


## Lies, Damned Lies, and Medical Science



Robyn Twomey/Redux

In 2001, rumors were circulating in Greek hospitals that surgery residents, eager to rack up scalpel time, were falsely diagnosing hapless Albanian immigrants with appendicitis. At the University of Ioannina medical school's teaching hospital, a newly minted doctor named Athina Tatsioni was discussing the rumors with colleagues when a professor who had overheard asked her if she'd like to try to prove whether they were true—he seemed to be almost daring her. She accepted the challenge and, with the professor's and other colleagues' help, eventually produced a formal study showing that, for whatever reason, the appendices removed from patients with Albanian names in six Greek hospitals were more than three times as likely to be perfectly healthy as those removed from patients with Greek names. "It was hard to find a journal willing to publish it, but we did," recalls Tatsioni. "I also discovered that I really liked research." Good thing, because the study had actually been a sort of audition. The professor, it turned out, had been putting together a team of exceptionally brash and curious young clinicians and Ph.D.s to join him in tackling an unusual and controversial agenda.

Last spring, I sat in on one of the team's weekly meetings on the medical school's campus, which is plunked crazily across a series of sharp hills. The building in which we met, like most at the school, had the look of a barracks and was festooned with political graffiti. But the group convened in a spacious conference room that would have been at home at a Silicon Valley start-up. Sprawled around a large table were Tatsioni and eight other youngish Greek researchers and physicians who, in contrast

to the pasty younger staff frequently seen in U.S. hospitals, looked like the casually glamorous cast of a television medical drama. The professor, a dapper and soft-spoken man named John Ioannidis, loosely presided.

One of the researchers, a biostatistician named Georgia Salanti, fired up a laptop and projector and started to take the group through a study she and a few colleagues were completing that asked this question: were drug companies manipulating published research to make their drugs look good? Salanti ticked off data that seemed to indicate they were, but the other team members almost immediately started interrupting. One noted that Salanti's study didn't address the fact that drug-company research wasn't measuring critically important "hard" outcomes for patients, such as survival versus death, and instead tended to measure "softer" outcomes, such as self-reported symptoms ("my chest doesn't hurt as much today"). Another pointed out that Salanti's study ignored the fact that when drug-company data seemed to show patients' health improving, the data often failed to show that the drug was responsible, or that the improvement was more than marginal.

Salanti remained poised, as if the grilling were par for the course, and gamely acknowledged that the suggestions were all good—but a single study can't prove everything, she said. Just as I was getting the sense that the data in drug studies were endlessly malleable, Ioannidis, who had mostly been listening, delivered what felt like a coup de grâce: wasn't it possible, he asked, that drug companies were carefully selecting the topics of their studies—for example, comparing their new drugs against those already known to be inferior to others on the market—so that they were ahead of the game even before the data juggling began? "Maybe sometimes it's the questions that are biased, not the answers," he said, flashing a friendly smile. Everyone nodded. Though the results of drug studies often make newspaper headlines, you have to wonder whether they prove anything at all. Indeed, given the breadth of the potential problems raised at the meeting, can *any* medical-research studies be trusted?

That question has been central to Ioannidis's career. He's what's known as a meta-researcher, and he's become one of the world's foremost experts on the credibility of medical research. He and his team have shown, again and again, and in many different ways, that much of what biomedical researchers conclude in published studies—conclusions that doctors keep in mind when they prescribe antibiotics or blood-pressure medication, or when they advise us to consume more fiber or less meat, or when they recommend surgery for heart disease or back pain—is misleading, exaggerated, and often flat-out wrong. He charges that as much as 90 percent of the published

medical information that doctors rely on is flawed. His work has been widely accepted by the medical community; it has been published in the field's top journals, where it is heavily cited; and he is a big draw at conferences. Given this exposure, and the fact that his work broadly targets everyone else's work in medicine, as well as everything that physicians do and all the health advice we get, Ioannidis may be one of the most influential scientists alive. Yet for all his influence, he worries that the field of medical research is so pervasively flawed, and so riddled with conflicts of interest, that it might be chronically resistant to change—or even to publicly admitting that there's a problem.

The city of Ioannina is a big college town a short drive from the ruins of a 20,000-seat amphitheater and a Zeusian sanctuary built at the site of the Dodona oracle. The oracle was said to have issued pronouncements to priests through the rustling of a sacred oak tree. Today, a different oak tree at the site provides visitors with a chance to try their own hands at extracting a prophecy. "I take all the researchers who visit me here, and almost every single one of them asks the tree the same question," Ioannidis tells me, as we contemplate the tree the day after the team's meeting. "'Will my research grant be approved?'" He chuckles, but Ioannidis (pronounced yo-NEE-dees) tends to laugh not so much in mirth as to soften the sting of his attack. And sure enough, he goes on to suggest that an obsession with winning funding has gone a long way toward weakening the reliability of medical research.

He first stumbled on the sorts of problems plaguing the field, he explains, as a young physician-researcher in the early 1990s at Harvard. At the time, he was interested in diagnosing rare diseases, for which a lack of case data can leave doctors with little to go on other than intuition and rules of thumb. But he noticed that doctors seemed to proceed in much the same manner even when it came to cancer, heart disease, and other common ailments. Where were the hard data that would back up their treatment decisions? There was plenty of published research, but much of it was remarkably unscientific, based largely on observations of a small number of cases. A new "evidence-based medicine" movement was just starting to gather force, and Ioannidis decided to throw himself into it, working first with prominent researchers at Tufts University and then taking positions at Johns Hopkins University and the National Institutes of Health. He was unusually well armed: he had been a math prodigy of near-celebrity status in high school in Greece, and had followed his parents, who were both physician-researchers, into medicine. Now he'd have a chance to combine math and medicine by applying rigorous statistical analysis to what seemed a surprisingly sloppy field. "I assumed that everything we physicians did was basically right, but now I was going to help verify it," he says. "All we'd have to do was

systematically review the evidence, trust what it told us, and then everything would be perfect.”

It didn't turn out that way. In poring over medical journals, he was struck by how many findings of all types were refuted by later findings. Of course, medical-science “never minds” are hardly secret. And they sometimes make headlines, as when in recent years large studies or growing consensus of researchers concluded that mammograms, colonoscopies, and PSA tests are far less useful cancer-detection tools than we had been told; or when widely prescribed antidepressants such as Prozac, Zoloft, and Paxil were revealed to be no more effective than a placebo for most cases of depression; or when we learned that staying out of the sun entirely can actually increase cancer risks; or when we were told that the advice to drink lots of water during intense exercise was potentially fatal; or when, last April, we were informed that taking fish oil, exercising, and doing puzzles doesn't really help fend off Alzheimer's disease, as long claimed. Peer-reviewed studies have come to opposite conclusions on whether using cell phones can cause brain cancer, whether sleeping more than eight hours a night is healthful or dangerous, whether taking aspirin every day is more likely to save your life or cut it short, and whether routine angioplasty works better than pills to unclog heart arteries.

But beyond the headlines, Ioannidis was shocked at the range and reach of the reversals he was seeing in everyday medical research. “Randomized controlled trials,” which compare how one group responds to a treatment against how an identical group fares without the treatment, had long been considered nearly unshakable evidence, but they, too, ended up being wrong some of the time. “I realized even our gold-standard research had a lot of problems,” he says. Baffled, he started looking for the specific ways in which studies were going wrong. And before long he discovered that the range of errors being committed was astonishing: from what questions researchers posed, to how they set up the studies, to which patients they recruited for the studies, to which measurements they took, to how they analyzed the data, to how they presented their results, to how particular studies came to be published in medical journals.

This array suggested a bigger, underlying dysfunction, and Ioannidis thought he knew what it was. “The studies were biased,” he says. “Sometimes they were overtly biased. Sometimes it was difficult to see the bias, but it was there.” Researchers headed into their studies wanting certain results—and, lo and behold, they were getting them. We think of the scientific process as being objective, rigorous, and even ruthless in separating out what is true from what we merely wish to be true, but in fact

it's easy to manipulate results, even unintentionally or unconsciously. "At every step in the process, there is room to distort results, a way to make a stronger claim or to select what is going to be concluded," says Ioannidis. "There is an intellectual conflict of interest that pressures researchers to find whatever it is that is most likely to get them funded."

Perhaps only a minority of researchers were succumbing to this bias, but their distorted findings were having an outsize effect on published research. To get funding and tenured positions, and often merely to stay afloat, researchers have to get their work published in well-regarded journals, where rejection rates can climb above 90 percent. Not surprisingly, the studies that tend to make the grade are those with eye-catching findings. But while coming up with eye-catching theories is relatively easy, getting reality to bear them out is another matter. The great majority collapse under the weight of contradictory data when studied rigorously. Imagine, though, that five different research teams test an interesting theory that's making the rounds, and four of the groups correctly prove the idea false, while the one less cautious group incorrectly "proves" it true through some combination of error, fluke, and clever selection of data. Guess whose findings your doctor ends up reading about in the journal, and you end up hearing about on the evening news? Researchers can sometimes win attention by refuting a prominent finding, which can help to at least raise doubts about results, but in general it is far more rewarding to add a new insight or exciting-sounding twist to existing research than to retest its basic premises—after all, simply re-proving someone else's results is unlikely to get you published, and attempting to undermine the work of respected colleagues can have ugly professional repercussions.

In the late 1990s, Ioannidis set up a base at the University of Ioannina. He pulled together his team, which remains largely intact today, and started chipping away at the problem in a series of papers that pointed out specific ways certain studies were getting misleading results. Other meta-researchers were also starting to spotlight disturbingly high rates of error in the medical literature. But Ioannidis wanted to get the big picture across, and to do so with solid data, clear reasoning, and good statistical analysis. The project dragged on, until finally he retreated to the tiny island of Siki-nos in the Aegean Sea, where he drew inspiration from the relatively primitive surroundings and the intellectual traditions they recalled. "A pervasive theme of ancient Greek literature is that you need to pursue the truth, no matter what the truth might be," he says. In 2005, he unleashed two papers that challenged the foundations of medical research.

He chose to publish one paper, fittingly, in the online journal *PLoS Medicine*, which is committed to running any methodologically sound article without regard to how “interesting” the results may be. In the paper, Ioannidis laid out a detailed mathematical proof that, assuming modest levels of researcher bias, typically imperfect research techniques, and the well-known tendency to focus on exciting rather than highly plausible theories, researchers will come up with wrong findings most of the time. Simply put, if you’re attracted to ideas that have a good chance of being wrong, and if you’re motivated to prove them right, and if you have a little wiggle room in how you assemble the evidence, you’ll probably succeed in proving wrong theories right. His model predicted, in different fields of medical research, rates of wrongness roughly corresponding to the observed rates at which findings were later convincingly refuted: 80 percent of non-randomized studies (by far the most common type) turn out to be wrong, as do 25 percent of supposedly gold-standard randomized trials, and as much as 10 percent of the platinum-standard large randomized trials. The article spelled out his belief that researchers were frequently manipulating data analyses, chasing career-advancing findings rather than good science, and even using the peer-review process—in which journals ask researchers to help decide which studies to publish—to suppress opposing views. “You can question some of the details of John’s calculations, but it’s hard to argue that the essential ideas aren’t absolutely correct,” says Doug Altman, an Oxford University researcher who directs the Centre for Statistics in Medicine.

Still, Ioannidis anticipated that the community might shrug off his findings: sure, a lot of dubious research makes it into journals, but we researchers and physicians know to ignore it and focus on the good stuff, so what’s the big deal? The other paper headed off that claim. He zoomed in on 49 of the most highly regarded research findings in medicine over the previous 13 years, as judged by the science community’s two standard measures: the papers had appeared in the journals most widely cited in research articles, and the 49 articles themselves were the most widely cited articles in these journals. These were articles that helped lead to the widespread popularity of treatments such as the use of hormone-replacement therapy for menopausal women, vitamin E to reduce the risk of heart disease, coronary stents to ward off heart attacks, and daily low-dose aspirin to control blood pressure and prevent heart attacks and strokes. Ioannidis was putting his contentions to the test not against run-of-the-mill research, or even merely well-accepted research, but against the absolute tip of the research pyramid. Of the 49 articles, 45 claimed to have uncovered effective interventions. Thirty-four of these claims had been retested, and 14 of these, or 41 percent, had been convincingly shown to be wrong or significantly exaggerated. If between a third and a half of the most acclaimed research in medicine was proving untrustwor-

thy, the scope and impact of the problem were undeniable. That article was published in the *Journal of the American Medical Association*.

Driving me back to campus in his smallish SUV—after insisting, as he apparently does with all his visitors, on showing me a nearby lake and the six monasteries situated on an islet within it—Ioannidis apologized profusely for running a yellow light, explaining with a laugh that he didn't trust the truck behind him to stop. Considering his willingness, even eagerness, to slap the face of the medical-research community, Ioannidis comes off as thoughtful, upbeat, and deeply civil. He's a careful listener, and his frequent grin and semi-apologetic chuckle can make the sharp prodding of his arguments seem almost good-natured. He is as quick, if not quicker, to question his own motives and competence as anyone else's. A neat and compact 45-year-old with a trim mustache, he presents as a sort of dashing nerd—Giancarlo Giannini with a bit of Mr. Bean.

The humility and graciousness seem to serve him well in getting across a message that is not easy to digest or, for that matter, believe: that even highly regarded researchers at prestigious institutions sometimes churn out attention-grabbing findings rather than findings likely to be right. But Ioannidis points out that obviously questionable findings cram the pages of top medical journals, not to mention the morning headlines. Consider, he says, the endless stream of results from nutritional studies in which researchers follow thousands of people for some number of years, tracking what they eat and what supplements they take, and how their health changes over the course of the study. "Then the researchers start asking, 'What did vitamin E do? What did vitamin C or D or A do? What changed with calorie intake, or protein or fat intake? What happened to cholesterol levels? Who got what type of cancer?'" he says. "They run everything through the mill, one at a time, and they start finding associations, and eventually conclude that vitamin X lowers the risk of cancer Y, or this food helps with the risk of that disease." In a single week this fall, Google's news page offered these headlines: "More Omega-3 Fats Didn't Aid Heart Patients"; "Fruits, Vegetables Cut Cancer Risk for Smokers"; "Soy May Ease Sleep Problems in Older Women"; and dozens of similar stories.

When a five-year study of 10,000 people finds that those who take more vitamin X are less likely to get cancer Y, you'd think you have pretty good reason to take more vitamin X, and physicians routinely pass these recommendations on to patients. But these studies often sharply conflict with one another. Studies have gone back and forth on the cancer-preventing powers of vitamins A, D, and E; on the heart-health benefits of eating fat and carbs; and even on the question of whether being over-

weight is more likely to extend or shorten your life. How should we choose among these dueling, high-profile nutritional findings? Ioannidis suggests a simple approach: ignore them all.

For starters, he explains, the odds are that in any large database of many nutritional and health factors, there will be a few apparent connections that are in fact merely flukes, not real health effects—it's a bit like combing through long, random strings of letters and claiming there's an important message in any words that happen to turn up. But even if a study managed to highlight a genuine health connection to some nutrient, you're unlikely to benefit much from taking more of it, because we consume thousands of nutrients that act together as a sort of network, and changing intake of just one of them is bound to cause ripples throughout the network that are far too complex for these studies to detect, and that may be as likely to harm you as help you. Even if changing that one factor does bring on the claimed improvement, there's still a good chance that it won't do you much good in the long run, because these studies rarely go on long enough to track the decades-long course of disease and ultimately death. Instead, they track easily measurable health "markers" such as cholesterol levels, blood pressure, and blood-sugar levels, and meta-experts have shown that changes in these markers often don't correlate as well with long-term health as we have been led to believe.

On the relatively rare occasions when a study does go on long enough to track mortality, the findings frequently upend those of the shorter studies. (For example, though the vast majority of studies of overweight individuals link excess weight to ill health, the longest of them haven't convincingly shown that overweight people are likely to die sooner, and a few of them have seemingly demonstrated that moderately overweight people are likely to live *longer*.) And these problems are aside from ubiquitous measurement errors (for example, people habitually misreport their diets in studies), routine misanalysis (researchers rely on complex software capable of juggling results in ways they don't always understand), and the less common, but serious, problem of outright fraud (which has been revealed, in confidential surveys, to be much more widespread than scientists like to acknowledge).

If a study somehow avoids every one of these problems and finds a real connection to long-term changes in health, you're still not guaranteed to benefit, because studies report average results that typically represent a vast range of individual outcomes. Should you be among the lucky minority that stands to benefit, don't expect a noticeable improvement in your health, because studies usually detect only modest effects that merely tend to whittle your chances of succumbing to a particular disease from



small to somewhat smaller. “The odds that anything useful will survive from any of these studies are poor,” says Ioannidis—dismissing in a breath a good chunk of the research into which we sink about \$100 billion a year in the United States alone.

And so it goes for all medical studies, he says. Indeed, nutritional studies aren’t the worst. Drug studies have the added corruptive force of financial conflict of interest. The exciting links between genes and various diseases and traits that are relentlessly hyped in the press for heralding miraculous around-the-corner treatments for everything from colon cancer to schizophrenia have in the past proved so vulnerable to error and distortion, Ioannidis has found, that in some cases you’d have done about as well by throwing darts at a chart of the genome. (These studies seem to have improved somewhat in recent years, but whether they will hold up or be useful in treatment are still open questions.) Vioxx, Zelnorm, and Baycol were among the widely prescribed drugs found to be safe and effective in large randomized controlled trials before the drugs were yanked from the market as unsafe or not so effective, or both.

“Often the claims made by studies are so extravagant that you can immediately cross them out without needing to know much about the specific problems with the studies,” Ioannidis says. But of course it’s that very extravagance of claim (one large randomized controlled trial even proved that secret prayer by unknown parties can save the lives of heart-surgery patients, while another proved that secret prayer can harm them) that helps gets these findings into journals and then into our treatments and lifestyles, especially when the claim builds on impressive-sounding evidence. “Even when the evidence shows that a particular research idea is wrong, if you have thousands of scientists who have invested their careers in it, they’ll continue to publish papers on it,” he says. “It’s like an epidemic, in the sense that they’re infected with these wrong ideas, and they’re spreading it to other researchers through journals.”

Though scientists and science journalists are constantly talking up the value of the peer-review process, researchers admit among themselves that biased, erroneous, and even blatantly fraudulent studies easily slip through it. *Nature*, the grande dame of science journals, stated in a 2006 editorial, “Scientists understand that peer review per se provides only a minimal assurance of quality, and that the public conception of peer review as a stamp of authentication is far from the truth.” What’s more, the peer-review process often pressures researchers to shy away from striking out in genuinely new directions, and instead to build on the findings of their colleagues (that is, their potential reviewers) in ways that only *seem* like breakthroughs—as with the exciting-sounding gene linkages (autism genes identified!) and nutritional findings (olive oil lowers blood pressure!) that are really just dubious and conflicting varia-

tions on a theme.

Most journal editors don't even claim to protect against the problems that plague these studies. University and government research overseers rarely step in to directly enforce research quality, and when they do, the science community goes ballistic over the outside interference. The ultimate protection against research error and bias is supposed to come from the way scientists constantly retest each other's results—except they don't. Only the most prominent findings are likely to be put to the test, because there's likely to be publication payoff in firming up the proof, or contradicting it.

But even for medicine's most influential studies, the evidence sometimes remains surprisingly narrow. Of those 45 super-cited studies that Ioannidis focused on, 11 had never been retested. Perhaps worse, Ioannidis found that even when a research error is outed, it typically persists for years or even decades. He looked at three prominent health studies from the 1980s and 1990s that were each later soundly refuted, and discovered that researchers continued to cite the original results as correct more often than as flawed—in one case for at least 12 years after the results were discredited.

Doctors may notice that their patients don't seem to fare as well with certain treatments as the literature would lead them to expect, but the field is appropriately conditioned to subjugate such anecdotal evidence to study findings. Yet much, perhaps even most, of what doctors do has never been formally put to the test in credible studies, given that the need to do so became obvious to the field only in the 1990s, leaving it playing catch-up with a century or more of non-evidence-based medicine, and contributing to Ioannidis's shockingly high estimate of the degree to which medical knowledge is flawed. That we're not routinely made seriously ill by this shortfall, he argues, is due largely to the fact that most medical interventions and advice don't address life-and-death situations, but rather aim to leave us marginally healthier or less unhealthy, so we usually neither gain nor risk all that much.

Medical research is not especially plagued with wrongness. Other meta-research experts have confirmed that similar issues distort research in all fields of science, from physics to economics (where the highly regarded economists J. Bradford DeLong and Kevin Lang once showed how a remarkably consistent paucity of strong evidence in published economics studies made it unlikely that *any* of them were right). And needless to say, things only get worse when it comes to the pop expertise that endlessly spews at us from diet, relationship, investment, and parenting gurus and pundits. But we expect more of scientists, and especially of medical scientists, given that

we believe we are staking our lives on their results. The public hardly recognizes how bad a bet this is. The medical community itself might still be largely oblivious to the scope of the problem, if Ioannidis hadn't forced a confrontation when he published his studies in 2005.

Ioannidis initially thought the community might come out fighting. Instead, it seemed relieved, as if it had been guiltily waiting for someone to blow the whistle, and eager to hear more. David Gorski, a surgeon and researcher at Detroit's Barbara Ann Karmanos Cancer Institute, noted in his prominent medical blog that when he presented Ioannidis's paper on highly cited research at a professional meeting, "not a single one of my surgical colleagues was the least bit surprised or disturbed by its findings." Ioannidis offers a theory for the relatively calm reception. "I think that people didn't feel I was only trying to provoke them, because I showed that it was a community problem, instead of pointing fingers at individual examples of bad research," he says. In a sense, he gave scientists an opportunity to cluck about the wrongness without having to acknowledge that they themselves succumb to it—it was something everyone else did.

To say that Ioannidis's work has been embraced would be an understatement. His *PLoS Medicine* paper is the most downloaded in the journal's history, and it's not even Ioannidis's most-cited work—that would be a paper he published in *Nature Genetics* on the problems with gene-link studies. Other researchers are eager to work with him: he has published papers with 1,328 different co-authors at 538 institutions in 43 countries, he says. Last year he received, by his estimate, invitations to speak at 1,000 conferences and institutions around the world, and he was accepting an average of about five invitations a month until a case last year of excessive-travel-induced vertigo led him to cut back. Even so, in the weeks before I visited him he had addressed an AIDS conference in San Francisco, the European Society for Clinical Investigation, Harvard's School of Public Health, and the medical schools at Stanford and Tufts.

The irony of his having achieved this sort of success by accusing the medical-research community of chasing after success is not lost on him, and he notes that it ought to raise the question of whether he himself might be pumping up his findings. "If I did a study and the results showed that in fact there wasn't really much bias in research, would I be willing to publish it?" he asks. "That would create a real psychological conflict for me." But his bigger worry, he says, is that while his fellow researchers seem to be getting the message, he hasn't necessarily forced anyone to do a better job. He fears he won't in the end have done much to improve anyone's health. "There

may not be fierce objections to what I'm saying," he explains. "But it's difficult to change the way that everyday doctors, patients, and healthy people think and behave."

As helter-skelter as the University of Ioannina Medical School campus looks, the hospital abutting it looks reassuringly stolid. Athina Tatsioni has offered to take me on a tour of the facility, but we make it only as far as the entrance when she is greeted—accosted, really—by a worried-looking older woman. Tatsioni, normally a bit reserved, is warm and animated with the woman, and the two have a brief but intense conversation before embracing and saying goodbye. Tatsioni explains to me that the woman and her husband were patients of hers years ago; now the husband has been admitted to the hospital with abdominal pains, and Tatsioni has promised she'll stop by his room later to say hello. Recalling the appendicitis story, I prod a bit, and she confesses she plans to do her own exam. She needs to be circumspect, though, so she won't appear to be second-guessing the other doctors.

Tatsioni doesn't so much fear that someone will carve out the man's healthy appendix. Rather, she's concerned that, like many patients, he'll end up with prescriptions for multiple drugs that will do little to help him, and may well harm him. "Usually what happens is that the doctor will ask for a suite of biochemical tests—liver fat, pancreas function, and so on," she tells me. "The tests could turn up something, but they're probably irrelevant. Just having a good talk with the patient and getting a close history is much more likely to tell me what's wrong." Of course, the doctors have all been trained to order these tests, she notes, and doing so is a lot quicker than a long bedside chat. They're also trained to ply the patient with whatever drugs might help whack any errant test numbers back into line. What they're not trained to do is to go back and look at the research papers that helped make these drugs the standard of care. "When you look the papers up, you often find the drugs didn't even work better than a placebo. And no one tested how they worked in combination with the other drugs," she says. "Just taking the patient off everything can improve their health right away." But not only is checking out the research another time-consuming task, patients often don't even *like* it when they're taken off their drugs, she explains; they find their prescriptions reassuring.

Later, Ioannidis tells me he makes a point of having several clinicians on his team. "Researchers and physicians often don't understand each other; they speak different languages," he says. Knowing that some of his researchers are spending more than half their time seeing patients makes him feel the team is better positioned to bridge that gap; their experience informs the team's research with firsthand knowledge, and

helps the team shape its papers in a way more likely to hit home with physicians. It's not that he envisions doctors making all their decisions based solely on solid evidence—there's simply too much complexity in patient treatment to pin down every situation with a great study. "Doctors need to rely on instinct and judgment to make choices," he says. "But these choices should be as informed as possible by the evidence. And if the evidence isn't good, doctors should know that, too. And so should patients."

In fact, the question of whether the problems with medical research should be broadcast to the public is a sticky one in the meta-research community. Already feeling that they're fighting to keep patients from turning to alternative medical treatments such as homeopathy, or misdiagnosing themselves on the Internet, or simply neglecting medical treatment altogether, many researchers and physicians aren't eager to provide even more reason to be skeptical of what doctors do—not to mention how public disenchantment with medicine could affect research funding. Ioannidis dismisses these concerns. "If we don't tell the public about these problems, then we're no better than nonscientists who falsely claim they can heal," he says. "If the drugs don't work and we're not sure how to treat something, why should we claim differently? Some fear that there may be less funding because we stop claiming we can prove we have miraculous treatments. But if we can't really provide those miracles, how long will we be able to fool the public anyway? The scientific enterprise is probably the most fantastic achievement in human history, but that doesn't mean we have a right to overstate what we're accomplishing."

We could solve much of the wrongness problem, Ioannidis says, if the world simply stopped expecting scientists to be right. That's because being wrong in science is fine, and even necessary—as long as scientists recognize that they blew it, report their mistake openly instead of disguising it as a success, and then move on to the next thing, until they come up with the very occasional genuine breakthrough. But as long as careers remain contingent on producing a stream of research that's dressed up to seem more right than it is, scientists will keep delivering exactly that.

"Science is a noble endeavor, but it's also a low-yield endeavor," he says. "I'm not sure that more than a very small percentage of medical research is ever likely to lead to major improvements in clinical outcomes and quality of life. We should be very comfortable with that fact."

*David H. Freedman is the author of *Wrong: Why Experts Keep Failing Us—And How to Know When Not to Trust Them*. He has been an Atlantic contributor since 1998.*