

Inductive risk and the contexts of communication

Keywords: Inductive risk; Values in science; Social epistemology; Neonicotinoid research; Public/private distinction; Communicative obligations

There have been precipitous declines in wild bee populations in many Northern Hemisphere countries (Stokstad 2012). We know that bees exposed in laboratory conditions to non-lethal quantities of neonicotinoids – a class of compounds used in insecticides – suffer from memory and navigation problems (Desneux, Decourtye and Delpuech 2007). Therefore, scientists have investigated whether wild bee populations might be negatively affected by neonicotinoid exposure. Two recent experimental studies claim to have shown such a link (Henry *et al* 2012; Whitehorn *et al* 2012). However, there are doubts over whether they provide high-quality evidence about real-world scenarios, for example because wild bee populations might be exposed to lower doses of insecticide than in experimental set-ups (Stokstad 2012). This debate is heated, because the political, economic and ecological stakes are high; arguably, a collapse of bee populations is worse than unnecessarily banning neonicotinoids – not only in ecological but in economic terms – but an unnecessary ban could seriously damage agricultural productivity.

On the one hand, it might seem that scientists should bear in mind the potential non-epistemic costs of failing to say that neonicotinoids are dangerous when in fact they are when deciding whether the evidence warrants that claim. Like many, I feel the pull of these concerns. In §1, I develop an account of scientists' communicative obligations, based on familiar arguments about “inductive risk”, which apparently justifies such an indirect role for non-epistemic values in scientific inference. On the other hand, allowing non-epistemic values to play this

role in scientific inference might seem problematic. Scientific research contributes to what Kitcher calls “public knowledge”, “that body of shared information on which people draw in pursuing their own ends” (Kitcher, 2011, p.85). Given that different people hold different values, a value-laden science may fail to contribute to “public” knowledge. I think this is a serious concern, which outweighs the considerations in favour of a value-laden science. Therefore, in §§2 and 3, draw on an unusual combination of Kant and Richard Jeffrey to argue that scientific inference aimed at *public* communication should not take account of non-epistemic concerns, thereby blunting the arguments in §1. §4 discusses how these arguments relate to scientists’ broader communicative obligations, including in neonicotinoid research, and to on-going debates over inductive risk and proper scientific inference. In conclusion I outline the broader implications of my arguments for understanding the “value free ideal” for science.

§1 Inductive risk and the Floating Standards Obligation

In 1953, Richard Rudner claimed that the scientist *qua* scientist “accepts or rejects hypotheses”, but no hypothesis is ever completely verified by the available evidence; therefore, decisions about acceptance must turn on whether the evidence is “sufficiently strong” (Rudner, 1953, p.2). More recently, Heather Douglas has set out a similar problem: all agents, including scientists, face choices about whether to make empirical claims which are not deductively implied by available evidence (Douglas, 2009, p.87). Both argue for a similar response to these problems. For Rudner, decisions about whether evidence is sufficiently strong are “a function of the *importance*, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis” (p.2, emphasis in original). Douglas argues that everyone, including scientists, has a moral responsibility to “consider the

consequences of error” (p.87) when making claims. Therefore, science is not value-free, in that “scientists should consider the potential social and ethical consequences of error in their work, they should weigh the importance of those consequences, and they should set burdens of proof accordingly” (p.87).

Rudner’s argument convinced many philosophers: for example, Hempel (1965) and Gaa (1977). More recently, following Douglas’s work, the “argument from inductive risk” has become commonplace, assumed in work by Kitcher (2011, 141-155) and Kukla (2012, 853-855) with discussions of its theoretical implications (Steel, 2010) and its practical implications (for “trust” in science (Wilholt, 2012) and model construction (Biddle and Winsberg, 2012)). Indeed, some now claim that her argument does not go far enough (Brown, forthcoming). In this paper, I will follow Rudner and Douglas in assuming that scientists face problems of “inductive risk”. I will, however, dispute their claims about how scientists must respond to these problems. To understand my proposals first it is necessary to clarify the problem, to present the strongest version of the Rudner/Douglas “floating standards” response, and to note just how radical that response is. These are the tasks of this section.

Inspired by Rudner and Douglas, I understand cases such as neonicotinoid research as follows. An expert (or group of experts) must decide whether or not to assert a claim – that neonicotinoids deplete wild bee populations – which is supported, but not deductively implied, by available evidence. In making this decision, she runs significant inductive risks (Hempel 1965): of a false positive – asserting a claim which is, in fact, false – and of a false negative – failing to assert a claim which is, in fact, true. Such a scientist requires (or can be seen as employing) an “epistemic standard for assertion”: i.e. a principle specifying how much evidence she should have in favour of the claim before asserting it. The “higher” the

standard – the more evidence required for warranted assertion – then the lower the risk of false positives, but the higher the risk of false negatives.¹ Which epistemic standards should scientists employ?

I follow Douglas in framing the problem of inductive risk in terms of assertion, rather than acceptance, for two reasons, related to Betz's insight that arguments from inductive risk are best understood in terms of what scientists morally, rather than logically, must do (Betz, 2013). First, a focus on "acceptance" snarls discussions of inductive risk in questions of whether cognitive attitudes should be sensitive to ethical considerations; assertion, by contrast, is clearly subject both to epistemic and ethical concerns. Second, a focus on assertion avoids a powerful response to inductive risk arguments. Many commentators, following Richard Jeffrey, agree that scientific practice involves establishing the degree of evidential support enjoyed by propositions, but deny that scientists do (or ought to) accept hypotheses outright, claiming that they (should) report degrees of evidential support (Jeffrey, 1956; Betz, 2013). Even if Jeffrey's response to Rudner is successful – which is controversial because scientists seem to face inductive risk problems in establishing evidence claims (Gaa 1977; Biddle and Winsberg 2010; Elliot, 2013) – it does not undermine ethical concerns about inductive risk and assertion. Claims like "given the evidence, it is extremely likely that neonicotinoids deplete bee populations" or "given the evidence, it is unclear that neonicotinoids deplete bee populations" do not go beyond the available evidence. However, we know that the former is likely to be heard as "neonicotinoids deplete bee populations" and the latter as "neonicotinoids do not deplete bee populations". The moral status of making a claim turns not only on what we say, but on how others (foreseeably) interpret what we say

¹ See John (2011, 502) for a slightly different account of "epistemic standards", which this account builds on.

(Saul, 2013). Understanding inductive risk in terms of assertion suggests that, workable or not, Jeffrey's proposal is of questionable moral significance.

Reframed in my terminology, Douglas proposes that when faced by problems of inductive risk, scientists should vary their epistemic standards for assertion in proportion to the expected consequences of different sorts of error. However, this proposal requires refinement, because not all consequences of assertions relate to the ethical status of those assertions in the same way.² For example, imagine that a government scientist knows that the claim “neonicotinoids deplete bee populations” is not currently well-enough warranted for policy-makers to act on it. However, she knows that if she reports this lack of evidence then industry interests will successfully twist her words to argue that neonicotinoids are clearly safe, a claim which could have disastrous consequences. This scientist is, undoubtedly, in a tricky position, but it seems strange to say that the fact that lobbyists will mendaciously twist her honest report to suit their own ends means that she should change her standards for reporting to policy-makers. Even if she is causally responsible for successful lobbying following her pronouncement, it would seem strange to say she is morally responsible. By contrast, it seems that she is not only causally but morally responsible for the consequences of her assertions which stem from her intended audience, policy makers, deferring to her claims.

A simple amendment to Douglas's proposal can, however, avoid these concerns. Often, scientists must set epistemic standards for assertions in situations where they can reasonably foresee that if they make a claim, then some intended audience will act on it (for example, if

² Elliott, 2011, raises a similar concern although the specific formulation below draws on unpublished work by Anthony Woodman. (See Steel and Whyte, 2012 and Elliott and MacKaughn, 2014, for other concerns about Douglas's work).

scientists say that neonicotinoids harm wild bee populations, then policy-makers will ban neonicotinoids); if they do not, hearers will not act on it (if scientists remain silent, policy-makers will not ban neonicotinoids). The amount of evidence which decision-makers should demand before acting on a claim – their “epistemic standard for acceptance” – should vary with the expected practical costs of false positives and false negatives. When we can identify an audience for scientific communication, we can also identify a “proper epistemic standard” for that audience’s acceptance of some claim. Indexing scientists’ communicative obligations to these standards, rather than to all foreseeable consequences of their assertions, suggests a refined version of Douglas’s proposal:

the “floating standards obligation” (FSO): scientists should consider their audience’s proper epistemic standards for acceptance when setting their own epistemic standards for assertion.³

In the rest of this paper, I will argue against the FSO as a general account of scientific communicative norms. To understand the proposal, however, first consider its application. In 2012 the UK’s Department of Environment, Food and Rural Affairs (DEFRA) instigated an analysis of published evidence on neonicotinoids and bees which concluded “while this assessment cannot exclude rare effects of neonicotinoids on bees in the field, it suggests that effects on bees do not occur under normal circumstances” (DEFRA 2013, 1). This report is interesting not for its conclusion, but because the authors assumed, without explicit justification, that to conclude that neonicotinoids harm bee populations would require very strong evidence, despite the obvious costs of “false negatives”. Interestingly, analyses of the

³ This reworked communicative obligation might be justified by a more general account of moral responsibility (as Douglas justifies her original proposal) or in some other way – such as by appeal to Grice’s “cooperative principle” to “make your contribution such as it is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged” (Grice 1975, 45). In this paper, I will not discuss the broader issue of how to justify communicative obligations more generally.

same literature by the European Food Standards Agency (EFSA) concluded that three commonly used neonicotinoids – clothianidin, imidacloprid and thiamethoxam – do pose significant risks to wild bee populations (EFSA Panel on Plant Protection Products and their Residues 2012; EFSA 2013a). Rather, they stressed that they too were careful not to extrapolate too far from the available evidence, but disagreed on the proper analysis of that data. The FSO implies that these scientists acted wrongly, because they failed to consider how the fact that “false negatives” would probably be practically more costly than “false positives” might affect policy-makers’ proper standards for acceptance and, hence, their own standards for assertion.

Even if this judgment seems plausible, note that the FSO has radical implications. The neonicotinoid researchers’ adherence to “high epistemic standards” is not mysterious or unusual. Although it is not true that *all* scientists *always* adopt “high standards” – in §3, I discuss clear counter-examples – “epistemic conservatism” seems a characteristic feature of much scientific practice. For example, consider how in statistical testing it is routine to reject the null hypothesis and accept the alternative hypothesis only when it is statistically significant according to a stringent “type 1” error of 0.05, or, more commonly, 0.01. This practice is highly institutionalised: statistical programmes routinely “black box” setting of “p values”; journals are reluctant to publish results which are not statistically significant; and so on.⁴ DEFRA and EFSA scientists used standards which are deeply embedded in scientific practice; they disagreed on their interpretation of the evidence, not on which epistemic risks were worth running.

⁴ See Ziliak and McCloskey (2007) for extremely thorough discussion of how significance tests are routinised.

This point relates back to older discussions of inductive risk. Isaac Levi (1960) responded to the Rudner/Jeffrey debate by arguing, *contra* Jeffrey, that scientists do “accept” hypotheses but, *contra* Rudner, that doing so did not require them to make non-epistemic value judgments. Rather, Levi claimed, scientists are guided by community-level “scientific standards of inference” (356); the scientist “*qua* scientist” does not make non-epistemic value judgments. If, as I argue, “high epistemic standards” are institutionalised, Levi has a point – maybe individual scientists do not *have to* appeal to non-epistemic values to resolve inductive risk problems – but this does not blunt the moral force of the FSO. Pirates might have a code of honour which determines which captives to kill. In describing these norms, we might define a “good pirate” (say, one who does not kill all the captives). But clearly those norms are themselves morally unacceptable. Similarly, taking the FSO seriously does not necessarily commit us to thinking that individual scientists – such as those at DEFRA or EFSA – acted in a morally culpable manner, but has a far more radical implication: that the institutions which govern scientific research systematically incentivise and reward morally problematic communication.

2. Publication and the limits of the Floating Standards Obligation

In this section and the next, I argue that there is no need for a radical overhaul of scientific institutions. In this section, by distinguishing between “private” and “public” communicative contexts, I argue that the FSO cannot govern much scientific assertion. In §3, I argue that in public contexts, there are good non-epistemic reasons why scientists should adopt fixed, high epistemic standards. (So, even if the reader was unconvinced that use of high standards is a characteristic feature of scientific research, I hope she will read on to learn why they should be.)

In §1, I outlined an argument for the FSO: i.e. when scientists make claims which go beyond the available evidence, they should consider their audience's epistemic standards for acceptance and set their epistemic standards for assertion accordingly. However, individuals can only be under an obligation if they can fulfil it: "ought" implies "can".⁵ Clearly, then, scientists can only be under the FSO if they can vary their epistemic standards for assertion in proportion to hearers' epistemic standards for acceptance. In some cases they can: a scientist working for a regulatory agency, such as DEFRA or EFSA, might know what paths of action are available to her hearer, a regulator, and this can be reflected in the epistemic standards she uses. However, the assertions of scientists acting in a regulatory agency, directed at some known set of policy-makers, seem distant from paradigm cases of scientific assertion: publication in scholarly journals. I shall now show that the nature of publication insulates such communicative contexts from the FSO.

Publication is, as the term's etymology suggests, a form of public communication, a speech act where we make claims to a public, rather than private, audience. In "private" communication, speakers aim to communicate to *ex-ante* known individuals. In "public" communication, by contrast, speakers communicate to *ex-ante* unknown audiences. In a phrase associated with Kant, in publishing we are "speaking to the world at large", rather than a circumscribed audience (Kant 1970; O'Neill, 1986). Note that the claim that publication involves addressing the "world at large" is compatible with someone who makes a public claim being able to make reasonable predictions about her likely *initial* audience ("there are only five other people in the world interested and competent enough to read what I

⁵ See Howard-Snyder, 1997 for a useful overview of the history and content of this principle.

have to say”), and, indeed, compatible with her having an intended audience (“of those five people, I want Jane to read my paper”). However, in virtue of the *permanence* of publication, there is always a possibility that *ex-ante* unidentifiable audiences will hear those claims in the future. It is important to clarify that the difference between private and public contexts of communication is not merely that it is harder to identify an audience in the latter context than in the former. Rather, the difference concerns the very nature of the speech act. In the context of “private” communication, it makes sense to think of the audience as a group of, at least in-principle, identifiable individuals, with identifiable needs, concerns, and so on. In the context of “public” communication, by contrast, speakers must make some assumptions about their audience, but they cannot, even in-principle, identify the needs, concerns and so on of all members of their audience.

The fact that scientific assertion – at least in its paradigm form – is a form of public communication creates problems for using the floating standards obligation to assess the propriety of such assertions. Even in “private communication”, it is possible that the very same assertion might be intended for an audience of more than one hearer, where different epistemic standards for acceptance might be appropriate for different hearers’ acceptance of the same hypothesis. As such, application of the FSO might be extremely difficult. However, in the context of “public communication”, the problem is of a different order: there is no identifiable set of individuals for a public communication, and, as such, there is, in principle, no way in which a scientist asserting a public claim can vary her standards in proportion to hearers’ epistemic standards for acceptance. Therefore, scientists cannot be under a moral obligation to govern their public assertions using floating standards.

These considerations are compatible with the FSO identifying important considerations about scientists' obligations in private communicative contexts, i.e. contexts where they are *ex-ante* aware of their audience. (Note that a slightly peculiar feature of the private/public distinction I borrow from Kant is that it implies that the communication of a publicly-funded body such as DEFRA to a public agency can be "private"!)

However, they pose a problem for understanding that obligation as concerning the responsibilities of scientists *qua* scientists. Before sketching a positive account of public communication in §3, I shall first clarify how my claims relate to debates over inductive risk more generally.

First, one might think that the arguments above do not show that Douglas was wrong that scientists should use "floating" standards, but, rather, that my apparently friendly amendment to her proposal was, in fact, misguided. Rather than claim, as I have proposed, that scientists should vary their epistemic standards for assertion in accordance with hearers' proper standards for acceptance, we might return to Douglas's own proposal: that they should vary their epistemic standards in accordance with all foreseeable consequences of error. If so, one might claim, scientists should just vary their standards according to the interests of those hearers whom they can identify *ex-ante*. However, even placing to one side the worry set out in §1, that such a principle confuses causal and moral responsibility, this response is doubly problematic. First, even if we can foresee who will hear our claims, as long as there is more than one identifiable hearer, it will be very difficult, probably impossible, to meet the "floating standards obligation". Consider, for example, scientists publishing a review of data on neonicotinoids and bees; even if they know that this topic is of interest solely to British, American and French policy-makers, given the different agricultural systems of these nations, it would be exceedingly hard to consider and balance *all* the foreseeable consequences of error. Second, even if scientists cannot reasonably foresee who might hear them, they can still

reasonably foresee that others might well hear them; therefore, they need to take account of this possibility. However, it is simply impossible for them to do so, when communications are public.⁶ Therefore, the problem with extending the FSO to public communication does not lie with the friendly amendment suggested in §1, but reflects a deeper problem with balancing risks of error when a communication has many potential audiences.

Second, my argument echoes a key, but often overlooked, aspect of Richard Jeffrey's argument against Rudner's proposal that scientists must vary their willingness to accept claims in proportion to the expected practical costs of false positives and false negatives. Jeffrey argued that because any hypothesis might be relevant to more than one decision, "it is certainly meaningless to talk of *the* cost of mistaken acceptance or rejection" (Jeffrey 1956, 422, emphasis in original). However, Jeffrey makes this claim as part of an argument where he assumes that either scientists accept hypotheses (in which case, they should vary their standards as Rudner suggests) or they simply report their degree of confidence. His claim about the impossibility of establishing "*the* cost" of mistaken acceptance or rejection is intended as a *reductio* of Rudner's argument, implying that scientists must (in some sense) report evidential probabilities. Discussion of Jeffrey's work has focused on showing that scientists cannot avoid making inductive leaps (Gaa, 1977). This is unfortunate, because it occludes important options in debate. Specifically, I have argued that (a version of) Jeffrey's concern about the multiplicity of potential uses of hypotheses implies that scientists simply *cannot* follow a moralised version of Rudner's recommendations, at least in public communication. However, in §1, I also argued that scientists who simply report the

⁶ Interestingly, Douglas herself suggests that Kevin Elliott's ethics of expertise, according to which experts are obliged to communicate that information which allows others to make informed choices, is problematic because it is unclear who experts' audiences are (Douglas, 2012).

probability of claims may also face moral problems, given how we know their claims will be interpreted. Jeffrey's argument is only partly successful: he is right that we cannot reasonably employ floating standards in public communication, but wrong to assume that this shows that we can ignore worries about inductive risk entirely.

3. Justifying high epistemic standards

Peter Lipton claimed that the problem of induction had both a descriptive aspect, adequately describing actual inductive practice, and a justificatory aspect, explaining how such practice is justified (Lipton 2004, Chap.1). Similarly, if we concede, as I do, that scientists face problems of “inductive risk”, we can ask both a descriptive question – concerning how they do solve those problems – and a normative question – concerning how they ought to solve them. In §1 I set out a possible answer to the normative question, that their assertions should be governed by the “floating standards obligation”. §2 showed that this conclusion might be limited, because the “floating standards obligation” is inapplicable in contexts of public communication. At the end of §1, however, I also suggested an answer to the descriptive problem: scientists tend to adopt “fixed high standards” (or, more accurately, there are strong institutional pressures on them to adopt such standards). Of course, this descriptive answer is questionable. However, I shall now argue that, regardless of what actually happens, public scientific assertion should be governed by fixed, high epistemic standards.

To introduce my argument in favour of “high epistemic standards”, first consider some problems with the FSO which were not discussed in the previous section, concerning problems of co-ordination. Torsten Wilholt (2013) has pointed out the following problem: it seems plausible that scientists face problems of inductive risk, and, as such, need some way

of setting the trade-off between false positives and false negatives. However, if each scientist were to set this trade-off in an idiosyncratic way, then scientists would face co-ordination problems in deciding when and whether to rely on others' results. According to Wilholt, fixed standards are far more efficient than floating standards, and perhaps even necessary for a functioning scientific community.

While Wilholt's argument identifies an important practical benefit of fixed standards for scientific assertion, it does not show why those standards must be "high", i.e. favour the avoidance of false positives over the avoidance of false negatives. Any standard would seem to serve the purpose of co-ordinating scientific work! One way in which to justify high standards would be by appeal to the distinctively epistemic goals of science. For example, we might argue that were scientists to adopt low standards to govern their public assertions, then it would be more likely that other scientists who accept their claims would base their research on falsehoods, thereby leading to the inclusion of a significant number of falsehoods in the corpus of scientific knowledge.⁷ Along similar lines, we might argue that there is an important relationship between the use of high epistemic standards and the production of "knowledge", given that beliefs generated in a manner which leads to many "false positives" may be the kinds of beliefs which, even if true, we cannot claim to know.⁸ Clearly, there are many interesting and important issues to be explored here. However, attempting to justify "high" standards by appeal to distinctively epistemic goods seems a mistake in the current context. After all, both DEFRA and EFSA scientists might have appealed to such epistemic values to justify their high epistemic standards, but use of these standards might still seem problematic on non-epistemic grounds. One way of reading the argument from inductive risk

⁷ I am grateful to Rune Nyrup for this point.

⁸ The final section of John (forthcoming) develops these points in greater detail.

is precisely as arguing that scientists' epistemic goals do not grant them a "moral exemption" from considering the practical consequences of inductive error. If so, it is unclear that appeal to truth or knowledge can serve as knockdown justifications for high standards in public communicative contexts.

Therefore, I suggest that a proper defence of use of high epistemic standards should, instead, appeal to non-epistemic goods which follow from scientists' use of such standards.

Specifically, I argue that we can build on Wilholt's work to argue that in communities where some people are uniquely well-qualified to collect, interpret and assess evidence bearing on hypotheses, there are good reasons why those individuals' "public" claims about those hypotheses should be governed by fixed, high epistemic standards. The first step in the argument extends Wilholt's concerns about efficiency beyond communication within the scientific community to consider non-experts' needs in their reliance on scientists. As Philip Nickel suggests, audiences' reasons to defer to scientists are not grounded on the scientist offering a personal guarantee of her competence and sincerity, but on the fact that scientists "are subjected to public scrutiny by experts applying stringent norms of evidence for assertions of that kind" (Nickel 2013, 215-16). From a hearer's perspective, it is clear why fixed standards (if not necessarily Nickel's "stringent" standards) are beneficial; it is easier for a hearer to know how to respond to scientists' public claims if she can reasonably assume that those claims meet a particular standard than if scientists' standards constantly vary. For example, if I know that the social institutions of science are such that scientists very rarely make claims unless they are very likely, then I can reasonably assume that some scientist's claim is very likely, whereas if scientists routinely change their standards, I must do more digging to discover precisely how well-supported some "public" claim is. Further gains also follow. For example, if we know that the institutions which govern public scientific assertion

tend to ensure that all such claims meet some standard, S, and policy-makers are committed to acting on claims which meet S, then we can more easily hold policy-makers to account than if scientists routinely change their standards.

Wilholt suggests that fixed standards generate efficiency gains within the scientific community; I suggest that they generate important efficiency gains across the broader community. Why, though, think that these considerations favour “high” standards?

Throughout this paper, I have assumed that individuals should vary their willingness to accept (i.e. act on) claims in proportion to the expected costs of acting on false positives and false negatives. I suggest that there is an “upper limit” to the proper epistemic standards for acceptance; for nearly all agents and nearly all claims, there is some degree-of-evidence such that those agents should accept those claims. If each member of an audience has good reasons to assume that the institutions which govern scientists’ assertions are such that scientists assert claims only when those claims are extremely unlikely to be false, then she can also reasonably assume that she should defer to those claims whatever her practical interests. If, by contrast, scientists were to adopt lower standards in making public claims, audience members would have to do more digging before deciding whether or not they – given their practical interests – should defer to those claims. Therefore, the same kinds of efficiency reasons which favour the institutionalisation of fixed standards also justify the institutionalisation of high standards, at least for public communication.⁹ The heterogeneity

⁹ These remarks relate to Edward Craig’s claim (1999) that the social role of the concept of “knowledge” is to identify “reliable informants”. I suggest that the institutions of scientific research ensure that scientists are a super-“reliable informant”: whatever a hearer’s practical interests, she has reason to defer to what they say.

of our practical interests provides us with reason to want there to be institutions which are above consideration of practical interests.

This is, of course, a very abstract way of framing matters. Furthermore, it leaves open important questions, which I return to below, of what scientists should do when they have good but not great evidence for policy-relevant claims, and, relatedly, how non-scientists should interpret scientists' silences. However, this abstract picture does capture real-life considerations. For example, were researchers on neonicotinoids, publishing in widely-distributed journals, to vary their epistemic standards in accordance to the (perceived) costs and benefits of policy-makers in one country acting on false positives and false negatives, then it would always be an open question whether policy-makers in a second country should accept their claims. When, by contrast, scientists' public claims about such matters are governed by fixed, high epistemic standards, we can reasonably assume that all policy-makers, whatever the country-specific issues involved, should defer to their testimony. Of course, this is not to say that there always will or must be a smooth path from scientists' public claims to others' acceptance of those claims. As the conflicting analyses of DEFRA and EFSA illustrate, shared epistemic standards are no guarantee of consensus! However, note how much harder it would be for non-experts to decide whom to trust in such cases of conflicting testimony if there was disagreement not only over interpreting evidence, but also over proper epistemic standards.

§4 Communicative obligations and the problems of institutionalisation

Kantians often argue that when we engage in “public reason” – when we speak (as if) to the world – we should be guided by different communicative norms than when we speak to

identifiable others – because forms of justification proper in “private” contexts (“I am your pastor, so listen to me”) will be improper in “public” contexts (O’Neill 1986). Discussion of these topics has tended to focus on political debate (Rawls 1993). I have argued that a similar distinction may be important to thinking through problems of inductive risk. In §1, I conceded that there is a *prima facie* plausible argument for the FSO. However, in §2 I argued that this obligation cannot be operative in “public communication”. §3 presented a positive account of the norms for public communication, in terms of what we might call the “high standards obligation”: scientists’ public assertions should be governed by fixed high epistemic standards. If I am right that, as a matter of fact, such standards are already institutionalised in much scientific practice, then this result may seem underwhelming. However, as I will now show, it raises important normative questions about the relationship between scientists’ private and public communication and their broader communicative obligations.

Even if there are good reasons why “public” communication should be governed by high standards, enforced through institutional mechanisms, such standards are clearly not unproblematic, for at least two reasons. First, I concede that it is plausible that “private” communication should be governed by (something like) the FSO. If, however, scientists are subject to (and/or have internalised) institutional pressures proper to “public communication”, then, plausibly, they will not vary their standards even when they should. As §1 suggested, maybe such a phenomenon is at play in DEFRA and EFSA’s reports on neonicotinoids: scientists unthinkingly appealed to norms proper to “public” communication to govern “private” communication. At the very least, it seems that a defence of high standards also needs to stress the importance of institutional norms and mechanisms which

allow and incentivise scientists to adopt “lower” epistemic standards in some “private” settings.

Second, perhaps more seriously, limiting scientists’ public assertions only to claims which meet high epistemic standards may leave them unable (properly) to say very much at all. In and of itself, this is not a problem: a certain kind of epistemic caution may seem to be a virtue of academic researchers in general. However, scientists may often be in a position where they are the only people aware that certain claims, although not well-enough established to warrant “public” assertion, are well-enough established to warrant action by others in the community. Remaining silent in such cases may seem an unacceptable abrogation of moral duty, and, given the complexities of gathering, interpreting and assessing evidence, scientists may often be in such situations. It seems, then, that a full account of “public communication” should hold that even if scientists are under the high standards obligation when they make certain sorts of public claims – with the full authority of science, as it were – they may also have further obligations to “speak out” about claims which are well-enough established to warrant action by some in the community, even when they are not well-enough established to warrant action by any rational agent.¹⁰

These obligations to “speak out” are particularly important, because there is a significant risk that policy-makers and members of the public will mis-interpret or mis-understand scientists’

¹⁰ Note here the interesting relationship to the “precautionary principle” in environmental and public health policy-making, which some authors (e.g. Sunstein, 2005) read as a reminder to policy-makers that a threat may be sufficiently well-warranted to justify action even if it is not sufficiently well-warranted to be “scientifically certain” of its existence. The proposals above suggest that as well as reminding policy-makers to beware of scientific reticence, maybe scientists should sometimes be less reticent. See John, 2010, for further comments on how the problem of inductive risk relates to interpreting the precautionary principle.

silence on various hypotheses as evidence that those hypotheses are not well-enough established to warrant action. Consider, for example, Lord de Mauley, the UK environment minister, who justified the UK government's vote against an EU-wide ban on neonicotinoids as follows: "having a healthy bee population is a top priority for us but we did not support the proposal because our scientific evidence doesn't support it" (quoted in Carrington, 2013). It seems plausible that de Mauley is confusing the claim that "scientific evidence" does not suffice to treat this claim as "scientifically proven" with the claim that "scientific evidence does not suffice to treat this claim as well-enough established for policy". It seems that, in virtue of their more general civic duties, scientists have an obligation to prevent and pre-empt such confusions through "speaking out".

Even if, as I have claimed, the "high standards obligation" should govern scientists' public claims, clearly this does not exhaust the ethics of scientific communication. Rather, we must also recognise scientists' obligations to employ floating standards in private contexts and their obligations to speak out in "quasi-private" contexts. I take the claim that scientists might be under such obligations to be (relatively) uncontroversial. The key issue, however, is how we might construct institutions which allow scientists to meet these obligations at the same time as ensuring that they might maintain high standards in public communication. I am no expert in institutional design, but note here two reasons to think that creating institutions which promote these goals is likely to be both practically and morally complex.

The first set of potential problems concerns the institutionalisation of the FSO in "private contexts". One problem here – already flagged above – is that even when scientists have a specific audience for their research, different members of that audience might have different proper standards for acceptance. It is unclear, for example, how to ensure that scientists at

EFSA employ the standards proper to “their” audience – EU policy-makers – given the widely different contexts of agricultural policy in different member states. A second problem arises because there may be actors who are not the intended audience of private communication, but who have a valid interest in being able to predict how scientists decide what to communicate in those contexts. For example, in relating her proposals to practice, Douglas (2009, Chapter 7) discusses the “inference guidelines” supplied by the US National Research Councils, which mandate inferences from evidence of chemicals’ toxicity in animals to claims about their toxicity in humans. In my terms, these guidelines seem to tell scientists to use “low epistemic standards” for assertion, as they recommend cross-species extrapolations which are known to be epistemically problematic. Given that NRC-funded scientists are typically communicating in a “private context” this practice may be in-line with the demands of the FSO. However, as Douglas notes, formulation of these guidelines was plagued by debate over the freedom scientists should have to change their testing practices on a case-by-case basis, because of fears that this would make the testing regime opaque and unpredictable (2009, 144). Clearly, other members of the community – such as industrial or charitable actors – do have some reasonable interest in transparent and predictable regulatory decisions. Therefore, even in private contexts such as regulatory agencies, there may be a difficult trade-off to be struck between moral sensitivity and broader social co-ordination. Even if we can justify institutions which allow for “low” standards, there may be restrictions on whether these standards should also be allowed to vary.

As well as these practical issues, institutional design may be morally complicated. One reason to be concerned about how scientists set their epistemic standards is that non-scientists often defer to their testimony. That is to say, scientists enjoy a kind of “epistemic authority” (Douglas 2009, 135). In turn, this power seems to generate responsibility: because others will

defer to scientists, scientists should be careful in how they trade-off false positives and false negatives. Why, though, do non-scientists defer to scientists? Above, I suggested that, at least in public communicative contexts, part of the answer lies in scientists' adherence to high epistemic standards. We defer to scientists' public assertions because we can reasonably assume that, whatever our interests, if scientists assert some claim, we should accept that claim (at least, as long as scientific institutions are working well). If so, scientists who vary their standards in private contexts or who speak out in public debate may be in a morally complicated situation, because they may be speaking with an authority which, properly speaking, they only enjoy in the "public" setting. Therefore, any account of how we should institutionalize scientists' broader communicative obligations, while retaining their commitment to "high epistemic standards" in public settings, will have to be alert to this risk of moral "passing off". Neither this nor the previous problem shows that we cannot create institutions which reflect the whole range of scientists' obligations, allowing, for example, that they might say one thing in Brussels, another thing in a journal, and a third thing in a newspaper editorial. What they do suggest, however is that constructing such institutions will be practically difficult, and that any set of institutions governing scientific communication might have significant moral costs.

In concluding this section's discussion of the complex normative problems raised by my arguments, it is useful to clarify how these worries relate to the "scope" of arguments from inductive risk. Sometimes, the argument from inductive risk is understood to imply a need for non-epistemic value judgments in all scientific work (Douglas, 2009, Chap.3).

Sometimes, it is understood more modestly as implying a need for non-epistemic value judgments when scientists act as "policy-advisors" (Steele 2012). The argument from inductive risk is, I suggest, attractive as an account of policy-advice (although note the

serious caveats above about how to institutionalise these concerns). However, because the argument is often framed in terms of how scientists should resolve a problem which arises in all scientific research – how to balance risks of false positives against risks of false negatives in inductive inference – it can be easy to think of “regulatory science” simply as a vivid example of a more general phenomenon, and that the norms proper for these cases are proper for all scientific research. I have argued that if we focus attention on contexts of communication, however, we can accept the force of the argument as an account of scientists’ obligations *qua* policy-advisors, but not *qua* scientists.

Note how this differs from an alternative strategy for limiting the argument from inductive risk to policy advice: that, even if the argument is relevant to some cases of “applied science”, it cannot be relevant to “theoretical science” because such research is often not directly relevant to any possible action (Levi 1960). My claims above do not rest on a distinction between types of research – I have argued that even when scientists are working in obviously practical fields, they have good reason to adopt high standards for public communication – but on types of communication.¹¹ In effect, I deny that there is a single answer to the normative problem of inductive risk; rather, it depends on audience.

¹¹ Furthermore, the proposed distinction between different forms of communication is preferable to Elliott’s similarly pluralistic suggestion that the propriety of scientists’ appeal to values depends on the particular “goals” prioritized in their context (see, for example, Elliott, 2013, p.381; Elliott and McKaughan, 2014). Elliott’s approach might seem to justify not appealing to non-epistemic values in, for example, journal articles if the “goals” of that activity are promoting truth, rather than aiding regulation. However, it is unclear why the fact that a scientist has a particular epistemic goal should grant her exemption from other moral considerations. What my argument does, then, is to “fill in” a non-epistemic justification for pursuing what might seem to be epistemic goals.

5. From freedom to neutrality; from ideal to second-best

The argument from inductive risk is often taken to show a problem for the “value free ideal” for science. Given the complexities around distinguishing different kinds and possible roles of value judgment in science, it can be unclear precisely what proponents of this ideal are committed to, but Gregor Betz’s recent definition – “the justification of scientific findings should not be based on non-epistemic (e.g. moral or political) values” (Betz 2013, 207) – captures the key idea. In conclusion, then, I will outline the implications of my arguments for the broader debate over the proper role of values in science.

I suggest that debates over “value freedom” often embody two confusions. First, as my comments on Rudner and Jeffrey at the end of §2 noted, many seem to assume that accepting that scientists do solve problems of inductive risk implies denying the value free ideal. Second, it may seem that defenders of the Value Free Ideal must ignore or downplay the complex relationships between (much) scientific inquiry and economic, social and political goals, in favour of a focus on the purely epistemic goals of inquiry. However, it is unclear that defenders of the Value Free Ideal must deny that scientists face problems of inductive risk or that much scientific inquiry is of great practical relevance. On the first point, as I noted in §1, we might concede that scientists do solve problems of inductive risk but deny that this involves appeal to non-epistemic values, as opposed to adherence to institutionalised standards. In turn, as §3 noted, but did not develop, use of “high standards” might be justified on purely “epistemic” grounds, as related to the generation of knowledge. On the second point, Betz’s own defence of value free science is motivated partly on the grounds that were scientists to appeal to non-epistemic values even “indirectly” in their work, they would

violate important democratic norms, according to which the people, rather than experts, should choose which values guide policy. I disagree with Betz's particular claims here – as scientists might take account of non-epistemic values but respect democratic norms if they were explicit on these value-judgments (Elliott, 2013) – but his general strategy raises an important point: the politically embedded nature of science may be a reason for, rather than against, value-free science.

This paper has developed both of these general thoughts in the following way. As I argued in §3, attempting to *justify* scientists' use of high epistemic standards by appeal solely to epistemic goods seems a weak response to the moral concerns raised by arguments from inductive risk. Although some think that it is important to show that an indirect role for values in science is compatible with a concern for epistemic values (Steel, 2010), the real challenge of the argument from inductive risk is, I suggest, that it makes us question the value of knowledge. It does so by reminding us that, for practical purposes, we might be better-off acting on not-known claims than only acting on known claims. However, I responded to this moral argument by a dual-level response, which distinguishes between the values which can be appealed to within a practice, and the values which we should use to justify having such a practice. At the first level – that of the practice of science – I have claimed, in-line with the Value Free Ideal, that scientists should not appeal to non-epistemic values in deciding which claims to make. However, this defence of “high epistemic standards”, which help to generate “knowledge”, does not itself appeal to the value of knowledge. Rather, I have argued for this practice in terms of how it allows for an efficient co-ordination of experts' claims and non-

experts' practical needs. That is to say, I have defended excluding non-epistemic values from science by appeal to non-epistemic values.¹²

At this point, some readers might be worried that these remarks, with their apparently strong distinction between epistemic and non-epistemic values and concerns, are in tension with an important strand in recent epistemology, according to which knowledge-ascriptions (perhaps even knowledge) are related to ascribers' or subjects' practical interests (Fantl and McGrath, 2012). For many epistemologists, it seems that knowledge does not require "high epistemic standards". However, note that the empirical data supporting claims of "pragmatic encroachment" are contestable (Gerken 2012). Furthermore, David Henderson (2011) has argued for a route from contextualist accounts of knowledge to the conclusion that scientific claims should be treated as known only when they meet high standards, on grounds similar to those above: that scientific communities are "*general-purpose source communities* – communities of inquirers having a social role of producing information of such a high epistemic quality that a somewhat indeterminate range of groups might freely draw on their results without hesitation." (87) Therefore, trends in contemporary epistemology which may seem to complicate my conclusions in fact lend support to the general thrust of my argument.

I have, then, argued that there is a version of the Value Free Ideal which is consistent with the claim that scientists solve problems of inductive risk, and which not merely recognises but is built upon an acknowledgment of the social, economic and political relevance of scientific

¹² Note then that there may be an interesting analogy here between scientific and legal contexts. In a recent paper, Enoch, Talia and Fisher (2012) have argued that courts' refusal to use statistical evidence might be understood in terms of the epistemic good of "sensitivity". However, as they also note, that we can redescribe courts' practices in this way leaves open a further justificatory question: why should courts care about this epistemic good, given that the exclusion of statistical evidence often seems to conflict with important aims of the legal system. They suggest, then, that "policy" considerations must be used to justify this practice. I suggest that a similar dual-level structure applies in the case of science.

research. However, it is unclear that this defence counts as a complete vindication of the Value Free Ideal for two reasons. First, my argument for “value free” science has turned on the importance of a certain form of “value neutrality” in public communications within societies characterised by value pluralism. I suspect that this may seem rather a weak argument to many who think that science ought to be value-free, who might hope for a more full-blooded commitment to epistemic values. Second, as I stressed in §4, the form of “value neutrality” I endorse in this paper is not unproblematic, but can create its own problems and difficulties. It is, as it were, not so much an “ideal” to be strived for, but the best available solution to a complex co-ordination problem, which still leaves many problems to be solved. The real lesson, then, may be that talk of the role of “values” in research requires supplementation by more discussion of the norms of communication.¹³

References

Betz, Gregor (2013), "In defence of the value free ideal." *European Journal for Philosophy of Science* 3: 207-220.

Biddle, Justin (Forthcoming), “Transient underdetermination and values in science.”

forthcoming in *Studies in History and Philosophy of Science Part A*

¹³ I am very grateful to the following people for discussion of the ideas raised in this paper: Anna Alexandrova, Shahar Avin, Marion Boulicault, Hasok Chang, Charlotte Goodburn, Tim Lewens, Emily McTernan, Onora O’Neill and Anthony Woodman. A previous version of this paper was presented at the Departmental Seminar, Department of History and Philosophy of Science, University of Cambridge, April 2013, and I benefitted from the discussion there. The comments by two anonymous referees for *Synthese* were unusually constructive and helpful. I am most grateful, however, to the several cohorts of undergraduate students who patiently sat through my lectures on the topic of inductive risk where I tried to articulate the concerns above.

Biddle, Justin, and Eric Winsberg (2010), "Value judgements and the estimation of uncertainty in climate modelling." pp127-97 in Jacob Busch and P.D. Magnus (eds) *New waves in philosophy of science*. London: Palgrave Macmillan.

Carrington, D (2013) "Bee-harming pesticides banned in Europe" *The Guardian* Mon, 29th April, 2013, at <http://www.theguardian.com/environment/2013/apr/29/bee-harming-pesticides-banned-europe> (Accessed 15th March 2014)

Churchman, C. West (1948), "Statistics, pragmatics, induction." *Philosophy of Science* 15(3): 249-268.

Craig, Edward (1999), *Knowledge and the state of nature*. Oxford: Clarendon Press.

Cranor, Carl (1993), *Regulating Toxic Substances*. Oxford: Oxford University Press.

Desneux, Nicols, Axel Decourtye and Jean-Marie Delpuech (2007) "The sublethal effects of pesticides on beneficial arthropods." *Annual Review of Entomology*. 52: 81-106

DEFRA (2013) *An assessment of key evidence about Neonicotinoids and bees*. London: Department for Environment, Food and Rural Affairs

Douglas, Heather E. (2000) "Inductive risk and values in science" *Philosophy of Science* 67 (4):559-579

Douglas, Heather E. (2009) *Science, policy, and the value-free ideal*. Pittsburgh, PA: University of Pittsburgh Press.

Douglas, Heather, E. (2012) “Book review of Kevin Elliott, *Is a little pollution good for you?*” *Philosophy of Science* 79(3): 425-28

Easterbrook P.J., Berlin J.A., Gopalan R., Matthews D.R. (1991), “Publication bias in clinical research.” *Lancet* 337(8746): 867-872.

EFSA Panel on Plant Protection Products and their Residues (PPR) (2012). “Scientific Opinion of the Panel on Plant Protection Products and their Residues on a request from the European Commission on the science behind the development of a risk assessment of plant protection products on bees (*Apis mellifera*, *Bombus* spp. and solitary bees)”. *The EFSA Journal* 2012; 10(5):2668

EFSA (2013) “Press release: EFSA identifies risks to bees from neonicotinoids” available at www.efsa.europa.eu/en/press/news/130116.htm

EFSA (2013b) “Guidance on the risk assessment of plant protection products on bees (*Apis mellifera*, *Bombus* spp. and solitary bees)”. *EFSA Journal* 2013;11(7):3295

Elliott, Kevin (2011) “Direct and Indirect Roles for Values in Science.” *Philosophy of Science* 78(2): 303-24

Elliott, Kevin (2011) *Is a little pollution good for you?* Oxford: Oxford University Press

Elliott, K (2013) “Douglas on Values: From Indirect Roles to Multiple Goals,” *Studies in History and Philosophy of Science* 44: 375-383.

Elliott, K and D McKaughan (2014) “Nonepistemic Values and the Multiple Goals of Science” *Philosophy of Science* 81(1): 1-21.

Enoch, David, Talia Fisher and Levi Spectre (2012) “Statistical Evidence, Sensitivity and the Legal Value of Knowledge.” *Philosophy and Public Affairs* 40(3): 197-224

Fantl, Jeremy, and Matthew McGrath (2010) *Knowledge in an uncertain world* Oxford: Oxford University Press

Flory, J and Philip Kitcher (2004). “Global Health and the Scientific Research Agenda.” *Philosophy and Public Affairs* 32 (1):36–65

Gaa, James (1977) “Moral autonomy and the rationality of science.” *Philosophy of Science* 44(4): 513-541.

Gerken, Mikkell (2012) “On the cognitive basis of knowledge ascriptions.” in Mikkell Gerken and Jessica Brown (eds) *Knowledge-ascriptions* Oxford: Oxford University Press

Grice, Paul (1975) "Logic and conversation." pp41-58 in Peter Cole and Jerry Morgan *Syntax and semantics 3: Speech acts* New York: Academic Press

Henderson, David (2011) "Gate-keeping contextualism." *Episteme* 8(1): 83-96

Henry, Mickael, Maxime Béguin, Fabrice Requier, Oriane Rollin, Jean-François Odoux, Pierrick Aupinel, Jean Aptel, Sylvie Tchamitchian and Axel Decourtye (2012) "A Common Pesticide Decreases Foraging Success and Survival in Honey Bees." *Science* 336(6079), 348-50

Hempel, Carl G. (1965) "Science and human values." pp81-96 in his *Aspects of scientific explanation*. New York: Free Press

Howard-Snyder, Frances (1997) "The Rejection of Objective Consequentialism." *Utilitas* 9(02): 241-248

Jeffrey, Richard (1956) "Valuation and acceptance of scientific hypotheses." *Philosophy of Science* 23(3) (1956): 237-246.

John, Stephen (2010) "In defence of bad science and irrational policies: an alternative account of the precautionary principle" *Ethical Theory and Moral Practice* 13(1): 3-18

John, Stephen (2010) "Expert testimony and epistemological free-riding." *The Philosophical Quarterly* 61(244): 496-517

John, Stephen (Forthcoming) “The example of the IPCC does not vindicate the Value Free Ideal: a reply to Gregor Betz” forthcoming in *European Journal for Philosophy of Science*

Kant, Immanuel (1970), “What is enlightenment?” in Hans Reiss (ed.) *Kant's Political Writings* Cambridge: Cambridge University Press

Kitcher, Philip (2011), *Science in a democratic society*. New York: Prometheus Books

Kukla, Rebecca (2012) ““Author TBD”: Radical Collaboration in Contemporary Biomedical Research.” *Philosophy of Science* 79(5): 845-858.

Levi, Isaac (1960). “Must the scientist make value judgments?” *The Journal of Philosophy* 57(11) (1960): 345-357.

Lipton, Peter (2004), *Inference to the best explanation*. London: Routledge

Nickel, Philip J. (2011) “Testimonial entitlement, norms of assertion and privacy.” *Episteme* 10(02): 207-217

O'Neill, Onora (1986) “The public use of reason.” *Political Theory* 14(4): 523-551.

Redmayne, Mike (2008) “Exploring the Proof Paradoxes.” *Legal Theory* 14: 281–309.

Rawls, John (1993) *Political liberalism* New York: Columbia University Press

Rudner, Richard (1953) "The scientist *qua* scientist makes value judgments." *Philosophy of Science* 20(1): 1-6

Saul, Jennifer (2013) *Lying, Misleading, and The Role of What is Said*. Oxford: Oxford University Press

Shrader-Frechette, Kristin (1991), *Risk and Rationality*. Berkeley: University of California Press

Steel, Daniel (2010) "Epistemic Values and the Argument from Inductive Risk." *Philosophy of Science* 77(1): 14-34.

Steel, D. and K. P. Whyte (2012). "Environmental Justice, Values, and Scientific Expertise," *Kennedy Institute of Ethics Journal* 22: 163-182.

Steele, Katie (2012) "The scientist *qua* policy advisor makes value judgments." *Philosophy of Science* 79(5): 893-904.

Stokstad, Erik (2012) "Field research on bees raises concern about low dose pesticides." *Science* 335 (6078): 1555

Sunstein, Cass R. (2005), *Laws of fear: Beyond the precautionary principle* Cambridge: Cambridge University Press

Trouwborst, Arie (2002) *Evolution and status of the precautionary principle in international law* The Hague: Kluwer Law International

United Nations General Assembly (1992) *Report of the United Nations conference on Environment and Development* New York: United Nations

Whitehorn, Penelope, Stephanie O'Connor, Felix Wackers, and Dave Goulson (2012) "Neonicotinoid pesticide reduces bumble bee colony growth and queen production" *Science* 336(6079): 351-352.

Wilholt, Torsten (2013) "Epistemic Trust in Science." *The British Journal for the Philosophy of Science* 64(2): 233-253.

Ziliak, Stephen and Deidre McCloskey (2007), *The Cult of statistical significance* Ann Arbor: University of Michigan Press