colleagues in compliance all the time.

But I want you to think of another role for ethics, one that actually says, you know, if a tool can actually improve patient care, or the health of populations, or foster a learning health care system, or support our aspirations in translational science, then maybe the thing to do is not stop or slow down. Maybe the thing is to begin studying it more rapidly, and begin to figure out ways that we can use it effectively.

And so you end up on the other side, there are a number of ethical duties as prescriptive. We identify them all the time-- reduce disparities, I assume no one disagrees with that. Foster and raise minimum standards, protect rights. As an homage to one of our colleagues-- there was a group of physicians and philosophers at Dartmouth in the '70s and '80s, the last of whom recently died. One of them, Bernie Gert, was a philosopher famous for saying, when it came to ethics-- "You know most of what I'm already going to tell you," he'd say. "And if I seem to say anything profound, you probably misunderstood me."

It gets tricky for us, however, when we try and do both at the same time. And so we have framed this idea, namely that progress is not unethical, it's an obligation. But we have a duty to somehow manage it in ways that support those core values. Progressive caution is the name that I think is useful to think of that.

Now I want to talk a little bit about software engineering. Anybody here write in Python or R? Thanks. When you do that, you are providing a service to yourself, to your colleagues. You're advancing biomedical knowledge in ways that most of your colleagues in the hospital have no idea about.

And yet, but for you at that level, in the bench world where you live, you worry about version control, provenance, attestation, curating. We have colleagues that know a lot about genomics, but they don't know so much about curating, right? When it comes to what you do with samples. We have tried to part-- I don't know if you've heard of Software Carpentry, it's not-for-profit group.

But it turns out, the most cited PLOS biology article in the last couple of years was by a group, they were called Best Practices for Scientific Computing. Very interesting, as the people who write the code that manages the data that's the foundation for translational science are a community unknown to most people who do it. And yet the ethical issues that arise in their work, I believe, bear on the current kerfuffle we're in related to reproducibility, and corroboration, and confirmation of results.

The philosophers in the room have been thinking about causation and confirmation for a very long time. That we've discovered that a lot of the research that we're paying for actually can't get replicated, has caused us to be concerned. I'm of the view, by the way, that we're not so much interested-- we ought not so much to be interested in reproducibility and replication, as corroboration. And it seems that we can't even do that.

I think therefore, it's a turtles all the way down issue, and we need to begin with our colleagues who are trying to write code, to make sense of the data generated by laboratories. And that leads, of course-- and while we now, in the world of bioresearch, talk about trust, and trust of our communities, and community engagement, our colleagues in information technology have been talking about trust for decades.

If I don't understand your code-- if I trusted you to come up with something that would work-- then something has gone wrong. Sometimes it can be very personal, sometimes it can be institutional, and sometimes it can be social. And so biomedical research is following the path of information technology, and discovering that at some very deep and profound level, the most important value on the table is not privacy, it might not be access, it might not be protecting human subjects. It might be something as simple as trust. Simple to utter, very difficult to earn, and to maintain.

I mention pareidolia and decision support-- that's just a fancy word for seeing bunnies in clouds. That is to say, pareidolia is where you look at something, and you impose on it a pattern-- that our brain tends to want to do-- that makes it look like something that was intentionally there, but which is not. So if you see a face on Mars, or a bunny in the clouds, or a pattern in your data, you need to make really sure-- given our really cool, powerful, and interest-generating patterns that data mining software generates-- that we're actually finding something that matters.

And I'm convinced that the writing of the software, along with education about these, is a crucial way to try and make clear that we're not seeing bunnies in clouds. By which, clouds, I mean, fluffy white ones in the sky, right? My favorite example of this is actually, it's a trick due to Wittgenstein. Forgive me the philosophical reference.

I'd like everybody to work with me now. I'm going to utter a series of positive integers, and I like you all, at the same time, to complete the series. Two, four, six, eight--

**DR. KENNETH GOODMAN:** Wrong. You see, it's 11. By which I mean, you increase by two three times, and then you increase by three. 11, 14-- and suddenly you see the series. We didn't know what the series was, and in some sense, if it's a complex set of phenomenon, you don't know what the series is until you get to the end of it. Well, that's too late to do us any good.

We have discovered that looking for patterns can be illusory. But it's also what we do, it's why identifying causal connections is so important as we collect data. But what we are so beguiled sometimes by our own tools that we forget these basic rules about following rules.

I used this example with Dr. Pullman earlier, it just occurred to me. I believe that we should carefully, and with great oversight, use animals for biomedical research. I'm not an anti-vivisectionist. Other people disagree with that position, and it's a reasonable position to take, I understand it. I just believe that it's mistaken. I believe that the work we do involving animals, especially with the controls in place, is important.

You would not, however, allow someone who opposed the use of animals to serve on the Institutional Animal Care and Use Committee-- which is looking at how to protect animals-- because they're opposed to it unutterably. There's no beginning, there's no end. They can't participate in a process that they believe is inherently wrong.

Fast forward to privacy. If I were of the belief that everybody needs to give their consent for every use of their information, that I would be of a belief that brands me as a privacy advocate, but as a privacy-- I want to use my words carefully-- purist, if you will. You might say zealot, or extremist, I don't want to say those words.

And that as we try and identify sources for translational science, as we try and get our heads around learning health care systems, and as we realize it's all information-intensive, we are now stuck with some very interesting problems. Namely, how do we use information collected for one purpose, for another? And how much consent do we need in the process? I want you to feel a vague discomfort about that for a while, and I'll return to it.

We have, for years, been well-served by the idea that data capture in the clinic is clinical data, and data collected during research is research data, and data collected by our colleagues in public health is surveillance data. Two of the three, the first two, actually require review by human subject, by an IRB, right? We make exceptions for public health data, although what someone in public health does for surveillance may be completely indistinguishable from what someone does and we call it research.

I move the view this distinction is no longer useful. And in fact, it is impeding the growth of knowledge that I think will be very important for growth of populations. I believe that if we do it right-- and we enjoy the trust of these people we keep saying that we're doing it for-- that in fact that distinction, when it's eliminated, will produce lots and lots of data which will, in fact, drive the biomedical research engine of the next several decades, if not longer. It's secondary use is a misnomer. Our colleagues in public health will tell us, what if we collected it for public health? Then using it for public health can't be secondary, right?

At what point are we not then gathering data in principle of use to public health? We want early warning systems for infection control in our hospitals. We collect data all the time from our communities for syndromic surveillance, and for bioterrorism, for that matter. We are constantly collecting data for all sorts of purposes.

And teasing out the reason that we gather a datum from one, as opposed to the other, I think, is increasingly not productive, not very useful scientifically, and I do not believe it is supported anymore by the way it used to be by the ethics community. That matters, because what it means is we're going to have a lot more data, and a lot more challenges for use of it-- for a secondary use-- it's very hard to give up.

What that means is-- this is the part where I'm uncomfortable even uttering it myself, it is provocative. And it's this, everybody in this room has benefited from other people's information being used. Everyone. If we could use your information to return to everyone else, does anybody want to say no to that? I'm not going to ask for a show of hands.

What you're thinking of is, well, but the literature shows that people actually like to be asked. The literature shows they like to know what is being put to their information. And all of that is true. But there's also a literature that shows that, in many respects, people assume that we-- namely, trusted agents in government, and academia, and public health-- are doing it already.

By the way, when I said trusted agents and government and academia, I was not being snide. We live in an era where our colleagues in government public health service do great work, they save lives by the millions. And the fact that they happen to do it because they work for a government is cause for some people to be to be concerned, strikes me as sad and inappropriate. I thank them for their service. As well should everyone.

The challenge now is this. If, in fact, we've benefited, then based on the old rule of what sort of goes around comes around, what reason can I offer, to say, I don't want you to use my information? I want to be able to opt out. And all of our institutions are going through great gyrations now, for the sake of translational science, to figure out how to create registries, how to create systems where patients can more or less nimbly opt out. What kind of consent can be presumed?

Or, as I-- I'm going to go back a section, a minute, to implied, or latent consent. I like the idea of latent consent. I am of the view that most ordinary people-- as has been shown by some interesting research-- already assume their information is being used for the health of populations and for research. They just assume that we can't tell easily that it's them, that it's Ken at 123 Elm Street. And by and large, it's true. We cannot tell easily that it's Ken at 123 Elm Street.

So if I want to opt out for the sake of public health surveillance, we would just say, no, you can't do such a thing. I was born I was died--- I was died. I died, I had a baby, I got Kuru, or Ebola. I don't get to say, I don't want anyone making a note of that, and having a number change in the counter. That's information about me, you need my permission to do that. What would we say about that?

We say that someone who argued that didn't understand vital statistics, going back, by the way, to the great John Graunt of London. If you've never done this, please Google Graunt, G-R-A-U-N-T. Actuarially, what caused people to die in London 1630-something-or-other? It's both telling, and it's really quite striking, about what caused people to die in London back in the day. The first actuarial tables used by insurance companies.

Now, I'm suggesting that for me to not want to contribute to that, in some acceptably de-identified way-- makes me akin to those people who would say, yeah, if I were dying, or my kid were, or my mom was, I'd sure hope to be able to transplant a kidney. But I think that's yucky, and I don't want to be a donor. Anyone who says I'd be happy to accept a kidney but don't want to be a donor is cheating.

Anybody who says, I really am glad you all got your kids vaccinated because I don't want get mine vaccinated, thank you very much-- I'm cheating, I become a free-rider. And so I want to explore with you, at least a little, the provocative idea of information free-riders. Namely that as we toil for our patients, and we toil for our communities, and we toil for translational science, I'm of the view that some of those ordinary people, at very low cost to them, can make it easy for us to study their data.

Now the last item there is, I want to pass, infrastructure support refusers. That almost sounds like something-- I have no political views whatsoever-- but I actually do enjoy paying my taxes. I do, I have a duty to do it. I enjoy paying taxes, because I contribute to civil society that way. I contribute to the excellence of institutions. When the Mayo Clinic gets an NIH grant, and does good work with it-- and my taxes paid for that grant-- I say yes, something's going right in civil society. I don't get to say, I don't want to pay taxes, but please make sure you share the results of that great research with me.

So it turns out that I think we have an opportunity here. Once again, I want to make sure you understand, this is not an anti-privacy argument. I think we can still do a really good job protecting privacy. In fact, as it is now, we don't do such a good job protecting privacy. Did the local newspaper have the ad from the company that the hackers took all the medical record numbers from?

I don't want to mention the name of it either, it's an insurance company, and somebody hacked their database, 80 million records. The largest data breach in the history of hominid evolution. Under HIPAA, of course, they're required to disclose this to everybody, so they've got full-page ads in newspapers around the country. Because they have to do it, right?

They didn't get medical information, they got medical record numbers, they got insurance information, maybe some credit cards, social security numbers, they were interested in identity theft. They were interested in being able to get data that medical records have. The hacker who hacks your medical record doesn't care that you have left toe boo-boos.

What the hacker who hacks your record wants to do is get enough information to try and either use your credit-- even credit card numbers are cheap, you my credit card number on the internet, on the black net for a couple of dollars. Medical records are worth more on the darknet, because it's got all these other information. There is a black market in medical records. Who'd have thunk it?

Should we be responsible for prophylaxis against all instances of that? I think it's an unfair burden to put on the research, and the health, and the nursing and medical communities. I also think that we have we have come to adopt two standards. One of them is this-- I hypothesize I'm not the only person who's said, I know someone is checking my clickstream, but I really just want to buy this thing now. Or I know when I stick my credit card in the machine that there's a bunch of metadata going with that, and someone's analyzing it.

I know that most of the mail I get is based on data-mining software that analyzed my purchasing patterns. Why am I utterly-- or why are we-- utterly sanguine about that? We shrug, and we say yes. But if someone says, "We'd like to use your de-identified data to try and reduce the incidence of pandemic flu this season," we get our knickers in a twist. How did that happen?

We are utterly blasé about banks-- nothing against banks-- credit card companies, love them dearly, financial corporations, I think we've rescued the economy since then, all of the things that involve financial transactions, we shrug at. But if someone wants to save a baby, we get we get all excited. I use saving babies not hyperbolically. The first example of systematic use of de-identified data, that I think is beautiful for this purpose-- well, it may be John Graunt, it may be John Snow, it may be any of the epidemiologists who systematically gathered data-- but I like the people who analyzed all the data from automobile accidents involving unrestrained children.

Ugly research, because what happens to a small person in a car that suddenly stops when it's going 60 miles an hour is not good. But the people who analyzed those data, without anybody's permission whatsoever, have come up with tools that saved the lives of thousands and thousands of children a year. Anybody have a problem with that?

Therefore, how do we have our cake and eat it too? That is to say, if Ken's proved too much, we start using all of these data willy-nilly, then we will have a problem with trust. So we have a bit of a disconnect there. We need to make sure that even if Ken is right about information free-riders, we need to do a good job with community engagement, and fostering the trust of the communities we purport to serve.

They need to understand what we do with that information. They need to understand that if somebody is able to, by very clever software, to be able to reconnect me with a particular data, which I'd rather they not have done. Where they're able to even take an anonymous tube of blood, as has been done, and connect it with a particular surname. If that sort of thing erodes confidence, we need to do a better job reminding people that, one, we can't guarantee that some hacker in the Aldebaran cluster is not going to somehow get into your records, any more than you can guarantee that for the bank.

But that's not a good reason for you to say you don't want to support the growth of biomedical knowledge. And we need to do that community engagement on a regular basis. In the same way, by the way-- I didn't intend to mention animal research-- but those of us who've talked to people about animal research and the ethics underlying it, also said our colleagues in the laboratories, find time to go talk to schoolkids about why you do research with animals. Find time to make sure that people understand why you're not torturing bunnies for fun.

Find time to make it clear that ordinary people, who are paying for this research, despite what they might have in the form of some moral discomfort-- which needs to be acknowledged-- understand the reasons why we justify the use of animals in research. You scientists need to do a better job communicating what you do to lay publics. I don't think they have, and we muddle through.

I think we have an opportunity here, as part of translational science, to do a better job of that. Do a better job of community engagement, and educating people about these very large databases that we use to improve decision support. It turns out that decision support ranges from alarms in hospitals, to drug interaction warnings in electronic health records, to diagnoses-- including a generation of differential diagnoses-- to prognostic scoring systems, including those will predict the likelihood of you surviving a hospitalization. In other words, these are data-intense operations.

Two, exception management. I like that-- it's related to incidental findings, or surprising scientific results-- that once we get a better handle on all of this, we will know what to make of incidental findings. If you're analyzing my genome, and you discover something that might be clinically actionable-- whatever that means, by the way-- but I'm not expecting you to tell me about it, well, we need a process in place. And that's some of the things that you're working on, right?

I'm of the view that once we wire all our electronic health records together, when it comes to incidental findings, you ain't seen nothing yet. We need to know when to break the glass, we need to know when there is an exception, and we need to know when some other value is more important than privacy, or privacy might be more important than some other value. We need to figure out how to take all of those data streams-- the monitoring, the surveillance, the research-- and recognized that compared to every other one, they're all secondary. And to applying this for all of these purposes.

Therefore, what we're all working on as well, then, is systems of trusted governance. Trust being the operational word there. That is to say, how can we put in place a system of checks and balances that are good enough to win and maintain the support of communities, even as we use people's information repeatedly, and without fine-grained consent?

So learning health care systems require that we share lots of data. The patients assume this. And it's needed for the kind of root practice, public health, and translational science that we all are embracing. Significantly-- I don't know how many patients know about learning health care systems-- but when you explain it to them, they say, yeah, sure. I mean, that's what I'm counting on you guys to do. That's why I like being patient at the Mayo Clinic, or the University of Miami Hospital. I trust you guys. We need to make sure that that's well-placed. I believe it is, but we need to continue to feed and water it.

OK, a couple of other practical big problems about this. One of them is-- for those of you familiar with meta analysis-- where the problem we have now, the reason why we created [clinicaltrials.gov](http://clinicaltrials.gov). We know of a lot of research. In fact, in the news, in the last couple of days, there have been some articles about how many research studies are still not being posted to clinical trials. In other words, there's research that's being conducted off the radar.

That's because when it ends up showing nothing, we're not embarrassed that we didn't accomplish anything, right? Or for whatever reason. That used to be the file drawer problem. Did a study, a large, well-designed study that didn't show much of anything. And we knew that editors, rumored to be human, don't like to publish negative or neutral results.

They all want positive results, therefore I won't even bother, and I put my beautiful manuscript, my large, well-designed trial manuscript in the drawer. That, among other things, had a problem of confounding many meta analyzes, which we rely on in many cases for resolving some scientific disputes. Or when there's a question about whether the randomized large-scale, double-blind, placebo-controlled trials are adequate to the task.

We have a file drawer problem that's built into the very wires-- or wireless networks-- that connect our institution, if in fact, we're not using that information. It's occulting information, which, if it's gathered carefully, can be quite useful. If we want a learning health system to make sense of all this big data, then-- as I promised you, we'll close with-- interoperability is a necessary condition, and a corrective to the reliance on flawed publication records, and other sorts of things. We need, for the sake of making it so our systems can talk to each other, to make sure-- and get the benefits of that-- that we eliminate those impediments.

And therefore, there's a bit of a syllogism. If interoperability improves outcomes, it becomes a duty to achieve it. If it reduces cost, increases safety, quality, and efficacy, then it becomes a duty to do it. If it fosters trust in our systems-- and you can complete that sentence for me. Unlike numbers, sometimes completing sentences are much constrained by syntax. We have time for a few questions.