

Nancy Cartwright

A Question of Nonsense

1. Introduction

As a philosopher of science I often feel compelled to teach about scientific realism. When I have done so, my best reconstruction of the issue formulates the central question this way: When are we rationally warranted in accepting a scientific theory? I have trouble teaching scientific realism however, and Edna's ideas, about picking and choosing (Ullmann-Margalit and Morgenbesser 1977, Ullmann-Margalit 2006a), and more lately the wonderful account of the interweaving that turns facts into evidence relative to the choice of other facts as evidence that we find in *Out of the Cave* (Ullmann-Margalit 2006b) have helped me figure out why. The reason is that the central question is, I think, nonsense. There are five major words in this question and every one of them is problematic. 'When are

- We
- Rationally
- Warranted
- In accepting
- A scientific theory?'

The issue of picking, or as I have been inclined to call it more graphically following Otto Neurath, *plumping*, has been with me for decades. As I shall explain, *Out of the Cave* has mattered to my recent efforts to develop a theory of evidence. Some people might assimilate evidence to warrant for acceptance, which I just said makes no sense. This assimilation, I argue, is a mistake. A theory of evidence asks, under what conditions does one fact speak for the truth of another? This endeavour is quite distinct from theory of knowledge, which is concerned with *warrant*, with justification for us. I am particularly concerned to keep them apart because I, unlike almost all other philosophers, do not believe in theory of knowledge, at least not

where I work: in the philosophy of science. There's no such thing as 'theory of scientific knowledge', no theory to address the question: When are we rationally warranted in accepting a scientific theory? That's because, as I said, I think there is not a single concept in this question that makes sense in real science.

So, let's have a look.

2. *What's Wrong with the Question?*

Scientific theory. For a long time philosophers of science have tried to axiomatize their favourite 'scientific theories', and in their own way scientists, especially physicists, have joined in this enterprise. Also for a good while now we in philosophy have debated the syntactic versus the semantic view of theories. Is a theory a set of claims that are supposed to be true of the world – in particular a small set of simple elegant general claims? Or is it instead a set of universal models that are supposed to map onto the world in some designated way? So we talk all the time about theory. But where in science is this special thing called 'theory'?

The last two decades of emphasis in Science Studies on scientific practice as opposed to scientific theory have underlined that there are a vast number of scientific practices that make up science, including classifying, experimenting, measuring, settling on standards, designing, and building machines – as in machine physics – and technology – like the laser or Lord Kelvin's Atlantic cable – constructing models and blueprints for specific cases, creating new substances and materials and new ways to change the world – which contemporary genetics witnesses repeatedly – developing new bits of mathematics, discovering, creating, and stabilising phenomena, calculating, inferring, making, refining and defending very specific concrete claims about specific situations and specific systems, providing a vast number of low level *ceteris paribus* laws, and so forth.

Scientific knowledge is reflected and encoded in all of these when carried out successfully. And, as I have argued for a long time, little of this knowledge – if any – is recorded, even in a very abstract way, in what gets labelled 'scientific theory' when we talk, say, about the semantic versus the syntactic view of theories. Some, though, is recorded in theories *as understood through practice and by practitioners*. It is not just that procedures and applications are necessary *in addition to* laws and theories. It

is the procedures and applications that give content to theoretical claims and laws. Without the procedures and applications, the claims are mere words, which is a well-known theme of Thomas Kuhn:

Consider, for a single example, the quite large and diverse community constituted by all physical scientists. Each member of that group today is taught the laws of, say, quantum mechanics, and most of them employ these laws at some point in their research or teaching. But they do not all learn the same applications of these laws, and they are not therefore all affected in the same ways by changes in quantum mechanical practice. On the road to professional specialization, a few physical scientists encounter only the basic principles of quantum mechanics. Others study in detail the paradigm applications of these principles to chemistry, still others to the physics of the solid state, and so on. What quantum mechanics means to each of them depends on what courses he has had, what texts he has read, and which journals he studies. [...] In short, though quantum mechanics (or Newtonian dynamics, or electromagnetic theory) is a paradigm for many scientific groups, it is not the same paradigm for all. (Kuhn 1962, 49–50)

These ideas are further developed in the works of Peter Galison that argues that different scientific groups imbed what can seem the same concepts in very different networks of practice and inference that thus embody different understandings. Communication between the groups with their different ‘thick’ understandings via stripped down ‘pidgin’ languages is not sufficient to do the tasks of any of the separate groups that speak them (Galison 1997).

I back up my scepticism about theory by a robust empiricism. The claims of science must be supported – in detail – by empirical facts. This support is witnessed by success in predicting and intervening precisely in the world. If so, what is supported then are claims as interpreted through the network of concrete assumptions and practices that afford the successful predictions and interventions. A *very* great number of claims are thus supported. But this wealth of very specific practice-interpreted claims are not in any way expressed or summarized by the axioms of theory, even if we allow for a great deal of adjustment in trying to get the summary to fit.

Who is we? Once it is acknowledged that,

- scientific knowledge is recorded in practices and machines as well as in words,
- the kinds of claims that are empirically supported are thickly interpreted through these practices, and
- there are very great differences in practice,

then the 'we' of 'When are *we* rationally justified...?' shatters into thousands of pieces. There are paradigms within paradigms within paradigms.

Who then is the 'we' of concern in the question 'When are we justified in accepting a scientific claim?' I want to point out that the 'we's' divide not only according to the network of paradigm practices and assumptions that provide a proper interpretation for the claim, but according to the purpose to which the claim will be put. I will tackle this issue simultaneously with the issue of acceptance and justification. But first – because I want to get you to think about *plumping*, or to use Edna's word, *picking* – I should like to discuss the rational of 'rational acceptance'.

Rational choice among thickly interpreted scientific claims: This we know is a vexed issue. From Otto Neurath, Thomas Kuhn, Paul Feyerabend, Imre Lakatos, and a great many others we have learned that there are no paradigm-independent standards from which to judge different paradigms. And even if we choose to operate always within a paradigm, what are taken to be the facts will always underdetermine any interesting scientific claims. To fix this problem there has been much talk of making decisions by resort to special 'virtues' that a theory or a hypothesis might have, like accuracy, consistency, scope, simplicity, production of novel predictions, and fruitfulness. The trouble is, these virtues are not up to the job. First, they are too vague to fix real choices; second, their verdicts often conflict; third, it is hard to see why if truth is the aim, they are virtues; fourth, the set seems arbitrary. Why focus on these rather than, for instance, Helen Longino's (1995) list that includes ontological heterogeneity, mutuality of interaction, applicability to human needs, and diffusion of power? So neither the facts nor the virtues can tell us what is the best choice to make.

Edna and Sidney Morgenbesser (1977) discussed issues like this with respect to rational choice theory and that is where I first learned about *picking*. In their case, rational choice was to be grounded in preferences. But, just as with facts and virtues in the case of arguments for scientific claims, sometimes preferences are not up to the job. They leave it open what the better choice is. Still we often must come to a decision. What do we do? We just decide. That's what Edna and Morgenbesser called *picking*.

Edna continued throughout her career to be concerned about picking and what happens in "the realm of decisions without preferences" (Ullmann-Margalit 2006a, 171). For instance:

One chooses for reasons; one *picks* when reasons cannot prevail. This happens when the alternatives are entirely symmetrical (or incommensurate). But reasons also fail to prevail when we come to the very end of the chain of reasons, when we run out of reasons altogether. If you choose to do X for reason A and, asked to justify A, you cite B and then you give C as your reason for B and so on, you eventually reach the very bottom, the substratum of all your reasons. (Ibid., italics added)

What Edna takes to be true for reasons to do X is equally true for arguments that speak for the truth of a scientific claim. The arguments require premises. And arguments for the premises require further premises. Eventually we choose not to pursue the arguments further. We rest content with where we stop, at least for the nonce. When it comes to science, it is widely acknowledged that there are no sure ways to build a case for our claims, and the same goes for the methods we might use to convince ourselves that one alternative is better than another. What seems hard to admit is the kind of voluntarism that follows in train, endorsed by Bas van Fraassen, Otto Neurath, me, and Edna when she speaks of decisions where “reasons cannot prevail” (2006a, 171). What we do then – all that’s left to do – is to pick, or as Otto Neurath put it: *we plump*.

So now let us think about acceptance and justification for it.

Acceptance. In a very nice 1953 paper, rejecting the value-free ideal and arguing that values *should* enter science that Heather Douglas has brought back to our attention, Richard Rudner begins his case with the claim that part of a scientist’s job is to accept or reject hypotheses. I disagree with Rudner’s starting premise, though not with his conclusion.

I do not think the common philosophical notion of ‘acceptance’ generally has any referent; and when acceptance goes, justification becomes essentially value infused, as it is throughout life. There are a very great many ways one can put a scientific claim to use. You can use it as an assumption in a blueprint to build a bridge or a laser or a back-to-work policy; as part of the basis for a new research endeavour; to predict the course of a cannonball or an accelerated particle at the large hadron collider in CERN; you can write it in a textbook – and that for school pupils, undergraduates, post graduates, ...; sign off on a consensus report about climate change or drug abuse policies that contains it; give a Nobel Prize to the person who envisioned it or evidenced it; vote in favour of an NSF grant for a young scientist to develop its ideas or to run an experiment to test it or to gather data to measure one of the parameters in it; and so forth.

These are all action descriptions with clear referents. But where is this thing called ‘acceptance’? Perhaps I have a peculiar phenomenology but I find nothing in either my inner or my outer life that corresponds to it. Nor do I understand it sociologically. Is it, as in Jeremiah 31.33, to ‘have the law inscribed in your inward heart’? I do not deny that special situations can create a reference in special cases. Marrano Jews perhaps could do and say exactly what Christians did even when talking most intimately among themselves and yet be in a state properly described as having the law of the God of Israel inscribed in their inward hearts. Perhaps here justification for writing something in our heart of hearts is an entirely personal matter. It may have no implications outside one’s own welfare and sense of self-respect. But these situations – where ‘acceptance’ has a clear meaning and where whatever it is one is supposed to be doing in ‘accepting’ has implications only for oneself – are not typical, and they are not salient for good scientific practice.

There is one sense of acceptance by particular communities that I do see instantiated – and that we must be wary of – because it implicates the welfare of others. That is, treating scientific claims like a well-tested product that you – the scientific community in question – put on a shelf in a warehouse for others to take off the shelf and put to use. My worries start when claims are put on warehouse shelves without clear instructions for use. As I have rehearsed: It is not strings of words or symbols with what Paul Grice would have called their ‘literal meaning’ that are tested in science. What the use of a claim in successful prediction, planning, and technology supports is that claim as interpreted through the dense web of techniques and practices which afford those successful predictions and uses. That empirical support can flow only to new predictions and uses which follow from the claim *as thus interpreted*.

Think about problem sets in textbooks. These both help provide not only content to the claims in the text but also clear paradigms of use. And as students get closer and closer to actual use as engineers, experimenters, lab technicians, doctors, options pricers, and the like, both the claims and the problems become more complex. These are cases where instructions for use are built into the product that is put on the warehouse shelves. Moreover, we will find a good many products that may at the loosest level of description look alike – they are on the shelves housing ‘Ohm’s law’ for instance – but these products are in fact highly differentiated.

The area I've been working in recently – evidence-based policy – is one where these lessons about thick interpretation, and the cautions which follow about putting scientific claims 'on the shelf' where they may be taken down to use in other 'thick' contexts, where they will in fact amount to a different claim, are largely ignored. There are a number of distinguished vetting agencies, which police evidence claims about policy efficacy: the Cochrane Collaboration for medicine and the Campbell Collaboration for social policy, for instance, or the US Department of Education's What Works Clearinghouse. These agencies check carefully that claims of policy efficacy are well supported by rigorous well-conducted studies, where an efficacy claim is a claim that the policy caused a targeted result in the study setting. Since these are causal claims, their gold standard for a top ranking – 1 in a scale of 1 to 4 say – is the randomised controlled trial (RCT).

So far so good. It is the next steps that are egregious, steps in which the label 'top ranked' slips from a specific efficacy claim to the policy itself. The very top rank, say 1**, is given when there are a number of well-supported efficacy claims about the same policy or a meta-analysis of studies about the same policy that treats the separate study populations as if they were drawn from one large sample. Policies associated with high rankings are then warehoused in special sites, like the What Works Clearing House or the seven What Works Centres on separate topics from crime to aging to education now being established in the UK; and policy makers mandated to pursue evidence-based policy are told to go to those sites to choose their policies if at all possible.

So here we have a case that may appropriately be labelled 'acceptance'. The vetting agency has *accepted* the claims it gives a high rank to (or probably better, for those who believe in degrees of acceptance, has accepted them to a high degree). The acceptance is witnessed by their recommending these claims for use by practitioners not able to judge the quality of the products themselves. But there is no clarity about:

- What the claims are.
- What the instructions for use are.
- How the uses to which they might be put can be warranted by the content of the claims.

The single efficacy claim ranked 1 on account of the good RCT has clear content: this policy produced the targeted effect in at least some units in the

study population. What about 1**? The methodology supports ‘This policy produced the targeted effect in some individuals in some number of different settings’. But the advice to policy consumers to choose policies with 1** rankings suggests that the claim is ‘This policy produces the targeted effect in most settings’. Or ‘This policy will produce the targeted effect in your setting’. Though sometimes there’s a warning label, as with the U.S. Department of Education website: ‘Don’t buy this policy if your setting is too different from the study settings’.

Here we might reasonably talk of ‘acceptance’ without asking about what is written in the inner heart. The warehouse vetters *accept* the 1 or 1** claims in the sense that they in good faith recommend them to non-expert users. But it is, I would judge, not justified acceptance, and not just because in the end the vetters have to plump. I am thus not concerned with *rational* justification, which is, I argue, anyway an empty notion, but with justification, simply justification.

Justification. My worries on this last issue follow the lead of Amartya Sen and Bernard Williams. In a 1983 article, Sen distinguished between the truth of scientific claims and the goodness of scientific accounts on the one hand and the goodness of actions on the other. The latter, being actions, are subject to moral scrutiny since all actions are subject to moral scrutiny. Sen says:

The problem here isn’t fearing that scientific action *might be* value-loaded, but fearing that it *might not*. Value-loading here is not so much a right as a duty. An action that is contrary to his or her own values [...] remains pernicious in terms of his or her own values, even if it happens to be related to science. (1983, 104, italics original)

So all those actions involving scientific claims that I described above when discussing acceptance: “You can use it ... to build a bridge; ... or for a new research endeavour; to predict the course of a cannonball; ... or write it in a textbook; ... sign off on a consensus report; ... give a Nobel Prize to the person who envisioned it;... vote in favour of a National Science Foundation grant to develop it...; and so forth” – all of these are actions and thus can call for justification – real justification. The justification for any particular use will involve genuine – quite ordinary – issues, which blend together questions of evidence, expertise, what can be taken as common knowledge, what is morally, socially, culturally acceptable, various benefits to various parties, what the costs are, what the costs of type 1 versus type 2 errors are,

and so forth: the usual ingredients of practical reasoning about what one ought to do.

Naturally issues of truth and evidence enter. But, recall Bernard Williams on morally thick descriptions such as *treachery*, *promise*, *brutality*, and *courage*: they are Janus-faced, looking to the world on one side and to issues of praise and blame on the other. As Williams puts it, they are “at the same time world-guided and action-guiding” (Williams 1986, 141). But they do not divide into two parts, glued together: matters of fact and matters of value. So too with actions involving scientific claims.

Consider: You are about to endorse a claim to a graduate student whom you know is readily influenced by you and is considering taking a position in a research group that uses this claim as a central pillar for its research. Before endorsing this claim in these circumstances, you should consider the evidence for it. You should also consider the abilities of the research team that propose to follow it up, the opinion of your colleagues about the evidence and what it shows, the talents of the student, the chances that she will end up with publishable papers even if the research program does not produce its promised results, and so forth. These issues will not separate nicely, as we might have hoped, to afford a two-stage deliberation: first wear your scientist’s hat to estimate the degree to which you are justified in ‘accepting’ the claim; then consider how justified you are in using a claim with that degree of warrant in the way proposed. Rather you must consider the issues all together in one fell swoop. And you should consider them. What you say to the student matters to her life, so you should take pains to ensure that what you do is justified. But that is not an exclusively scientific enterprise.

Interim conclusion. Returning at last to my opening concern, the central question of scientific realism: ‘When are we rationally justified in accepting a scientific theory?’ I conclude from these considerations that this question does not make sense. And providing rational justification for acceptance is not something that scientists should attempt to engage in.

Evidence and evaluation. BUT, there is something clearly at the core of science that this activity can be mistaken for, which is: producing evidence for scientific claims and producing reasons that it *is* evidence – just the kind of thing that Edna explores so beautifully in *Out of the Cave*. I advocate a stark theory of evidence, the Argument Theory. An empirical claim *e* is evidence for an empirical hypothesis *h* just in case *e* is an essential premise in a sound argument for *h*, that is, a valid argument with true premises.

My colleague Julian Reiss defends an alternative theory of evidence. He dislikes the Argument Theory because, he claims, it is not a theory of evidence *for us* (Personal communication). He is here criticising me on my very own grounds – hoist on my own petard. I regularly argue that metaphysics and methods must march hand-in-hand. It is a poor analysis of what something is – say, evidence, causality, the electron, malaria – if the analysis leaves it mysterious why our best methods for diagnosing whether something is evidence or an electron or a case of malaria should work.

Sherrilyn Roush has this worry too:

I don't see how we could use [the definition from the Argument Theory] to figure out when we can use a claim as evidence then, because it doesn't give us any guidance. Even if this tells us the conditions that obtain when something is evidence, it doesn't tell us anything about how to tell when those conditions do apply. (Personal communication, August 2012)

This objection seems to me misplaced. To figure out whether *e* is evidence for *h*, the Argument Theory guides you to look for good arguments connecting *e* and *h*. Of course it doesn't tell you how to tell if an argument is good. But that's not in its job description. Coming up with an argument is part of the ordinary normal science job of scientific discovery. To check that it is valid, perhaps one needs a good logician or a good mathematician. To tell if the premises are true, we employ the normal methods available in the paradigm in which we work for assessing the kinds of claims the premises make.

There is another version of the objection though that can seem to have more bite. We want, Roush urges, "some guidelines for dealing with our lack of the kinds of Arguments we seek. How confident should we be, given the very partial information that we have, that there is a sound, valid argument from [*e*] to [*h*], and *why*?" (Personal communication, August 2012)

My first remark on this is that there is no *should* about it. Following the point of view I have been developing here, there is no such thing as 'rational warrant for accepting' a scientific claim, even if we restrict attention to claims thickly interpreted by the practices in which they are imbedded. Roush herself, like many philosophers engaged in debates about statistical inference or confirmation theory, does not like the idea of acceptance either. And for good reason. Recall that at the conclusion of my discussion of rational choice, I urged, with Neurath, van Fraassen, and Edna in cases where preferences won't suffice, that in the end, the individual must plump. Here is what Roush thinks about that:

This is where a Bayesian has a strength because, since you never have complete evidence from which the conclusion follows deductively, you never plunge but only apportion confidence to the strength of your evidence, and acting on that confidence/degree of belief retains awareness of the imperfection of your epistemic situation by the fact that you (the rational subject) will put less stake on a [hypothesis] that you have less evidence for. (Personal communication, August 2012)

This looks again like a case of being hoist on my own petard since I am always urging, especially in advice about evidence-based policy, that when we act, we should generally hedge our bets heavily. Roush puts very directly the objection that a very great many people will have to the Argument Theory of evidence: that it provides no purchase for degrees of support.

I agree that it does not. But I do not take that to be a problem. I have argued that we should not be trying to answer the question ‘When are we warranted in accepting a scientific claim?’ Equally we should not be trying to answer the question, ‘What degree of confidence are we rationally warranted to have in a scientific claim?’ If the other premises of a valid argument from e to h are true, e partly supports h . If they are not, e does not partly support h . There’s no fact of nature, among the premises, which does the heavy lifting. They must all be true or no lifting is done at all.

There is no such thing as ‘the degree of confidence I am entitled to have in h .’ As a scientist, if I want to assess the truth of a novel hypothesis that is not already a part of the canon, I gather evidence for it and I reason about it. If I want to assess whether a result I have obtained – e – really is evidence for the hypothesis – h , I gather more evidence and do more reasoning. So: I have measured the deflection of a particle from a straight line trajectory in an electromagnetic field. Is my result, call it e again, evidence about h , the charge of the particle? To address that question, I turn to Maxwell’s electromagnetic theory; I double-check that my measurements of the electromagnetic field strength are correct as well as my identification of the particle’s mass; I check that the apparatus is working as I expect; and I go back over my calculations to see that I have made no misstep. Perhaps I also check that my understanding is right that Maxwell’s theory coupled with Newton’s does imply that the deflection in the trajectory is a function of the particle’s charge. That is how I defend my claim that the deflection is evidence for the charge. Then in the end, if the situation calls for me to use my result as the correct charge, I plump. My plumping may be justified and it may not be, given the situation and the use to which I am putting my result.

But that is an ordinary matter of practical/moral justification that we might face with respect to any act we perform.

Maybe I do not have to plump for h . Perhaps we are about to do a cost-benefit analysis and all it takes is my estimate of the chance that h is true. That is the happier situation that Roush envisages. And it is the reply that Richard Jeffrey (1956) gave to Richard Rudner (1953). Rudner argued that values necessarily enter science at the stage when the scientist accepts a hypothesis: If there are foreseeable heavily negative consequences from accepting the hypothesis if it is false, the scientist must have a far higher degree of confidence before accepting it. In Sen's terms, it seems that Rudner saw 'accepting' as an action that the scientist undertakes in a specific situation and as an action, it may call for justification – real moral justification. Jeffrey replied that it was not the job of the scientist to accept or reject claims but merely to put probabilities on them.

I agree with the first part of Jeffrey's claim. It is not the job of the scientist to accept claims since this notion of acceptance does not have any proper application in science. But I reject the second part. Suppose Jeffrey is talking about assigning objective probabilities for the hypothesis (if there are any such things). In that case the scientist is just back in the normal science project of investigating an empirical hypothesis – that $P(h) = X$ – and there is no more call for *accepting* this empirical claim than for accepting h . Though of course there may be many actions that the scientist is called upon to engage in that involve this claim in important ways; and as actions, they may call for justification... again, real practical justification.

There is another problem that arises from the Argument Theory, which I urge more sympathy with than with its failure to accord degrees of confidence we are entitled to have. This is the problem of recognizing that something is or is not evidence for a hypothesis. I urge more sympathy for this problem because whether something is evidence or not can play a role both in assessing the truth of the hypothesis and in justifying us in many of the activities we might undertake involving that hypothesis, like assuming it in building a bridge or funding more research for a program growing out of it. The problem arises because on the Argument Theory, evidence is a 3-place relation, e is evidence for h relative to a good argument for h in which e plays an essential role. This means that whether a given fact is evidence or not is conditional on the other premises, each of which in turn is itself evidence on the Argument Theory – but only relative to e 's being evidence. If e is

not evidence, then they are not either (or at least that is not secured by the argument in view), but e cannot be evidence unless they are.

This is a problem for us, but not I think an objection. It is what evidencing looks like in real cases. I first saw this clearly in grappling with Edna's *Out of the Cave*. There we see that different hypotheses about the Dead Sea Scrolls and the nearby settlement at Qumran are supported by different complex narratives. Two important features of these narratives stand out. First is how compelling they are; should they be true, they almost force their conclusions – as in a good argument. Second, from the point of view of one narrative, most of the facts adduced as evidence in the others are irrelevant – they only make a difference to the case *given* the story we are already in the midst of. If the Qumran site was the setting of a military fortress then the fact that it is two kilometres from the place where the scrolls were located has no bearing at all on whether the scrolls are Essene or not.

One can assimilate this to a well-known fact that we teach undergraduates in our philosophy of science classes under the heading 'Duhem-Quine holism'. Single hypotheses are seldom testable in isolation. Generally it takes a whole set of claims to deduce an observable consequence and should that consequence not obtain, it can be the fault of one of the 'auxiliary' hypotheses rather than of the hypothesis we hope to test. Put this way, this seems to be a fact about *hypotheses* and how secure they can be. Quite possibly I have been slower than others to see that this way of looking at Duhem-Quine holism can be turned upside down. It is equally about what is and what is not *evidence*. The deduced consequence is evidence only on condition the 'auxiliaries' are true. Immediately this makes evidence a 3-place relation, contrary to many accounts, from Peter Achinstein's in which the essential criterion is that e be explanatorily relevant to h , which Achinstein seems to treat as 2-place, to various probabilistic accounts that build from 2-place conditional probabilities, like $P(ehl)$, $P(hle)$, $P(h'le)$ for some alternative $h' \neq h$, and $P(elh')$.

But I did not start by turning Duhem-Quine wholism on its head. I started with *Out of the Cave* and its thick narratives of interwoven strands each of which leads to the hypothesis only because the others do so as well; where without the rest of the story, the separate pieces have no bearing on the final conclusion – they are just dangling odd facts that may or may not be of interest on their own to someone or for some purpose, some purpose other than speaking to the truth of the hypothesis. What I took away from Edna's

book then were two related insights. The status of any one fact as evidence depends on the evidential status of a host of other facts, and taken all together what counts as evidence must be able to shape itself into a good argument: if the rest of the facts aren't true that could flesh out the narrative to compel the conclusion, the facts we have gathered do not bear one way or another on the truth of that conclusion.

3. Conclusion

Whether we are scientists or not, or 'acting in the role of scientist' or not, we are responsible for what we say and what we do, so we may need justification for that because what we say and do can affect our own welfare and the welfare of others. Apart from that what we believe is a matter that hardly invites justification, unless God is looking into our inner hearts to judge us by what is found there. So do not worry about 'warrant for belief'.

But sometimes we are in the business of assessing the truth of a claim. In that case, evidence can help: facts that speak for the truth of a claim. Evidence for a claim, I have maintained, are facts that play an essential role in a good argument for that claim. And the better the evidence for the premises of that argument and the better the arguments for the premises in the premises, the better off we are at our enterprise. But eventually we have to pick, to plump.

I say that there is no single 'we' in 'we accept' a scientific claim nor is there any sense to 'accept'. Okay then, one may object, why should non-experts aim to do anything about anthropogenic climate change? The standard answer is because the scientific community accepts that it is real and that it will be a disaster, which uses just those concepts I warn against.

I think they should because there is good evidence that anthropogenic climate change is real and that it will lead to disaster; because most of us reading this are able to do some thing or other, if only to vote, to improve matters; and because we are most of us, *ceteris paribus*, not justified in not thus acting. This verdict comes in part from plumping for the truth of the claim that the Intergovernmental Panel on Climate Change (IPCC) is likely to have got it right. And in part because the causality claim is in a pidgin we all can speak and even within that pidgin and the thin meanings of the concepts there, disaster ensues from business as usual. What specifically to do about building seawalls in Brighton or maintaining green taxes on energy providers in the UK is, though, a matter that can only be seriously investigated

using thick concepts supported by networks of small communities speaking their own thick languages.

Edna has taught us that sometimes we need to plump, something many would prefer to avoid because we feel justified when decisions are dictated by rules of rationality. Plumping is an action, though, and can affect others. So where we come down on the likelihood that the IPCC has got it right can call for serious justification.

University of Durham

References

- Galison, Peter. 1997. *Image and Logic: A Material Culture of Microphysics*. Chicago: University of Chicago Press.
- Jeffrey, Richard. 1956. "Valuation and Acceptance of Scientific Hypotheses," *Philosophy of Science* 23 (3): 237-246.
- Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Longino, Helen. 1995. "Gender, Politics, and the Theoretical Virtues," *Synthese* 104 (3): 383-397.
- Rudner, Richard. 1953. "The Scientist qua Scientist Makes Value Judgment," *Philosophy of Science* 20:1-6.
- Sen, Amartya. 1983. "Accounts, Actions and Values: Objectivity of Social Science," in Christopher Lloyd, ed., *Social Theory and Political Practice: Wolfson College Lectures 1981*. Oxford: Oxford University Press.
- Ullmann-Margalit, Edna and Sidney Morgenbesser. 1977. "Picking and Choosing," *Social Research* 44 (4): 757-785.
- Ullmann-Margalit, Edna. 2006a. "Big Decisions: Opting, Converting, Drifting," *Royal Institute of Philosophy Supplement* 58:157-172.
- Ullmann-Margalit, Edna. 2006b. *Out of the Cave: A Philosophical Inquiry into the Dead Sea Scrolls Research*. Cambridge: Harvard University Press.
- Williams, Bernard. 1986. *Ethics and the Limits of Philosophy*. Cambridge: Harvard University Press.