How many scientists does it take to change a paradigm?

New ideas to explain scientific observations are everywhere—we just need to learn how to see them

Neil R. Smalheiser

People who read about scientific discoveries might get the misleading impression that scientific research produces a few rare breakthroughs—once or twice per century—and a large body of ‘merely incremental’ studies. In reality, however, breakthrough discoveries are reported on a weekly basis, and one can cite many fields just in biology—brain imaging, non-coding RNAs and stem cell biology, to name a few—that have undergone paradigm shifts within the past decade.

The truly surprising thing about discovery is not just that it happens at a regular pace, but that most significant discoveries occurred only after the scientific community had already accepted another explanation. It is not merely the accrual of new data that leads to a breakthrough, but a willingness to acknowledge that a problem that is already ‘solved’ might require an entirely different explanation. In the case of breakthroughs or paradigm shifts, this new explanation might seem far-fetched or nonsensical and not even worthy of serious consideration. It is as if new ideas are sitting right in front of everyone, but in their blind spots so that only those who use their peripheral vision can see them.

Scientists do not all share any single method or way of working. Yet they tend to share certain prevalent attitudes: they accept ‘facts’ and ‘obvious’ explanations only provisionally, at arm’s length, as it were; they not only imagine alternatives, but—almost as a reflex—ask themselves what alternative explanations are possible.

When teaching students, it is a challenge to convey this critical attitude towards seemingly obvious explanations. In the spring semester of 2009, I offered a

seminar entitled *The Process of Scientific Discovery* to Honours undergraduate students at the University of Illinois-Chicago in the USA. I originally planned to cover aspects of discovery such as the impact of funding agencies, the importance of mentoring and hypothesis-driven as opposed to data-driven research. As the semester progressed, however, my sessions moved towards ‘teaching moments’ drawn from everyday life, which forced the students to look at familiar things in unfamiliar ways. These served as metaphors for certain aspects of the process by which scientists discover new paradigms.

For the first seven weeks of the spring semester, the class read *Everyday Practice of Science* by Frederick Grinnell [1]. During the discussion of the first chapter, one of the students noted that Grinnell referred to a scientist generically as ‘she’ rather than ‘he’ or the neutral ‘he or she’. This use is unusual and made her vaguely uneasy: she wondered whether the author was making a sexist point. Before considering her hypothesis, I asked the class to make a list of assumptions that they took for granted when reading the chapter, together with the possible explanations for the use of ‘she’ in the first chapter, no matter how far-fetched or unlikely they might seem.

For example, one might assume that Frederick Grinnell or ‘Fred’ is from a culture similar to our own. How would we interpret his behaviour and outlook if we knew that Fred came from an exotic foreign land? Another assumption is that Fred is male; how would we view the remark if we discover that Frederick is short for Fredericka? We have equally assumed that Fred, as

The truly surprising thing about discovery is [...] that most significant discoveries occurred only after the scientific community had already accepted another explanation

with most humans, wants us to like him. Instead, perhaps he is being intentionally provocative in order to get our attention or move us out of our comfort zone. Perhaps he planted ‘she’ as a deliberate example for us to discuss, as he does later in the second chapter, in which he deliberately hides a strange item in plain sight within one of the illustrations in order to make a point about observing anomalies. Perhaps the book was written not by Fred but by a ghost writer? Perhaps the ‘she’ was a typo?

Looking for patterns throughout the book, and in Fred’s other writing, might persuade us to discard some of the possible explanations: does ‘she’ appear just once? Does Fred use other unusual or provocative turns of phrase? Does Fred discuss gender bias or sexism explicitly? Has anyone written or complained about him? Of course, one could ask Fred directly what he meant, although without knowing him personally, it would be difficult to know how to interpret his answer or whether to take his remarks at face value. Notwithstanding the answer, the exercise is an important lesson about considering and weighing all possible explanations.

Arguably, the most prominent term used in science studies is the notion of a ‘paradigm’. I use this term with reluctance, as it is extraordinarily ambiguous. For example, it could simply refer to

a specific type of experimental design: a randomized, placebo-controlled clinical trial could be considered a paradigm. In the context of science studies, however, it most often refers to the idea of large-scale leaps in scientific world views, as promoted by Thomas Kuhn in *The Structure of Scientific Revolutions* [2]. Kuhn's notion of a paradigm can lead one to believe—erroneously in my opinion—that paradigm shifts are the opposite of practical, everyday scientific problem-solving.

A paradigm is recognized by the set of assumptions that an observer might not realize he or she is making...

Instead, I propose here a definition of 'paradigm' that emphasizes not the nature of the problem, the type of discovery or the scope of its implications, but rather the psychology of the scientist. A scientist viewing a problem or phenomenon resides within a paradigm when he or she does not notice, and cannot imagine, that an alternative way of looking at things needs to be considered seriously. Importantly, a paradigm is not a viewpoint, model, interpretation, hypothesis or conclusion. A paradigm is not the object that is viewed but the lenses through which it is viewed. A paradigm is recognized by the set of assumptions that an observer might not realize he or she is making, but which imply many automatic expectations and simultaneously prevent the observer from seeing the issue in any other fashion.

For example, the teacher–student paradigm feels natural and obvious, yet it is merely set up by habit and tradition. It implies lectures, assignments, grades, ways of addressing the professor and so on, all of which could be done differently, if we had merely thought to consider alternatives. What feels most natural in a paradigm is often the most arbitrary. When we have a birthday, we expect to have a cake with candles, yet there is no natural relationship at all between birthdays, cakes and candles. In fact, when something is arbitrary or conventional yet feels entirely natural, that is an important clue that a paradigm is present.

It is certainly natural for people to colour their observations according to their expectations: "To a man with a hammer, everything looks like a nail," as Mark Twain put it. However, this is a pitfall that scientists

(and doctors) must try hard to avoid. When I was a first-year medical student at Albert Einstein College of Medicine in New York City, we took a class on how to approach patients. As part of this course, we attended a session in which a psychiatrist interviewed a 'normal, healthy old person' in order to understand better the lives and perspectives of the elderly.

A man came in, and the psychiatrist began to ask him some benign questions. After about 10 minutes, however, the man began to pause before answering; then his answers became terse; then he said he did not feel well, excused himself and abruptly left the room. The psychiatrist continued to lecture to the students for another half-hour, analysing and interpreting the halting responses in terms of the emotional conflicts that the man was experiencing. 'Repression', 'emotional blocks', and 'reaction formation' were some of the terms bandied about.

However, unbeknown to the class, the man had collapsed just on the other side of the classroom door. Two cardiologists happened to be walking by and instantly realized the man was having an acute heart attack. They instituted CPR on the spot, but the man died within a few minutes.

The psychiatrist had been told that the man was healthy, and thus interpreted everything that he saw in psychological terms. It never entered his mind that the man might have been dying in front of his eyes. The cardiologists saw a man having a heart attack, and it never entered their minds that the man might have had psychological issues.

The movie *The Sixth Sense* [3] resonated particularly well with my students and served as a platform for discussing attitudes that are helpful for scientific investigation, such as "keep an open mind", "reality is much stranger than you can imagine" and "our conclusions are always provisional at best". Best of all, *The Sixth Sense* demonstrates the tension that exists between different scientific paradigms in a clear and beautiful way. When Haley Joel Osment says, "I see dead people," does he actually see ghosts? Or is he hallucinating?

It is important to emphasize that these are not merely different viewpoints, or different ways of defining terms. If we argued about which mountain is higher, Everest or K2, we might disagree about

which kind of evidence is more reliable, but we would fundamentally agree on the notion of measurement. By contrast, in *The Sixth Sense*, the same evidence used by one paradigm to support its assertion is used with equal strength by the other paradigm as evidence in its favour. In the movie, Bruce Willis plays a psychologist who assumes that Osment must be a troubled youth. However, the fact that he says he sees ghosts is also evidence in favour of the existence of ghosts, if you do not reject out of hand the possibility of their existence. These two explanations are incommensurate. One cannot simply weigh all of the evidence because each side rejects the type of evidence that the other side accepts, and regards the alternative explanation not merely as wrong but as ridiculous or nonsensical. It is in this sense that a paradigm represents a failure of imagination—each side cannot imagine that the other explanation could possibly be true, or at least, plausible enough to warrant serious consideration.

The failure of imagination means that each side fails to notice or to seek 'objective' evidence that would favour one explanation over the other. For example, during the episodes when Osment saw ghosts, the thermostat in the room fell precipitously and he could see his own breath. This certainly would seem to constitute objective evidence to favour the ghost explanation, and the fact that his mother had noticed that the heating in her apartment was erratic suggests that the temperature change was not simply another imagined symptom. But the mother assumed that the problem was in the heating system and did not even conceive that this might be linked to ghosts—so the 'objective' evidence certainly was not compelling or even suggestive on its own.

...when scientists reach a conclusion, it is merely a place to pause and rest for a moment, not a final destination

Osment did succeed eventually in convincing his mother that he saw ghosts, and he did it in the same way that any scientist would convince his colleagues: namely, he produced evidence that made perfect sense in the context of one, and only one, explanation. First, he told his mother a secret that he said her dead mother had

told him. This secret was about an incident that had occurred before he was born, and presumably she had never spoken of it, so there was no obvious way that he could have learned about it. Next, he told her that the grandmother had heard her say “every day” when standing near her grave. Again, the mother had presumably visited the grave alone and had not told anyone about the visit or about what was said. So, the mother was eventually convinced that Osment must have spoken with the dead grandmother after all. No other explanation seemed to fit all the facts.

Is this the end of the story? We, the audience, realize that it is possible that Osment had merely guessed about the incidents, heard them second-hand from another relative or (as with professional psychics) might have retold his anecdotes whilst looking for validation from his mother. The evidence seems compelling only because these alternatives seem even less likely. It is in this same sense that when scientists reach a conclusion, it is merely a place to pause and rest for a moment, not a final destination.

Near the end of the course, I gave a pop-quiz asking each student to give a ‘yes’ or ‘no’ answer, plus a short one-sentence explanation, to the following question: Donald Trump seems to be a wealthy businessman. He dresses like one, he has a TV show in which he acts like one, he gives seminars on wealth building and so on. Everything we know about him says that he is wealthy as a direct result of his business activities. On the basis of this evidence, are we justified in concluding that he is, in fact, a wealthy businessman?

Even when looked at carefully, not every anomaly is attractive enough or ‘ripe’ enough to be pursued when first noticed

About half the class said that yes, if all the evidence points in one direction, that suffices. About half the class said ‘no’, the stated evidence is circumstantial and we do not know, for example, what his bank balance is or whether he has more debt than equity. All the evidence we know about points in one direction, but we might not know all the facts.

How do we know whether or not we know all the facts? Again, it is a matter of

imagination. Let us review a few possible alternatives. Maybe his wealth comes from inheritance rather than business acumen; or from silent partners; or from drug running. Maybe he is dangerously over-extended and living on borrowed money; maybe his wealth is more apparent than real. Maybe Trump Casinos made up the role of Donald Trump as its symbol, the way McDonald’s made up the role of Ronald McDonald?

Several students complained that this was a ridiculous question. Yet I had posed this just after Bernard Madoff’s arrest was blanketing the news. Madoff was known as a billionaire investor genius for decades and had even served as the head of the Securities and Exchange Commission. As it turned out, his money was obtained by a massive Ponzi scheme. Why was Madoff able to succeed for so long? Because it was inconceivable that such a famous public figure could be a common con man and the people around him could not imagine the possibility that his livelihood needed to be scrutinized.

To this point, I have emphasized the benefits of paying attention to anomalous, strange or unwelcome observations. Yet paradoxically, scientists often make progress by (provisionally) putting aside anomalous or apparently negative findings that seem to invalidate or distract from their hypothesis. When Rita Levi-Montalcini was assaying the neurite-promoting effects of tumour tissue, she had predicted that this was a property of tumours and was devastated to find that normal tissue had the same effects. Only by ‘ignoring’ this apparent failure could she move forward to characterize nerve growth factor and eventually understand its biology [4].

Another classic example is Huntington disease—a genetic disorder in which an inherited alteration in the gene that encodes a protein, huntingtin, leads to toxicity within certain types of neuron and causes a progressive movement disorder associated with cognitive decline and psychiatric symptoms. Clinicians observed that the offspring of Huntington disease patients sometimes showed symptoms at an earlier age than their parents, and this phenomenon, called ‘genetic anticipation’, could affect successive generations at earlier and earlier ages of onset. This observation was met with scepticism and sometimes ridicule, as everything that was known about genetics at the time indicated that genes do not change across generations.

Ascertainment bias was suggested as a much more probable explanation; in other words, once a patient is diagnosed with Huntington disease, their doctors will look at their offspring much more closely and will thus tend to identify the onset of symptoms at an earlier age. Eventually, once the detailed genetics of the disease were understood at the molecular level, it was shown that the structure of the altered huntingtin gene does change. Genetic anticipation is now an accepted phenomenon.

...in fact, schools teach a lot about how to test hypotheses but little about how to find good hypotheses in the first place

What does this teach us about discovery? Even when looked at carefully, not every anomaly is attractive enough or ‘ripe’ enough to be pursued when first noticed. The biologists who identified the structure of the abnormal huntingtin gene did eventually explain genetic anticipation, although they set aside the puzzling clinical observations and proceeded pragmatically according to their (wrong) initial best-guess as to the genetics. The important thing is to move forward.

Finally, let us consider the case of Grigori Perelman, an outstanding mathematician who solved the Poincaré Conjecture a few years ago. He did not tell anyone he was working on the problem, lest their ‘helpful advice’ discourage him; he posted his historic proof online, bypassing peer-reviewed journals altogether; he turned down both the Fields Medal and a million dollar prize; and he has refused professorial posts at prestigious universities. Having made a deliberate decision to eschew the external incentives associated with science as a career, his choices have been written off as examples of eccentric anti-social behaviour. I suggest, however, that he might have simply recognized that the usual rules for success and the usual reward structure of the scientific community can create roadblocks, which had to be avoided if he was to solve a supposedly unsolvable problem.

If we cannot imagine new paradigms, then how can they ever be perceived, much less tested? It should be clear by now that the ‘process of scientific discovery’ can proceed by many different paths.

However, here is one cognitive exercise that can be applied to almost any situation. (i) Notice a phenomenon, even if (especially if) it is familiar and regarded as a solved problem; regard it as if it is new and strange. In particular, look hard for anomalous and strange aspects of the phenomenon that are ignored by scientists in the field. (ii) Look for the hidden assumptions that guide scientists' thinking about the phenomenon, and ask what kinds of explanation would be possible if the assumptions were false (or reversed). (iii) Make a list of possible alternative explanations, no matter how unlikely they seem to be. (iv) Ask if one of these explanations has particular appeal (for example, if it is the most elegant theoretically; if it can generalize to new domains; and if it would have great practical impact). (v) Ask what kind of evidence would allow one to favour that hypothesis over the others, and carry out experiments to test the hypothesis.

The process just outlined is not something that is taught in graduate school; in fact, schools teach a lot about how to test hypotheses but little about how to find good hypotheses in the first place. Consequently, this cognitive exercise is not often carried out within the brain of an individual scientist. Yet this creative tension happens naturally when investigators from two different fields, who have different assumptions, methods and ways of

working, meet to discuss a particular problem. This is one reason why new paradigms so often emerge in the cross-fertilization of different disciplines.

There are of course other, more systematic ways of searching for hypotheses by bringing together seemingly unrelated evidence. The Arrowsmith two-node search strategy [5], for instance, is based on distinct searches of the biomedical literature to retrieve articles on two different areas of science that have not been studied in relation to each other, but that the investigator suspects might be related in some fashion. The software identifies common words or phrases, which might point to meaningful links between them. This is but one example of 'literature-based discovery' as a heuristic technique [6], and in turn, is part of the larger data-driven approach of 'text mining' or 'data mining', which looks for unusual, new or unexpected patterns within large amounts of observational data. Regardless of whether one follows hypothesis-driven or data-driven models of investigation, let us teach our students to repeat the mantra: 'odd is good'!

CONFLICT OF INTEREST

The author declares that he has no conflict of interest.

REFERENCES

1. Grinnell F (2009) *Everyday Practice of Science: Where Intuition and Passion Meet*

Objectivity and Logic. Oxford, UK: Oxford University Press

2. Kuhn TS (1996) *The Structure of Scientific Revolutions*. Chicago, Illinois, USA: University of Chicago Press
3. *The Sixth Sense* (1999) Directed by M. Night Shyamalan. Pennsylvania, USA: A Kennedy/Marshall/Barry Mendel Production, Hollywood Pictures, Spyglass Entertainment
4. Levi-Montalcini R (1989) *In Praise of Imperfection: My Life and Work*. New York, New York, USA: Basic Books
5. Smalheiser NR, Torvik VI, Zhou W (2009) Arrowsmith two-node search interface: a tutorial on finding meaningful links between two disparate sets of articles in MEDLINE. *Comput Methods Programs Biomed* **94**: 190–197
6. Bruza P, Weeber M (2008) *Literature-based Discovery*. Berlin, Germany: Springer

**Neil Smalheiser is Associate Professor in Psychiatry at the University of Illinois at Chicago, USA.
E-mail: neils@uic.edu**

EMBO reports (2013) **14**, 861–865; published online 6 September 2013; doi:10.1038/embor.2013.125