

In Conversation With... Mark L. Graber, MD

January 1, 2016

In Conversation With... Mark L. Graber, MD. PSNet [internet]. 2016.

<https://psnet.ahrq.gov/perspective/conversation-mark-l-graber-md>

Editor's note: *Dr. Graber is a Senior Fellow at RTI International and Professor Emeritus of Medicine at the State University of New York at Stony Brook. He is a national leader in the field of patient safety and has pioneered efforts to address diagnostic errors. He founded the [Society to Improve Diagnosis in Medicine](#) and the journal, [Diagnosis](#). We spoke with him about the recent National Academy of Medicine (formerly Institute of Medicine [IOM]) [Improving Diagnosis in Health Care](#) report, and about diagnostic errors more generally.*

Dr. Robert M. Wachter: What got you interested in [diagnostic errors](#)?

Dr. Mark L. Graber: I'm a nephrologist by training and background, and in nephrology we get consults all the time on things that turn out not to be real diseases. A patient may appear to have hyperkalemia, but really it was just a case of hemolysis. Or their labs may suggest hyponatremia, but it was really that their lipid levels were too high, or an IV was running when the sample was drawn. Those seem like silly things, but clinicians were treating patients for disorders that they didn't really have. That was an introduction to diagnostic errors.

More importantly, I was the Chief of Medicine at a VA Hospital, and in our peer review program, we were seeing all sorts of diagnostic errors that kept recurring year after year after year. For example, diagnosis of lung cancer cases that should have been made in 3 or 4 months but stretched out over 9 or 12 or 15 months, or mistakes reading X-rays in the emergency department because there was no radiologist on call. It struck me that this was a patient safety problem, just like all the others we were dealing with at that time but one that hadn't been discussed.

RW: What were your hypotheses about why it wasn't having the same resonance in the patient safety field and being discussed as much as some of the other kinds of harms?

MG: There are many possible explanations. One is that diagnostic errors are not quite as obvious as a fall or a medication error. They play out over time and are hard to define. There is often no agreement on whether an error occurred or not. Some people view it as just an acceptable variation in practice that it took them 4 months to diagnose something that somebody else might have diagnosed in 2 months.

It's my own belief that the major factor was the perception that we were doing okay in regard to diagnosis, and I think that's been the assumption for many decades. It's very understandable because I think we are. We get it right more than 90% of the time, and that's really amazing given the complexity of the task and the fact that there is uncertainty at every step.

RW: Although the data from malpractice cases and other sources show that even if we get it right 90% or 95% of the time, that's a huge burden of harm from diagnostic errors.

MG: You're absolutely right. The best estimate from autopsy studies is that there are 40,000 to 80,000 deaths a year from diagnostic error. Every physician knows about malpractice cases—maybe they've been unfortunate enough to be involved in one—but they generally assume it's the other guy who's not as good or not as careful, or it's the other hospital where things are not as well organized. So there is a tremendous amount of overconfidence in medicine today. Part of the reason is that we don't have autopsies anymore, and our patients that were harmed might just go elsewhere; we never really hear from them again. There are way too many patients being harmed from diagnostic error, and the problem needs to be addressed. If you ask physicians if they can remember their last diagnostic error, many will say that they cannot. If you do the math, the average physician will not see a fatal case of diagnostic error except *maybe* once in a lifetime. So those very egregious cases don't come to mind very readily, and the less harmful errors are much more common, but physicians tend to be unaware of these for various reasons.

RW: Do you think they are not seeing it because they're not happening or they're not seeing it because the patient goes some other place and the error becomes manifest but it's off their radar screen?

MG: It's a little bit of both. It's certainly true that many physicians will never hear about the diagnostic error that their patient encountered. For example, if I don't diagnose something correctly in clinic, that patient may go on into the hospital or may be transferred somewhere else, and I may not even know the outcome of their care. I've heard there are even HIPAA concerns about informing a physician earlier upstream about subsequent care because that patient is no longer under your direct care. But that's one factor; you just don't know what happens to all your patients.

Another factor is that we tend to rationalize away some of the mistakes that are made. We say, well that was a really difficult case, or I was distracted at the time. And life is full; we move on to the next case. We don't really dwell on the errors that we do become aware of.

RW: Let's talk a little bit about politics. It seems like you got very interested in this issue 10 to 15 years ago. You and colleagues began to chart a strategy to get this on the radar screen, and ultimately that led to the recent [IOM report](#). How strategic were you in the beginning? Did you say we need to do this, this, and this? Or was it kind of doing one thing leading to the next thing leading to the next thing?

MG: The progress in accepting diagnostic error as a major safety concern has been very slow. It wasn't because we were getting pushback; it was because people didn't recognize it as a problem—and that's still the case. Most health care organizations still don't recognize diagnostic error as a priority, and many physicians don't believe that it's an important problem in their practice. Our Society [SIDM] thought if we could convince the IOM to undertake a major report on diagnostic error, that would be the most effective way to raise awareness across all the stakeholders.

RW: As you began mapping out the steps that might be required to get more attention and more resources—in retrospect, I look back and I see the development of a society. I see the development of a journal. I see a research agenda. Then, ultimately I see the Institute of Medicine getting involved. How did that all happen?

MG: The first thing was finding people who shared our interest in diagnostic error and were willing to have a national conference annually. The conferences were outstanding in developing the passion, energy, and shared vision that we needed to work on this. But after 3 or 4 years of conferences, we realized that we get together once a year and we reaffirm that this is a priority, but then all this time passes between conferences when nothing is happening.

That led to the formation in 2011 of the [Society to Improve Diagnosis in Medicine](#) (SIDM), through which we dedicated ourselves to working on this every day instead of just once a year. For many years, we were pondering what could we do to bring attention to this—to get people to wake up to this issue. Given the success of the first IOM report, *To Err Is Human*, we were hoping that a follow-up report might do the trick. As you know from your talks at our conferences and from your own [writings](#), diagnostic error was mentioned in the original report, but only a few times. There were no major conclusions or recommendations, and the topic kind of disappeared from public view. The decision to request a second report from the Institute of Medicine was definitely strategic, and we are hoping that it will bring the attention that the problem needs. This new report is the missing chapter from *To Err Is Human*.

RW: Most people don't know how an IOM report gets generated. Is it just a matter of you writing them a letter saying we think this would be a good idea, or somehow the funding has to materialize? How does that work?

MG: We didn't know either, but after a few calls to the Institute of Medicine we learned. Most of the reports are requested by Congress and are funded through that mechanism, but they do some reports at the request of organizations and societies. You have to petition the Institute of Medicine, requests are reviewed by their board, and if approved, you then need to come up with the funding, roughly \$1.5 million. So we got to work and were successful in helping the IOM raise the money to get this report started. The report involved the work of some 30 individuals over a period of 18 months, with the final report issued in September 2015.

RW: Tell us about the report and what you were hoping to achieve from it. Maybe briefly summarize what you think the key take-home messages were from the report.

MG: The report is a landmark document that hopefully will be cited for years to come. It's a wonderful compilation of all the evidence that's available in regard to diagnostic error. The problem of diagnostic error was analyzed by an incredibly talented and knowledgeable committee that paid full attention to every aspect of this. It breaks new ground in a couple of areas. For one, it's focused very squarely on the patient. The patient is part of a new definition of diagnostic error that's provided in the report. Diagnosis is something that's done *for* the patient and the patient needs to be a partner in this process. That's an important statement that was missing from the earlier definitions.

Another major advance in the report is the framework that it presents about the diagnostic process. Diagnosis starts with a history and physical examination, then you order some tests, come up with some ideas, maybe get some consults, and ultimately you have a diagnosis. But the report very clearly lays out what those steps are and portrays this as a process. The importance of that is that health care organizations are very familiar with process improvement. Defining diagnosis as a process allows them to apply all the things they know about how to improve processes but in this case, apply that to diagnosis. How can we do a better job getting our patients access to care? How can we do a better job making sure that all the information that we need is at our fingertips? How can we do a better job making sure all the lab tests reach the right people who need to know? Having that framework is very valuable to improving care and also to research because it clearly identifies the gaps. Another major feature of the framework is that it illustrates the dynamic interplay between the actions of the individuals involved, and the environment and system in which those actions take place. Being able to apply concepts from human factors engineering will be one of the most productive approaches, we believe, to help reduce the risk of diagnostic errors.

RW: Given that the initial IOM report catalyzed the patient safety movement and did it with a big visible, visceral analogy—the jumbo jet a day analogy—how much pressure did you and the group feel to come up with a similar headline?

MG: That was a topic of very active discussion. Many people felt that we needed a jumbo jet kind of a number to grab people's attention. Other people felt that the numbers weren't really solid enough to pick the right one. All the numbers we have were based on research studies. For example the 40,000 to 80,000 deaths a year was just a back-of-the-envelope calculation based on autopsy studies. [Hardeep Singh](#) did a [study](#) of chart reviews and concluded that 1 in 20 ambulatory patients will experience a diagnostic error every year. [Secret shopper studies](#) show that 10% to 15% of the diagnoses rendered in classic presentations are incorrect. Although we have these various numbers from research, it's absolutely true that we really don't know the rate of diagnostic error in actual practice. Not a single health care organization in the country or in the world is measuring the incidence of error in their own practices right now. There is no one number; that was the conclusion of the committee. But they did agree that the evidence was sufficient to say that it's very likely that each of us will have a diagnostic error in our lifetime. That's a very important statement that each of us can relate to—that we are all at risk. There is a 50% chance that you will have three diagnostic errors in your lifetime if you do the math. We are hoping that that statement will be the equivalent of the jumbo jet impact factor that was introduced in *To Err Is Human*.

RW: Another element of these reports is a handful of high impact recommendations—if the system just did a, b, c, d, we would make this problem much better. The report contains a lot of recommendations. What handful of recommendations do you think would be the most impactful?

MG: Let's start with education. It's fascinating to me that there is no course on diagnosis in medical school. In most programs you learn diagnosis by watching the faculty and your mentors, and it's interesting how well that works out. We all turn out to be pretty good doctors. But the question is can we be better? If there were some training on how cognition works, how easily it goes astray, and how you might be able to prevent it from going astray, that might lead to better diagnosticians. This should be interprofessional training because so many other members of the professional team are involved in making diagnoses:

nurses, dentists, and optometrists, for example.

Other recommendations in the report are really addressed to everyone. There are recommendations for patients. We want them to be more engaged in the diagnostic process if they're so inclined and willing to do so. There is growing evidence that engaged patients have better outcomes. To the extent that they can become more involved in their own diagnosis, it might be safer. The report outlines steps for physicians to take in terms of reflective practice. The report asks health care organizations to get involved in terms of facilitating communication and making sure laboratory tests are appropriately relayed to the right people. We are hopeful that each of the parties involved in diagnosis—and it's a long list of stakeholders that could have an impact—could take a look at the report and seriously start to consider what they could do.

RW: If you watched Watson beat the Jeopardy champions a few years ago, you couldn't help but say, "I know how this is going to get fixed. Computers will take over the process of diagnosis for doctors." Where did that come out in the committee's deliberations, and what are your thoughts about that?

MG: I have to tell you one of the most pleasurable things I did this year was to read [your book](#) on the digital doctor. There were so many moments of truth in there. My experience is a little different than yours because I was absolutely in love with my electronic medical record. I practiced in the VA and it was a huge help for me. It saved me time. I think it made me a better doctor. And every month it got better. New features kept being added in that I liked and that I used. I think that is the future of medicine. We will ultimately get to decision support and medical record systems that really do improve diagnosis. The electronic medical record is so much a part of diagnosis at this point that the two are practically inseparable. This is where we document our thoughts, find our information, and communicate with others.

But you're absolutely right that there are so many unfulfilled promises right now. The electronic record has become a wall that's led to separations that I hope can be replaced, but I'm not sure. I used to love to go down and talk to the pathologists about a case and look at the slides together. That doesn't happen anymore because they just send their report back to me in the medical record and that conversation doesn't occur. It's the same for radiologists. The radiologists know so much more about a film than they put in their report, and you used to get that information when you went down to discuss a case with them. But you will not get it in just reading the report. The medical record will hopefully fulfill many of its promises someday, but we're clearly not there now.

RW: How about the more specific issue of the computer as diagnostician? You're pointing out the degree to which the computer can facilitate information flow or get in the way of social relationships, and all of those are very big deals. But in some ways the promise held out over the last 30 years was ultimately that we wouldn't need the doctors to be all that smart because the computers would make the diagnosis for them.

MG: I think that computers will never be as good at diagnosis as a physician, given all the information that has to be integrated and the interaction with the patient that's so important in getting the information that you need. I don't ever see computers replacing us. But they can play an incredibly important role in preventing some kinds of diagnostic errors. In a [famous paper](#), Georges Bordage analyzed diagnostic error and concluded that the most common reason was that you just didn't think of the right diagnosis. We saw that a lot in the [research study](#) I did of diagnostic errors. Physicians settled on a diagnosis that explained all

the facts in front of them. And they were happy. They were moving on to the next patient. They didn't really consider other possibilities. Hardeep Singh had a [study](#) of diagnostic errors in which 80% of the time when a diagnostic error had occurred, no differential diagnosis was present on the patient's chart. That is where computers and decision support can help us. There are, according to some estimates, 8,000 or 10,000 different diseases out there, but the average primary care doctor will only see 300 or 400 a year. How good am I going to be at diagnosing a disease that I may have only seen once in my lifetime, or I may have never seen? That's where decision support systems can help us create an appropriate differential diagnosis and make sure we don't miss something that we otherwise would have.

RW: But you don't see a day in the either not so distant or moderately distant future where all of the data go into the electronic hopper and it spits out a diagnosis and says, "You have lupus"?

MG: I think that is a valid question that needs to be studied. We've had decision support products for 3 or 4 decades now that could help physicians come up with a differential diagnosis. By and large, they just haven't been used. So if you ask physicians why—they know that these are available and can help them—they say I don't need that and I don't have time to do it. So if you're relying on the physicians to use these products, they may not.

That's a problem for me in regard to diagnostic error because so many diagnostic errors arise from cases where the physician was sure of their diagnosis. They didn't feel that they had to go look something up or go to a computer decision support system to obtain a better differential diagnosis. The systems that push out and answer at you instead of requiring you to go in there and pull it out might be more effective in preventing diagnostic error, and that's something that we really do need to look at through research.

RW: I'm assuming that there is a bell-shaped curve of people's diagnostic abilities. Whether the goal was to give them feedback or to make an assessment about whether they're good enough at it, you'd want to build in some layer where you're testing people on their diagnostic abilities and giving them some feedback. How are we doing on that and what do you think the future directions are?

MG: That is a great question. Right now, we're nowhere. If you pass the board, you're good enough. But we don't really look at the variability from physician to physician in terms of diagnostic ability. You're absolutely right there has to be a spectrum. There is some data; for example, optometrists looking at diabetic eyes don't all agree on what grade diabetic retinopathy it is. There is variation there and in all sorts of areas that have been looked at. And we're not the same physician from one day to the next. I may be more distracted one day or more tired or ill, and I won't do as good a job as I might do on some other day. This problem of variability really does need to be addressed at some point to try to narrow those bands so that we all perform at a slightly higher level. How you assess that is a completely open question that people are just starting to think about.

Right now most of the board exams are based on assessing your knowledge of the content in that area. They don't really examine your ability to differentiate one condition from another to generate a differential diagnosis, to excel at clinical reasoning. I know the boards are very interested in this. There was a meeting last month hosted by the [National Board of Medical Examiners](#) starting to ask how can they do a better job assessing the cognitive aspects of clinical reasoning. The [American Board of Medical Specialties](#) is asking

exactly the same question.

In addition to assessing whether we know the content, we need to assess the calibration of a physician. If I'm sick, I want to see somebody who's really good at whatever condition I have but who also knows their limits, knows what they know, knows what they don't know, knows when they need to slow down, knows who they need to get a consult on. Those are all higher levels of cognition that need to be and I think could easily be assessed. We could start in medical school to ask students, how well do you think you're going to do on this test? Or how likely are you to get this next question right on the exam? If we got in the habit of knowing how well calibrated we were, we might end up better-calibrated physicians in the long haul.

RW: A couple of strategies have been used in the past. I wonder whether they need to be exhumed. One of them is the old yellow pen, the test on the boards and in other places where you went down a diagnostic tree and it told you got it right or didn't get it right, and then you moved on finally making it to a diagnosis. The second is the oral exam, which I think is still done in the UK but is long gone in medicine in the US. What do you think about those two directions?

MG: The oral exams are definitely a wonderful way to go. You get so much more information from watching how a clinician analyzes a case than you do from how they perform on a written test, and that's irreplaceable. So there may be a movement back toward oral exams. But the other point you bring up touches on feedback a little bit and how important that is in becoming an expert. If you read the literature on expertise, that's how experts become that way. They practice like crazy but they get really good feedback on how they're doing, what they're not doing right, and the areas in which they need to focus. In medical school we get some feedback, we're always asking for more, and we always feel like we're not getting enough. But once you get out into practice that kind of feedback is scarce. It's almost a cultural taboo. If you discover a diagnostic error that I made, it's an awkward conversation for you to call me up and say, "Hey Mark, you blew it on that case." But we need to get past that. We need to get much better feedback if we want to excel at this, and we need to create a culture where it's okay or I give you permission. I say, "Bob, please, if you become aware of a diagnosis that I made that changed, I'd like to know about that. It will help me be a better doctor." I'd like to see health care organizations set up ways to facilitate physicians finding out what the outcome was on their patients, whether they got it right or whether they got it wrong. I'd like to see autopsies come back. I don't know whether that's a wild dream or not, but there was nothing like an autopsy to sensitize you to the limitations of our skills and our science. There is no lesson more powerful than seeing your own mistake on the autopsy table and nothing's replaced that. Also I'm a huge fan of the AHRQ WebM&M conferences, and there are many, many valuable lessons in there that we could all learn from.

RW: Tell us a little more about the reaction to the National Academy of Medicine report. Was it what you expected or has it been surprising?

MG: I would say 95% of the feedback I've received is positive. People are saying you got it right. You identified the main problems. You outlined some important recommendations, this is terrific, and we're going to get started. Other people wanted to see more detail and more granularity in the report. For example, the report includes very little about what laboratories could do or what radiology could do. I think that's a valid criticism. But looking at IOM reports, they are already very long. This report is 300 or 400

pages. To get to that level of granularity would double the size of the typical report. What I've learned is that IOM reports tend to stay at a very high level, and it's up to the stakeholders to start analyzing what's in there and taking appropriate action in their domains.

RW: It sounds like you've gotten good feedback on it. Are you getting a sense that there is momentum growing among stakeholders to do just that?

MG: Absolutely. The IOM report will open a lot of doors for us, and it already has. We're seeing growing interest on the part of health care organizations and physicians to talk about this, and our own Society is growing—SIDM. And we are very happy about a new coalition we're putting together. It's called [The Coalition to Improve Diagnosis](#), and it includes professional organizations like the [American Board of Internal Medicine](#) and the American Association of Nurse Practitioners. These societies will be working together to come up with collective actions that could be taken to improve diagnosis, but also agreeing to individually take on some projects for their members and in their domain that could help to improve the reliability of diagnosis. We see that as one of the major ways that the IOM report could actually lead to the changes in practice that will be needed to improve the safety in diagnosis.