

# *Hanging On To The Edges: Morale is high (since I gave up hope)*

If at first you don't succeed, try, try again. Then quit.  
There's no point in being a damn fool about it.  
- W. C. Fields

In April 2015, Richard Horton wrote as follows: 'The case against science is straightforward: much of the scientific literature, perhaps half, may simply be untrue'. Horton goes on to provide a worrying charge-sheet: scientists typically leap to generalisations from overly small samples, and are abetted by the establishment in doing so; they pursue dubious trends for extra-scientific reasons; statistical inference is poor and formulaic; data are sifted to support predictions; predictions are altered to retrofit data ("These are our values, and if you don't like them...we have others!"); researchers are driven by the maximization of their own status metrics; there are often blatant conflicts of interest; universities behave like sweat shops for making more, rather than more credible, scientific outputs; scientific journals are for-profit entities that want to attract attention to their brand, not reveal the truth about the universe. The consequence is a scientific literature much of which we should be very careful about trusting.

Just what we needed, you may be thinking. Some anti-science nut, chaining together half-truths and conspiracy theories in order to undermine the case for more public investment in science, evidence-based public policy, or the teaching of evolution. We need to be out there defending the enlightenment and its progeny against this kind of flat-earth knavery, which has a very definite agenda of its own. But Richard Horton is most certainly not an anti-science nut. He's the editor of *The Lancet*, one of the top medical journals in the world. What he is reporting on in this particular editorial is a symposium involving the major funders of biomedical research, as well as some of the most senior individuals in the field, to consider 'the idea that something has gone fundamentally wrong with one of our greatest human creations'<sup>1</sup>. That's our real problem, you see: it's not just the barbarians outside the gates saying that the empire is decadent and corrupt. Increasingly, there is unease among the citizens inside the gates too<sup>2</sup>. This unease has been preoccupying me. I don't just mean that I have been reviewing my own working practices to understand how they could be more robust, though I have. I mean something deeper: it has been affecting my morale, my motivation to carry on.

I attempted to go into science out of a very naïve, very pure and rather spiritual sense of love. I was always interested in the arts too, and in fact I worked in the arts for a few years. But I fell in love with

---

<sup>1</sup> Horton, R. (2015). Offline: What is medicine's 5 sigma? *The Lancet* 385: 1380. By the way, the title of this essay comes from that of a show by Powder Keg theatre company, a show that was aptly enough about searching for something to cling to amidst mess and uncertainty.

<sup>2</sup> Non-exhaustive list of key references: Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine* 2, 696–0701; Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science* 22: 1359–1366; Prinz, F., Schlange, T., & Asadullah, K. (2011). Believe it or not: how much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery* 10: 712; Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science* 349: aac4716; Smaldino, P.E. & R. McElreath (2016). The natural selection of bad science. *Royal Society Open Science* 3: 160384.

science through the enchanting writings of authors such as Richard Dawkins and E. O. Wilson<sup>3</sup>. I can hardly describe the exhilaration of learning about science from these masters: not just that it made the drugs and the computers work, not that it added to the size of the economy, but the sweeping intellectual and even aesthetic case that underpinned it. Science: an unbounded golden web of elegant theory, beautiful experimentation, and the best of the human potential. A shibboleth that makes us different from, and, frankly, better than, creationists on one side, and post-modernists on the other. The science/everything else division became for me, I now see, the division between the sacred and the profane, remade by these great writers in a new and astonishing way. Wilson, in *Consilience*, stated very clearly that science is a qualitatively distinct kind of activity from other expressions of human belief. Other belief systems may serve 'psychological functions', he concedes, but science is revolutionary in its ability to discover truth. The enlightenment is a singularity, and science is a new phase of human life.

This is why the current problems in science are so unsettling. To discover that the revolutionary sacred activity probably misses truth at least as often as it hits, not just through bad luck but through systematically stupid and bad behaviour; to discover that all kinds of 'psychological functions' such as confirmation bias, protection of fiefdom, the quest for status, exaggeration of a case in order to market a product, and so on, are deeply embedded in the one institution supposed to be different; what this adds up to is discovering that the ordinary, disappointing regularities of the profane are right there in the heart of the sacred. This poses the question 'How can I carry on?'. Although various theories posit income-maximisation or cultural conformity as prime movers of human behaviour, my personal experience is rather different: people, including me, want to do things that they could readily justify to a jury of their peers (including the jury within). The opportunity cost of trying to do science rather than something else is very big. You have to feel convinced it is worth it. You need to know that the things you believe and promote have some validity. You need to have done due diligence enough to be sure that it isn't all some kind of delusion, quackery, or racket.

I have two questions today. The first is really the warm-up: how can we simultaneously accept the evidence that the actual practices of science are flawed, and its products often wrong; and yet hang on to the assurance that science is a special kind of activity whose long-term arc bends towards the truth? The second is the small matter of how, having answered the first question, we can best live.

## §

The first question turns out to be surprisingly easy. In order for the long-term arc of science to bend toward the truth, science does not have to be perfect. It only has generate a force that is on average very slightly stronger than the forces that hold human knowledge back. That's how science can be both very similar to other kinds of human activities (shamanism, rhetoric, product marketing and what not), and also revolutionarily different. Those other activities all generate a velocity slightly less than that needed for epistemic escape; science generates a velocity that is at least sometimes slightly greater. A small difference, with big consequences.

Let us spell this out with an example: powered flight. What was revolutionary about early powered aircraft is not that they were efficient. They weren't. In fact, there were terrible. The vast majority of the energy they generated was wasted as heat. Of the energy they did manage to generate as motive force, very little was converted into lift. So the point was not that they were very good. They were about the worst devices for powered flight you could come up with, *except for all the other devices that had been tried out in the history of humanity*. For most of those earlier devices, the lift they produced was insufficient to exceed the pull of gravity. The early powered aircraft were only

---

<sup>3</sup> For example, Dawkins, R. (1998). *Unweaving the Rainbow* (Boston: Houghton Mifflin) and Wilson, E.O. (1998) *Consilience: The Unity of Knowledge* (New York: Knopf).

incrementally different, perhaps, but the increment was a consequential one: it was the increment that reversed the sign of difference between gravity and lift, not by much, and not always, but enough for something unprecedented to happen. And once the sign was reversed, once the planes took flight, their design could be gradually improved by the cumulative tinkering that characterises human culture.

Let us return to science. It is not that the people, or even the institutions, that characterise science as a profession are so very different from any other body of people or institutions. It's that somehow, the interaction of those people with those institutions has led to a slow accretion of better understanding of the world over long passages of time. Much of the energy is wasted: the ideas and the claims in any individual publication or even career mostly turn out to be nugatory. But the resultant of all the chaotic motion is a ratchet of gradually better understanding of the processes of the world. The good stuff is just slightly more likely than the bad stuff to be generated and retained, on average. The improving arc is more perceptible the further away you stand: close up, you only see the individual sparks flying off in all directions, mostly not the right one. Only from afar do you see that there is a bit more energy going in one direction than in the others. Just as in the powered flight example, once a science has achieved some kind of lift off, its efficiency can be cumulatively improved. We should be putting as much energy into reforming methods (improving the efficiency of plane design) as we do into individual studies (going for a fly around). The current debates within the science community, the so-called 'replication crisis', should really be seen as discussions about how best to do this, not repudiations of the whole scientific enterprise.

This view of science leads very naturally to seeing scienceness as a continuum. The best cases for the revolutionary nature of science can be made from physics, from chemistry, and from certain parts of basic biology. It's no accident that the best cases made by Dawkins and Wilson come from those areas of science, whereas the day-to-day reality of my working life comes from the study of behaviour and society, where the situation is rather less decisive. If lift routinely and decisively exceeds gravity in physics and chemistry, then the two forces are much more nearly equal in the social and behavioural sciences. Where the two forces are about equal, there is a lot of scope of bump along with bad ideas persisting too long; multiple incompatible views being held simultaneously; fads that appear and vanish like the morning mist; and rhetoric, ideology, and social influence determining the disposition of the field. That's why professional disputes are often so prolonged and so bitter in social and behavioural fields: because, as it were, the stakes are so low. Still, we have to hold on to the hope that even in these fields, the arc towards the truth proves a bit stronger than the will to power in the very long term.

## §

The problem of how to live, as a scientist, is the following one. You need faith in order to be able to do the work. Faith that what you are doing is sensible and worthwhile. Faith that you have the right methods and design. Faith that the patterns you see could be real patterns. Faith that the way you have analysed your data is a sensible way. Faith that the arguments you make are good arguments, and important ones. You need faith in all these things because the whole process is genuinely difficult, and very slow; you are constantly knocked off course by obstacles and distractions; peer reviewers can be quite gratuitously unpleasant, as well as sloppy; rejection is designed in to the process; and employment conditions are often less than ideal. So if you do not have sufficient faith in what you are doing, you will quite sensibly walk away.

Yet science is a system of organized scepticism. Faith is the one thing you should never have. The view of science I have sketched in the last few pages suggests that, rationally, you should make the pessimistic meta-induction: the specifics of the thing you are working on will probably not turn out to

be as you believe them to be; your results will probably not replicate; your methods will appear naïve and flawed to a future generation; and the world will not turn out to be quite as you contend. It's a hard thing to pull off this trick: enough conviction to get up and go to work every morning, enough scepticism to remain a good scientist.

The way most successful scientists solve this is very simple: they have faith in their own work, and relentless scepticism about everyone else's<sup>4</sup>. It's much like the large majority of car drivers who apparently believe their driving to be better than average. But we are all grown-ups here: we should know that we can't all be better than average. The lesson I take from the replication crisis is not that there are few bad apples in science who should be pilloried. It's that we've *all* been doing bad science, probably still are, in myriad banal ways that are so habitual that we don't even realise their significance. The very fact we can't see anything wrong with our own practices is precisely the point. Though we understand that others might fudge the theory, hypothesise after the results are known, exercise researcher degrees of freedom, or torture the analysis in search of the 'significant' *p*-value, it doesn't feel like we ourselves do it. But we must concede there are thousands of tiny judgements involved in the writing of every single paper, the analysis of every dataset, and these are not recorded in any ledger anywhere. And the thing about self-deception is that you are always the last person to know.

For some reason I have a particular terror that the work I have done will not hold up. I shouldn't worry, since I described the data I collected honestly, I made the best suggestions I could, and what the universe decides to do next is not in my control. In fact, some of the things I have done have been replicated multiple times, not just exact replication, but also 'generalizing replication', the extension of the principle of the finding to a new study system or an alternative methodology. My balance sheet where replication is applicable is looking respectably healthy: some things looking very repeatable, some where the dust has not settled, one or two looking like they were probably flukes. Mostly, though, I just do not know yet, and will not know for a long time, if ever, what value the observations and their interpretation had. You can't hope for much better than this, so I have nothing to be ashamed of, really. Yet something keeps me awake at night. Why?

Part of the reason is to do with the usual human concerns: self-consciousness about reputation, status and apparent competence. I am the first scientist in my family (when a teacher suggested at a parent-teacher evening that I should consider further study of mathematics, my father queried whether they were discussing the correct child). I have never felt quite at ease in the august and self-confident circles in which I can nowadays move, so there is more than a hint of the imposter's fear of being found out. I think of my dream-friend Franz Kafka. The key to understanding Joseph K.'s odd passivity when his accusers arrive at his door in *The Trial* is that somewhere in his heart, he already suspected he was guilty. In some diffuse way, he was expecting it. He just didn't yet know what it was he was guilty of<sup>5</sup>.

But there is more to it than this. When you analyse a dataset, there is a lot of tedious merging and cleaning and preliminary analysis. Then at a certain point, you try an analysis or two, and suddenly see a pattern. That moment is psychologically completely and utterly compelling. You instantly grasp why that is the pattern that makes sense, that made sense all along. It jumps out at you with a concrete and immanent reality. It's very hard to tell yourself in that moment that it is most likely to be a fluke, that you may have tortured the dataset until it confessed to something, that your prediction has shifted and you are rewarding yourself after the fact. The evidence of the senses is not easily gainsaid by the caveats of the frontal cortex. The experience of getting a 'significant' result, then writing it up, is completely engrossing. It's almost like a visual illusion: the fact that people are constantly telling

---

<sup>4</sup> See *How my theory explains everything: And can make you happier, healthier and wealthier*.

<sup>5</sup> He still didn't know as he was being repeatedly stabbed to death in a quarry. This is possibly an analogy I should not extend too far.

you it is your brain playing a trick on you does not make it any easier to see the lines as the same length. Later, when someone else, or the very next experiment you do, fails to reproduce your finding, it is as unnerving as the discovery that a conversation you had yesterday was in fact a hallucination. But I saw it! The *p*-value was there! Oh please, let me not be mad!

## §

When middle-aged scientists go through replication crises, some turn to drink, some to drugs, some to Bayesian statistics. I have been turning to Buddhism. I can't say I have got very far (aren't there a lot of numbered lists?). Nonetheless, there are some elements of Buddhist thought that seem like they could come in handy in getting papers written and equanimity preserved.

At the heart of Buddhist thought, as far as I understand it, lies a network of linked ideas. The first, expressed in the first noble truth, is that living is synonymous with suffering, or at least, exposure to the possibility of suffering (*dukkha*). This is because, in living, we crave and are attached to—indeed are fuelled by—worldly things that are by their nature impermanent and likely to fail us or fade away. These worldly things include pleasure, material goods, status and renown, but also beliefs and habits. To the extent that we condition our happiness on obtaining or maintaining these things, we are locked into a cycle of endless living-suffering (*samsara*), because pleasures always fade, beliefs turn out to be wrong, and status is never enough.

This is the bad news, but the good news follows: once we recognise the reality of suffering, and its causes, we see that it can also cease, and that there is an available route to liberating ourselves from it. This route requires nothing other than enlightenment of our minds. We achieve the liberation not by satisfying our attachments and cravings, which would after all just bind us further into *samsara*, but by living according to the eight-fold (or just possibly three-fold, or  $8 \times 3 = 24$ -fold) noble path. This path consists of right speech, right action, right livelihood, right effort, right mindfulness, right concentration, right view, and right resolve. Basically, trying to live mindfully, wisely and well. The noble path leads directly to the release from the compulsory cycle of desire, striving, attachment and suffering known as *nirvana* (often translated in US English as 'tenure').

The noble path is not just asceticism. The Buddha certainly rejected living by trying to fulfil one's hedonic cravings as unprofitable and painful; but he rejected a life of extreme austerity and self-mortification too, on the same grounds. After all, attachment to austere ritual, to self-denial, is just another form of attachment. Instead he proposed the middle way: in the world, of the world, yet trying to navigate it nobly.

Now before I completely lose any credibility I had, I would invite you to consider the above doctrine not as supernatural or even religious, but as a set of rules of thumb for living worked out over centuries by thoughtful members of a smart species of ape<sup>6</sup>. A species with no single mental governor, but whose mind consists of a noisy parliament of different and perfectly explicable motivations—for resources, for reputation, for sex, for power, for avoiding danger—which together conspire to produce powerful patterns of habit and thought, patterns that can in the long term produce interpersonal and intrapersonal problems. But this same species of ape also, for extraordinary but not supernatural biological reasons, possesses a surprising capacity for offline reflection and reasoning, a capacity that can be used to calm fractious disputes and reach wise compromises, including, critically, disputes and compromises within the parliament of the mind. The first noble truth reminds us that our very real and natural motivations can make us disappointed or miserable and cause us problems; the noble

---

<sup>6</sup> This is the approach of Jonathan Haidt's 2006 book *The Happiness Hypothesis: Finding Truth in Ancient Wisdom* (New York: Basic Books).

path reminds us that we have mental resources to deal with these problems, and it's a good idea to practice using them.

Let's apply some of these ideas back to the life of the scientist. To be a scientist is synonymous with suffering, or at least exposure to the possibility of suffering. The reasons for this are part banal—we all want the big grant, the big promotion, our papers to be accepted in the selective journals, but we can't all achieve this. Others will succeed and mostly we will fail. There are more metaphysical reasons, too: the beliefs and hypotheses to which we devote thousands of difficult hours of our lives will more often than not turn out to be wrong. We won't know this for a very long time, perhaps never, and, individually, will have very little control over how it turns out. So we must live with doubt and uncertainty about the aspects of our lives that are more important to us than anything else. And what we achieve always, in my experience, falls short of what we hoped to achieve: as E. O. Wilson put it, all scientists are 'are children of Tantalus, frustrated by the failure to grasp that which seems within reach'<sup>7</sup>. Thus, if we base our well-being on getting what we crave, or attachment to what we have done before, we can never be really satisfied, for we are trying to hold on to a will o'the wisp.

Faced with this dilemma, two courses suggest themselves. The first is equivalent to the life of hedonic gratification: in the long run, no-one knows who is going to be right, and I won't be around to find out anyway, so I will just make as good a career for myself as I can. Consider researchers of type A. They make a big name for themselves with their seminal *Hypothetical Attachment Theory (HAT)*, or whatever. They make the positive case for *HAT* in big idea piece after big idea piece. They churn out empirical studies, and present them in the best pro-*HAT* light the peer reviewers will let them get away with. The studies are as good as they need to be, but no better. They aggressively confront journal editors who reject their papers. When peer-reviewing, they recommend 'reject' if the authors do not cite enough *HAT* references, and that is nearly always. They know who their rivals are and make sure to rubbish their grants so that they do not get funded. As the evidence accumulates that *HAT* is probably old hat, they dismiss the criticisms as ill-founded or personally motivated. They defend their dung-hills. They see the exchange of academic views as a social game to be won. It is because senior people are type-A researchers that science sometimes seems to advance, in the famous paraphrase of Max Planck, one funeral at a time. It is obvious that type-A researchers suffer from grasping attachment: to status, power and their habits of thought. They are locked into the *samsara* of revise and resubmit, busy-work, dogma and self-promotion.

Now consider researchers of type B. All too aware of the possibility of error and luck, they are desperately cautious about what they will publish. They always want to take more time over everything. They need more data, another replication, before they will show their results to the world. They always need to do more reading, in case there is something they have missed. And when they do write, there are caveats on their caveats: their papers will not come down strongly for any position, all too aware that to do so might lead to saying something that could turn out to be wrong. The type-B approach is also problematic. Science could not work if everyone followed it. The dynamism of science emerges from it being a vigorous and passionate conversation, with people willingly sticking their heads above the parapet with interesting data and the strongest possible advocacy of particular ideas. And, less obviously, researchers of type B are suffering from grasping attachment too, like the ascetics criticized by the Buddha. In fact, it's another form of egotism. They crave a kind of certainty and definitiveness that we can never really have; they are too attached to their own self-image, their personal comfort and their rituals of scholarship to be prepared to let the ideas and the data speak for themselves.

You can see where I am going here: towards the middle way, and a path to the cessation of suffering. The middle way is where we are quite prepared to put out what we have done, including strong and

---

<sup>7</sup> Wilson, E.O. (1998) *Consilience: The Unity of Knowledge* (New York: Knopf), p. x.

principled, even passionate, arguments for what we think it means theoretically; yet on the other hand we are open to changing our minds at a moment's notice; we encourage alternative views and welcome those whose starting point is different from our own; we are quite prepared to say when we were wrong, and patient to say why if we still think we might be right. Really it comes down to humility and openness: Open sharing of our data, openness about what operations and analyses we have performed, openness to sticking a preprint out there that turns out to be ill-informed, openness to alternative views, openness to trying to see things a different way. The eight-fold noble path (right speech, right action, right livelihood, right effort, right mindfulness, right concentration, right view, and right resolve), perhaps enriched for our time with *not using so many disposable drinks containers* (that really drives me nuts) is a pretty good recipe for making contemporary science not just more effective, but also a nicer place to live. We need to provide both the incentives and the social norms that encourage this kind of nobility, and discourage its opposite.

## §

What, then, of *nirvana*? This is often translated as 'release from the cycle of existence'. I have always struggled a little with this. Many of my very favourite things depend rather heavily on me existing. So why would I follow a path whose best-case scenario is self-obliteration? This question, in the general case, is rather above my eschatological pay-grade, but we can make a useful translocation of the *nirvana* concept for science. In fact, there is a surprising link to Karl Popper, of all people.

In contemplating the human capacity for reasoning, and hence science, Popper made the following point (and in doing so revealed himself as a better philosopher than zoologist). When most animals hold a false belief, it can lead to their deaths (think of a deer with a false belief about which species are dangerous predators, or a false belief about which food is safe to eat). So the stock of beliefs is only improved by the cycle of birth (which introduces variation in beliefs), and death (which disposes of the false ones). The miracle of being human rather than being some other kind of animal is that the deaths of our ideas can become decoupled from our own deaths. We can represent ideas symbolically, then debate, converse, test, adjudicate, modify, falsify, and eventually reject them, all in relative safety. We can, as Popper put it, 'let our false theories die in our stead.'<sup>8</sup> This opens up the possibility of an adaptive evolution of ideas, with a generation time much faster than our biological generations, but an eventual survival that could be much longer: ideas take on a life of their own. Science, perhaps above all else, is the commitment to fostering this artificial life: ideas proliferating, mutating, recombining, dying and becoming immortal in the rich, distributed ecology of the scientific literature.

What does this have to do with *nirvana*? The following: what is it that a scientist can most sincerely hope for? What most can sustain him in feeling that it was all worthwhile? It is not being right. It is not being the cleverest. It is that the ideas to which he devoted his life ultimately released him from their cycle of birth and death. They began to evolve on their own, in ways he could not predict, through other people, through artefacts and dusty books, through conversations in coffee shops, through different technologies or social institutions, through a shy student's first moment of enchantment. This possibility is profoundly and existentially comforting though, paradoxically, the person for whom it is comforting vanishes from the picture.

This is what I think. It's a long old game. If I stick at it well enough, I will cease to matter. I will sooner or later disappear to the hills, another wind-lined, inconclusive, Pennine man; increasingly vague; stalwart of the choral society; my worldly heft gradually diminishing through ordinary biological processes. But maybe I'll know that, somewhere in the world, these ideas that I have cherished, these ideas will be dying—and living—in my stead.

---

<sup>8</sup> Magee, B. (2010). *Popper*. London: Fontana, p. 64.