Book Review: The Structure Of Scientific Revolutions

Posted on January 8, 2019 by Scott Alexander

I

When I hear scientists talk about Thomas Kuhn, he sounds very reasonable. Scientists have theories that guide their work. Sometimes they run into things their theories can't explain. Then some genius develops a new theory, and scientists are guided by that one. So the cycle repeats, knowledge gained with every step.

When I hear philosophers talk about Thomas Kuhn, he sounds like a madman. There is no such thing as ground-level truth! Only theory! No objective sense-data! Only theory! No basis for accepting or rejecting any theory over any other! Only theory! No scientists! Only theories, wearing lab coats and fake beards, hoping nobody will notice the charade!

I decided to read Kuhn's *The Structure Of Scientific Revolutions* in order to understand this better. Having finished, I have come to a conclusion: yup, I can see why this book causes so much confusion.

At first Kuhn's thesis appears simple, maybe even obvious. I found myself worrying at times that he was knocking down a straw man, although of course we have to read the history of philosophy backwards and remember that Kuhn may already be in the water supply, so to speak. He argues against a simplistic view of science in which it is merely the gradual accumulation of facts. So Aristotle discovered a few true facts, Galileo added a few more on, then Newton discovered a few more, and now we have very many facts indeed.

In this model, good science cannot disagree with other good science. You're either wrong – as various pseudoscientists and failed scientists have been throughout history, positing false ideas like "the brain is only there to cool the blood" or "the sun orbits the earth". Or you're right, your ideas are enshrined in the Sacristry Of Settled Science, and your facts join the accumulated store that passes through the ages.

Simple-version-of-Kuhn says this isn't true. Science isn't just facts. It's *paradigms* – whole ways of looking at the world. Without a paradigm, scientists wouldn't know what facts to gather, how to collect them, or what to do with them once they had them. With a paradigm, scientists gather and process facts in the ways the paradigm suggests ("normal science"). Eventually, this process runs into a hitch – apparent contradictions, or things that don't quite fit predictions, or just a giant ugly mess of epicycles. Some genius develops a new paradigm ("paradigm shift" or "scientific revolution"). Then the process begins again. Facts can be accumulated within a paradigm. And many of the facts accumulated in one par-

adigm can survive, with only slight translation effort, into a new paradigm. But scientific progress is the story of one relatively-successful and genuinely-scientific effort giving way to a different and contradictory relatively-successful and genuinely-scientific effort. It's the story of scientists constantly tossing out one another's work and beginning anew.

This gets awkward because paradigms look a lot like facts. The atomic theory – the current paradigm in a lot of chemistry – looks a lot like the fact "everything is made of atoms and molecules". But this is only the iceberg's tip. Once you have atomic theory, chemistry starts looking a lot different. Your first question when confronted with an unknown chemical is "what is the molecular structure?" and you have pretty good ideas for how to figure this out. You are not particularly interested in the surface appearance of chemicals, since you know that iron and silver can look alike but are totally different elements; you may be much more interested in the weight ratio at which two chemicals react (which might seem to the uninitiated like a pretty random and silly thing to care about). If confronted with a gas, you might ask things like "which gas is it?" as opposed to thinking all gases are the same thing, or wondering what it would even mean for two gases to be different. You can even think things like "this is a mixture of two different types of gas" without agonizing about how a perfectly uniform substance can be a mixture of anything. If someone asks you "How noble and close to God would say this chemical sample is?" you can tell them that this is not really a legitimate chemical question, unless you mean "noble" in the sense of the noble gases. If someone tells you a certain chemical is toxic because toxicity is a funda-

mental property of its essence, you can tell them that no, it probably has to do with some reaction it causes or fails to cause with chemicals in the body. And if someone tells you that a certain chemical has changed into a different chemical because it got colder, you can tell them that cold might have done something to it, it might even have caused it to react with the air or something, but chemicals don't change into other chemicals in a fundamental way just because of the temperature. None of these things are obvious. All of them are hard-won discoveries.

A field without paradigms looks like the STEM supremacist's stereotype of philosophy. There are all kinds of different schools – Kantians, Aristotelians, Lockeans – who all disagree with each other. There may be progress within a school – some Aristotelian may come up with a really cool new Aristotelian way to look at bioethics, and all the other Aristotelians may agree that it's great – but the field as a whole does not progress. People will talk past one another; the Aristotelian can go on all day about the telos of the embryo, but the utilitarian is just going to ask what the hell a telos is, why anyone would think embryos have one, and how many utils the embryo is bringing people. "Debates" between the Aristotelian and the utilitarian may not be literally impossible, but they are going to have to go all the way to first principles, in a way that never works. Kuhn interestingly dismisses these areas as "the fields where people write books" – if you want to say anything, you might as well address it to a popular audience for all the good other people's pre-existing knowledge will do you, and you may have to spend hundreds of pages explaining your entire system from the ground up. He throws all the social sciences in this bin – you may

read Freud, Skinner, and Beck instead of Aristotle, Locke, and Kant, but it's the same situation.

A real science is one where everyone agrees on a single paradigm. Newtonianism and Einsteinianism are the same kind of things as Aristotelianism and utilitarianism; but in 1850, everybody believed the former, and in 1950, the latter.

I got confused by this – is Aristotelian philosophy a science? Would it be one if the Aristotelians forced every non-Aristotelian philosopher out of the academy, so that 100% of philosophers fell in line behind Aristotle? I *think* Kuhn's answer to this is that it's telling that Aristotelians haven't been able to do this (at least not lately); either Aristotle's theories are too weak, or philosophy too intractable. But all physicists unite behind Einstein in a way that all philosophers cannot behind Aristotle. Because of this, all physicists mean more or less the same thing when they talk about "space" and "time", and they can work together on explaining these concepts without constantly arguing to each other about what they mean or whether they're the right way to think about things at all (and a Newtonian and Einsteinian would *not* be able to do this with each other, any more than an Aristotelian and utilitarian).

So how does science settle on a single paradigm when other fields can't? Is this the part where we admit it's because science has objective truth so you can just settle questions with experiments?

This is very much not that part. Kuhn doesn't think it's anywhere near that simple, for a few reasons.

First, there is rarely a single experiment that one paradigm fails and another passes. Rather, there are dozens of experiments. One paradigm does better on some, the other paradigm does better on others, and everyone argues over which ones should or shouldn't count.

For example, one might try to test the Copernican vs. Ptolemaic worldviews by observing the parallax of the fixed stars over the course of a year. Copernicus predicts it should be visible; Ptolemy predicts it shouldn't be. It isn't, which means either the Earth is fixed and unmoving, or the stars are unutterably unimaginably immensely impossibly far away. Nobody expected the stars to be that far away, so advantage Ptolemy. Meanwhile, the Copernicans posit far-off stars in order to save their paradigm. What looked like a test to select one paradigm or the other has turned into a wedge pushing the two paradigms even further apart.

What looks like a decisive victory to one side may look like random noise to another. Did you know weird technologically advanced artifacts are sometimes found encased in rocks that our current understanding of geology says are millions of years old? Creationists have no trouble explaining those – the rocks are much younger, and the artifacts were probably planted by nephilim. Evolutionists have no idea how to explain those, and default to things like "the artifacts are hoaxes" or "the miners were really careless and a screw slipped from their pocket into the rock vein while they were

mining". I'm an evolutionist and I agree the artifacts are probably hoaxes or mistakes, even when there is no particular evidence that they are. Meanwhile, probably creationists say that some fossil or other incompatible with creationism is a hoax or a mistake. But that means the "find something predicted by one paradigm but not the other, and then the failed theory comes crashing down" over-simplification doesn't work. Find something predicted by one paradigm but not the other, and often the proponents of the disadvantaged paradigm can – and should – just shrug and say "whatever".

In 1870, flat-earther Samuel Rowbotham performed a series of experiments to show the Earth could not be a globe. In the most famous, he placed several flags miles apart along a perfectly straight canal. Then he looked through a telescope and was able to see all of them in a row, even though the furthest should have been hidden by the Earth's curvature. Having done so, he concluded the Earth was flat, and the spherical-earth paradigm debunked. Alfred Wallace (more famous for pre-empting Darwin on evolution) took up the challenge, and showed that the bending of light rays by atmospheric refraction explained Rowbotham's result. It turns out that light rays curve downward at a rate equal to the curvature of the Earth's surface! Luckily for Wallace, refraction was already a known phenomenon; if not, it would have been the same kind of wedge-between-paradigms as the Copernicans having to change the distance to the fixed stars.

It is all nice and well to say "Sure, it *looks* like your paradigm is right, but once we adjust for this new idea about the distance to the stars / the refraction of light, the evidence actually supports

my paradigm". But the supporters of old paradigms can do that too! The Ptolemaics are rightly mocked for adding epicycle after epicycle until their system gave the right result. But to a hostile observer, positing refraction effects that exactly counterbalance the curvature of the Earth sure looks like adding epicycles. At some point a new paradigm will win out, and its "epicycles" will look like perfectly reasonable adjustments for reality's surprising amount of detail. And the old paradigm will lose, and its "epicycles" will look like obvious kludges to cover up that it never really worked. Before that happens… well, good luck.

Second, two paradigms may not even address or care about the same questions.

Let's go back to utilitarianism vs. Aristotelianism. Many people associate utilitarianism with the trolley problem, which is indeed a good way to think about some of the issues involved. It might be tempting for a utilitarian to think of Aristotelian ethics as having some different answer to the trolley problem. Maybe it does, I don't know. But Aristotle doesn't talk about how he would solve whatever the 4th-century BC equivalent of the trolley problem was. He talks more about "what is the true meaning of justice?" and stuff like that. While you can twist Aristotle into having an opinion on trolleys, he's not really optimizing for that. And while you can make utilitarianism have some idea what the true meaning of justice is, it's not really optimized for that either.

An Aristotelian can say their paradigm is best, because it does a great job explicating all the little types and subtypes of justice. A

utilitarian can say *their* paradigm is best, because it does a great job telling you how to act in various contrived moral dilemmas.

It's actually even worse than this. The closest thing I can think of to an ancient Greek moral dilemma is the story of Antigone. Antigone's uncle declares that her traitorous dead brother may not be buried with the proper rites. Antigone is torn between her duty to obey her uncle, and her desire to honor her dead brother. Utilitarianism is… not really designed for this sort of moral dilemma. Is ignoring her family squabbles and trying to cure typhus an option? No?

But then utilitarianism's problems are deeper than just "comes to a different conclusion than ancient Greek morals would have". The utilitarian's job isn't to change the ancient Greek's mind about the answer to a certain problem. It's to convince him to stop caring about basically all the problems he cares about, and care about different problems instead.

Third, two paradigms may disagree on what kind of answers are allowed, or what counts as solving a problem.

Kuhn talks about the 17th century "dormitive potency" discourse. Aristotle tended to explain phenomena by appealing to essences; trees grew because it was "in their nature" to grow. Descartes gets a bad rap for inventing dualism, but this is undeserved – what he was really doing was inventing the concept of "matter" as we understand it, a what-you-see-is-what-you-get kind of stuff with no hidden essences, which responds mechanically to forces (and

once you have this idea, you naturally need some other kind of stuff to be the mind). With Cartesian matter firmly in place, everyone made fun of Aristotle for thinking he had "solved" the "why do trees grow?" question by answering "because it is in their nature", and this climaxed with the playwright Moliere portraying a buffoonish doctor who claimed to have discovered how opium put people to sleep – it was because it had a dormitive potency!

In Aristotle's view of matter, saying "because it's their essence" successfully answers questions like "why do trees grow?". The Cartesian paradigm forbade this kind of answer, and so many previously "solved" problems like why trees grow became mysterious again – a step backwards, sort of. For Descartes, you were only allowed to answer questions if you could explain how purely-mechanical matter smashing against other purely-mechanical matter in a billiard-ball-like way could produce an effect; a more virtuous and Descartes-aware doctor explained opium's properties by saying opium corpuscles must have a sandpaper-like shape that smooths the neurons!

Then Newton discovered gravity and caused an uproar. Gravity posits no corpuscles jostling other corpuscles. It sounds almost Aristotelian: "It is the nature of matter to attract other matter". Newton was denounced as trying to smuggle occultism into science. How much do you discount a theory for having occult elements? If some conception of quantum theory predicts the data beautifully, but says matter behaves differently depending on whether someone's watching it or not, is that okay? What if it says that a certain electron has a 50% chance of being in a certain

place, full stop, and there is no conceivable explanation for which of the two possibilities is realized, and you're not even allowed to ask the question? What if my explanation for dark matter is "invisible gremlins"? How do you figure out when you need to relax your assumptions about what counts as science, versus when somebody is just cheating?

A less dramatic example: Lavoisier's theory of combustion boasts an ability to explain why some substances gain weight when burned; they are absorbing oxygen from the air. A brilliant example of an anomaly explained, which proves the superiority of combustion theory to other paradigms that cannot account for the phenomenon? No – "things shouldn't randomly gain weight" comes to us as a principle of the chemical revolution of which Lavoisier was a part:

> In the seventeenth century, [an explanation of weight gain] seemed unnecessary to most chemists. If chemical reactions could alter the volume, color, and texture of the ingredients, why should they not alter weight as well? Weight was not always taken to be the measure of quantity of matter. Besides, weight-gain on roasting remained an isolated phenomenon. Most natural bodies (eg wood) lose weight on roasting as the phlogiston theory was later to say they should.

In previous paradigms, weight gain wasn't even an anomaly to be explained. It was just a perfectly okay thing that might happen. It's

only within the constellation of new methods and rules we learned around Lavoisier's time, that Lavoisier's theories solved anything at all.

So how *do* scientists ever switch paradigms?

Kuhn thinks it's kind of an ugly process. It starts with exasperation; the old paradigm is clearly inadequate. Progress is stagnating.

> Awareness [of the inadequacy of geocentric astronomy] did come. By the thirteenth century Alfonso X could proclaim that if God had consulted him when creating the universe, he would have received good advice. In the sixteenth century, Copernicus' coworker, Domenico da Novara, held that no system so cumbersome and inaccurate as the Ptolemaic had become could possibly be true of nature. And Copernicus himself wrote in the Preface to the *De Revolutionibus* that the astronomical tradition he inherited had finally created only a monster.

Then someone proposes a new paradigm. In its original form, it is woefully underspecified, bad at matching reality, and only beats the old paradigm in a few test cases. For whatever reason, a few people jump on board. Sometimes the new paradigm is simply more mathematically elegant, more beautiful. Other times it's petty things, like a Frenchman invented the old paradigm and a German the new one, and you're German. Sometimes it's just that there's

nothing better. These people gradually expand the new paradigm to cover more and more cases. At some point, the new paradigm explains things a little better than the old paradigm. Some of its predictions are spookily good. The old paradigm is never conclusively debunked. But the new paradigm now has enough advantages that more and more people hop on the bandwagon. Gradually the old paradigm becomes a laughingstock, people forget the context in which it ever made sense, and it is remembered only as a bunch of jokes about dormitive potency.

But now that it's been adopted and expanded and reached the zenith of its power, *this* is the point at which we can admit it's objectively better, right?

For a better treatment of this question than I can give, see Samzdat's Science Cannot Count To Red. But my impression is that Kuhn is not really willing to say this. I think he is of the "all models are wrong, some are useful" camp, thinks of paradigms as models, and would be willing to admit a new paradigm may be more useful than an old one.

Can we separate the fact around which a paradigm is based (like "the Earth orbits the sun") from the paradigm itself (being a collection of definitions of eg "planet" and "orbit", ways of thinking, mathematical methods, and rules for what kind of science will and won't be accepted)? And then say the earth factually orbits the sun, and the paradigm is just a useful tool that shouldn't be judged objectively? I think Kuhn's answer is that facts cannot be paradigm-independent. A medieval would not hear "the Earth orbits

the sun" and hear the same claim we hear (albeit, in his view wrong). He would, for example, interpret it to mean the Earth was set in a slowly-turning crystal sphere with the sun at its center. Then he might ask – where does the sphere intersect the Earth? How come we can't see it? Is Marco Polo going to try to travel to China and then hit a huge invisible wall halfway across the Himalayas? And what about gravity? My understanding is the Ptolemaics didn't believe in gravity as we understand it at all. They believed objects had a natural tendency to seek the center of the universe. So if the sun is more central, why isn't everything falling into the sun? To a medieval the statement "the Earth orbits the sun" has a bunch of common-sense disproofs everywhere you look. It's only when attached to the rest of the Copernican paradigm that it starts to make sense.

This impresses me less than it impresses Kuhn. I would say "if you have many false beliefs, then true statements may be confusing in that they seem to imply false statements – but true statements are still objectively true". Perhaps I am misunderstanding Kuhn's argument here; the above is an amalgam of various things and not something Kuhn says outright in the book. But whatever his argument, Kuhn is not really willing to say that there are definite paradigm-independent objective facts, at least not without a lot of caveats.

So where *is* the point at which we admit some things are objectively true and that's what this whole enterprise rests on?

Kuhn only barely touches on this, in the last page of the book:

> Anyone who has followed the argument this far will nevertheless feel the need to ask why the evolutionary process should work. What must nature, including man, be like in order that science be possible at all? Why should scientific communities be able to reach a firm consensus unattainable in other fields? Why should consensus endure across one paradigm change after another? And why should paradigm change invariably produce an instrument more perfect in any sense than those known before? From one point of view those questions, excepting the first, have already been answered. But from another they are as open as they were when this essay began. It is not only the scientific community that must be special. The world of which that community is a part must also possess quite special characteristics, and we are no closer than we were at the start to knowing what these must be. That problem—What must the world be like in order that man may know it?—was not, however, created by this essay. On the contrary, it is as old as science itself, and it remains unanswered. But it need not be answered in this place.

II

A lot of the examples above are mine, not Kuhn's. Some of them even come from philosophy or other nonscientific fields. Shouldn't I have used the book's own examples?

Yes. But one of my big complaints about this book is that, for a purported description of How Science Everywhere Is Always Practiced, it really just gives five examples. Ptolemy/Copernicus on astronomy. Alchemy/Dalton on chemistry. Phlogiston/Lavoisier on combustion. Aristotle/Galileo/Newton/Einstein on motion. And ???/Franklin/Coulomb on electricity.

It doesn't explain any of the examples. If you don't already know what Coulomb's contribution to electricity is and what previous ideas he overturned, you're out of luck. And don't try looking it up in a book either. Kuhn says that all the books have been written by people so engrossed in the current paradigm that they unconsciously jam past scientists into it, removing all evidence of paradigm shift. This made parts of the book a little beyond my level, since my knowledge of Coulomb begins and ends with "one amp times one second".

Even saying Kuhn has five examples is giving him too much credit. He usually brings in one of his five per point he's trying to make, meaning that you never get a really full view of how any of the five examples exactly fit into his system.

And all five examples are from physics. Kuhn says at the beginning that he wished he had time to talk about how his system fits biology, but he doesn't. He's unsure whether any of the social sciences are sciences at all, and nothing else even gets mentioned. This means we have to figure out how Kuhn's theory fits everything from scattershot looks at the history of electricity and astronomy and a

few other things. This is pretty hard. For example, consider three scientific papers I've looked at on this blog recently:

- Cipriani, Ioannidis, et al perform a meta-analysis of antidepressant effect sizes and find that although almost all of them seem to work, amitriptyline works best.
- Ceballos, Ehrlich, et al calculate whether more species have become extinct recently than would be expected based on historical background rates; after finding almost 500 extinctions since 1900, they conclude they definitely have.
- Terrell et al examine contributions to open source projects and find that men are more likely to be accepted than women when adjusted for some measure of competence they believe is appropriate, suggesting a gender bias.

What paradigm is each of these working from?

You could argue that the antidepressant study is working off of the "biological psychiatry" paradigm, a venerable collection of assumptions that can be profitably contrasted with other paradigms like psychoanalysis. But couldn't a Hippocratic four-humors physician of a thousand years ago done the same thing? A meta-analysis of the effect sizes of various kinds of leeches for depression? Sure, leeches are different from antidepressants, but it doesn't look like the belief in biological psychiatry is affecting anything about the research other than the topic. And although the topic is certainly important, Kuhn led me to expect something more profound than that. Maybe the paradigm is evidence-based-medicine itself, the

practice of doing RCTs and meta-analyses on things? I think this is a stronger case, but a paradigm completely divorced from the content of what it's studying is exactly the sort of weird thing that makes me wish Kuhn had included more than five examples.

As for the extinction paper, surely it can be attributed to some chain of thought starting with Cuvier's catastrophism, passing through Lyell, and continuing on to the current day, based on the idea that the world has changed dramatically over its history and new species can arise and old ones disappear. But is that "the" paradigm of biology, or ecology, or whatever field Ceballos and Lyell are working in? Doesn't it also depend on the idea of species, a different paradigm starting with Linnaeus and developed by zoologists over the ensuing centuries? It look like it dips into a bunch of different paradigms, but is not wholly within any.

And the open source paper? Is "feminism" a paradigm? But surely this is no different than what would be done to investigate racist biases in open source. Or some right-winger looking for anti-Christian biases in open source. Is the paradigm just "looking for biases in things?"

What about my favorite trivial example, looking both ways when you cross the street so you don't get hit by a bus? Is it based on a paradigm of motorized transportation? Does it use assumptions like "buses exist" and "roads are there to be crossed"? Was there a paradigm shift between the bad old days of looking one way before crossing, and the exciting new development of looking both ways

before crossing? Is this really that much more of a stretch than calling looking for biases in things a paradigm?

Outside the five examples Kuhn gives from the physical sciences, identifying paradigms seems pretty hard – or maybe too easy. Is it all fractal? Are there overarching paradigms like atomic theory, and then lower-level paradigms like organic chemistry, and then tiny subsubparadigms like "how we deal with this one organic compound"? Does every scientific experiment use lots of different paradigms from different traditions and different levels? This is the kind of thing I wish Kuhn's book answered instead of just talking about Coulumb and Copernicus over and over again.

III

In conclusion, all of this is about predictive coding.

It's the same thing. Perception getting guided equally by top-down expectations and bottom-up evidence. Oh, I know what you're thinking. "There goes Scott again, seeing predictive coding in everything". And yes. But also, Kuhn does everything short of come out and say "When you guys get around to inventing predictive coding, make sure to notice that's what I was getting at this whole time."

Don't believe me? From the chapter *Anomaly And The Emergence Of Scientific Discovery* (my emphasis, and for "anomaly", read "surprisal"):

The characteristics common to the three examples above are characteristic of all discoveries from which new sorts of phenomena emerge. Those characteristics include: the previous awareness of anomaly, the gradual and simultaneous emergence of both observational and conceptual recognition, and the consequent change of paradigm categories and procedures often accompanied by resistance. **There is even evidence that these same characteristics are built into the nature of the perceptual process itself.** In a psychological experiment that deserves to be far better known outside the trade, Bruner and Postman asked experimental subjects to identify on short and controlled exposure a series of playing cards. Many of the cards were normal, but some were made anomalous, e.g., a red six of spades and a black four of hearts. Each experimental run was constituted by the display of a single card to a single subject in a series of gradually increased exposures. After each exposure the subject was asked what he had seen, and the run was terminated by two successive correct identifications.

Even on the shortest exposures many subjects identified most of the cards, and after a small increase all the subjects identified them all. For the normal cards these identifications were usually correct, but the anomalous cards were almost always identified, without apparent hesitation or puzzlement, as normal. The black four of hearts might, for example, be identified as the four of either spades or hearts. Without any awareness of trouble, it was immediately fitted to one of the conceptual categories prepared by prior experi-

ence. One would not even like to say that the subjects had seen something different from what they identified. With a further increase of exposure to the anomalous cards, subjects did begin to hesitate and to display awareness of anomaly. Exposed, for example, to the red six of spades, some would say: That's the six of spades, but there's something wrong with it— the black has a red border. Further increase of exposure resulted in still more hesitation and confusion until finally, and sometimes quite suddenly, most subjects would produce the correct identification without hesitation. Moreover, after doing this with two or three of the anomalous cards, they would have little further difficulty with the others. A few subjects, however, were never able to make the requisite adjustment of their categories. Even at forty times the average exposure required to recognize normal cards for what they were, more than 10 per cent of the anomalous cards were not correctly identified. And the subjects who then failed often experienced acute personal distress. One of them exclaimed: "I can't make the suit out, whatever it is. It didn't even look like a card that time. I don't know what color it is now or whether it's a spade or a heart. I'm not even sure now what a spade looks like. My God!" In the next section we shall occasionally see scientists behaving this way too.

Either as a metaphor or **because it reflects the nature of the mind**, that psychological experiment provides a wonderfully simple and cogent schema for the process of scientific discovery.

And from *Revolutions As Changes Of World-View*:

> **Surveying the rich experimental literature from which these examples are drawn makes one suspect that something like a paradigm is prerequisite to perception itself. What a man sees depends both upon what he looks at and also upon what his previous visual-conceptual experience has taught him to see. In the absence of such training there can only be, in William James's phrase, "a bloomin' buzzin' confusion." In recent years several of those concerned with the history of science have found the sorts of experiments described above immensely suggestive.**

If you can read those paragraphs and honestly still think I'm just just irrationally reading predictive coding into a perfectly innocent book, I have nothing to say to you.

I think this is my best answer to the whole "is Kuhn denying an objective reality" issue. If Kuhn and the predictive coding people are grasping at the same thing from different angles, then both shed some light on each other. I think I understand the way that predictive coding balances the importance of pre-existing structures and categories with a preserved belief in objectivity. If Kuhn is trying to extend the predictive coding model of the brain processing information to the way the scientific community as a whole processes it, then maybe we can import the same balance and not worry about it as much.