1 UNDERSTANDING SCIENCE

> It is necessary to get behind someone,
> before you can stab them in the back.
>
> Sir Humphrey Appleby
> *Yes, Prime Minister* (BBC), 1987

We want to teach you how to overthrow a scientific theory.

That might sound a little “anti-science”, but actually you’ll be doing scientists a favour. We learn something when bad ideas are exposed. Science often progresses by supporting the reigning ideas, but at other times it has been necessary to storm the castle and install a new monarch. That’s how many great scientists rose to fame. *Vive la révolution!*

But you’ve got to do it right, and that’s what this book is about. Revolutions fail for attacking the wrong target, following the wrong tactics, and underestimating the old order. Scientific theories are ideas about the natural world. They claim to know what the universe is like and how it behaves. This tells you how to dethrone a scientific idea: take up the weapon of *observations* and aim squarely at its *predictions*. Show that it can’t handle the truth. And be ready with your new monarch when the throne is vacant.

To do all that, you must know your enemy. These wise words from Sun Tzu (or, if you prefer, Rage Against the Machine) are very relevant here: before you can launch a scientific revolution, you need to know the facts, and you need to know the ruling theory and its predictions. Theories aren’t installed on the scientific throne by accident, so do your homework.

This book will hand you the facts, point you in the direction of the castle walls, and wish you the very best of luck. In particular, we’ll be looking at the biggest scientific target of them all.

This book is about the universe.

It's about how we observe the universe, either with our naked eyes, or with the many telescopes that now survey the heavens, sensitive to radiation our eyes cannot perceive. But more than that, this book is about how we *understand* the workings of the universe, from its fundamental properties to its largest features. It's about how we put the pieces together.

Current scientific orthodoxy paints a picture of the cosmos that has been built up from many centuries of observation, experimentation, and hard thinking. Great minds throughout scientific history have laid the groundwork, carefully studying the basic rules of motion, space, time, atoms, light, and gravity, to provide the mathematical tools we need to comprehend the changing heavens. Today, *cosmology* – the study of the universe as a whole – is hailed as a paradigm of scientific success.

But what a strange picture! Many find modern cosmology completely unbelievable. The universe, we are told, was born almost 14 billion years ago in a hot and fiery event, cheekily named the *big bang*. At its beginning, everything was compressed into a point of infinite density and infinite temperature. In the aftermath, the universe is *expanding*, but it's not expanding *into* anything. Space itself is stretching.[1] Today, the galaxies we observe in the night sky all appear to be moving away from us. A vast sea of galaxies, stars, and planets fills this expanding space, but because light only moves so fast, most of this universe will be forever beyond the reach of our telescopes, over the *horizon*.

What about the stuff in the universe? Compiling an inventory would appear to be straightforward, if painstaking: just add up all of the stars, planets, and gas clouds that inhabit galaxies and the spaces between the galaxies. But cosmologists say that there is more to the universe than the stuff that we can see. Much, much more. A *dark side* of the universe, which we cannot touch or feel, dominates its energy budget and controls its expansion.

Firstly, modern cosmology tells us that there is *dark matter*. This stuff pervades every galaxy, holding stars in their orbits with its gravitational pull. But dark matter emits no light of its own, and so remains unseen by our telescopes. Stars illuminate the heavens, but dark matter accounts for more than 85% of the mass in the universe. The atoms that make up you and me, stars and planets, are little more than frosting on the cosmic cake.

And then modern cosmology tells us about *dark energy*, a substance as pervasive yet more elusive than dark matter. The case for dark energy was made only in the past few decades. We are told that this substance governs the dynamics of the universe on its largest scales, causing the expansion to accelerate, and driving us towards a cold, dark, dead future.

Why would anyone believe all of that?

A quick internet search turns up plenty of websites, blogs, and videos decrying modern cosmology as wrong, illogical, or even a conspiracy of the scientific establishment that suppresses voices of criticism. Modern cosmology, they claim, is a sham, purposefully distorted and hyped in the hunt for funding. Cosmologists are little more than a self-serving cabal, crushing all opposition.

Maybe, dear reader, you are one of these revolutionary voices, wanting to put science right. Maybe you have ideas about the laws of physics and how they impact our view of stars and galaxies. Maybe you have tried to engage with established astronomers and cosmologists to express your ideas and explain why their view is misguided, but have received a cold shoulder. Why are academics, locked up in their ivory towers, so sure they are right?

Our goal is to explain how physicists, astronomers, and cosmologists developed their picture of how the universe behaves, why they talk about it the way they do, and to tell you what you need to do to confront their strange ideas and begin a revolution. We'll help you build a strategy to battle modern science on a more even playing field, and to ensure that your voice is heard amongst the scientific din.

Just What Is Science?

Warning: the following discussion is very physics-o-centric!

To an outsider, science can be a difficult beast to understand. The media – and especially health advertisements – often tell us

"Science says ..." and "Scientists have discovered that ...", but science is not a single, monolithic enterprise. The scientific community consists of many thousands of individuals who often specialize in a narrow set of fields. Some scientists design experiments, some perform observations, and others wrestle with abstract mathematical theories. All spend far too much time in front of a computer. But what is the goal of science?

We begin with an important point: scientists try to predict the future.

If you are not familiar with the workings of science, this might seem a little strange. A flick through popular science magazines such as *New Scientist* or *Scientific American* will reveal stories that focus on *big* scientific questions such as "What is spacetime *really*?" and "What is quantum mechanics *really* telling us about the universe?" But we can't attack these deeper, foundational issues without some help.

In particular, it will help if we can bring these lofty questions down to a practical level. This is the part of science that plays "what if" games, constructing possible physical scenarios and teasing out implications. What if particles of light (*photons*) possessed a tiny amount of mass? What if a cloud of matter collapsed under its own gravity? What if I heat some hydrogen to 10 million degrees? Answering such questions requires more than a vivid imagination: we need our ideas to be translated into the language of mathematics. Sometimes, entirely new mathematical ideas need to be discovered and developed.

The goal of this precision is to connect our ideas to data. Can our new idea account for existing observations of the universe? And, just as importantly, are there any future observations that we could make that would provide further evidence for or against our idea? Can we get one step ahead of nature?

Take gravity as an example. In the 1680s, Isaac Newton published his incredibly successful theory of gravity. With one simple law, he explained how apples fall and how the planets move. Using Newton's law, Edmund Halley was able to predict the future motion of the comet that now bears his name. In 1705, he calculated that it would return in 1758. Sure enough, on Christmas day, it was spotted by a German farmer. Sadly, neither Newton nor Halley was alive to see it.

However, in the mid 1800s, Newton's theory was struggling. Astronomers had discovered that the innermost planet, Mercury, was orbiting slightly out of place, as if pulled by an unseen planet near the Sun. Some even claimed to have observed this newest member of the Solar System, which had been dubbed "Vulcan". Other astronomers, however, could not confirm this sighting. As evidence evaporated and Vulcan consistently failed to turn up where it was predicted to be, this mysterious shortcoming of Newton deepened into a crisis.

In the early 1900s, Einstein proposed his radical new theory of gravity – called the *general theory of relativity* – in which space and time themselves warp, stretch, and wobble. While Einstein's prediction of the orbit of Mercury is only slightly different from that of Newton, that was enough to beautifully align theory with observation. The planet Vulcan was banished to the scientific scrap heap.

Einstein's explanation of Mercury's orbit is impressive, but, like Newton's explanation of the motions of the planets, it comes after the data. We knew about the orbit of Mercury before Einstein proposed his theory. This is sometimes called a "post-diction".

Is there anything wrong with post-diction? We certainly can't discard all the evidence we found before a theory was proposed. Our scientific results would be swayed by something as contingent as what historical order we human beings happened to discover some idea or perform some experiment. That could depend on all sorts of irrelevant factors, like whether Thelma the Theorist took a few days off, or Xavier the Experimenter had a particularly good breakfast.

In principle, prediction and post-diction carry equal weight. But in practice we want to know whether a theory explains the data *naturally*, rather than being glued together from makeshift bits and pieces. Sometimes we can judge this by directly examining the assumptions that underlie the theory. But it is not always easy to tell. Predictions dispel this worry: you can't cook up a theory just to explain data if you don't have the data yet. If a theory correctly predicts the result of an experiment that we haven't done yet, then that is impressive.

So, when a new theory is proposed, we start asking "what if" questions. With Einstein's theory in hand, we have a whole new theoretical universe to explore. We look for new opportunities to test whether these ideas are correct. Einstein predicted that gravity would bend the path of light rays moving near massive objects. Famously, this effect was observed by the British astronomer Sir Arthur Eddington during a solar eclipse in 1919, confirming general relativity's predictions and propelling Einstein to further international fame.

Einstein's theory continues to make successful predictions. In 2015, a hundred years after Einstein's announcement of his new theory, scientists confirmed a hugely important prediction of general relativity: gravitational waves. Space and time can ripple. The discovery of these feeble vibrations, typically swamped by the everyday groans and grumbles of life on Earth, required half a century of effort to build an extraordinarily sensitive detector called the Laser Interferometer Gravitational-Wave Observatory (or LIGO for short). The results were spectacular, with the first signal revealing the merging of two black holes 3 billion years ago in the distant universe. LIGO has opened up a new window on the cosmos.

While Einstein's name is synonymous with scientific genius, you don't need to venture far into the outskirts of the internet to find many people who object to his ideas. Some play the man, rather than the ball, accusing him and the scientific community of outright fraud. Relativity is obviously crazy, they say, but it allows fat-cat scientists to keep feeding off the public purse. Others will decry the "logic" of relativity, often voicing a dislike of the notion of curved space and time, and even accusing the scientific establishment of wilful blindness to their unrecognized genius.

But science holds onto general relativity, not because of hero worship of Einstein, or because we are part of a secret conspiracy. Rather, we use his theory because it works. Physicists dream of proving Einstein wrong; we just haven't been able to do it. We are devising new ways to draw out predictions, and building new experiments to test those predictions.

As we said at the beginning of the chapter, the reigning monarchs of science didn't get there by accident. But they are always vulnerable, because every prediction is a chance to fail. So, what do you need if you want to revolutionize science? A new monarch. You need a model!

Just What Is a Model?

The word *model* has several meanings in the English language, and this can lead to some confusion when talking about a "scientific model". Anarchic comedian Alexei Sayle once said, "my girlfriend's a model. She's an Airfix kit of a Stuka dive bomber!"

We can understand the most important thing about a scientific model by thinking of a model house. Everything in the model is to scale, with one-twentieth size windows, doors, rooms, cupboards, and more. The useful thing about this model is that we can use it to answer questions about the real house. Suppose you want to know whether you can rearrange the living room to incorporate that new sofa you've had your eye on. You can answer this question with the model. If we make a one-twentieth scale model of the new sofa, then we can easily rearrange the model room to see if everything fits. For an accurate model, if the model sofa fits into the model house, then the real sofa would fit into the real house.

This is the crucial feature of a model: using the right translation, we can turn a problem in the real world (will the sofa fit in my living room?) into a problem in the model (will the model sofa fit in the model living room?). We then solve the problem in the model. If the model is an accurate representation of reality, then we have also solved the problem for the real world.

In the case of a model house, the translation between the model and reality is simple: it's just 20 times smaller. For a scientific model, the mathematical framework can be more complicated, but the crucial feature is the same: we can translate a question about the real world into a question about the model. Because we can relate between the two, we can make predictions. We can ask questions such as “what if I performed such-and-such experiment?”

Let’s take another look at Newton’s model for gravity. (We’re physicists. We like Newton!) We can express his idea in words: gravity will produce a force between two masses, whose magnitude is proportional to the product of the two masses, and inversely proportional to the square of the distance between the masses. That’s interesting, but not much use to a working scientist. To a scientist, the useful form of Newton’s law of gravity looks like this:

$$\vec{F} = -G\frac{M_1\,M_2}{r^2}\,\hat{r}$$

If you are not a fan of mathematics, and if this equation looks like little more than gobbledygook, don’t worry too much. We can look at this like a machine, where we input two values for the masses, M_1 and M_2, and the distance between them, r, and this machine returns the gravitational force between them. The other number in the equation is G, which is known as Newton’s gravitational constant. It scales the numbers so the result has the correct unit (which, for force, is the *newton*). Finally, $\hat{r}$ (“r” with a little hat) is known as a *unit vector*; it tells you that the force pulls the masses towards each other. But what can you do with this bit of mathematics?

We turn to Newton’s laws of motion. We can state the idea in words as “forces cause objects to change their speed and direction of travel”. But as we have noted, it’s the mathematical version of the law that allows us to make precise predictions:

$$\vec{F} = m\,\vec{a}$$

This equation might be familiar from high-school physics; F is the force, a is the acceleration, and m is the mass. Combining these equations, we can start with information about the position and velocity (which encodes speed and direction) of the objects in the system at a particular time, and transform it into a prediction about the future of the system. For example, if we know where all the Solar System's planets are today, and how fast and in what direction they are moving, we can calculate where they will be at any future time.

The point of all physical models – Newton's, Einstein's, and anyone else's – is that we can ask questions about the universe. Given where I saw the planet Mercury last night, where will I see it tonight? By how much will the path of a light ray bend as it passes close to the Sun? We can ask Newton's model, and we can ask Einstein's model, and then we can actually look at the universe to see if either is correct.

The lesson is that if you are going to revolutionize science, you need a mathematical model. Words will not do. As scientists, we regularly get emails and letters espousing new ideas about the cosmos, from theories about fundamental particles to new interpretations of galaxy redshifts and the expansion of the universe. Surprisingly often, the author confesses that they are unable to express these ideas mathematically. I'm sure my idea is correct, they say, I just need some help working out the mathematics. To a scientist, and particularly to a physicist, this is a bit like saying "I have a great idea for a symphony; I just need some help with the musical notes" or "I'm sure I could do brain surgery; I just need some pointers on where to start cutting."

For a physicist, you don't really *have* a theory until you can think about it clearly enough to put it in mathematical form. Without precise predictions, it is too easy to fool yourself into thinking that the data is consistent with your idea. We need to predict measurements and observations, so that we can hold this mathematical model up to nature.[2]

What Makes a *Good* Scientific Model?

What does a scientist want in a scientific model? We have emphasized that your model must present a precise, quantitative picture of the universe, one that allows us to predict the results of experiments. But this is not the only criterion that scientists use. Historians and philosophers of science, by studying how scientists actually argue for and against theories, have proposed sets of *theoretical virtues*, that is, traits of a good scientific idea.

Not everyone agrees about all of the virtues, of course, but there is a common core that scientists will recognize. We will look at a recent list of twelve theoretical values (TVs) compiled by historian Mike Keas.[3] His list is helpfully comprehensive: while the twelve values overlap somewhat, each pinpoints something important about good scientific theories.

The first three relate to how your theory handles the evidence.

TV1. Evidential accuracy: your theory accounts for or fits the data well.

TV2. Causal adequacy: your theory posits causes that account for the effects we see in the data.

TV3. Explanatory depth: your theory applies to a wide range of scenarios.

Clearly, if your theory is correct, or at least approximately correct, then it should explain the data (TV1). All the data! Cherry-picking – focusing on the results that your mathematical model can describe, while ignoring those where it fails – is a scientific sin. This is a sure road to being ignored by the scientific community.

But scientists want more from a theory than this. The theory that the continents can move over the surface of the Earth explains why they appear to fit together like a jigsaw puzzle. But when it was first proposed, this theory was rightly criticized because it lacked causal adequacy (TV2): it didn't tell us *how* the continents moved. Frankly, no one had much of an idea of how something as large as a continent *could* slide around the Earth's surface. The theory of plate tectonics added the all-important details.

But the theory of plate tectonics does even more. It has implications for a wide range of facts about the Earth's surface: how mountains form, how lava comes to the surface in volcanos, and the origin of earthquakes along fault lines. Scientists prefer broad theories that explain a lot about the universe (TV3).

The next three virtues are about how your theory hangs together.

TV4. Internal consistency: your theory does not contradict itself.

TV5. Internal coherence: your theory's various parts fit together neatly and naturally, with no internal tension or tacked-on assumptions.

TV6. Universal coherence: your theory sits well with other warranted principles.

Obviously, if one part of your theory contradicts another part, then it has self-destructed (TV4). But more generally, a theory can fail to hang together in a convincing way. It may need too many ad hoc bits and pieces, tacked on for no good reason. The paradigmatic example is *epicycles*: when the ancient Greeks observed that the planets don't move in perfectly circular orbits, Claudius Ptolemy proposed that the planets move along "circles on circles". There is no deep or natural reason for these epicycles. They explain the data, but in a clumsy way. Even Copernicus's model, which correctly placed the Sun at the centre of the Solar System, needed epicycles. A more coherent model of the Solar System would await the work of Johannes Kepler and Isaac Newton.

Also, if your theory violates established principles, then scientists are suspicious. Suppose your theory, despite fitting all the experimental facts, fails to conserve energy in certain circumstances. This is reason to worry since conservation of energy is a time-tested principle. We aren't going to discard, or even grant a few exceptions, to this principle on a whim.

The next three virtues are aesthetic.

TV7. Beauty: your theory strikes scientists as beautiful.

TV8. Simplicity: your theory explains the same facts with fewer starting assumptions.

TV9. Unification: your theory explains more *kinds* of facts than rivals, relative to its starting assumptions.

TV7 might surprise you, but a long line of important physicists can be marshalled in support. The physicist Paul Dirac went as far as saying, "it is more important to have beauty in one's equations than to have them fit experiment." Henri Poincaré spoke of “the intimate beauty which comes from the harmonious order of its parts and which a pure intelligence can grasp”. In a letter to Albert Einstein, Werner Heisenberg stated the following:

> *That these interrelationships display, in all their mathematical abstraction, an incredible degree of simplicity, is a gift we can only accept humbly. Not even Plato could have believed them to be so beautiful. . . .*
>
> *You must have felt this too: the almost frightening simplicity and wholeness of the relationships which nature suddenly spreads out before us and for which none of us was in the least prepared.*[4]

Why would scientists – hard-nosed, no-nonsense, just-the-facts people, supposedly – be concerned with something as subjective and nebulous as beauty?

One important reason is found in physics as we know it. Many of the theories of modern physics, when one can fluently speak the appropriate mathematical language, are strikingly elegant, symmetric, and ingenious. Einstein’s general theory of relativity is a great example. The central insight – gravity is the warped geometry of space and time – is a piece of creative brilliance, arguably the greatest single theoretical insight in the history of physics. The connection of this idea to the beautiful mathematics of curved spaces seems almost inevitable; in the appropriate language, the theory is stated very simply, even poetically. The range of observed data that the theory explains is enormous. Within one concise equation, sitting innocently on the blackboard, one feels the weight of a million worlds to be explored – expanding universes, spinning black holes, slowing clocks, bending light rays, gravitational waves, and even the possibility of time travel.

But what exactly does a physicist mean when they say that a theory is beautiful? Most physicists, we think, would say that they don’t exactly know.[5] As with all experiences of beauty, it is vivid and immediate, but not easy to describe. The other theoretical virtues try to explain a bit more.

One common feature of mathematical beauty is simplicity. The scientific enterprise generates an extraordinary amount of data every day. A single night at a large telescope will easily fill a computer hard drive, and that's just astronomy. One of the lovely things about beautiful theories is how succinctly they can be stated: a few postulated kinds of entities and laws are all you need to explain a mountain of data. It's like cracking a code or solving a puzzle.

Another important feature is *unity*. Beautiful scientific theories give a sense of how the whole of nature fits together into one grand picture. Ptolemy, in the second century AD, glimpsed something wonderful in his theory of the Solar System:[6]

> *I know that I am mortal by nature and ephemeral, but when I trace at my pleasure the windings to and fro of the heavenly bodies, I no longer touch earth with my feet. I stand in the presence of Zeus himself and take my fill of ambrosia.*

When a theory unifies our view of nature, it explains many seemingly disconnected facts. It explains how nature holds together. (Part of the reason for the beauty of our scientific laws, of course, is that they describe our beautiful universe.)

The final set of theoretical virtues concern how a theory fares over time, as it is examined, extended, and utilized by the scientific community.

TV10. Durability: your theory has survived testing by new experiments and new data.

TV11. Fruitfulness: your theory has pointed to new discoveries, such as successful novel prediction and unification.

TV12. Applicability: your theory has led to the development of new technology.

If your theory has really peeked into the inner workings of nature, then it should continue to illuminate the way the universe works. It should be able to explain new data that we gather (TV10).

As we discussed above, the advantage of prediction – theory first, observation second – over post-diction is that we can be sure that the theory wasn't jerry-rigged to explain the data. It is particularly impressive if (TV11) a new theory suggests an experiment or observation we hadn't even thought to make. Physicists love this: it gives us something to do.

One particular form of evidence for a new theory of nature is whether it can be put to technological use (TV12). This isn't true of all theories, of course: working out how a galaxy makes its stars isn't going to lead to star-making factories any time soon. But if we really understand how electrons and gravity work, then this should help us design and build devices that make use of this knowledge. We understand electrons well enough to build CRT televisions, and we understand gravity well enough to put satellites in orbit.

We note again that while these criteria are overlapping, each captures something important about what we expect from a good scientific theory. Our best theories fulfil most of them. But if your theory is all numerical coincidences and no physical insight, if it lacks unifying principles, if it ignores huge swathes of data, if it is made from an odd assortment of unconnected assumptions, if it is mathematically mundane, if it doesn't seem to produce new predictions or new ways of looking at the physical world, then scientists will have a sneaking suspicion that there's something missing.

The Real Process of Science Part 1: Publishing

So, now you have your theory. You've thought through its foundations, its implications, and its applications. You have a clear vision of the real nature of the universe. You're sure it ticks off enough of the theoretical virtues to be taken seriously. You're ready to explain your discovery to anyone who will listen. How do you make scientists pay attention?

Some advice: writing letters or emails to prominent scientists probably won't work. They receive too many of those. If well-known scientists responded to all the supposedly revolutionary ideas that landed in their inbox, they'd have time for little else.

A scientist – even a famous one – might answer a concise question if you're polite, especially if you open with, "I read your book/paper/article and I have a query." But your glorious 400-page unsolicited Word document, with seven different fonts, Microsoft Paint diagrams, and ALL CAPS, is headed straight for the trash. We haven't studied this phenomenon systematically but, in our experience, there is a remarkably strong correlation between how utterly bonkers an idea is and how many fonts are used by **its** *defender*. Even too much **bold** and *italics* is a red flag. Don't do this. OR UNLEASH YOUR INNER CAPS LOCK. Or use multiple colours. Or write a webpage that is one enormous paragraph. Good typesetting is subtle. If you really want your document to look scientific, learn to use a program called LaTeX.

Raging in internet chat rooms and blog comments won't achieve much either. In fact, scientists are unlikely to hear of your idea even if you write a book about it; new theories in physics are rarely first published in a book.

Popular culture often talks about the scientific method as an idealistic "wash, rinse, and repeat" procedure: you have an idea, test the idea in an experiment, then accept or reject your idea and start again with a new idea or experiment. Real science isn't quite this neat. Different fields approach nature differently. But there is at least one key common element to the practice of science: publication.

Publications in a scientific journal are the currency of science. On the résumé of a scientist, there will be a list of their publications and the journals in which they appear. Why are journal publications so highly regarded? Because this is the *first step* of peer review by the broader research community. To get published in an established journal, a manuscript is submitted to an editor. The editor takes the first glance at a paper, just to check that it looks like, well, a paper: is it laid out appropriately, does it present a background, approach, methods, results, and discussion, and does it look like it makes at least a little bit of sense? For articles that overcome this low bar, the editor will then seek comment from external referees.

If it passes their review, sometimes after a few revisions, then it is published.

Scientists have a choice about which journal they send their paper to. We look for a good reputation, proper editorial oversight, and peer review by other scientists. The hierarchy of journals is sometimes ranked by something called *impact factor*, which measures how many times (on average) a paper in that journal is cited over a particular time period. It's an awful statistic, strongly skewed by individual papers that are very influential and heavily cited, shining glory onto other articles that share the same pages. This discussion of journals, impact factors, and gaming will cause some of our colleagues' blood to boil. Scientists don't particularly appreciate reducing their work down to a blunt statistic. Still, impact factor is easy to calculate.

High-impact journals such as *Nature* and *Science* cover all of science and reject most papers that they receive. Well, you can't be exclusive without excluding. They require submitted work to be highly novel and innovative. This doesn't mean that it is right; in some ways, it means that the work is riskier.

Then there are the "bread-and-butter" journals in a scientific field. In astronomy, for example, most papers are submitted to journals such as the *Astronomical Journal*, *Astrophysical Journal*, *Monthly Notices of the Royal Astronomical Society*, and *Astronomy & Astrophysics*. These journals accept papers from all over the world. While their refereeing is robust enough to see many papers rejected, they aren't necessarily looking for the latest sensation. As well as these larger journals, smaller journals can be found attached to national societies or observatories. These provide a place for less impactful research, conference proceedings, technical reports, or student projects with an interesting result.

The reward for publishing in a higher impact journal is a more impressive CV for a scientist and a boost for their university up the rankings. As a result, it is not unheard of for a paper to be submitted initially to the highest impact journal, only to be rejected and be resubmitted to a lower ranking journal, and so on.

The lesson is that merely getting your paper published isn't as important as *where* it is published. Scientists will take note of where your paper has appeared, as this will signal the level of refereeing and editorial review it has received. If you claim to have revolutionized cosmology, waving your published paper as evidence, it had better be published in a high-ranking physics or astronomy journal. Appearing in the *Bulgarian Journal of Basket Weaving* will not garner much attention.

The Real Process of Science Part 2: Peer Review

As part of its assessment at a journal, a submitted paper will be sent out to *referees*: experts in the field who provide detailed comments on the paper, in particular identifying any glaring mistakes. Referees also identify whether the work is interesting, significant, and potentially of use to the scientific community. Papers must pass this interrogation to be published.

But don't overestimate the importance of a thumbs-up from a referee. A passing grade does not mean that the paper is correct, or has been accepted as immutable scientific orthodoxy. What it means is that some scientists have judged that it is not obviously wrong and probably of use to other researchers.

And don't think that the assessment of an article for a journal by referees is all there is to peer review. It is merely *the first step* in assessment by the scientific community. Peer review continues long after an article has been published, in the continual assessment of whether the idea is interesting, accounts for new data, and spurs new ideas. The most definitive indicator of scientific impact, craved by scientists and by university bean-counters equally, is the *citation*.

When scientists list their journal papers, they usually also list the number of times their article has been cited in someone else's work. Why? Because this means that others in the scientific community have studied their work and found it useful. Oscar Wilde's quip that "there is only one thing in life worse than being talked about, and that is not being talked about" is true for scientific papers. The worst thing that can happen to a scientist is that nobody takes notice of your work. At best, no-one has realized its importance *yet*. This can happen: in 1967, Steven Weinberg published a theory in particle physics for which he would later win the Nobel Prize.[7] But in the first four years after publication, it was only cited twice. Once his colleagues realized its importance, however, this soon changed. It has been cited, on average, four times a week for the last 50 years, for a total of 11,000 citations.

But let's be honest. A lack of citations is more likely to be a sign that it was not an interesting or good piece of work.

Citations are an important measure of scientific impact. Universities and funding bodies are continually assessing the research activities and output of their staff. They are glued to international university rankings like a football fanatic glued to the league table; positioning in the table brings renown and prestige, and with it, better chances at securing research funding and foreign students. Universities and funding bodies want impact to be immediate.

Scientists, however, tend to have mixed feelings about citations. Those of us engaged in so-called *blue-sky* fields, which study fundamental questions about the structure of matter and the universe at large, know that it can take a long time for some research to be appreciated and result in scientific impact. Like Weinberg's paper, plenty of good scientific work in astronomy and cosmology, such as the prediction of gravitational lensing or gravitational waves, or the initial observational clues of dark matter, waited for many decades before its importance was realized.

At this point, we tip our hats to our mathematician cousins. Physics is grateful for the constant flow of new mathematical ideas that allow us to understand the workings of the universe. However, advances in mathematics may lie around for centuries before being noticed for what they are. This means that mathematicians can die unrecognized in their lifetime.

The Real Process of Science Part 3: Presentation

As well as publishing material in recognized journals, scientists spread the word about their work by presenting at conferences and workshops. These are quite varied occasions, some with a handful of people, some with hundreds, some covering broad areas of astronomy and cosmology, others focused on very specific topics.

For example, these words are being written on a Boeing-737, winging its way to Cairns in northern Queensland, Australia. The authors are heading for “Diving in the Dark”, a conference about the dark side of the universe, both dark matter and dark energy. While at the airport, we discovered that Cairns is hosting a second astronomical meeting this week, specifically focused upon the centre of the Milky Way galaxy. At each meeting, astronomers will present, discuss, and argue about the latest research, share new ideas, and forge new collaborations. Walking away from such meetings, an astronomer will have a feel for the latest research in their area, have gained new ideas of where to take their research and, importantly, have promoted their own ideas to the community.

No matter how lucid your writing, you just can’t beat a good audio-visual presentation for conveying the big picture of your idea. Talking about your latest research at major meetings makes it much more likely that other researchers will read your papers. The goal is not just appreciation: you want them to run with your idea, thinking of new implications and new ways of testing those implications. When they write their next paper, they may even give you one of those treasured citations.

As with publications, presentations should be at significant conferences that highlight the latest data and results. Wowing the monthly meeting of your local Quidditch society doesn’t do much for your theory’s credentials. When you’re a young researcher, you may only get 5 or 10 minutes, so have an elevator-pitch version of your ideas ready!

Does this sound like a lot of hard work? It is. And that’s the way it should be. Revolutions shouldn’t be easy!

Most importantly, be prepared to fail. In fact, accepting that you are wrong and moving on is a key attribute of a good scientist. Sticking to your guns when all the data says that your idea is wrong is the stuff of pseudoscience and irrelevance.

Admittedly, the evidence against a theory isn't always completely definitive, so keeping an option on the table can be appropriate. But, as evidence mounts and the scientific community moves on, ideas that can't keep up are continually being discarded.

Does Science Want a Revolution?

Isn't this all a bit optimistic? It assumes that scientists are, at least on average, rational, reasonable, and unbiased, immune to cognitive distortions such as groupthink, confirmation bias, and lack-of-coffee. But scientists are human beings, too. Maybe they're just doing whatever it takes to get another research grant or win the approval of their peers.

In particular, if scientists *wanted to* suppress new ideas, then the mechanisms we have described – journals, editors, peer review, and citations – seem to provide an all-too-easy way to do it: you tread down an idea by preventing it from being published. Have reviewers reject it or, even more effectively, have the editors of the journal reject it before it is sent out for review. Without the minimal stamp-of-approval from a journal, most of the time-starved scientific community won't consider it. And so, an idea – even a good one – can be suppressed by the scientific illuminati.

Now, we can't discuss every supposed case of scientific censorship, nor do we want to defend everything that every scientist or journal editor has ever done. We cannot deny that the history of science is littered with theories that took too long to die, cherished by scientists who fought hard to ignore the evidence that the time for their pet theory had passed. This is not a new phenomenon, as a great scientist of the early twentieth century, Max Planck, was reputed to say, science progresses "one funeral at a time".

We're formulating a game plan here, not settling every grievance. What's important is that while science, like every human activity, isn't perfect, we can make these imperfections work for us.

Firstly, as a scientific revolutionary, you must realize that cranks exist, as do nuts, loons, and buffoons. In addition, there are well-meaning amateurs without the breadth of knowledge or mathematical ability to contribute meaningfully to the scientific enterprise. Explaining to each of them why they are wrong would be a colossal waste of time. More scientific and would-be scientific work is published every day than any person can read and digest. If you want to be noticed, you'll have to do *something* to distinguish yourself from people that scientists have good reason to ignore.

Secondly, most scientists acknowledge an obligation to communicate their results to the public. For example, scientists who study the natural world will want to inform people (and their elected representatives) of threats to the environment or the survival of a particular species. Many scientific advances have potential technological applications, which will need to be explained to potential investors and customers. Astronomy, in particular, is driven by the human thirst for knowledge, rather than being directly aimed at new technology. We are mostly supported by government research grants, so having the tax-paying public share our curiosity about the universe is in our best interests.

If scientists expect their ideas to be respected by the layperson, we cannot act like a secret society. We must show openness to new ideas, or, at the very least, not tamper with the scientific ecosystem: reason and evidence, peer review and citations. If their taxes pay our salary, we *owe* (to some extent) the public an explanation of why we prefer our ideas to theirs.

Thirdly, what if my new ideas threaten the old order and their lucrative research grants? Scientists, so the story goes, are up in their ivory towers, with their cash and their students, wanting little to change. Surely great ideas are continually suppressed by the "Establishment" to stop them from upsetting this cushy apple cart?

But this story doesn't hold together: you don't get new grant money for doing old science. When a scientist applies to a funding agency for a grant to support their research, they must propose to do something *new*. Similarly, PhD students need a novel project to work on, something that they can make their own, that no one has done before. Scientists are constantly on the lookout for a *hook*, a new idea or method, which they can present in a grant application and say, “look at this new thing that I can do!” So, show us how your new scientific idea could *contribute* to our case for more funding.

Finally, if you ask a young scientist why they chose this career path, or ask an older scientist what the most satisfying moment of their career was, they will often point to the excitement of a new discovery or insight. In the words of physicist Ed Hinds, “those of us engaged in scientific research generally do it because we can’t help it – because Nature is the biggest and most complicated jumbo holiday crossword puzzle you have ever seen.” No one becomes a scientist so that they can plod along behind the establishment. Scientific revolutions get the Nobel Prize.

These forces – earning public respect, winning funding, and advancing knowledge – keep science from stagnating and keep us open to new ideas. This is an important lesson for the cosmic revolutionary: scientists don’t want just any old revolution. They want one that offers new ideas and new directions, that creates opportunities and deepens understanding.

So *that’s* your hook. What’s your great insight, and how will it help us understand this marvellous universe? Give us something that can be held up to nature, firmly founded in mathematics. Give us a deep insight into an old puzzle. Uncover an unseen simplicity in our data. Point the way to new experiments that will test your idea. If you can do this, revolution is at hand.

How to Read This Book

This book examines the universe in discrete pieces, considering key observations that tell us something deep about the workings of the cosmos. We will endeavour to present these observations in as *raw* a form as possible, undigested and untouched by theory. This isn’t completely achievable, unfortunately. But in astronomy we can at least focus on the question: what have our telescopes recorded?

We can report, for example, that the planets appear to move with respect to the fixed stars, and that some stars get brighter and dimmer at regular intervals. When we use a telescope that can measure the temperature of a light source, we pick up a strong signal all across the sky at almost three degrees above absolute zero. The interpretation of these facts, that is, the story of the universe that ties them all together, is up to you and your revolutionary idea.

Our aim is to provide the facts that need to be explained. We will give the most up-to-date versions of the key observations. This is important: we have noticed that many would-be revolutionaries focus too much on the *first* evidence that was presented for a certain claim. For example, they spend their energy critiquing Edwin Hubble's observations in the 1920s of a sample of 30 galaxies, completely ignoring the millions of galaxies we have observed since then.

Throughout, we will give references in the endnotes to the scientific literature regarding both observations and theoretical interpretation. While papers in scientific journals are often hidden behind a paywall, many can be found for free at *arxiv.org* with a quick search. Most scientists, if you ask them succinctly and politely, would be happy to email you a copy of one of their papers.

Remember: if your scientific ideas can account for all of the observations presented in this book and provide future predictions that can be held up to nature, then scientific legitimacy awaits. Alas, fame, fortune, and a lucrative book deal are not guaranteed.

Good luck.