

Expert Judgement without Values: Credences not Inductive Risks
Rivkah Hatchwell
David Papineau

[To be published in D. Pritchard and M. Farina (eds) *Expertise*, Oxford University Press]

Many contemporary philosophers of science think it is proper for ethical considerations to play a role in deciding scientific claims. We disagree, not least because this is likely to bring science into disrepute.

Plan

1. Initial Meta-Ethical Observations
2. Evidence not Expedience
3. Rudner and “Inductive Risk”
4. Whose Values?
5. Which Choices?
6. The Current Debate
7. Credences versus Significance Tests
8. Rudnerian Calculations
9. Higher-Order Judgements
10. The Problem of the Priors
11. Against Categorical Beliefs
12. Racing to the Bottom

1. Initial Meta-Ethical Observations

We shall be arguing that values should play no role in deciding scientific claims. But let us start by distancing ourselves from a more general notion of scientific “value freedom”.

It is often said that science should be “value-free”. As it is usually meant, we do not think this dictum is defensible. That is, we do not think that scientific claims should be free of evaluative implications. In our view, such implications are both inevitable and desirable.

The dictum that science should be “value-free” seems to us to be premised on outmoded meta-ethical views. We have in mind vulgar relativist views on which there is some unbridgeable gap between facts and evaluations. On such views, any descriptive claim leaves it quite open what evaluations should follow. Different people are free to adopt different evaluative responses to the facts. Given this, the business of science is solely to determine descriptive information, not to prescribe any evaluative add-ons. Not only will science be exceeding its authority if it adds evaluative glosses to its claims, but it will run the risk of alienating audiences with different evaluative reactions.

This view of values might still be common in certain scientific circles, but from a meta-ethical perspective it is nearly a century out of date. Contemporary meta-ethical views are generally agreed in rejecting relativism. This covers both realist and non-realist meta-ethical stances. While these options disagree about the import of ethical claims, they generally

agree in their resistance to relativism. Thus moral realists hold that the natural facts in and of themselves metaphysically entail evaluative claims, and perhaps also conceptually entail them. And most moral non-realists, while denying that moral attitudes answer to substantial moral facts, still incline to the non-relativist view that humans will generally be drawn to similar evaluative conclusions given the same descriptive facts. On all these views, then, we can expect given factual claims to have agreed evaluative consequences.

There is of course much more to be said here. We shall content ourselves with pointing to the oddity of relativism. Is it really optional how you should react evaluatively to news that, say, an infectious agent is causing children to go blind, or that a breakdown in medical supplies is resulting in painful deaths? These are not unusual examples. Many factual descriptions of real-life situations have obvious ethical implications. In consequence, any search for “value-free” terminology that is somehow insulated from evaluative connotations is doomed to failure. Given that many descriptions of natural facts automatically dictate evaluative judgements, science cannot possibly be “value-free”. What we want from science is the facts, not the facts shorn—*per impossibile*—of evaluative significance.¹

Even if people agree generally about values, this does not mean that they will always agree about how to weigh them against each other. In rejecting relativism, we take it that humans will generally agree about which features of a situation have positive value and which negative value. Suffering and death are bad, health, prosperity and justice are good. But this leaves plenty of room for even fully informed people to disagree about the best course of action, all things considered. For example, some might attach more weight to prosperity, others to justice. Agreement on *pro tanto* values does not therefore ensure a universally shared calculus for decisions. Even reasonable people in full possession of the facts can diverge in their choices of political and practical policies. This point will be central to what follows.

2. Evidence not Expedience

We say it is futile for science to aspire to “value-freedom”. This does not mean, however, that it is acceptable to adopt scientific claims *because* of your values. Scientific claims may carry inevitable evaluative consequences. But this by no means implies that our scientific beliefs should be tailored to pre-existing values, rather than to evidence for their truth.

After all, why would anybody want to adopt scientific claims in the absence of good evidence for them? Whatever values are in play, it would seem to be in everybody’s interests to believe truly if they can. If you harbour false views of the world, you will be ill-placed to select effective means to your ends and your projects will fail. Tailoring your beliefs to pre-existing values rather than good evidence does not seem a sensible strategy for anybody to adopt.

Of course, some people might have an interest in *persuading others* to adopt false beliefs. If I am rich, I might promulgate the theory that high marginal taxes have a negative impact on

¹ In line with this, we do not take “thick” concepts (Anderson 2002, Kirchin 2013) or “mixed” concepts (Alexandrova 2018) to be incompatible with good science.

the overall economy, in the hope this will stop my own taxes being raised. Again, if I grow crops, I might try to persuade the public that neonicotinoids don't poison bees, so as to lessen the danger of a ban on my using them. Even so, it would be silly for me to believe these things myself if the evidence doesn't support them. As before, it will only reduce the chance of my being able to select the appropriate means to my ends. In particular, if I care at all about the general economy, or the bees, as well as about my income and my crops, then I will be ill-advised to embrace my own propaganda, for it will only lead me to underestimate the dangers to some of the things I value, and so to advocate policies ill-suited my overall interests. (And if I don't care at all about the general economy or the bees, then why should I want to fool myself about the impact my selfish interests will have on them?)

As a matter of fact, of course, people with sectional interests do often delude themselves in just this way. In addition to promulgating factual views that favour policies that serve their own interests, they tend to believe those views themselves. It is an interesting question why this happens. After all, it is not only epistemically but also prudentially irrational, for the reasons just given. Still, people are not always rational. Perhaps an aversion to "cognitive dissonance" is operative: rather than accept that your different ends are in conflict, you come to persuade yourself, against the evidence, that they are in truth jointly achievable. Or perhaps a more simple form of self-deception is responsible: if you are going to espouse propaganda that serves your sectional interests, it is psychologically more comfortable, and indeed more effective, to believe in it yourself.

But we can leave these psychological conundrums for others. The nether regions of the human mind are not our present subject. Whatever irrational processes sometimes influence humans, we take it to be uncontroversial that it's bad to induce others to form unevicenced beliefs in order to further your sectional interests, and it's not made any better if you end up duping yourself into those beliefs too.

Surprisingly, an influential strain in contemporary philosophy of science comes close to arguing that it is the duty of scientists to do just these things.

3. Rudner and "Inductive Risk"

The line of thought goes back to the 1950s. In "The Scientist Qua Scientist Makes Value Judgements" (1953), Richard Rudner argued that existing evidence can never conclusively prove or falsify any scientific claim. As he saw it, there will always be some inductive risk² involved in definitely accepting or rejecting any claim on the basis of finite evidence. In the face of this risk, said Rudner, scientists ought to compare the relative dangers of falsely accepting or rejecting the claim. To the extent that false acceptance will have significant adverse consequences, they should demand stronger evidence before endorsing it. And the

² As it happens, the now-standard phrase "inductive risk" does not appear in Rudner's original paper. It was introduced by Carl Hempel in "Science and Human Values" (1965) when discussing commitment to universal generalizations.

same applies in reverse. If mistaken rejection will lead to costly results, they should wait on strong negative evidence before rejecting the claim.³

Let us illustrate with neonicotinoids and the bees. Suppose that our hypothesis is:

(N) Neonicotinoid spraying reduces bee populations by more than 20% each year.

And suppose that neonicotinoids will be banned if and only if this is accepted.

So if we reject this hypothesis when it is true, we will carry on spraying neonicotinoids and the bee population will be inadvertently depleted. On the other hand, if we accept the hypothesis when it is false, we will stop spraying and crop yields will decrease to no good effect. Suppose now that scientists agree that a significant fall-off in the bee population would be a worse outcome than reduced crop yields. According to Rudner's analysis, they should therefore make it easier to accept that neonicotinoids significantly affect bee populations. They should be ready to embrace N on weaker evidence than they'd require to reject it. This would have the happy result of reducing the risk of continuing to spray neonicotinoids when it will be hurting the bees.⁴

4. Whose Values?

One obvious problem with Rudner's recommendation is: "Whose evaluations of the costs and benefits should go into the calculation?"

Rudner's suggestion is that the scientists should look to their own values to tell them when the evidence is sufficient to warrant accepting or rejecting a hypothesis. In our example, the idea was that, given that scientists will value bee populations more than crop yields, they should be relatively quick to accept that neonicotinoids significantly affect bee populations. Still, there is no reason why everyone should share these values with the scientists. Rural communities, or indeed the general public, might well have different evaluative priorities. They might attach higher value than the scientists to food production, and so be inclined to demand more evidence before banning neonicotinoid spraying on the grounds that it has been shown to significantly affect bee populations.

Relying on the values of scientists would thus seem to be in danger of pre-empting decisions that should properly be subject to democratic debate. At bottom, the Rudnerian line means that neonicotinoids will be banned more readily if the scientists prioritise bees over crops more than the wider public do. However, put like that, it seems clearly wrong. Why ever should the scientists' values carry more weight in these matters than those of wider public?

³ Rudner in fact only claims that scientists *will* compare these dangers, not that they *should*. In line with most subsequent literature, we shall skip over this. Given that a Rudnerian appeal to values is not the only possible response to inductive risk (see Levi, 1967, John, 2015a), the philosophically pressing issue is whether it is the best response. The actual practice of scientists is a secondary matter.

⁴ See John (2015a) for further discussion of this example.

What about modifying the Rudnerian line so that the scientists don't use their own values in deciding how to set evidential thresholds, but rather those of the democratically answerable politicians they are advising?⁵

That strikes us as better, but not a lot better. If the values determining evidential thresholds are those of the democratic majority, then at least the policies that come to adopted will reflect those majority values. However, there still seems something wrong with allowing even majority values to influence decisions on scientific hypotheses. After all, not everybody shares those majority values, and such dissenters might reasonably take exception to scientific experts categorically asserting claims on relatively weak evidence, just because those claims align with majority preferences.

5. Which Choices?

Our main concern in this paper will be with the way that Rudnerian choices are sensitive to *whose values* are in play, in the way just explained. But before proceeding we would like to observe that Rudnerian choices are also sensitive to *which choices* are being addressed.

At bottom, the Rudnerians are making a choice aimed at certain non-epistemic ends—in our example, food production and bee preservation. Their eventual aim is to maximise expected value over those ends. True, it is not a simple expected value maximization calculation, because their framework requires them to find a route to these ends that goes via a choice of an evidential threshold and subsequent formation of a categorical judgement. Even so, if we ignore the bells and whistles, the Rudnerians are in effect making a choice aimed at certain non-epistemic ends.

One consequence is that, in a different context where different ends are being pursued, Rudnerians might end up setting a different evidential threshold for some hypothesis, and so might reach a different categorical judgement, even given the same evidence. To illustrate, imagine that we are still interested in:

(N) Neonicotinoids reduce bee populations by more than 20% each year.

But now imagine a context where we are concerned, not with bee conservation and crop yields, but about the best way to deal with an infestation of bees that is disrupting a school. Should we use neonicotinoids or some more proven but more expensive insecticide? Assuming it would be much worse to fail to deal with the infestation than to spend a bit more money, Rudnerians will now have reason to set a higher threshold for accepting N, the better to guard against using the cheaper neonicotinoids when they are in truth ineffective. So they might end up categorically rejecting N on the basis of evidence which would have led them to accept it when they were concerned with bee conservation and crop yields.

In the last section we pointed out that the Rudnerian setting of evidential thresholds is sensitive to whose values are at issue: the thresholds favoured by ecologists will be different

⁵ This is advocated by Heather Douglas (2000, 2009). She argues that evidential thresholds ought to be determined, not by the values of scientists themselves, but by democratic values generated by involving the public in science.

from those favoured by agriculturalists. We now see that, even when we are dealing with one fixed set of values, the setting of thresholds is also sensitive to which choices are being addressed.

It might have been thought that it is the responsibility of science to issue judgements that all people can trust to inform whichever choices they need to make. Rudnerian science, however, turns out not to be able to fulfil this responsibility. The categorical judgements it delivers are inevitably tailored to specific values and specific choices.

6. The Current Debate

The debate about values and inductive risk has had a had a new lease of life over the past couple of decades, no doubt due to the increased awareness among philosophers of the political implications of scientific claims. A surprising number of philosophers endorse the Rudnerian line that values should influence commitment to scientific claims, among others Philip Kitcher, Heather Douglas, Katie Steele, Eric Winsberg, and Torsten Wilholt.⁶

We are going to argue against the Rudnerian line. Our objection will not be to his proposal for setting evidential thresholds by reference to values, but to the whole idea that we need such thresholds at all. We think that Rudner goes wrong from the start, in assuming that science needs to overcome “inductive risks”. Once this assumption is made, the intrusion of values into scientific theory choice becomes hard to avoid. Still, as we shall explain in the next section, we don’t need to start where Rudner does. A better account of science is available, one on which worries about the role of values fall away.

We should say at this point that we will not be able to address all the issues raised in the contemporary debate about values and inductive risk. These are many and varied. For a start, different authors locate the role of values at different stages of the scientific process. Thus Douglas (2000) argues that they are needed at intermediate levels of research, as for example when a tumour is categorised as “cancerous” or not. Others have suggested we should be concerned, not with the claims scientists themselves believe, but rather with those that they communicate publicly (John 2015a). Yet others have focused on particular cases, especially involving climate change. In this last context some have advocated a system of intervals to represent uncertainty (Betz 2013) while others (John 2015b, Frank 2017) have responded that this will not eliminate the need to address inductive risks.

Rather than attempting to engage with all these detailed concerns, our focus in this paper will be on more foundational matters.

7. Credences versus Significance Tests

Our view, as we said, is that Rudner goes wrong from the start in the way he thinks about the scientific enterprise. He assumes without argument that the job of science is to reach *categorical judgements* about the truth or falsity of scientific claims. However, it is contentious that we should think of science in this way. The alternative would be for science

⁶ See Kitcher 2001, Douglas 2000 2009, Wilholt 2009, Steele 2012, Winsberg 2012.

to advise us on what *credences* (equivalently *degree of belief*, or *subjective probabilities*) in scientific claims are appropriate given the evidence. After all, Rudner himself insists that existing evidence will never compel a definite decision on any scientific claim. Given this, why assume that science must issue in definite categorical judgements⁷, rather than indicating appropriate degrees of belief?

This response to Rudner's line is of long standing. Rudner's original paper in 1953 soon prompted a response from the early Bayesian Richard Jeffrey, who observed that there would be no need to invoke values to fix evidential thresholds for categorical commitments if scientists simply eschewed such commitments (Jeffrey 1956). Instead they could just report the degree of belief warranted by the evidence. If scientists didn't try to bridge the "inductive risk gap" by setting evidential thresholds, Jeffrey observed, they wouldn't need evaluate the costs and benefits of doing so. As Jeffrey saw it, instead of forming categorical judgements, "the activity proper to the scientist is the assignment of probabilities" (1956 237).

At the time he wrote, Jeffrey's Bayesian line was fighting against the tide of scientific orthodoxy. For better or worse, scientists were committed to arriving at categorical conclusions, and they standardly appealed to the apparatus of *significance testing* to facilitate this.

We take it that this apparatus will be familiar, but it will be convenient to rehearse the basics briefly. The central idea of significance testing is that, when presented with some hypothesis and some evidence supporting it, scientists should assess how likely it is that we would have found supporting evidence of that strength if the hypothesis were not true. If statistical analysis shows that this probability is small enough, say less than 5%, or less than 1%, then we should conclude that the hypothesis is indeed true. The probability required is called the *significance level*, or the *p-value*, and the rationale for this procedure is that it means we will only rarely end up endorsing a hypothesis when it is in fact false—5% of the time, say, or 1%, depending on what *significance level* is chosen.

The Rudnerian line on value-ladenness goes hand-in-hand with this methodology of significance testing. Indeed Rudner and many others in the ongoing debate explicitly refer to significance levels as what they mean when they say thresholds for acceptance should be sensitive to the costs of inductive mistakes. As they see it, if there is a high cost to rejecting the neonicotinoid hypothesis N when it is actually true, then we should choose a relatively undemanding significance level, with a 5% p-value or higher, rather than 1%, so as to lessen the risk of wrongly going on spraying.

Still, this affinity with the methodology of significance tests lends no real support to the Rudnerian approach. This is because the whole logic of significance testing is itself deeply flawed. When properly examined, the idea that we should categorically accept a hypothesis

⁷ How exactly should we understand "categorical judgement" in the present context? Given that the attitude is supposed to be a response to non-zero inductive risks, it can scarcely be required to exclude any possibility of error. What it does require, however, is unclear. However, we shall say little about this issue here (though see section 11 below). Since we deny that categorical judgements play any serious intellectual role in science or policy making, we regard the issue as a problem for our opponents rather than ourselves.

H whenever we observe a significant result makes little sense. Sure, if we do that we'll only end up accepting H *when it is false* one time in 20 (or 100, or whatever). But that's no reason at all to think H is likely to be true—which one might have thought would be a minimum requirement for categorically accepting it.

To see the point, imagine that scientists are indefatigable assessors of unlikely hypotheses, and that only one in 100 of the hypotheses that they test are in fact true. Now focus on the 99 of every 100 hypotheses that are false. We can expect about 5 of them to show a significant result at the 5% level, and about 1 to do so even at the 1% level. After all, that is precisely what “significance level” means—how often you will get a significant result if your hypothesis is false. So, even if we take it for granted that true hypotheses will always generate significant results themselves, a significant result at the 5% level will only warrant about a 1/6 credence in the hypothesis under test—since we will get a significant result from false hypotheses 5 times for every one from a true hypothesis—and even at the 1% level only a 1/2 credence will be warranted. This would seem a terrible basis for the categorical commitment to the hypothesis that is mandated by the logic of significance testing.⁸

Much of the motivation for the traditional “frequentist”⁹ approach to statistical inference lay in the way that the alternative “Bayesian” credences depend on “prior probabilities”, that is, initial pre-testing degrees of belief in the hypotheses under test. For example, the toy analysis of the preceding paragraph hinged on the prior assumption that the hypothesis under test has a one in a 100 probability of being true. Since this assumption preceded the gathering of any evidence, it would seem to lack any proper empirical backing. Reliance on such prior probabilities thus seems in tension with a commitment to evidence-based science. The attraction of frequentist significance testing was that it seemed to offer a way of by-passing any dependence on arbitrary initial ideas.

We shall come back to the supposed arbitrariness of prior probabilities in section 10 below. Still, whatever view we take on that issue, we trust that it is already clear that the methodology of significance testing is itself fatally flawed. The toy analysis above was not meant to be fanciful. Science is currently beset by a “replication crisis”. In areas that rely on statistical evidence, surveys indicate that most published results accepted on the basis of significant evidence turn out not to be replicable in repeat studies (Open Science Collaboration 2015, Baker 2016). No doubt a number of different factors have contributed to this, but a growing consensus now takes the view that the main cause has simply been the misplaced faith in significance testing. As we have seen, failures of replication are exactly what we would expect if the preponderance of hypotheses subject to test are in fact false. From this perspective, the replication crisis is simply the chickens of significance testing coming home to roost.

⁸ For more on this critique, see Papineau 2018, Bird 2021.

⁹ From a philosophical point of view, there is little to recommend “frequentist” as a description of Neyman-Pearson significance testing. The distinguishing feature of that methodology is that it relies solely on objective probabilities, not that it favours a frequency interpretation of those objective probabilities over alternative propensity or chance interpretations. Still, this usage is now so well-established that there seems little point in resisting it.

8. Rudnerian Calculations

The last section outlined the natural affinity between the Rudnerian view of science and the traditional methodology of significance testing. Both are entirely focused on categorical judgements and do not assign any significant role to credences in scientific practice.

Somewhat curiously, however, the Rudnerian approach to theory choice itself hinges essentially, if implicitly, on credences. It needs them for its calculations about how high or low to set evidential thresholds. In this section we would like to digress briefly to explain this. The point is not crucial to our overall argument, but will be worth rehearsing, if only because it adds to the case that there is no cogent way of avoiding credences in science.

Rudnerians argue that we need to assess the possible costs and benefits of different evidential thresholds. However, Rudnerians are rarely specific about the calculations that this might involve. To make things definite, let us assume, in line with much of the Rudnerian literature, that the evidential thresholds in question are p-values. For example, to stick with the conventional options, should we set the significance level at 5% or 1%? The former will give us a 5% and the latter a 1% chance of false positives (taking neonicotinoids to be toxic when they aren't, for example). Then there is also the issue of false negatives. To figure out the chances of these we will need to estimate the *power* of the tests at issue. Suppose for the sake of the argument that we are able to do this. This will then tell us how much higher the chance of false negatives (not taking neonicotinoids to be benign when they are) will be with a 1% than a 5% p-value.

Let us thus suppose that we have so determined the relative chances of false positives and negatives for both 5% and 1% p-values. The crucial Rudnerian idea is then that we need to add in *evaluations* of the costs of false positive and negatives to get from these chances to a decision about evidential thresholds. As the Rudnerians see it, the ecologists will be more worried about false negatives than the agriculturalists, and vice versa about false positives, and so the ecologists will tend to favour the 5% significance level.

Well, the ecologists might be *more inclined* to the higher p-value than the agriculturalists. But will they actually opt for it? If they are rational, that will have to depend on a further factor, namely their *credence* in N—what probability do they attach neonicotinoids being toxic?¹⁰ After all, even the ecologists will presumably attach *some* value to maintaining food production. And so they won't want to threaten to reduce food production without some reasonable expectation of compensating benefit. But, for all that's been said so far, that could well be precisely the upshot of opting for the 5% significance level. To make the point vivid, suppose the ecologists only had a one in a million credence that neonicotinoids are toxic. Then it would be very odd for them to opt for the 5% rather than 1% significance level.

¹⁰ Should this be the prior or posterior credence in N? We leave it to the Rudnerians to tell us. After all, it's their calculation and they certainly need some credence for N to carry it out. They're probably better off invoking their *prior* credence in N—the idea would thus be to choose the significance level before gathering evidence for or against N. Alternatively, they could in principle get the evidence, update their credence in N on that basis, and then use that *posterior* credence to choose the significance level for deciding whether or not to accept N—It seems particularly weird, though, to design a significance test to take you from the evidence to a categorical judgement when you have already used that same evidence to update your degree of belief.

That would multiply the chance of stopping spraying and reducing crop yield fivefold, even though such a cessation is almost certain to be pointless.

All right, we could imagine extreme ecological values—no amount of wheat is worth even one bee’s life—that would lead ecologists to favour the higher p-value even if their credence in neonicotinoid toxicity was miniscule. But the general point should be clear. If you attach any value to food production, you already need a credence about the toxicity of nicotinoids to assess whether the increased threat to food production posed by the 5% p-value is worth it.¹¹

9. Higher-Order Judgements

As we said, Bayesians like Richard Jeffrey point out that science can avoid all the Rudnerian contortions straightforwardly enough by simply avoiding categorical judgements and sticking to quantitative probabilities. A long-standing objection to Jeffrey, however, is that categorical judgements will inevitably re-emerge at a higher level. When the scientists say that some hypothesis *P is probable to such-and-such a degree*, won’t that itself be a categorical judgement about a probability, and so won’t it reintroduce all the issues of inductive risk?

This anti-Bayesian objection strikes us as muddling up subjective and objective probabilities.

We take *subjective probabilities* (or *credences*, or *degrees of belief*) to characterise the strength of a person’s commitment to a proposition, the extent to which they expect it to be true. If your subjective probability that Trump will win the 2024 presidential election is 0.4, say, then you expect it to that degree, as manifested in the choices you make, for example by hedging proportionately more against his losing than his winning. (And we take *rational subjective probabilities/credences/degrees of belief* to be the extent to which a person *ought* to expect the relevant proposition to be true.)

By contrast, we take *objective probabilities* to characterise the tendency for a certain kind of outcome to happen in a certain kind of situation, such as the tendency for fair coins to come down heads when tossed, or for radium atoms to decay within 1000 years. The nature of objective probabilities is much debated, but they are clearly distinct from subjective degrees of belief. Radium atoms would still have a certain objective tendency to decay within 1000 years even if no humans or any other creatures had ever thought about the matter.

¹¹ Here’s a table to make the issue clear. False negatives in the left-hand column, false positives in the right. Let us suppose, for the sake of the argument, that opting for a 1% p-value rather than 5% will reduce the power of the test from 90% to 80%. Environmentalists will be more worried about the left-hand false negatives than the agriculturalists, and vice versa for false positives. But are the environmentalists worried enough to make it worth choosing 5% and so substantially increasing the chance of a false positive and reducing crop yields unnecessarily? This must depend, not just on how much they are concerned about the bees, but also on their overall credence for Toxic versus Benign.

	Toxic	Benign
5% p-value	Pr(Allow) = 0.10	Pr(Ban) = 0.05
1% p-value	Pr(Allow) = 0.20	Pr(Ban) = 0.01

Now, objective probabilities are part of the *subject* matter of many scientific claims. For example, we might be interested in:

(M) The objective probability of lifetime breast cancer for women with the BRCA1 gene is over 50%.

On the other hand, subjective probabilities are a matter of *attitudes to* scientific claims. Someone believes a scientific claim to a certain degree—25%, 50%, 90%, whatever. The scientific claim in question might have objective probability as part of its subject matter, as in M, or it might not.

Rational subjective probabilities are not entirely unrelated to objective ones. There is an obvious and natural connection between them—David Lewis called it the “Principal Principle” (1980). This connection applies in situations which do involve objective probabilities, like coin tosses and radium atoms. The connection is then that, to the extent we know the objective probability of some outcome, we ought to set our subjective degree of belief equal to that objective probability. For example, if you think some coin’s landing heads has an objective probability of $\frac{1}{2}$, then you ought to have a credence of $\frac{1}{2}$ in its landing heads.

However, this kind of constraint on subjective probabilities has no obvious application to the kind of case we have been considering, where we attach a subjective probability, not to the claim that a coin will land heads, or that a radioactive atom will decay, but rather to a general scientific hypothesis about the toxicity of insecticides, say, or the risks associated with the BRCA1 gene. Even if objective probabilities enter into to subject matter of such hypotheses, as in the latter example, this doesn’t at all mean that the hypothesis itself will have an objective probability. It is not as if there is some physical mechanism that systematically ensures that a certain proportion of such scientific hypotheses are true, in the way that the physics of coin tossing is responsible for a certain proportion of coins landing heads. So we don’t here have the same kind of constraint on credences that arises with coin tosses and the like. Our credences might be directed at claims *about* objective probabilities, but those credences won’t themselves be *reflections* of any objective probabilities. Rather those credences will simply express the confidence that scientists have in the relevant hypotheses, given their background knowledge and the existing evidence.

The objection mounted against Jeffrey was that, even if credences of the form *P is probable to such-and-such a degree* do not commit categorically to non-probabilistic facts, they at least involve a categorical judgement about a probability, and this in itself is enough to reintroduce all the issues of inductive risk. We can now see what is wrong with this objection. Someone who expresses a certain degree of belief in a hypothesis isn’t thereby making a categorical claim about some objective probability. There are just expressing their confidence in the hypothesis, and that’s it.

The point is clearest when the hypothesis at issue isn’t itself about objective probabilities. If I say I am 80% confident that *mad cow disease is caused by prions*, I am not ascribing a probability to anything, categorically or otherwise. The sole proposition in play is the non-

probabilistic *mad cow disease is caused by prions*, and the 80% only characterises my attitude *to* that proposition, not any probabilistic aspect *in* my subject matter. The same point applies even when we're dealing with a hypothesis about objective probabilities. I am 80% confident that M—the objective probability of lifetime breast cancer for women with the BRCA1 gene is over 0.5. Again, the 80% only characterises my attitude *to* M. I don't make any judgement, categorical or otherwise, *about* anything being 80% probable.¹²

In short, when we say scientists should express credences, not make categorical judgements, we don't mean that scientists should make categorical judgements about probabilities. Rather they should simply express their confidence that P is true, grounded in their background knowledge and their evidence. Nothing in that involves any categorical higher-order judgements about anything being probable.¹³

10. The Problem of the Priors

The last section argued that we should understand scientists' credences in hypotheses, not as judgements about objective probabilities, but simply as expressions of their informed expectations that the hypotheses in question are true. However, this now returns us to the "problem of the priors" mentioned earlier. For Bayesians, rational credences are a function of the empirical evidence and prior credences. A rational subject will start with some initial degree of belief in a hypothesis H. When evidence E is found, they will increase (or decrease) their degree of belief in H to the extent H implied a higher (lower) objective probability for E than the alternative hypotheses. While this updating procedure hinges in part on objective probabilities, it also depends on the prior degree of belief the subject starts with. And it is unclear what basis this prior credence could have, beyond the subject's initial intuition, hunches or prejudices.

On the face of things, this seems a serious worry about Jeffrey's Bayesian line. Do we really want to stop scientists making definite statements and instead restrict them to statements like "Given my evidentially ungrounded initial assessment of H's credibility, plus the evidence that has now come in, I'd say that the odds in favour of H are 75%"? Not only would the initial credences seem to lack any support, but they are also likely to vary between scientists, with the result that the scientific community will end up sending mixed messages in cases where clear guidance is needed.

Still, what is the alternative? When the evidence leaves room for reasonable experts to disagree, is it really better for them to speak with one categorical voice, just because that

¹² Sometimes credences about objective probabilities will be expressed in the form of "confidence intervals". For example, we might have a 95% degree of belief that the mean of the underlying probability distribution lies within a certain distance of some observed statistic. (Note, however, that if confidence intervals are so understood as expressing degrees of belief, they too will depend on prior subjective probabilities, and cannot, as is widely but erroneously supposed, simply be read off from the statistic and the knowledge that, say, it is normally distributed.)

¹³ We find it surprising that this response to the higher-order objection to Jeffrey is not more widely accepted. Perhaps part of the reason is that Jeffrey's own comments about the issue at the end of his original paper are unclear and confusing. We would attribute this to his maintaining, along with other Bayesians of the time, that objective probabilities should be entirely eliminated in favour of subjective probabilities. (We would like to thank Richard Bradley for advice about Jeffrey's early views.)

opinion supports some policy that they or some favoured section of the population would choose?

Fortunately, this dilemma is not inescapable. In our view, the so-called “problem of the priors” is overblown. Even if scientists were to turn away from Runderian dogmatism, and embrace the Bayesian alternative, there is no reason why their expression of credences should not command general respect.

An initial point is that the prior credences of scientists won’t just be random guesses, but informed by their education and experience, even if not by newly generated empirical evidence. Worries about prior credences are often amplified by citing archaic medical procedures that rested on nothing but the folklore of senior physicians. Traditional medical science, however, is arguably a poor basis for condemning all prior scientific judgement. As a general rule, educated scientists in empirically grounded disciplines will have a reasonable sense of which hypotheses are likely to be true, and of which should be discounted until they are backed up by substantial supporting evidence.

A second point relates to the way credences change once empirical evidence is uncovered. New evidence will tend to wash out differences in priors and push credences in the same direction. For example, when Barry Marshall and Robin Warren first proposed that bacteria were largely responsible for stomach ulcers, most medical scientists thought the idea absurd (Marshall and Warren 1983). However this scepticism was soon enough eliminated as more evidence came in. In truth, there’s no reason why the Bayesian judgements of informed scientists shouldn’t display reassuring conformity even in cases where they start off with significantly divergent priors.

In particular, this consensus will often amount to de facto certainty. Back in the 1950s, at the time of Rudner’s original article, the implicit presupposition was that a claim was only firmly established if it was *logically entailed* by the evidence. Anything short of that meant that the issue remained underdetermined. This is the rationale for Rudner’s insistence on unavoidable “inductive risk”. Categorical commitment always incurs a real chance of error. More recent philosophy of science, however, has recognized that the demand for logically compelling evidence sets the bar far too high. Yes, the empirical evidence is never logically compelling, but it is often sufficient to rule out all but the most fanciful alternatives. And in those cases even Bayesians will agree that scientists can issue unqualified endorsement of hypotheses. One can imagine fanciful alternatives to the claims that *COVID-19 is caused by a virus*, that *smoking causes cancer*, or indeed that *bacteria affect stomach ulcers*—perhaps an evil demon wants to deceive us, perhaps an alliance of world governments is conspiring against us—but, if these are the only alternatives consistent with the evidence, Bayesians will view the probability of the orthodox theories as indefinitely close to one, and expect scientists to endorse them accordingly.

So the “problem of the priors” does not inevitably condemn Bayesian scientists to destructive disagreement. For a start, they will be able fully to endorse scientific claims in those many cases where the evidence leaves no real room for doubt. Second, even when the evidence is less than conclusive, it will tend to push their degrees of belief in the same direction. And, finally, even when the evidence has not yet produced consensus, the

credences of scientists will generally express educated judgements rather than arbitrary guesses.

11. Against Categorical Beliefs

In our view, as we have said, the whole debate about inductive risk starts in the wrong place. The real issue is not the best way to arrive at categorical beliefs in the face of such risks, but whether we need categorical beliefs at all. All the problems of inductive risk only arise if we aspire to categorical beliefs.

So what would go wrong if we didn't bother with categorical beliefs at all and dealt only in credences? While a great deal of mainstream epistemology is concerned with the formation of categorical beliefs, it contains surprisingly little discussion of the need for them.¹⁴

One thought commonly aired, though not necessarily developed in detail, is that probabilistic calculations quickly become too complicated for ordinary humans, and so we need categorical beliefs to render our everyday decision-making tractable (see for example Sturgeon 2020).

Maybe so. Still, devices designed to simplify decision-making would seem inappropriate for scientifically informed policy making, even if they are necessary for ordinary citizens in everyday contexts. When we are dealing with serious issues which do not need to be decided in a hurry, what reason is there not to figure out which options will really maximise expected utility given the evidentially supported credences? As a rule we can expect policy makers to be perfectly capable of doing calculations and weighing up costs and benefits properly. Communicating scientific results via categorical beliefs rather than credences will only create room for policy makers to arrive at sub-optimal decisions. Maybe the cost of categorical beliefs is inevitable in many everyday contexts, but that is no argument for importing them into serious policy-based decision-making.

A rather different rationale for categorical beliefs is offered by our King's College London colleague David Owens. He observes that, given the way humans are psychologically constituted, we can't engage emotionally with the world without forming categorical beliefs. You can't *blame* somebody for their bad behaviour, or be *angry* with them, or *pleased* at their success, or *happy* because Spurs won . . . if you only think it's *likely* these things happened. We need fully to believe before we can react emotionally. And since our emotional engagement with the world is central to our lives, argues Owens, we cannot do without categorical beliefs (Owens 2017).

We do not want to take issue with Owens' argument. But, once more, we do not regard it as relevant to serious decision-making. Let us grant that we need categorical beliefs to form emotional reactions. That is no argument for thinking that policy makers need them when evaluating the pros and cons of different options. If optimal decisions are best arrived at on

¹⁴ Within mainstream epistemology, the debate about "pragmatic encroachment" covers many of the points we have been discussing. Unfortunately it would take us too far afield to engage with this debate here.

the basis of informed credences, basing them on categorical beliefs instead can only degrade decision-making.¹⁵

12. Racing to the Bottom

In much of the literature on Rudner's proposal, the dangers of *manipulation* and *wilful misunderstanding* are standardly cited as reasons for scientists making categorical pronouncements rather than communicating credences. Rudner's supporters argue that vested interests and anti-science factions will latch onto any hint of scientific uncertainty or doubt to undermine the scientific case for sound policies. They will declare that the evidence is equivocal, that there is no scientific proof of insecticide harms, vaccine safety, global warming, and so on (Douglas 2000 2009, John, 2015a).

The solution, according to Kevin Elliott (2011) is for scientists to adopt a "no passing the buck principle". It's no good their passing on their credences to the politicians and policy makers and then leaving it to them to make rational decisions. Their failure to make categorical pronouncements will only be exploited by bad actors. Instead the scientists should grasp the nettle themselves, identify the right course of action, and promulgate as definitively established those claims that support it.¹⁶

Perhaps this is indeed the pass we have come to in the modern world of fake news and social media bubbles. Maybe there is no alternative to pre-empting debate in the way Elliot urges. If that is so, however, a great deal will have been lost. Let us conclude by pointing out the downsides of Elliot's recommendation, in order to emphasise the importance of seeking another way.

We will do well to remember that the bad actors in our story are not necessarily evil. In the first instance, all that distinguishes them from the consensus of right-thinking scientists is that they have different priorities. They might place more weight on their immediate household incomes than the long-term future of the planet. Or perhaps they prioritise generation of profit over environmental impact. In themselves these divergences are not pathological. They are the kinds of differences in values and interests that we can expect among reasonable people, and there seems no immediate reason why democratic processes should not be able to accommodate and resolve them.

Rudner's supporters might respond that it is not these diverging priorities themselves that are now being invoked to justify categorical scientific pronouncements. Rather it is the consequent sin of manipulating the facts, of seizing on uncertainty to make it seem as if there is no evidence that difficult measures are needed. We entirely agree that this is a sin

¹⁵ This point perhaps deserves further discussion. What if the options at issue themselves involve *acting out of an emotion*? For example, Clayton Littlejohn (2017) has argued that criminal courts cannot convict without being in a position to *blame*, and so in effect face a choice between freeing the defendant and *convicting-on-categorical-belief*. Our response is that in such cases categorical beliefs will feature among the *ends* aimed at, but that the choice of *means* to those ends should still be guided by rational credences. Littlejohn (personal communication) does not disagree.

¹⁶ Douglas similarly holds that scientists have a moral responsibility to promulgate value-based categorical judgements (Douglas, 2000 563).

that should be condemned. Democratic processes will not be able to find the right way of resolving conflicting priorities if they are beset by misinformation about the relevant facts.

Still, is the solution for the scientists to get their manipulation in first?¹⁷ To our mind, the cure seems just as bad as the problem. At bottom, Elliot's "no passing the buck" principle responds to the threat of fake news by urging we manipulate the facts in the opposite direction. Rather than honestly communicating the evidential situation, Elliot and other Rudnerians want scientists to present partially evidenced claims as conclusively proven, whenever that will support the policies that right-thinking citizens favour.

The real danger with this strategy is that it is likely to set us off on a race to the bottom. Perhaps it will have commendable upshots in particular cases, ensuring that we arrive at policies that would be democratically favoured in the absence of propaganda from powerful bad actors. Still, it threatens a deeper cost. If the Rudnerian strategy comes to be adopted as standard practice, it is all too likely to undermine trust in science. Those who are unsure of the scientists' preferred policies will have every reason to suspect that the scientists are overegging the evidence.

It has been put to us that this kind of mistrust would be misplaced if the scientists were guided, not by their own values, but by those of the democratic majority, in the way suggested earlier. But this misses the point. The whole rationale for "not passing the buck" is to outmanoeuvre those who do not share the majority's priorities. Evidential uncertainties are to be concealed from them, in the interests of bolstering the case for favoured policies. They are to be told that things are beyond doubt, when they aren't. Once this becomes standard practice, the danger is that the intended audience will turn against science. Why should they credit official scientific pronouncements, if they are designed precisely to outwit them? "We've had enough of experts" would seem an entirely rational response.

Democratic decisions need to be informed by accurate scientific information, including information about partial credences. Policy makers can then take into account differences in values and interests and use the science to see how best to resolve these. But our institutions will struggle to do this if all sides are wielding propaganda to support their preferred policies. In the end, allowing evaluations to shape scientific decisions only threatens to devalue the coin of science.

Bibliography

Alexandrova, A. 2018 "Can the Science of Well-Being be Objective?" *British Journal for the Philosophy of Science* 69: 421-445

Anderson, E. 2002 "Situated Knowledge and the Interplay of Value Judgments and Evidence in Scientific Inquiry" in Gärdenfors, P., Kijania-Placek, K. and Woléński, J. eds *In the Scope of Logic, Methodology and Philosophy of Science vol 2* 497–517. Springer

¹⁷ "Let's get our retaliation in first" urged captain Willie John McBride in his team talk before the first test of the 1974 Lions rugby tour of South Africa.

- Baker, M. (2016) "1,500 Scientists Lift the Lid on Reproducibility" *Nature* 26: 452–4.
- Betz, G. (2013) "In Defence of the Value Free Ideal" *European Journal for Philosophy of Science* 3: 207–20.
- Bird, A. (2021) "Understanding the Replication Crisis as a Base Rate Fallacy" *British Journal for the Philosophy of Science* 72: 965–93.
- Douglas, H. (2000) "Inductive Risk and Values in Science" *Philosophy of Science* 67: 559–79.
- Douglas, H. (2009). *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh Press.
- Elliott, K. (2011). *Is a Little Pollution Good for You?* Oxford University Press.
- Elliott, K. & Richards, T. (2017) "Introduction" in Elliott, K. and Richards, T. eds *Exploring Inductive Risk: Case Studies of Values in Science*. Oxford University Press.
- Frank, D. M. (2017) "Making Uncertainties Explicit: The Jeffreyan Value-Free Ideal and Its Limits" in Elliott, K. and Richards, T. eds *Exploring Inductive Risk: Case Studies of Values in Science*. Oxford University Press.
- Hempel, C. (1965) "Science and Human Values" in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science* 81–96. The Free Press.
- Jeffrey, R. C. (1956) "Valuation and Acceptance of Scientific Hypotheses" *Philosophy of Science* 23: 237–46.
- John, S. (2015a) "Inductive Risk and the Contexts of Communication" *Synthese* 192: 79–96.
- John, S. (2015b) "The Example of the IPCC does not Vindicate the Value Free Ideal: a Reply to Gregor Betz" *European Journal for Philosophy of Science* 5: 1–13.
- Kirchin, S. ed 2013 *Thick Concepts* Oxford University Press.
- Kitcher, P. (2001) *Science, Truth, and Democracy*. Oxford University Press.
- Levi, I. (1967). *Gambling with Truth*. Knopf.
- Lewis, D. (1980) "A Subjectivist's Guide to Objective Chance" in Jeffrey, R. ed *Studies in Inductive Logic and Probability Volume II* 263–93. University of California Press.
- Littlejohn, C. (2020) "Truth, Knowledge, and the Standard of Proof in Criminal Law" *Synthese* 197: 5253–86.
- Marshall, B. and Warren, J. (1983) "Unidentified Curved Bacilli on Gastric Epithelium in Active Chronic Gastritis" *The Lancet* 321: 1273–5.

Open Science Collaboration. (2015) "Estimating the Reproducibility of Psychological Science" *Science* 349(6251).

Owens, D. (2017) *Normativity and Control*. Oxford University Press.

Papineau, D. (2018) Thomas Bayes and the Crisis in Science. *Times Literary Supplement website*.

<https://www.davidpapineau.co.uk/uploads/1/8/5/5/18551740/bayestlspapineau.pdf>

Rudner, R. (1953) "The Scientist Qua Scientist Makes Value Judgements" *Philosophy of Science* 20: 1-6

Steele, K. (2012) "The Scientist qua Policy Advisor makes Value Judgments" *Philosophy of Science* 79: 893–904.

Sturgeon, S. (2020) *The Rational Mind*. Oxford University Press.

Wilholt, T. (2009) "Bias and Values in Scientific Research" *Studies in History and Philosophy of Science Part A* 40: 92–101

Winsberg, E. (2012) "Values and Uncertainties in the Predictions of Global Climate Models" *Kennedy Institute of Ethics Journal* 22: 111–137.