Uncertainty Attitudes as Values in Science

Joe Roussos

April 6, 2022

There is now a large literature on values in science, discussing whether and how science can be objective while realistically acknowledging and managing the impact of values in the production of scientific information. In this paper, I am concerned with what counts as a value in this literature. I argue that we ought to consider scientists' attitudes to uncertainty as values. I will be concerned with inductive risk, and the claim I make is a conditional one: if you are concerned about inductive risk in a particular part of science, then that concern should include uncertainty attitudes alongside the more commonly considered moral, social, or political values.[1]

Here is a decision-focussed presentation of the argument from inductive risk (IR), and some major entries in the debate about it. The argument goes like this: Science involves decisions—about questions, methods, analyses, representations, uncertainty management, and whether to accept hypotheses, to name but a few. Decisions are a function of beliefs and desires, or to use language more familiar in this context, evidence

[1]For brevity, I will simply refer to "moral values", meaning that term to encompass the wide range of values discussed in this literature.

and values. The value-free idealist hopes that these decisions can be made using only harmless epistemic and cognitive values. Rudner (1953) argues that the decision to accept or reject a hypothesis must involve moral evaluations of the badness of potential errors. Jeffrey (1956) argues that there is not one set of consequences for the scientist to consider but many, one for each application of the hypothesis. Thus, scientists should avoid acceptance decisions and merely report their probabilities for the hypotheses in question. Douglas (2009) responds that scientists face crucial decision points prior to the formation of the probabilities which guide the acceptance decision, so that Jeffrey's response is insufficient to remove values from science.

I focus on decisions because the core move of this essay is to note that decisions in fact depend on more than just beliefs and desires, they also depend on the decision maker's attitudes to uncertainty. The most familiar attitude to uncertainty is risk aversion, which I will discuss below alongside its cousin ambiguity aversion. I argue that these are evaluative attitudes, that they are not plausibly epistemic or cognitive values, and that they raise the same worries that motivate for the value-free ideal.

I will be concerned with two motivations for value-free science. The first reason is that moral values are thought to interfere with the pursuit of the core epistemic and cognitive goals of science: true theories, offering explanations and understanding of the world. If scientists are concerned with moral values, their scientific products might deviate from the truth. One way that this concern is presented is in terms of "wishful thinking": a scientist whose concern for wellbeing informs their choice of methods, or acceptance of theories, might come to hold beliefs based on what they desire the world to be like rather than based on what it is really like. This argument has been associated with very strong forms of value-freedom, such as a complete absence of moral evaluations in the core practice of science.

The second reason is democratic. In order to be democratic, decisions must respond appropriately to the values of the people (represented by policymakers, political decision makers, etc.). Scientific input is required in policy, as we want our policies to respond to how the world really is. If science is value-laden then it "bakes in" some value judgements which are not those of the people, and these can influence which policy decisions are made. This subverts democratic control, as scientists' values dilute or supplant the values of the people. The result is what is sometimes called "liberal epistemic division of labour"—a picture of science-based policy according to which scientists provide the facts, and politicians provide the values (Brown, 2009). Here, the relevant ideal is freedom from values which would interfere with democracy.

So, my main claim is that, insofar as moral values are a problem for these reasons, so are attitudes to uncertainty. Or, if these are not problems per se but rather something to be managed properly, then so too are uncertainty attitudes.

1 Values in the inductive risk literature

The values that are discussed in IR debates have a particular form. They take as their objects concrete outcomes—states of the world described without reference to probabilities. This is contrast with attitudes to uncertainty, which take as their object the state of uncertainty under which a decision is made.

Allow me to illustrate. In Rudner's (1953, 2) classic article, he argues that scientists "must make the decision that the evidence is sufficiently strong or that the probability is sufficiently high to warrant the acceptance of the hypothesis" under consideration. He presents examples which illustrate that this judgement of sufficiency is a moral matter. The first example contrasts a case in which "the hypothesis under consideration were to the effect that a toxic ingredient of a drug was not present in lethal quantity" with one where the "hypothesis stated that, on the basis of a sample, a certain lot of machine stamped belt buckles was not defective". He notes that the former would require higher confidence, in virtue of the high ethical stakes. We are, I take it, invited to imagine that the scientists must consider the harms to innocent medicine takers, weighing the badness of their illness or death against the potential benefits of treating the disease. This is in contrast to hapless customers whose trousers don't quite stay up.

Or, to take a more recent example, consider Winsberg, Oreskes and Lloyd's (2020) discussion of the science of extreme weather event attribution. Here, scientists attempt to determine whether a catastrophic event like Hurricane Harvey was due to or made worse by climate change. There is an ongoing methodological debate in this field between the so-called storyline approach and the more dominant fraction of attributable risk approach. Defenders of the storyline approach are "concerned that the risk-based approach will falsely fail to attribute the extreme event to climate change. [They are] concerned that this approach has a propensity to underestimate harm." Harm is here a straightforwardly moral matter: harm to people and society due to climate change and extreme weather events. The risk-based folks are "concerned about the risk of overstatement of human effects... [i.e.,] about making too many false positive errors, or overstating the role of climate change." These concerns are practical: "time and money might be spent preparing for events that will not occur" (Winsberg et al., 2020, 145–46).

The object of evaluation in each case is a state of the world which occurs after the decision and which are described without reference to the probabilities governing the decision. They are possibilities described in terms of the morally relevant facts, such as the harms of extreme weather events or the damage to the reputation of the scientists.

To make this more precise, let me present IR formally in terms of a stylised example of a decision to accept an hypothesis. (The presentation is Bayesian, so I use "accept"/"reject" rather than the more familiar frequentist framing in terms of rejecting a null hypothesis.) Table 1 displays such an acceptance decision. Acceptance decisions are often framed in terms of the threshold of probability θ which is required to accept the hypothesis. In this framing, the object of interest is the value of θ, the idea being that only when $P(H) > \theta$ should the scientist accept H. The argument from IR says, roughly, that θ depends on (non-epistemic) evaluations of the consequences of that decision—represented by the utilities of outcomes in the decision table. E.g., $u(FP)$ represents the value of a false positive.

Table 1: Acceptance decision

	$P(H)$	$P(\neg H)$
Accept H	$u(TP)$	$u(FP)$
Reject H	$u(FN)$	$u(TN)$

I want to present the situation a little differently. Suppose that we have two scientists, who face a similar decision about accepting hypothesis H. They have the same evidence, which they have evaluated identically, and thus they assign the same probability to H, say 0.6. They are faced with the decision to accept H, following Douglas (2009), let us suppose that they have first applied their epistemic and cognitive values. Nonetheless, a gap remains, which must be bridged by their evaluations of the consequences of error. Here, they differ; let us imagine that Scientist 1 takes a false positive to be neutral while Scientist 2 takes it to be moderately bad. The result is that the first scientist accepts hypothesis H and the second rejects it. Tables 2 and 3 display the decisions faced by each scientist and we can read off the third column that Scientist 1 accepts H while Scientist 2 rejects H.

The decision is guided by expected utility considerations: Scientist 1's expected utility of accepting is higher than that of rejecting. The point of the example decision tables is simply to make clear the nature of the evaluations and the they are playing. The evaluation is of a consequence—a fully-specified state of the world which results if H is true (false) and if the scientist accepts (rejects). (This, not coincidentally, is how Jeffrey (1956) describes the consequences of choices.)

Table 2: Scientist 1's acceptance decision

	$P(H)$=0.6	$P(\neg H)$=0.4	
Accept H	2	0	$EU(A) = 1.2$
Reject H	1	1	$EU(R) = 1$

Table 3: Scientist 2's acceptance decision

	$P(H)$=0.6	$P(\neg H)$=0.4	
Accept H	2	-1	$EU(A) = 0.8$
Reject H	1	1	$EU(R) = 1$

2 Uncertainty attitudes and scientific decisions

The decision theoretic underpinnings of the above example are highly idealised. In particular, they ignore the fact that many actual agents have attitudes towards uncertainty itself. In simple terms, an uncertainty attitude is a liking of or aversion to uncertainty itself. An aversion to uncertainty manifests as a preference for making decisions with less uncertainty over making decisions with more uncertainty; more subtly, such attitudes are measured via willingness to trade material consequences in exchange for a reduction of the uncertainty associated with making a decision. As I use the term here, "uncertainty attitude" is an umbrella category, which encompasses risk attitudes and ambiguity attitudes.

My argument for taking them seriously in the values and science debate is very simple: Many people have these uncertainty attitudes, and we should expect the same to be true of scientists.[2] These attitudes make a difference to the decisions people make—as I demonstrate below. They are plausibly rational.[3] So, a rational theory of scientific inference should account for them.

I begin with attitudes to risk. A decision-maker who is risk averse will prefer to receive 100 euros for sure than to place a wager which has 50% chance of paying 0 and 50% chance of paying 200. These bets have the same expected value in euros, so the decision maker's preference for the sure 100 is explained by their distaste for risk when it comes to getting euros. "Risk" here means the kind of uncertainty present in the wager: there are widely spread out outcomes which occur with known probabilities. The simple orthodox decision theory that I used above, expected utility theory, has room for only a limited kind of attitude to uncertainty. This is a form of risk aversion which can be captured in the shape of the agent's utility function: agents who are risk averse in euros

[2]For empirical evidence and discussion of trends, see (Di Mauro and Maffioletti, 2004) and (Trautmann and van de Kuilen, 2015)

[3]As evidenced by the normative models of rational choice which account for them, e.g., (Buchak, 2013) and (Stefánsson and Bradley, 2019) for risk, and (Gilboa and Schmeidler, 1989) and (Bradley, 2017) for ambiguity

are described with utility functions that are concave in euros. This means that they get more utility from the first 100 euros than they do from the second 100.

This way of representing risk averse behaviour can capture some real behaviour amongst people (and, I presume, scientists). But it is severely limited. The use of concave utility functions conflates two psychologically distinct phenomena: decreasing marginal utility of a good and an aversion to risk (Stefánsson and Bradley, 2019). Expected utility theory also has no room for agents who are risk averse in utility itself—who would prefer a sure 100 units of the good over a 50-50 gamble of 0 and 200 units. Nevertheless, this is both possible and plausibly rational. In fact, there is significant evidence that agents have risk attitudes which cannot be represented via concave utility functions; indeed, such agents cannot be represented as maximising expected utility at all.[4] It is this kind of non-EU attitude that I refer to as "risk aversion" in the remainder of this essay.

For non-EU agents with uncertainty attitudes, these attitudes are an additional ingredient in their decision making, beyond their beliefs and desires. Agents who are risk averse are not merely responding to the expected value of their decisions, they are also responding to the fact that the outcomes are distributed a certain way, and that the component outcomes have the specific probabilities they do. Where the risk neutral expected utility maximiser regards all ways of getting an expected 100 utils as equivalent, the risk sensitive decision maker regards the sure 100 as different from the 50–50 bet on 0 and 200, and as different from the 1/3–2/3 bet on 0 and 150, and so on. Philosophers have defended this as a rational pattern of preference an offered models of such choices, e.g., the risk-weighted expected utility theory (REU) developed by Buchak (2013).

[4]See fn. 2.

Buchak's model contains a risk preference function r for each agent. Risk averse agents have convex risk functions, e.g., $r(p) = p^2$, while risk neutral agents have $r(p) = p$. This function modifies the probabilities, before they are combined with the agent's utilities in an expected value-type calculation. The details of the calculations don't matter; what does is that agents with identical probabilities and utilities can reach different decisions because they have different risk attitudes, represented by different r functions.

In canonical examples of IR, we suppose that scientists have some probabilities to hand and that they are deciding whether to accept a hypothesis on the basis of them—as in Table 2. This is a straightforward example of a decision under risk. We then say that the scientist makes moral evaluations of the outcomes of error. Scientists, like other people, are plausibly risk averse when it comes to decisions about such values. So, imagine two scientists each confronted with the decision in Table 2, and suppose that they have the same evidence, same priors, and have arrived at the same probability for H. Suppose also that they make the same IR assessments: they identify all of the same consequences, and evaluate them identically. They agree, in other words, on exactly how bad each kind of error would be, and on how important true positives and true negatives are. It is nevertheless possible for these two scientists to make different inferences about H, because they differ in their taste for risk. So, since some scientific decisions involve risk, and since we should expect at least some scientists to have non-EU risk attitudes, we should expect these attitudes to be influencing the scientific information they produce.

The same thing can happen for attitudes to ambiguity. Ambiguity refers to a kind of uncertainty where we lack the information required to estimate probabilities precisely. Imagine a game based on drawing balls from an urn. Urn 1 has 50 red and 50 black. If it is black, there is no payment. Now consider Urn 2, which also has 100 balls in some unknown combination of red and black. A ball is drawn at random from the urn, and if it is red then the player is paid out 200. An agent is offered a choice: they can take draw from Urn 1 or from Urn 2. Knowing nothing about the distribution of balls in Urn 2, the decision-maker might as well take them to be equally likely, but importantly they don't know that this is the case—unlike with Urn 1. A decision-maker who is *ambiguity averse* will prefer the bet on Urn 1, with known 50–50 odds, to a similar bet on Urn 2, with unknown odds.

Situations like this are sometimes represented with imprecise probabilities, e.g., a range of probabilities for H like 0.7–0.9. Decisions under ambiguity are controversial and, since most work on ambiguity has been done by economists and descriptive decision theorists, there is less in the way of normative theorising here. Nonetheless, philosophers like Bradley (2017) have defended the rational permissibility of attitudes to ambiguity and offered decision rules for different degrees of ambiguity aversion and ambiguity seeking. An example is the Alpha Maximin rule, according to which decision makers value options at a mixture of the worst and best expected utility, relative to the range of probabilities they entertain. (Alpha is the parameter controlling how much weight the worst-case EU gets.) Again, the precise details don't matter. What does is that with identical sets of probabilities and identical utilities can reach different decisions because they have different ambiguity attitudes, represented by different α parameters.

Many actual scientific decisions are involve ambiguity; i.e., they are more like betting on Urn 2 than Urn 1. For example, recall Douglas' (2000, 571) case of of scientists characterising unclear evidence. There, toxicologists examined rat liver slides under microscopes, looked for abnormalities, and classified them as benign or malignant. This is a skilful judgement, relying on experience and tacit knowledge. The classification of borderline cases is, Douglas argues, subject to IR. But it is not best characterised as a "decision under risk" in the technical sense described above since there no clear probabilities in play.

Scientists, like other people, plausibly have a range of attitudes to ambiguity. Consider a pair of scientists who face the same choice of whether to accept H. Let us suppose that there are different lines of evidence, perhaps gathered through different methods. Using each line of evidence, the scientists can assess the probability of H. Suppose again that their epistemic assessments of the evidence are identical. However, they don't know how the combine the different lines of evidence. So, they represent their uncertainty about H using the set of probabilities supported by the evidence. This is a form of imprecise probability, a common framework for representing ambiguity which is less demanding than the precise probabilistic framework used above. We again add the constraint that they make the same IR assessments: they identify all of the same consequences, and evaluate them identically. Nevertheless, one might accept H and the other reject it, because they differ in the attitudes to ambiguity. So, since some scientific decisions involve ambiguity, and since we should expect at least some scientists to have non-EU ambiguity attitudes, we should expect these attitudes to be influencing the scientific information they produce.

2.1 Is this new?

At this point, the reader might wonder whether this is really a new observation. Isn't the whole discussion about inductive *risks*, and attitudes thereto? Here it is worth noting that the English word "risk" is ambiguous between several meanings. It can mean an unwanted event, as in “the risk of getting cancer”. It can mean the probabilities of unwanted events, as in “the risk that a smoker’s life is shortened is 50%.” In risk analysis, it almost always means the expectation values of unwanted events, as in “the risk of smoking is 12.5 life-years lost on average.” Or, it can refer to spreads of outcomes over possibilities, as when I described the sure 100 euros as less risky than the 50–50 gamble on 0 or 200 euros. The IR discussion focusses on the value of the state of the world in which one has made the inductive error. This is a use of risk in the sense of bad outcome. But as we have seen, that is not what matters for non-EU uncertainty sensitive agents.

There are, to be sure, scattered references to risk aversion in the literature on values and science, for example the single unexplained mention in the SEP article on Objectivity by Reiss and Sprenger (2017). But it is not a central topic. It appears nowhere in Elliott and Richards (2017) anthology on inductive risk, not even in the wide-ranging typology of risk by Biddle et al. (2017). Nor does it appear in Biddle’s earlier (2013) “state of the field” review. To some extent this is not surprising, as much of the action in the philosophy of risk and ambiguity attitudes is recent. When uncertainty attitudes appear, they do so in the context of recent discussions of science-driven policy decisions. For example, Bradley and Steele (2015) and Winsberg (2018) discuss the importance of the permissibility of uncertainty attitudes when discussing how policymakers can respond to the scientific uncertainty in climate science. They do not make the move I am making here, of recognising that those very scientists make IR decisions which themselves could permissibly have been affected by their attitudes to uncertainty.

3 Uncertainty attitudes matter

Now, drawing on Douglas (2000) and Winsberg (2018) I suggest that we see scientific information as the result of a series of decisions, possibly made by different scientists, with multiple points for IR-considerations to enter. We have just seen that uncertainty attitudes can alter how scientists make IR-sensitive decisions. Since any scientific product is the result of a sequence of such decisions, uncertainty attitudes could have a significant cumulative effect, leading scientists (or groups of scientists) who differ only in their attitudes to uncertainty to arrive at quite different conclusions at the end of the sequence.

So, they make a difference. Does this difference matter? One way we can think about this is to consider directly what kind of evaluation is happening when scientist's attitude to uncertainty affects their decisions. On an illuminating account due to Bradley and Stefánsson, uncertainty attitudes involve conative attitudes towards chances (Stefánsson and Bradley, 2015; Bradley, 2016; Stefánsson and Bradley, 2019). For example, risk seeking agents enjoy *facing risks*: for an experience mountain climber "there is an optimal region of risk, where the chances of death or injury are high enough to require courage of the climber but not so high as to make the activity foolish" (Stefánsson and Bradley, 2015, 605). It is not merely the concrete outcomes which are evaluated, but the chances of those outcomes. To take a more morally loaded example, one might think that scarce medical resources should be distributed by lottery because there is *value in having a chance* of receiving the treatment, even if one doesn't get it in the end. Clearly, evaluating chances-of-outcomes is related to an evaluation of the outcomes themselves—it is the harm of death that creates the thrill for the climber, and the benefit of life-saved that creates the value of the lottery. So, an uncertainty attitude in a decision situation depends on the underlying evaluation of the consequences, though it is not reducible to them (as the risk seeking mountain climber demonstrates). In other words, scientists' uncertainty attitudes are a distinct avenue through which their moral values affect their decisions.

Now, this is debated, so I will also offer two more direct considerations. First, I think there is a clear prima facie case that attitudes are a challenge to objectivity in the same way that moral values are. Consider the two cases side by side. Suppose that a scientist, Karim, wants to make use of hypothesis H in his work. To establish whether he should rely on H, he reads two scientific papers. Suppose that they emerged from the same lab and so the authors had the same prior empirical beliefs, collected exactly the same data, and evaluated it identically. Nevertheless, the authors reach different conclusions. Karim discovers that the difference is due to how they morally evaluated the badness of the potential inferential errors. To many, this dependency seems bad for Karim's project of learning facts about the world. Consider a variation of the story where the scientists also report, honestly let us suppose, identical moral values. They are upfront about their IR reasoning, and it is clear that they identified all the same outcomes and evaluated them identically. Nevertheless, the authors reach different conclusions. Karim investigates, and discovers that the difference is due to a difference in their risk attitudes: one is moderately risk averse, the other moderately risk seeking. If you are concerned by the first story, I find it hard to see how you can avoid being concerned by this variation.

A second concern is especially pressing for those whose preferred response to the value-ladenness of science is a form of democratic procedural solution in which stakeholder values are incorporated into science. The idea behind such approaches, as I understand them, is that scientists by eliciting and incorporating stakeholder values scientists come to act on behalf of the non-expert users of science, so that the science reflects what they would conclude if they were in the scientist's epistemic position. For example, discussing uncertainty management in climate adaptation work, Parker and Lusk (2019, 1647) write: "if choices must be made, they could be made in light of the [IR] preferences of the user or client: if it would be particularly bad for the user's purposes for uncertainty to be underestimated, then the provider might select those methodological options that will deliver a broader uncertainty estimate." The key point once again is that decisions aren't a only function of evidence and values. These are filtered through attitudes to uncertainty when a decision is made. So if one attempts to incorporate stakeholder values into science while neglecting stakeholder uncertainty attitudes, this will lead to decisions which still diverge from how stakeholders would make them if they were in the scientist's epistemic position.

4 Potential objections and replies

There are several ways one could respond to my claim that we should worry about uncertainty attitudes whenever we worry about inductive risk.

A first objection might be that there is an obvious solution: each attitude-type discussed comes with a "neutral" variant, risk neutrality and ambiguity neutrality. Surely these are the attitudes required for value-neutral science. However, this would be to be misled by the names. There is nothing "neutral" about risk neutrality, in the sense typically meant be value-free idealists. Risk neutrality is just one of many evaluative stances; the one according to which (100kr for sure) is equal in value to (0kr, 0.5 ; 200kr, 0.5). The strong value-free ideal is that scientists should make no moral evaluations whatsoever, and weaker forms distinguish particular kinds of problematic judgements, such as those which interfere with democracy. Clearly risk neutrality is not suitable for strong value-freedom. Now, one might argue that there is a risk attitude which does not interfere with democracy, and further argue that it is risk neutrality. But linguistic coincidence cannot supply that argument.

A second objection might be that uncertainty attitudes are harmless, as the values they encode are one of the kosher varieties, either cognitive or epistemic values. Let's begin with cognitive values. As I have used the term, these are properties of theories and models, which facilitate scientific cognition (following Douglas, 2009, 93-94). That is to say that they facilitate thinking with and understanding these scientific object, for agents with our cognitive capabilities. For example, simplicity "is a cognitive value because complex theories are more difficult to work with, and the full implications of complex theories are harder to unpack" (Douglas, 2009, 93). Uncertainty attitudes don't seem to fit the type: they aren't properties of theories or models, and achieving the preferred valence doesn't facilitate cognition. The uncertainty averse agent favours decision situations in which the harms and benefits are "clumped up" rather than "spread out" across the possibilities that they are aware of. But the benefit isn't cognitive. The risk averse agent has no trouble reasoning about spread out outcomes—indeed, they do so as part of their decision-making.

Perhaps then uncertainty attitudes are of the pure epistemic type. These, recall, constitute the truth-seeking mission of science. Whereas cognitive values can have little to do with truth seeking or truth preservation (Laudan, 2004), epistemic values are tightly focused on these aims. On this categorisation, the epistemic is a small category containing values such as predictive accuracy, internal consistency, and empirical adequacy. The initial plausibility of including uncertainty attitudes is therefore quite low. Uncertainty attitudes are towards decision situations and patterns of outcomes of choices, not to states of knowledge. We can see this most clearly in the famous proofs that risk averse agents, who appear to value certainty, can sometimes rationally turn down free information (Wakker, 1988; Buchak, 2010; Campbell-Moore and Salow, 2020).

But we can do better than a mere argument by elimination of alternatives, even if we don't accept Bradley and Stefánsson's account of the evaluative basis of uncertainty attitudes. There are independent reasons to think that uncertainty attitudes are a moral matter, which come to us from the ethics of social decision-making. There is now a small literature on whether the uncertainty attitudes of social decision makers matter to the ethics of their decisions. For example, in a recent manuscript Buchak (ms) argues that specific attitudes are morally required: social decision-makers should be risk averse but ambiguity neutral, unless they know the attitudes of every person on whose behalf they decide, in which case they should defer to those. Stefánsson (ming), without relying on his values-of-chances account, argues that social decision makers ought to be more risk seeking than individuals would be if they were making the decisions themselves. Rowe and Voorhoeve (2018) argue that ambiguity aversion is permissible for social decision makers, and that this fact supports a form of egalitarianism. Now recall that, on the democratic motivation for the value-free ideal, an important use of science is that it informs policy. Here, scientists are involved in decisions made on behalf of others, and so the above gives us an angle on the moral status of their attitudes. The conclusion I want to draw from this is that, at minimum, uncertainty attitudes matter to moral decision making. So, the uncertainty attitudes of scientists may interfere with democratic decision making, just as their moral attitudes might.

5 Conclusion

Where does this leave us? It might be that there is a a morally correct set of uncertainty attitudes for scientists to take. (They might even be the neutral attitudes.) Or there might be a set of attitudes which do not interfere with democracy. Or we could follow the turn in the values and science debate towards stakeholder engagement, and insist that science ought to be based on the "right values" in a procedural sense: values supplied to them as the output of a consultative procedure with some relevant group of users of the scientific outputs.

All of these are live options in the values and science debate. Moreover, they occur because of the ethical implications of uncertainty attitudes. This establishes my conditional claim: if you are concerned about inductive risk in a particular part of science, then that concern should include uncertainty attitudes alongside the more commonly considered moral, social, or political values.

References

Biddle, J. (2013, March). State of the field: Transient underdetermination and values in science. *Studies in History and Philosophy of Science Part A 44*(1), 124–133.

Biddle, J. B., Q. Kukla, K. C. Elliott, and T. Richards (2017). The geography of epistemic risk. In *Exploring Inductive Risk: Case Studies of Values in Science*. Oxford: Oxford University Press.

Bradley, R. (2016). Ellsberg's Paradox and the Value of Chances. *Economics and Philosophy 32*(02), 231–248.

Bradley, R. (2017). *Decision Theory with a Human Face.* Cambridge University Press.

Bradley, R. and K. Steele (2015). Making Climate Decisions. *Philosophy Compass 10*(11), 799–810.

Brown, M. B. (2009). *Science in Democracy: Expertise, Institutions, and Representation.* Cambridge, MA: MIT Press.

Buchak, L. (2010). Instrumental rationality, epistemic rationality, and evidence gathering. *Philosophical Perspectives 24*(1), 85–120.

Buchak, L. (2013). *Risk and Rationality.* Oxford University Press.

Buchak, L. (ms). Risk and Ambiguity in Ethical Decision Making.

Campbell-Moore, C. and B. Salow (2020, July). Avoiding Risk and Avoiding Evidence. *Australasian Journal of Philosophy 98*(3), 495–515.

Di Mauro, C. and A. Maffioletti (2004, March). Attitudes to risk and attitudes to uncertainty: Experimental evidence. *Applied Economics 36*(4), 357–372.

Douglas, H. (2000, December). Inductive Risk and Values in Science. *Philosophy of Science 67*(4), 559–579.

Douglas, H. (2009). *Science, Policy and the Value-Free Ideal.* Pittsburgh: University of Pittsburgh Press.

Elliott, K. and T. Richards (2017). *Exploring Inductive Risk: Case Studies of Values in Science*. Oxford: Oxford University Press.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.

Jeffrey, R. C. (1956). Valuation and Acceptance of Scientific Hypotheses. *Philosophy of Science 23*(3), 237–246.

Laudan, L. (2004). The Epistemic, the Cognitive, and the Social. In P. Machamer and G. Wolters (Eds.), *Science, Values, and Objectivity*, pp. 14–23. Pittsburgh: University of Pittsburgh Press.

Parker, W. S. and G. Lusk (2019). INCORPORATING USER VALUES INTO CLIMATE SERVICES. *American Meteorological Society 100*(9), 1643–1651.

Reiss, J. and J. Sprenger (2017). Scientific Objectivity. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Metaphysics Research Lab, Stanford University.

Rowe, T. and A. Voorhoeve (2018). Egalitarianism under Severe Uncertainty. *Philosophy and Public Affairs 46*(3), 239–68.

Rudner, R. (1953). The Scientist Qua Scientist Makes Value Judgments. *Philosophy of Science 20*(1), 1–6.

Stefánsson, H. O. (forthcoming). Longtermism and social risk taking. In Barrett, H. Greaves, and D. Thorstad (Eds.), *Essays on Longtermism*. Oxford: Oxford University Press.

Stefánsson, H. O. and R. Bradley (2015). How valuable are chances? *Philosophy of Science 82*, 602–625.

Stefánsson, H. O. and R. Bradley (2019). What is risk aversion? *The British Journal for the Philosophy of Science 70*, 77–102.

Trautmann, S. and G. van de Kuilen (2015). Ambiguity Attitudes. In G. Keren and G. Wu (Eds.), *The Wiley Blackwell Handbook of Judgment and Decision Making*, pp. 89–116. John Wiley & Sons.

Wakker, P. (1988). Nonexpected utility as aversion of information. *Journal of Behavioural Decisiom Making 1*, 169–175.

Winsberg, E. (2018). *Philosophy and Climate Science*. Cambridge: Cambridge University Press.

Winsberg, E., N. Oreskes, and E. A. Lloyd (2020). Severe weather event attribution: Why values won't go away. *Studies in History and Philosophy of Science Part A 84*, 142–149.