
Objectivity and Intellectual Humility in Scientific Research: They're Harder Than You Think

NANCY CARTWRIGHT *  & FARON RAY **

*Department of Philosophy, Durham University, Stockton Road, Durham DH1 3LE, UK and Department of Philosophy, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA. Email: nancy.cartwright@durham.ac.uk

**Department of Philosophy, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA. Email: fray@ucsd.edu

We begin from the assumption that where scientific research will predictably be used to affect things of moral significance in the world, you have a special duty, a duty of care, to 'get it right'. This, we argue, requires a special kind of objectivity, 'objectivity to be found'. What is it that's to be found? In any kind of scientific endeavour, you should make all reasonable efforts to find the right methods to get the right results to serve the purposes at stake and neither exaggerate nor underestimate the credibility of what you have done. That, we take it, is what in this context constitutes objectivity and intellectual humility. But where your results will affect the world, you have a more demanding duty: a duty to 'get it right' about the purposes the endeavour should serve. Often the most morally significant purposes are those that 'go without saying' and because they are not said, we can too easily overlook them, sometimes at the cost even of human lives. We illustrate this with the example of the Vajont dam design and the flawed modelling that resulted in the Hillsborough football disaster.

Introduction

This discussion aims to set you thinking about objectivity in science research, because all too often research objectivity is far more demanding than we think. I have in mind particularly research that is likely to be used to intervene in the world. In this

case, it is not only good science that calls for objectivity: you have a *duty of care* to be objective. This generates extra demands that are generally not well articulated – though they are very real indeed.

What constitutes objectivity in scientific research? The last decade has seen a great deal of work on this in philosophy of science and in science studies and long lists of suggestions are on offer. These include freedom from values, or judgement, or subjectivity; use of warranted procedures; neutrality; impartiality; independence; convergence; consistency; reliability; dependability; generalizability; transferability; accountability; transparency; rigour; auditability; reproducibility (Hyde 2023).

This hodgepodge of proposals has led some philosophers to urge getting rid of the notion of objectivity altogether. Ian Hacking, for example, urges: ‘Let’s not talk about objectivity’ (2015: 19); instead ‘Let’s get down to work on cases’ (Hacking 2015: 29). According to Hacking, one should say what one means in the case at hand, e.g. ‘use only rigorous methods like randomized controlled trials’ or ‘do not rely on expert judgment’ rather than issue the empty injunction ‘be objective’.

I will argue that this is a mistake. We need the empty injunction for scientific research. That’s because the kind of objectivity we want in science is what Eleonora Montuschi has labelled ‘objectivity to be found’. In this article I will first explain and defend ‘objectivity to be found’ and then point out that in those cases where objectivity is a duty of care, there is an extra burden in what you need to find. We all recognize that good research requires finding the right research questions and the right methods to address them. However, where there’s a duty of care, you also need to identify the right purposes that the research should serve. I will illustrate with two examples. The first developed from work with Eleanora Montuschi on ideas she first published with Pierluigi Barrotta – the Vajont dam disaster (Montuschi and Barrotta 2018a). The second is an example of Faron Ray’s from some joint work we’ve done on objectivity – the Hillsborough football stadium disaster. Finally, I will note the importance of intellectual humility for securing objectivity and point out how difficult the appropriate exercise of intellectual humility is to achieve. It is all too easy for an expert researcher to *presume* that the methods and theories that they have spent years practising and honing are good enough for a job in their area of expertise.

Why We Need ‘Objectivity’

To explain why objectivity – and especially objectivity to be found – should not be done away with, I will begin with a brief excursus through some philosophical terrain that may be unfamiliar to you. This is in order to make sense of why the demands of objectivity are open-ended and why the duty of care can be so hard to discharge.

Following suggestions from the wonderful work of Lorraine Daston and Peter Galison (2007) on the history of objectivity, I suggest that the reason the list of what constitutes objectivity is so long and varied is that ‘objective’ is what the mid-twentieth-century Oxford philosopher JL Austin (1962) called a ‘trouser word’.^a A trouser word is a word that gets its meaning in a setting from what it denies there.

One of Austin's prime examples is the word 'real', which he tells us 'is not to contribute positively to the characterization of anything but to exclude possible ways of being not real – and these ways are both numerous for particular kinds of things, and liable to be quite different for things of different kinds' (Austin 1962: 70). As a course guide from the University of Reading explains, 'Possible ways of not being real really are numerous. We can contrast real limbs with artificial limbs, real teeth with false teeth, real cream with synthetic cream, real diamonds with paste diamonds, real bullets with dummy bullets, real ducks with decoy ducks, real cars with toy cars, real horses with pictures of horses, etc., almost ad infinitum' (Preston 2018). I propose that objectivity is like that. As Austin says of 'real', 'it does not have one single, specifiable, always-the-same meaning ... Nor does it have a large number of different meanings – it is not ambiguous, even "systematically".'

Objectivity then is a trouser word. This makes sense once we recognize that objectivity is not an end in itself. We value objectivity because we take it to be an aid in securing something else that we value in scientific research – 'getting it right'. What Daston and Galison's history suggests is that what objectivity calls for in a given setting depends on what threats to getting it right are perceived to loom in that setting as well as what courses of action are thought likely to diminish those threats. Consider activist research. Whether researchers are working for a political cause or for a company with vested interests (say an oil company or a pharmaceutical), being deeply imbedded in and committed to a point of view can incline them to overlook or misjudge the strength of evidence that does not support the interests of their cause. So particular commitments and points of view can raise a threat to getting it right. On the other hand, standpoint theory argues that researchers can never leave behind assumptions and inclinations from their social-economic-political-cultural background, which raises the same kinds of threats to getting it right. So, standpoint theorists urge that the way to ameliorate these threats is not to try to get rid of perspective, aiming impossibly for a 'view from nowhere' in research, but rather to diversify research teams to include a variety of perspectives in designing and carrying out the research.

The point of this excursus through trouser words and threats is to answer Ian Hacking and others like him who urge eliminating the loose concept of objectivity and substituting it with something more precise that we actually mean. We should not eliminate 'objectivity' and express our demand more precisely because what, more precisely, is required will vary from case to case depending on the threats to getting it right that loom in that context. And those threats are often not at all apparent at the start. Often you can only recognize some of them once you are deep into the research itself. The threats, and the methods you should take to avert them, have to be *found*, and not just in advance, you have to watch out for them throughout the research.

This puts a heavy burden on the researcher. It is much easier to satisfy a requirement if what you have to do is spelled out precisely. But there will be many cases where no matter how much effort is put into identifying threats ahead of time and saying just what is expected in the research to avert them, unexpected threats will

arise in the course of the research. You are not being sufficiently responsible if you do not take due care to identify and deal with these as they turn up.

By way of analogy, compare EU *regulations* which, once passed, are directly effective exactly as written in every member state, with EU *directives* that set a goal that all EU countries must achieve but leave it up to individual countries to devise their own laws on how to reach that goal. The reason for directives is that the way to get the same effect of a desired change expressed into law in one country may differ from another as different countries have entirely different legislative landscapes, judicial bodies and methods of enforcement. It seems Hacking wants us to be passing regulations that dictate exactly what should happen. Yet, what we need for getting it right in science is the open-ended directive: do what it takes in your case to be objective there.

Of course, just as with directives versus regulations, in general the use of more abstract, less concrete concepts, guidance and instructions can make interpretation and application difficult – murky – and maybe prohibitively expensive (e.g. small firms needing expert legal advice to ensure that what they do complies with financial guidelines or guidelines on negligence and duties of care or extra research resources directed to checking out possible threats to see how damaging they might be). And it can seem too open to interpretation: you never know if you really are doing the right thing. That is of course a constant danger in science; it is a good part of why doing good research by being objective is so hard.

Of course, you cannot be expected to identify everything correctly. But that is not what duty requires. What you must do is exert due diligence. This English expression I understand is like the Spanish concepts *diligencia debida* and *diligencia de cuidado*. Individuals/corporations are potentially liable to *negligencia* (negligence) when they fail to exercise good care to fulfil the duty. You are not expected to have perfect insight, but rather to do what a reasonable person with an appropriate range of expertise and training would judge necessary. This is like the idea in English tort law: ‘There must be, and is, some *general conception* of relations giving rise to a duty of care, of which the particular cases found in the books are *but instances* ...’. The ‘reasonable person’ in English law is sometimes described as ‘the man in the street’, or ‘the man in the Clapham omnibus’: ‘Such a man taking a ticket to see a cricket match at Lord’s would know quite well that he was not going to be encased in a steel frame to protect him from the one in a million chance of a cricket ball dropping on his head’ (Cartwright et al. 2022).

Now we can begin to see why achieving objectivity is so hard. Being objective involves taking due effort to get it right: to ask just the right questions to serve the purposes at hand and adopt the right methods to answer them, or ones that are right enough for those purposes. But, contrary to Hacking’s suggestion, seldom can it be specified in advance and by some external agent or norm or convention what this should consist of in any given setting. Of course, your training and expertise will guide you – though sometimes this is just what blinds you, which is a topic I will turn to later. In the meantime, the point I want to underline is that scientific research is expected to be objective, and part of the reason that that expectation is so demanding

is that, far from being told what to do, researchers have to *find* what it is that being objective requires in the setting. I call this the problem of *objectivity to be found*, an idea originated by the philosopher of social science Eleonora Montuschi in joint papers with Pierluigi Barrotta (2018a, 2018b) and developed in our recent multiply authored book, *The Tangle of Science* (Cartwright et al. 2022).

That's not the end of the difficulties however. Matters are even harder when research will predictably be used to intervene in the world. In this case there is an extra burden, a duty of care, to get it right, since getting it right involves, beyond all I have just said, taking due diligence to identify what the *right purposes* are in the setting. Finding what the right purposes are is critical since what the purposes are affects what are the right research questions to ask and what must be done to address them. I will illustrate this, developing first an example Montuschi originally used to illustrate this – the Vajont dam disaster – and second, an example of Faron Ray's from our joint paper, *Modelling Objectively* (Cartwright and Ray 2023) – the Hillsborough football tragedy.

The Vajont Dam

In 1960 the tallest dam in the world was completed – the Vajont dam in the Dolomites, 100 km north of Venice. The dam was a fantastic engineering achievement. On 9 October 1963, a huge landslide went into the reservoir and the dam stood. Indeed, from this point of view it was an engineering triumph: it withstood a force eight times what it was designed for. But a tsunami of 50 million cubic metres of water overtopped the dam in a wave 250 m high. Virtually the entire reservoir blew over the dam destroying the village of Longarone down the Piave valley and severely damaging other villages, killing over 2,000 people.

What went wrong? There was some outright misconduct. Information suggesting matters might be problematic was ignored and not passed on to everyone who should have been apprised of it, and generally the force of this information, though it might not otherwise have seemed impressive, was underestimated given the stakes. The engineer in place for the last few years before the disaster was eventually sentenced to six years in prison.^b What I want to illustrate is not points of outright misconduct and illegal behaviour though, but rather what looks to be a failure of due diligence to be objective.

In this case, wrong models were developed and relied on and wrong purposes held centre stage. Most of the effort focused on the formidable engineering challenges of building the dam. For instance, a system of 146 equations with as many unknown variables was perfected and solved – with the aim of building a dam that would stand and generate electricity. Safety purposes were not taken seriously enough in developing and assessing various models given the stakes.

Beyond conscious and deliberate misconduct, how did this happen? A lot has been said about this since the disaster. For illustration, I shall focus on three things. All three look to involve a failure of *due intellectual humility*. For purposes here I

take intellectual humility to be a virtue that sails between the Charybdis of presumptive overconfidence and the Scylla of unreflective excessive doubt. ‘Presumptive’ and ‘unreflective’ matter. There are many reasons one can be excessively over- or underconfident. You could simply be mistaken or incompetent or deliberately overinflate or underestimate. The first might well be an honest mistake, the second could be either an epistemic or a moral failure and the third is generally a moral failure. As I use the term, what makes over- or underconfidence a failure of *intellectual humility* is the failure to reflect sufficiently on what you are doing, to just *presume* you’ve got it right, or presume you probably have it wrong. In all three cases there seems to have been insufficient attention to the all-important purpose of securing the safety of the local communities, given how high the stakes were. The dam designers presumed that the great efforts they were making were sufficient to build the kind of dam needed. In doing so, too much importance was given to the purpose of building a dam that would stand and generate electricity.

The first problem I want to note is that it was very over-optimistically assumed that the movement of large landslide masses in the region of the dam could be managed by elevating the level in a careful manner and drawing down when troubles appear. Let us look at a brief history of this.

At first filling in 1960 to a height of 170 m above river, a 2 km long crack appeared in the nearby land mass suggesting a landslide had occurred. Fill continued in 1960 until, at a fill height of 180 m on 4 November, 700,000 cubic metres of material slid into the lake in 10 minutes. The level was dropped to 135 m by December, reducing movement of the sliding landmasses from 8 cm to 1 mm/day. In October 1961, a by-pass tunnel was constructed as an aid in case of landslides.

A second filling took place slowly in 1961 and into 1962. In November 1962 at a fill height of 235 m, the land velocities increased to 1.2 cm/day. So, a second draw down began and in April, 1963 at 185 m, velocities reduced to approximately zero.

So, a *third filling* was undertaken. By early September 1963 at a fill height of 245 m, velocities increased to 3.5 cm/day. Again, a draw down was commenced to bring rates of creep back under control, but land movement velocities kept slowly increasing; at some places 20 cm/day was recorded. By 9 October, 235 m was reached. That’s when the catastrophe occurred.

As an AGU landslide blog explains, ‘The experiences gained from the 2nd filling and subsequent draw-down confirmed to the engineers that control of the landslide was possible by altering the level of the reservoir’ (Wolter 2016). But even without hindsight, given the safety stakes, the confidence in this certainly looks hubristic.

The second problem had to do with modelling land movements and landslides. As one geologist explains: ‘The continuous rejection of the worst case scenario by authorities and the electric power company running the dam, was in part based on a lack of understanding of large mass movements at the time’ (Bressan 2011). For more than three years, land movements were monitored and geologists studied the slope.

Some geologists warned of a deep-seated landslide. Others proposed superficial sliding planes, able to cause only small landslides. Small landslides, as happened in 1960, were always expected. So, as I noted, in 1961 the construction of a by-pass tunnel was started in case the reservoir became partially obstructed by a landslide. At that time calculations, based on a small model of the entire reservoir, suggested that a (small) landslide into the lake could generate a 30 m high wave. Technicians suggested not exceeding a water height in the reservoir that was surpassed in 1963 by 10 m.

Here is what Alessandro Franci et al. (2020: 1) noted about this in *Engineering Geology*:

Reduced-scale experiments . . . two years before the Vajont disaster were carried out with a material not representative of the actual rockslide behavior and [did not consider] the simultaneous failure of the whole landslide body.

And here is the diagnosis offered by Franci et al. (2020: 1):

Prediction of multi-hazard slope stability events requires an informed and judicious choice of the possible scenarios. . . . An incorrect definition of landslide conditions . . . can lead to inaccurate predictions . . . and wrong engineering and risk management decisions.

The dam developers were not modelling objectively given the possible huge moral costs. They made the very optimistic presumption that the modelling was good enough. But the reality was that they didn't really know how to model landslides. Of course, you can't be faulted for not knowing something that nobody has figured out yet. The assumption that they could proceed as they did without better models could have been intentional overconfidence in order to get the dam built and get the electricity generation started, and thus genuine misconduct. Or maybe it wasn't intentional – they merely *presumed* it was enough without the amount of reflection due given the stakes. In that case it would clearly have been a failure – a really disastrous failure – of the intellectual humility due where lives are seriously at stake.

The third problem was highlighted by Montuschi in a paper with Pierluigi Barrotta (2018b), 'Expertise, relevance and types of knowledge'. Here, Montuschi and Barrotta argue that there was too much reliance on general knowledge and too little attention to local knowledge. The engineers were steeped in general knowledge; they were real experts in their field. But, Barrotta and Montuschi urge, the engineers did not make serious enough efforts to integrate their well-established and general knowledge with a miscellany of less rigorously established local knowledge, both knowledge of local facts and knowledge held by local inhabitants.

For a start on local facts, there is the very name of the mountain that looms over the reservoir, 'Monte Toc', which, Barrotta and Montuschi note, is short for 'pa-toc' – spoiled, rotten, damaged. Other crucial facts included 'deep fractures, tremors and loud noises coming from Mount Toc . . . , a 50 million cubic meter slide in a nearby artificial lake that killed one man, and mounting evidence that a much

bigger, much older landslide could have set itself in motion at any time' (Montuschi and Barrotta 2018a: 391).

This was further exacerbated by the sources of much of this knowledge, as they explain:

Part of the reason for ... [the neglect of local knowledge] ... is that some of this local knowledge was held by local people (the inhabitants of the valley, peasants and mountaineers in the Vajont). This knowledge was not 'scientific', it was not formalized in a textbook, nor was it discovered by scientific method and expressed in sophisticated geological classifications. Nonetheless, it could count as knowledge. In fact, the mountaineers' system of beliefs was warranted by a secular and detailed acquaintance with the slopes of the valley. However, at least partly because it was a type of knowledge formulated in such a way that did not command assent and credibility, the overall role of local knowledge in building relevance to 'the case at hand' was by and large overlooked. This was a mistake. (Montuschi and Barrotta 2018a: 392)

It was not only a mistake, they argue, it was a failure of objectivity – objectivity to be found. The dam designers did not find all the knowledge that could reasonably inform their decisions. It is equally a failure of intellectual humility: the engineers presumed without enough due reflection that their methods and knowledge were the right ones to employ.

The Hillsborough Football Tragedy

Here is what Faron Ray has to tell us about this disaster (Cartwright and Ray 2023):

On April 15th 1989, 96 supporters were killed and 766 injured at Hillsborough Football Stadium in Sheffield, South Yorkshire in the UK when a fatal crush occurred in the stadium's enclosed pens. South Yorkshire Police were de facto in charge of crowd safety and thus had both a moral and legal duty of care to ensure that those entering the ground were not exposed to unreasonable levels of risk. They failed to fulfil this duty.

What went wrong? One could explain the failure by tracing the chain of mistakes made by various officers on the day. First, officers stationed outside the ground lost control of crowds as they waited to enter, leading to the onset of a dangerous crush by Leppings Lane turnstiles. Next, to ease the crush outside a senior officer requested for an exit-gate (Gate C) to be opened in order to allow maximum flow into the ground. Finally, Chief Superintendent David Duckenfield granted the officer's request, leading crowds of supporters to be funnelled down the natural channel that existed from Gate C to one of the already full enclosed pens.

Whilst accurate, however, this story only goes so far as an explanation for South Yorkshire Police's failure to fulfil their duty of care, for the failures of the South Yorkshire Police as an institution were not limited to the actions taken by individual officers that day. Rather, they extended to the manner in which South Yorkshire Police set about preparing to act. That is, they extended to the modelling conducted by South Yorkshire Police. Indeed, it was precisely that the police were working with bad models that led to the catastrophic mistakes they made that day.

So, what exactly was so bad about the modelling carried out by South Yorkshire Police? They had not estimated certain factors, sure, but then they had not estimated lots of things, many of which turned out to be inconsequential. What was it about the pre-match modelling by South Yorkshire Police that proved so fatal that day? We suggest that South Yorkshire Police's failure to fulfil their duty of care was in part caused by their failure to find the right purposes for their modelling. How so? First, it seems clear that the police conceived of their role at the ground to be first and foremost a disciplinary one. This was wrong. As the only party capable of ensuring crowd safety at the ground, the duty to do so naturally fell upon them. Their failure to adequately focus on the full nature of their role thus led to their identifying the wrong purposes of their pre-match preparation; pre-match preparation which consequently concentrated on the narrow mechanics of crowd control rather than the much more complex task of crowd safety.

In sum, it was South Yorkshire Police's failure to find the right purposes for their pre-match preparation that constituted their failure to model objectively and this, in turn, helps explain their failure to fulfil their duty of care.

Again, as in the case of the Vajont dam, it looks likely that this failure of objectivity was due to presumptive overconfidence in the police view of their role and in their safety methods.

In sum: these two examples are meant to illustrate how easy it is for objectivity to go astray and especially to highlight the harms that can be generated from focusing on too narrow a set of purposes, and in consequence not asking the right questions and not choosing the right methods. They are also meant to point to a particular cause of failures of objectivity that I hypothesize is frequently operative: a failure of intellectual humility. Failures of intellectual humility, I suggest, are a ready threat to 'getting it right' – to focusing on the right purposes, asking the right questions and employing the right methods. Noting this specific kind of cause matters when it comes to thinking about how to better promote objectivity in research, which I turn to next.

What Can be Done to Promote Due Intellectual Humility in Scientific Research and Thereby Diminish the Threats to Objectivity?

Before we think about the question in the title of this section, it is worth turning to a prior question: who, or what, is to blame for failures of intellectual humility in science research? The natural answer seems to be ‘the researchers who designed and carried out the research’. This seems right in cases where the researchers are consciously and/or deliberately arrogant or, conversely, over-diffident. But recall that when I introduced the concept of intellectual humility, I pointed your attention to what we do *without* notice or deliberate intent, using the label ‘presumptive’ for this.

I focus on this because presuming is absolutely essential in science. Each new scientific endeavour is built on an unimaginably vast tangle of previous work. Some of this previous work that is seen as especially salient may attract special scrutiny, but we couldn’t proceed without massive unreflective presumption. You cannot constantly be reviewing the methods, facts, theories and models that you will employ. You have to take for granted an immense body of knowledge and practice, and what you take for granted will depend very much on what (sub)discipline you are in, how you have been trained and whom you take seriously.

This already points to the limits of blaming individual researchers for failures of intellectual humility. Whilst failures of intellectual humility do sometimes arise from the culpable arrogance or diffidence of particular researchers, more often than not they take the form of *systems failures*. Often, the presumptive under- or overconfidence characteristic of failures of intellectual humility arise from the very norms, habits and practices of scientific institutions and communities themselves. In these cases, an individualistic analysis breaks down. We must look, instead, to the features of scientific institutions and communities that give rise to such presumptive under- or overconfidence. In this closing section, I look at two such features. The first goes by the name of *silozation* whilst the latter is often referred to as the problem of *integration*.

Scientific knowledge is now so complex that mastering what it takes to make progress in any special problem area takes years of intense training and fine honing of knowledge and skills. That makes the knowledge and skills employed in other problem areas opaque. We usually know little or nothing about the methods and background knowledge employed in other areas. That makes this knowledge and these methods inaccessible to us and, if they are offered, we are quite often suspicious of them because we are not in a position to judge whether the knowledge is well-established and rigorous and whether the methods can deliver what is promised. This feature of current scientific knowledge production is known as the problem of ‘silozation.’ Here is Harvard economist Dani Rodrik hinting at the problem of silozation and its effects in the context of economics:

Because economists share a language and method, they are prone to disregard, or deprecate, noneconomists’ points of view. Critics are not taken seriously – what is your model? where is the evidence? – unless they’re willing to follow the rules of engagement. (Rodrik 2015: 80)

Similarly, economist Robert Skidelsky writes of the way in which work in economics can become blind to the possibilities lying outside a certain set of methods and assumptions:

Neoclassical economics has developed a peculiar method . . . and the use of any other method is not regarded as economics . . . Models based on this method allow for only a limited range of possibilities. Events which might occur outside this range are not picked up on economists' radar screens. (Skidelsky 2021: 1)

Unfortunately, the silo-ization of economics and its consequent detachment from the other social sciences, such as sociology, gives rise to mutual distrust and suspicion between these different communities. This phenomenon of mutual distrust is well-documented. Indeed, as Skidelsky writes, 'economists and sociologists . . . each "view the other through a glass darkly"' (Skidelsky 2021: 97). Here is Rodrik speaking of the attitudes he found amongst philosophers, sociologists and historians to the discipline of economics whilst working at the Princeton Institute for Advanced Study:

[There was] a strong undercurrent of suspicion toward economics. To them, economists either stated the obvious or greatly overreached by applying simple frameworks to complex social phenomena . . . [T]he few economists around were treated as the idiots savants of social science: good with math and statistics, but not much use otherwise. (Rodrik 2015: xii)

Unfortunately, disrespect is a two-way street. Rodrik goes on:

The irony is that I had seen this kind of attitude before – in reverse. Hang around a bunch of economists and see what they say about sociology or anthropology! To economists, other social scientists are soft, undisciplined, verbose, insufficiently empirical, or (alternatively) inadequately versed in the pitfalls of empirical analysis. Economists know how to think and get results, while others go around in circles. (Rodrik 2015: xii)

This disturbing pattern of mutually reinforcing distrust and suspicion is what Faron Ray has referred to as a *hubristic feedback loop*:

Unfortunately, here enters what I will call a *hubristic feedback loop*, for such criticisms, made by sociologists, inspire a similarly hubristic response from economists. Indeed, it is precisely this hubristic attitude of sociologists that has the potential to help further entrench the unwarranted prejudice of economists *against* the sociologists' qualitative methods – 'You criticise our methods but you don't even understand them! How can you and your methods be trusted if you exhibit no interest in seriously engaging with the insights we have to offer?' (Ray 2022)

Hence, from the silo-ization of scientific disciplines there can arise a presumptive dismissal by researchers of the methods and assumptions present within other disciplines. Again, this is a failure of intellectual humility, understood not as a failure of particular individuals but as an emergent feature of the norms, habits and practices within scientific institutions and communities.

Finally, in addition to silo-ization there is also the well-known problem of *integration*. The language, world view, methods and models of different disciplines do not slot easily – if at all – into each other. For instance, whilst sociology likes *complexity*, economics likes *simplicity*. These different sensibilities give rise to entirely different approaches to studying social phenomena. Here is sociologist Kieran Healy, for instance, writing (disparagingly) about sociology’s own taste for complexity, where he refers to lovers of complexity as ‘connoisseurs’:

Connoisseurs call for the contemplation of complexity almost for its own sake or remind everyone that things are subtler than they seem ... Connoisseurship gets its aesthetic bite from the easy insinuation that the person trying to simplify things is a bit less sophisticated a thinker than the person pointing out that things are more complicated. (Healy 2017)

Contrast Healy’s remarks about the sociologist’s hankering for complexity with how MIT economist Jonathan Gruber speaks of one of the defining features of economics. Economics, as Gruber makes clear, likes *simplicity*:

We’re never going to get it perfectly, but you’ll be amazed at how the small a number of assumptions we need to explain an enormous number of things. (Gruber 2012)

These two approaches, and the models and methods that they give rise to, do not slot easily together. Indeed, as distinguished economist Paul Samuelson notes, the simplicity and rigor of economics gives it a world view all of its own:

[Economists’] map of the world differs from that of the layman. Perhaps our map will never be a best seller. But a discipline like economics has a logic and a validity of its own. (Samuelson 1962: 18)

This last point is particularly important, for the internal logic of economics with its emphasis on formalization gives rise to what Rodrik (2015: 80) refers to as the ‘strange paradox’ of economics; the very emphasis on formalization that allows economists to state their assumptions clearly and thus *challenge* the assumptions of their peers tends to foster an inability of economists to appreciate or even see challenges coming from outside. The very feature that makes economics sensitive to criticism from *inside* the discipline makes it insensitive to criticism coming from *outside* the discipline.

We have seen the same kind of insensitivity to external input and criticism in the engineers’ dismissal of local knowledge in the case of the Vajont dam disaster. Recall what Barrotta and Montuschi claimed about the local knowledge held by local people:

This knowledge was not ‘scientific’, it was not formalized in a textbook, nor was it discovered by scientific method and expressed in sophisticated geological classifications. Nonetheless, it could count as knowledge. In fact, the mountaineers’ system of beliefs was warranted by a secular and detailed acquaintance with the slopes of the valley. However, at least partly because it was a type of knowledge formulated in such a way that did not command assent and credibility, the overall role of local knowledge in building relevance to ‘the case at hand’ was by and large overlooked. (Montuschi and Barrotta 2018b: 6)

These problems, of silo-ization and a lack of integration, point to the conclusion that securing the right kinds and degrees of intellectual humility and promoting objectivity is not best done at the individual level. It is entirely natural, indeed necessary, for researchers who are experts in an area in which they see a problem to presume that the viewpoint, knowledge and methods they employ are right enough to solve that problem. It is incredibly difficult from within a given area of expertise to spot difficulties that might seem fairly transparent from another viewpoint. This suggests that securing better chances for genuinely objective research in areas where research is likely to affect the world is not primarily an individual-level undertaking: it needs to be an institutional one. It is too easy for research to be biased/prejudiced/presumptive without notice or intention. Objective research is demanding. Too demanding, I urge, to put onto individual researchers. Objective modelling is an institutional problem.

I end by leaving you with a crucial research question: what can we do to help develop institutions that support modelling objectively? One place to start may be to provide a better sketch of the mechanisms, or causal pathways, through which failures of intellectual humility arise in scientific institutions. I have pointed to two sources from which such failures can arise: silo-ization and a lack of integration. But there are likely more, and further research should be done about the different ways in which the norms, habits and practices within scientific institutions may lead to failures of intellectual humility. With this in hand, we might then be in a better position to pose more pointed questions and devise more practical solutions. For instance, what kinds of interdisciplinarity might work to block the kinds of hubristic dismissal of others’ methods and assumptions that would make particular research projects more objective? Are there ways in which researchers can better *export* the insights of their own disciplines into other disciplines? These are important questions, and they deserve serious attention.

Notes

- a. From what is probably no longer an acceptable expression: ‘Who wears the trousers in the family’.
- b. For instance, this is what a lengthy report on risk concealment says: ‘After the accident, an investigation commission stated that the main cause of the disaster was “bureaucratic inefficiency, muddled withholding of alarming information, and buck-passing among top-officials” . . . Four years later, the court found 11 executives of ENEL/SADE and government officials guilty’ (Chernov and Sornette 2016).

References

- Austin JL** (1962) *Sense and Sensibilia*. Oxford: Clarendon.
- Bressan D** (2011) October 9, 1963: Vajont. *historyofgeology*. Available at <https://historyofgeology.wordpress.com/2011/10/09/october-9-1963-vajont/> (accessed May 2023).
- Cartwright N, Hardie J, Montuschi E, Soleiman M and Thresher AC** (2022) *The Tangle of Science*. Oxford: Oxford University Press.
- Cartwright N and Ray F** (2023) Modelling objectively. In Erdbeer RM, Hagenmeyer V and Stierstorfer K (eds) *The Modelling of Energy Transition. Cultures – Visions – Narratives*. London: Palgrave Macmillan.
- Chernov D and Sornette D** (2016) Examples of risk information concealment practice. In *Man-made Catastrophes and Risk Information Concealment*. Springer, pp. 9–245, doi: [10.1007/978-3-319-24301-6_2](https://doi.org/10.1007/978-3-319-24301-6_2), PMID: PMC7175960.
- Daston L and Galison P** (2007) *Objectivity*. New York: Zone Books.
- Franci A, Cremonesi M, Perego U, Oñate E and Crosta G** (2020) 3D simulation of Vajont disaster. Part 2: Multi-failure scenarios. *Engineering Geology* **279**, 105856.
- Gruber J** (2012) Lec 1 – MIT 14.01SC Principles of Microeconomics. Available at <https://www.youtube.com/watch?v=Vss3nofHpZI> (accessed May 2023).
- Hacking I** (2015) Let's not talk about objectivity. In Tsou JY, Richardson A and Padovani F (eds), *Objectivity in Science*, pp. 19–33. New York: Springer.
- Healy K** (2017) Fuck nuance. *Sociological Theory* **35**, 118–127.
- Hyde BVE** (2023) *Objectivity in Ideological Research*. Undergraduate dissertation, Philosophy Department, Durham University.
- Montuschi E and Barrotta P** (2018a) The dam project: Who are the experts? A philosophical lesson from the Vajont disaster. In Barrotta P and Scarafile G (eds), *Science & Democracy*. John Benjamins.
- Montuschi E and Barrotta P** (2018b) Expertise, relevance and types of knowledge. *Social Epistemology* **32**, 1–10.
- Preston J** (2018) Course notes on Austin's Sense and Sensibilia. Department of Philosophy, University of Reading. Available at <https://www.reading.ac.uk/AcaDepts/ld/Philos/jmp/TheoryofKnowledge/Austin.htm> (accessed May 2023).
- Ray F** (2022) *Notes on Intellectual Humility and Scientific Institutions* (unpublished).
- Rodrik D** (2015) *Economics Rules*. New York: Norton.
- Samuelson P** (1962) Economists and the history of ideas. *The American Economic Review* **52**, 1–18.
- Skidelsky R** (2021) *What's Wrong with Economics?* London: Yale University Press.
- Wolter A** (2016) The Vajont Slide: A new event chronology and the importance of geomorphology. AGU Blogosphere. Available at <https://blogs.agu.org/landslideblog/2016/02/25/the-vajont-slide-1/> (accessed May 2023).

About the Authors

Nancy Cartwright is a philosopher of natural and social science with special interests in modelling, causality, objectivity and evidence (especially for predicting policy outcomes). She is a Professor of Philosophy and Director of the Centre for Humanities Engaging Science and Society at Durham University in the UK and a Distinguished Professor at the University of California at San Diego in the US. She is

a fellow of the British Academy and the UK Academy of Social Science and a member of the American Academy of Arts and Sciences, of the Leopoldina and of the Academia Europa and is a former MacArthur Fellow.

Faron Ray is a PhD student in philosophy at the University of California, San Diego. He works in moral and political philosophy as well as in the philosophy of science, with special interests in causation, explanation and the philosophy of the social sciences.