

On the Mitigation of Inductive Risk

GABRIELE CONTESSA

Department of Philosophy
Carleton University
1125 Colonel By Drive
Ottawa, ON, K1S 5B6
Canada

gabriele.contessa@carleton.ca

The last couple of decades have witnessed a renewed interest in the notion of inductive risk among philosophers of science. However, while it is possible to find a number of suggestions about the mitigation of inductive risk (i.e., its assessment and management) in the literature, so far these suggestions have been mostly relegated to vague marginal remarks. This paper aims to lay the groundwork for a more systematic discussion of the mitigation of inductive risk. In particular, I consider two approaches to the mitigation of inductive risk—the individualistic approach, which maintains that individual scientists are primarily responsible for the mitigation of inductive risk, and the socialized approach, according to which the responsibility for the mitigation of inductive risk should be more broadly distributed across the scientific community or, even more broadly, across society. I review some of the argument for and against the two approaches and introduce two new problems for the individualistic approach, which I call the problem of precautionary cascades and the problem of exogenous inductive risk, and I argue that a socialized approach might alleviate each of these problems.

1. INTRODUCTION

The last couple of decades have witnessed a renewed interest in the notion of inductive risk among philosophers of science. While the label ‘inductive risk’ is due to Carl Hempel (1965), the notion is often credited to Richard Rudner (1953). Rudner argued that, since no scientific hypothesis is ever conclusively verified or falsified, whenever scientists decide to accept or reject a hypothesis, they should take into account the non-epistemic (e.g., moral, social, and political) consequences of accepting a false hypothesis or rejecting a true one. If Rudner’s argument is sound, then even choices that seem to be purely epistemic (such as the choice of an epistemic standard of acceptance or rejection of a certain hypothesis) presuppose non-epistemic value judgments.

Consider, for example, the case of glyphosate, which is a widely used broad-spectrum herbicide that is suspected of being a human carcinogen (IARC 2015). If Rudner is right, then the epistemic decision to accept the hypothesis that glyphosate causes cancer should take into account the potential consequences of error. If, on the one hand, researchers wrongly reject the hypothesis glyphosate causes cancer, then its continued use would result in a higher incidence of certain forms of cancer with all the negative human and social consequences that that entails. If, on the other hand, researchers wrongly accept the hypothesis that glyphosate causes cancer, then their decision might result in the overregulation of (or even a ban on) glyphosate-based pesticides.

Heather Douglas, who is primarily responsible for the current revival of interest in the notion of inductive risk, has expanded and strengthened Rudner’s original argument in a number of ways (see, in particular, (Douglas 2009)). Let me mention three here. First, Douglas has emphasized the increasingly prominent role played by science in the policymaking process over the past century and

with it the increasing relevance of inductive risk (see, in particular, (Douglas 2009: Ch 2)). Second, Douglas has argued that researchers and advisors have a moral duty to consider the possible consequences of their scientific decisions (see, in particular, (Douglas 2009: Ch 4)). Third, Douglas has maintained that inductive risk is not circumscribed to the decision to accept or reject a hypothesis, but it also affects other epistemic decisions upstream the research process. For example, she maintains that certain methodological decisions should also take into account non-epistemic values (see, in particular, (Douglas 2000)).

While most of the literature so far as focused on whether inductive risk is a reason to abandon the value-free ideal of science or on how inductive risk affects different areas of science,¹ very little attention has been paid so far to *the mitigation of inductive risk*—i.e., to how inductive risk should be assessed and managed at various stages of the research and advice process. While it is possible to find a number of suggestions about the mitigation of inductive risk in the literature, so far these suggestions have been mostly relegated to vague marginal remarks. This paper aims to lay the groundwork for a more systematic discussion of the mitigation of inductive risk.

For our purposes, it is convenient to distinguish two broad approaches to the mitigation of inductive risk, which, for the sake of convenience, I call, respectively, *the individualistic approach* and *the socialized approach*. On the individualistic approach, individual scientists are primarily responsible for the mitigation of inductive risk. On the socialized approach, the responsibility for the mitigation of inductive risk is more broadly distributed across the scientific community or, even more broadly, across society.

Given the brevity of these descriptions, a few remarks are in order. First, the individualistic and the socialized approaches are best understood as the two extremes of a spectrum of possible approaches to the mitigation of inductive risk. For our purposes, it is convenient distinguishing between what we might call *the strictly individualistic approach*, according to which individual scientists are *solely* responsible for the mitigation of inductive risk, and a *broadly individualistic approach*, individual scientists are *primarily* (though not necessarily solely) responsible for the mitigation of inductive risk. As far as I can see, no one accepts the strictly individualistic approach and (as I discuss in §2) for good reasons. Even Douglas, who, due to her emphasis on the moral responsibilities of scientists (see, in particular, (Douglas 2009, chap. 4)), might appear sympathetic to an individualistic approach, clearly rejects what I call the strictly individualistic approach (see, e.g., (Douglas 2009, Ch 8) and (Douglas 2018)). Other contributors to the literature on inductive risk are even more explicit about their support for a socialized approach to the mitigation of inductive risk (see, e.g., (Elliott 2017: 97–99) or (Biddle and Kukla 2017)).

Second, while “the strictly individualistic approach” labels a clearly defined view, the other two labels are better understood as umbrella term that cover a number of possible views about how to distribute the mitigation of inductive risk among scientists. The difference between the broadly individualistic approach and the socialized approach is a matter of degree—it is a matter of how to balance individual and collective responsibilities for the mitigation of inductive risk. The socialized approach does not deny that individual scientists should play a role in the mitigation of inductive risk. Even on a socialized approach, individual scientists do face decisions that require them to assess the relevant inductive risks. However, the socialized approach denies that, individual scientists are solely or even primarily responsible for the mitigation of inductive risk.

Third, the distinction between the individualistic and the socialized approaches ignores another important dimension of disagreement about the mitigation of inductive risk, which concerns which values should be used in the mitigation of inductive risk. Should scientists employ their own personal

¹ See, e.g., the debate between Douglas and Gregor Betz in (Elliott and Steel 2016) and the contributions to (Elliott and Richards 2017).

values in the mitigation of inductive risk?² Or should they, instead, employ the values that are prevalent in their society?² While it might be tempting to believe that the individualistic approach relies on the individual values of scientists, this need not be the case.³ Therefore, the question of *who* should be responsible for the mitigation of inductive risk is largely independent from the question of *which* values they should use. While the latter question is very important, I do not attempt to address it in this paper.

Finally, the labels for the two approaches are somewhat misleading, as the question of *who* should be responsible for the mitigation of inductive risk is only one of the questions to which the various versions of the two approaches give different answers. Another equally important question concerns *where* in the research and advice process inductive risk should be mitigated. The individualistic approach presupposes that the mitigation of inductive risk should always happen *in situ* (i.e., at the stage of the research and advice process where the individual scientist faces a certain inductive risk). While the version of the socialized approach I sketch here suggests that there are, at least, two *loci* where the responsibility for the mitigation of inductive risk should be more broadly distributed among various actors. The first is at the stage in which a scientific community sets (or revises) its field-specific epistemic and methodological standards (e.g., the choice of a conventional level of statistical significance appropriate for a certain scientific field). The second is at the advisory stage, when the researchers are called on to advise policymakers on specific issues.

The plan for the paper is as follows. In the next section (§2), I briefly discuss some of the advantages and disadvantages of the individualistic approach. In the following two sections, I introduce two problems for the individualistic approach, which I call, respectively, *the problem of precautionary cascades* (§3) and *the problem of exogenous inductive risk* (§4). In the final section (§5), I conclude with some general reflections on the mitigation of inductive risk and some brief and tentative remarks about the mitigation of inductive risk during the COVID-19 pandemic.

2. THE INDIVIDUALISTIC APPROACH: PRO AND CONTRA

As I mentioned in the previous section, the most important question is not so much *whether* we should adopt a socialized approach to the mitigation of inductive risk as much as *which* version of the socialized approach we should adopt. This is because the strictly individualistic approach seems to suffer from a number of serious problems. In this section, I briefly mention three of them.

The first is that individual scientists might make value judgements unconsciously or unreflectively, which is more likely to lead to poor value judgments (as in the case in which unconscious prejudices cloud the judgment of scientists). A socialized approach to the mitigation of inductive risk would likely contribute to more explicit and reflective value judgements being made, as it would require the participants to openly discuss and justify their reasons for their epistemic decisions, including the role, if any, that values play in those decisions.⁴

The second problem is that scientists upstream in the research process might find it harder to assess and manage inductive risks, as the further these decisions are from the point of practical application,

² See, e.g., (Elliott 2017) and (Schroeder forthcoming)

³ For example, it seems possible to maintain that individual scientists are solely responsible for the mitigation of inductive risk but that they should employ values that are widely shared in their society in the management of inductive risk.

⁴ This is a contention that is often made by supporters of deliberative decision-making processes and, in particular, deliberative democracy. In general, supporters of the socialized approach would seem to be able rely on many of the arguments that are offered in favor of deliberative democracy (see, e.g., (Landemore 2017)).

the less clear it is what their non-epistemic consequences of them might be.⁵ For example, it is unclear what the non-epistemic consequences of incorrectly accepting the hypothesis that glyphosate causes cancer are without knowing what possible measures different countries would adopt in response to the acceptance of that hypothesis. Would all countries ban glyphosate? Would some countries only regulate its use? This seems to suggest that scientists upstream in the research process might not always be in the best position to try to mitigate the relevant inductive risks.

The third problem is that, as a matter of fact, scientific communities already partly adopt a socialized approach to the mitigation of inductive risk. This is because specific scientific communities already manage inductive risk collectively in some important respects. As I mentioned above, one way in which they do so is by setting and revising their field-specific epistemic and methodological standards (see, e.g., (Biddle and Kukla 2017))). While the process by which scientific communities collectively manage inductive risk can be likely improved, the notion that inductive risk should be managed solely (or even primarily) by individual scientists (as the strictly individualistic approach claims) seems to disregard the irreducibly social dimension of the scientific process.⁶

However, while the strictly individualistic approach might be both prescriptively and descriptively inadequate, even supporters of the socialized approach must concede that, to some extent, individual scientists will have to shoulder some of the responsibility for mitigating inductive risk (if only because it is sometimes the most expedient way to mitigate inductive risks). As I mentioned above, the crucial question, therefore, is how to best balance individual and collective responsibilities in the mitigation of inductive risk and how to develop better norms for distributing those responsibilities among different actors, including individual scientists, scientific communities, more or less formal scientific institutions, advisory committees, stakeholders, and policymakers.⁷ In this sense, it is still fruitful to discuss the limitations of the individualistic approach, as it might help us develop better models and procedures for the mitigation of inductive risk. In this spirit, the next two sections discuss two specific problems for a (broadly) individualistic approach and sketch how versions of the socialized approach might help preventing them.

3. THE PROBLEM OF PRECAUTIONARY CASCADES

In this section, I discuss what I call the problem of precautionary cascades. If, as I assume here, Douglas is by and large correct, then researchers and advisors might have to rely on non-epistemic value judgments at a number of steps in the course of the research and advice process. The problem of precautionary cascades arises from the fact that, in certain contexts, scientists are likely to make similar value judgments at the different steps and that the epistemic outcomes of each of these decisions might accumulate and lead to what I call a *precautionary cascade*.

To illustrate this point, consider a case in which the relevant value judgments might be relatively unambiguous, such as that of aluminium compounds and Alzheimer's Disease. The relevant hypothesis, which I shall call *the aluminium hypothesis*, is that aluminium causes Alzheimer's Disease.⁸ Given that aluminium compounds are widely used in personal care products (such as commercial deodorants and antiperspirants), the use of these products might expose the population to an increased risk of Alzheimer's Disease. Let me assume, if only for the sake of the argument, that the moral, social,

⁵ This point is similar to a point made by Richard Jeffrey (1956) in his response to Rudner's original argument.

⁶ On the irreducibly social nature of science, see, e.g., (Longino 1990).

⁷ I would like to thank one of the anonymous reviewers for suggesting this way to formulate this point.

⁸ For a review of the evidence for and against the Aluminium Hypothesis, see (Klotz et al. 2017). For a more critical view, see (Lidsky 2014).

and economic costs of a ban on the use of aluminium compounds in personal care products would be negligible or, at least, that the costs would be so small that the potential non-epistemic consequences of wrongly rejecting the aluminium hypothesis are unanimously judged to be far worse than the potential consequences of wrongly accepting it. In these circumstances, it is reasonable to expect that, if disinterested researchers and advisors are individually responsible for the relevant value judgments, the vast majority of them would come to a similar conclusion and try to err on the side of caution—i.e., on the side of wrongly accepting the aluminium hypothesis rather than on the side of wrongly rejecting it. However, if, at each juncture in the research and advice process, the experts rely on similar value judgments without taking into account the epistemic adjustments made at previous stages in the research process due to similar value judgments, then the epistemic consequences of each of these decisions are likely to accumulate, effectively lowering the epistemic standards for accepting the aluminium hypothesis to a level that might be excessively low. When this occurs, we face what I call a *precautionary cascade*.

Let me elaborate this point. Given that it would be obviously unethical to run a randomized controlled trial of a potentially neurotoxic substance on humans, researchers rely on two main sources of evidence about the safety of aluminium compounds—i.e., animal studies and population studies. For the sake of simplicity, let me assume that inductive risks have to be assessed and managed at two distinct stages—the research stage and the advice stage. At the research stage, researchers conducting individual animal or population studies face a number of unforced epistemic choices about methods, materials, data management, data analysis, etc. that might affect the ultimate outcomes of their studies. As Douglas has persuasively argued (see, in particular, (Douglas 2000)), researchers have a duty to take into account the potential non-epistemic consequences of each of these decisions. For example, researchers can use a wide variety of methods to examine the neural tissue samples from experimental rabbits for abnormalities (including using different stains, different light sources, and different types of microscopes) and some of these methods might be more likely to lead to the classification of ambiguous samples as abnormal than others.

Now, in cases like the one we are discussing, it is likely that, insofar as researchers take the potential non-epistemic consequences of their epistemic choices into account, all (disinterested) researchers will tend to err on the side of caution and make choices that are more likely to result in the acceptance than in the rejection of the aluminium hypothesis. For example, they might use a method that are more likely to lead to ambiguous samples being classified as abnormal than methods that are more likely to classify them as normal. This might result in a systematic bias in the outcomes of the vast majority of studies, which is more likely to lead to the eventual acceptance of the aluminium hypothesis.

Proponents of inductive risk have persuasively argued that this sort of bias not only not detrimental, but it is actually beneficial. However, the problem is that advisors at the advice stage of the process are likely to make similar value judgments as they review, evaluate, and aggregate the evidence from the individual studies in order to decide whether the totality of the evidence available warrants accepting the aluminium hypothesis. The problem is that, if the advisors also decide to err on the side of caution when reviewing and aggregating the evidence from the research stage and if they do so without taking into account how researchers in the previous stage also adjusted their epistemic standards in light of non-epistemic considerations, this will result in a precautionary cascade—the epistemic effects of the individual epistemic choices will accumulate and, possibly, tip the scales in favour of accepting the aluminium hypothesis.

Prima facie, this might not seem to be a particularly serious problem. After all, I have assumed that the costs of wrongly accepting the aluminium hypothesis are negligible. However, this positive assessment seems to underestimate the danger of precautionary cascades. Let me focus on three problems here. The first problem is that precautionary cascades result in the collective adoption of

epistemic standards that are looser than the standards each individual scientist would adopt, which, in itself, seems problematic. The second problem is that, if the community unintentionally adopts excessively loose epistemic standards for accepting the aluminium hypothesis, then it might become practically impossible to prove aluminium compounds to be safe. In this sense, a precautionary cascade might lead to a situation in which we set a “guilty until proven innocent” standard while, at the same time, creating conditions that make it practically impossible to prove aluminium “innocent.” After all, even the randomized control trials that are part of the approval process for prescription medications often fail to reveal serious but rare adverse effects and these adverse effects are often only discovered when the drug has been in use for some time in the general patient population.⁹ The third and most important problem is that precautionary cascades can also occur when the costs and the benefits of each option are not as lopsided as in the case I have discussed. All that is required for a precautionary cascade is that, independently from and unbeknownst to one another, a large proportion of the scientists involved along the research and advisory process adjust their epistemic standards in the same direction on the basis of their individual and uncoordinated assessments of the inductive risks relevant to the case.

Supporters of the individualistic approach might argue that precautionary cascades can be avoided if individual scientists are transparent about their value judgments and how they affect their epistemic decisions. If researchers upstream in the research and advice process communicate clearly where and how non-epistemic values informed their epistemic choices, then researchers and advisors downstream in the process can keep track of the epistemic adjustments made at each of the previous stages and take them into account in their decisions, thus avoiding a precautionary cascade.

However, while transparency is undoubtedly part of the solution, it is not clear that it is always sufficient to deal with the problem of precautionary cascades. First of all, each individual study requires researchers to make myriad unforced epistemic choices that might affect its outcome. It would likely be not only exceedingly impracticable but actually unhelpful to keep track of and reporting each of these choices and how they were informed by practical and normative considerations. It is difficult to see how advisors who are reviewing dozens of studies at the advisory stage would not be simply overwhelmed by such an overabundance of details. Second, in an era of increasing specialization, it is unlikely that any scientist can fully grasp the consequences of each of these choices and the motivations behind them when they are not within their narrow area of specialization. For example, an epidemiologist is unlikely to understand all of the subtleties of conducting experiments on lab animals and a biologist is unlikely to understand all of the subtleties of conducting a population study. Finally, as Richard Jeffrey (1956) pointed out in his response to Rudner’s original argument, the notion of the non-epistemic consequences of an epistemic decision is ill-defined. While Jeffrey seemed to take this as a reason for thinking that trying to fathom the possible non-epistemic consequences of our epistemic decisions is a fool’s errand, it is more likely that this means that inductive risk has to be comprehensively (re)assessed and managed at the advice stage, as it is usually at that stage that the non-epistemic consequences of the advisors’ decisions are clearest.

While my aim here is to shed light on the problem of precautionary cascades rather than to try to offer a solution to it, let me briefly mention how a socialized approach might reduce the risk of precautionary cascades. A possible solution to the problem of precautionary cascades would be to include researchers from as many relevant sub-fields as possible in the review of the evidence at the advisory stage and to promote an open discussion of how to best (re)assess and manage the relevant inductive risks at that stage. To be clear, on the socialized approach, the individual scientists involved

⁹ This is due, among other things, to the fact that most randomized control trials are relatively small and that they include many exclusion criteria and that as a result the trial population is not representative of the general population that will use the drug. For a discussion of these issues in relation to inductive risk, see (Stegenga 2017).

in the research process still bear some responsibility for mitigating the inductive risks they face to the best of their abilities and they should do so as explicitly and transparently as possible. However, it is the advisory committee that is primarily responsible for ensuring that the steps taken by the individual scientists upstream in the process do not result in a precautionary cascade.

Admittedly, this proposal is likely only going to reduce the risk of precautionary cascades, not avoid them completely. Moreover, as it is presented here, it is just a sketch. A number of questions would need to be answered before trying to implement it. For example, is it advisable for the advisory committee to involve the authors of the studies that are being reviewed? On the one hand, they are likely to be the foremost experts on the subject and they are best positioned to walk the other members of the advisory committee through their decision-making process. On the other hand, their involvement might inhibit an honest assessment of the quality of the studies that are being reviewed. One possible solution is for the advisory committee to consult with the original researchers while making their own decisions at arm's length from the original researchers. However, it is beyond the scope of this paper to develop a detailed approach to reducing the risk of precautionary cascades and more work needs to be done to design effective processes in contexts in which precautionary cascades are particularly likely to occur.

4. THE PROBLEM OF EXOGENOUS INDUCTIVE RISK

In this section, I discuss what I call the problem of exogenous inductive risk. In order to illustrate the problem, consider again the case of glyphosate briefly discussed in §1. Glyphosate is a widely used broad-spectrum herbicide. In its 2015 review of the available evidence, World Health Organization's International Agency for Cancer Research classified glyphosate as 'probably carcinogenic to humans' (IARC 2015, 112). However, the European Food Safety Authority (EFSA) has concluded that 'glyphosate is unlikely to pose a carcinogenic hazard to humans' (EFSA 2015) and glyphosate is still approved for use by both the EFSA and the Environmental Protection Agency in the United States.¹⁰

Given the widespread use of glyphosate-based pesticides for agricultural and horticultural purposes, this is clearly a case in which inductive risk plays a major role. However, while many of the examples discussed in the literature on inductive risk presuppose that it is clear both what the non-epistemic consequence of certain epistemic decisions are and how to weight them against one another, in real-world cases, the task of assessing and weighing the potential non-epistemic consequences of one's epistemic decisions is much more difficult and complex. In the real world, it is often even a challenge to clearly identify what the potential consequences of error might be. For example, in order to assess the potential consequences of error in the glyphosate case, researchers and advisors would have to answer a number of complex empirical questions, which include questions such as: 'If glyphosate is carcinogenic, how many excess cases of cancer can be attributed to its use?', 'What is the survival rate for the cancers supposedly caused by glyphosate?', 'What treatments do currently exist for those cancers?', 'How effective and expensive are those treatments?', 'Are the alleged carcinogenic effects of glyphosate limited to agricultural workers or do they affect consumers as well?', 'If it mainly affects workers, are there measures that can be taken to reduce their exposure to the product?', 'How effective are the alternatives to glyphosate-based herbicides?', 'How would the use of safer herbicide affect different crop yields?', 'Might bans in richer countries affect food production in poorer countries?', 'What effect would a ban have on domestic and international food prices?', 'Might a ban

¹⁰ Similarly, the Joint Meeting of the Food and Agriculture Organization and the World Health Organization has concluded that it is 'unlikely to pose a carcinogenic risk to humans from exposure through diet' (JMPR 2016, 2). For a discussion of the apparent disagreement between the IARC and these other agencies, see (Tarazona et al. 2017).

on glyphosate contribute to hunger and starvation, especially in countries where a large portion of the population already suffers from hunger and food insecurity?⁷.

While this list of questions is far from exhaustive, it suffices to illustrate what I call *the problem of exogenous inductive risk*. The problem is that, in complex cases such as this one, the assessment of inductive risk requires relying on resources from a large variety of disciplines ranging from the biomedical sciences to the agricultural sciences, to economics. Since the range of resources required ranges well beyond the area of expertise of any individual scientist, this poses a problem for the individualistic approach.

The problem of exogenous inductive risk gives rise to two additional problems. The first is what we might call *the problem of second-order inductive risk*—in order to accept an answer to each of the questions on the list above also requires taking into consideration the inductive risks of doing so. In order to decide whether to accept a certain hypothesis about, say, the mortality rate of non-Hodgkin's lymphoma or the likelihood that a ban on glyphosate would cause hunger and starvation in poorer countries, the scientist who is trying to assess the inductive risk of accepting the hypothesis that glyphosate causes cancer faces inductive risks that are specific to the purpose at hand and these inductive risks are best assessed and managed by the experts of the relevant fields rather than by the scientist who is an expert on the carcinogenicity of glyphosate. For example, oncologists usually face a different sort of inductive risk when deciding whether to accept a certain estimate of the survival rate for non-Hodgkin's lymphoma from the inductive risk relevant to this case. Their acceptance of a certain estimate should take into account how that estimate affects, among other things, the screening and treatment of non-Hodgkin's lymphoma.¹¹ Accepting an overestimate would be much less appropriate in that context than in the context we are currently considering.

The second problem is what we might call *the problem of inductive risk salience*. Researchers who are experts in, say, the carcinogenicity of glyphosate are unlikely to be able to determine what all of the questions relevant to the management of inductive risk are. The danger is that, left to their own devices, the researchers and advisors working in a certain field would only ask the sorts of questions that are most salient to them, while ignoring other important questions that are not equally salient to them but are no less relevant (and would be salient to experts from other disciplines). A particularly important manifestation of this phenomenon is what we might call *asymmetric salience*, which occurs when the more direct and clear non-epistemic consequences of one's decisions are more salient to them than less direct and clear consequences of it. For example, the risk of excess deaths is clear and direct in the scenario in which the experts wrongly reject the hypothesis that glyphosate causes cancer, but it might be much less clear and direct in the scenario in which they wrongly reject it. For example, it might be difficult for an epidemiologist to realize that a ban on glyphosate in a rich country might affect food prices in a food insecure country, thereby causing deaths in that country. In cases such as this, no individual researcher seems to be in a good position to properly assess the inductive risks they face, and this might lead to the mismanagement of inductive risk.

While I do not aspire to offer a detailed solution to these problems, one possible solution is to involve experts from a number of different relevant disciplines as well as the relevant stakeholders in the assessment and management of inductive risk at the advisory stage instead of entrusting the mitigation of inductive risk at that stage solely to those who are experts on the questions that bear directly on the relevant policy decisions. One problem with this solution is that the very nature of the problem suggests that it is not always easy to determine what the relevant disciplines might be. However, in cases in which the stakes are particularly high or in which exogenous inductive risk is particularly likely to arise, a prudent approach would be to include in the advisory committee experts from a wide variety of relevant disciplines (including, possibly, philosophers of science) as well as

¹¹ For an excellent discussion of similar issues in the context of breast cancer prevention, see (Plutynski 2017).

stakeholders. Admittedly, this proposal is, too, just a sketch and it is likely only going to alleviate the problems discussed above rather than avoiding them altogether.

5. CONCLUSION: HOW SHOULD INDUCTIVE RISK BE MITIGATED?

This paper aimed to lay the groundwork for a more systematic exploration of different approaches to the mitigation of inductive risk. After distinguishing between individualistic and socialized approaches and reviewing some of considerations for and against individualistic approaches, I have introduced two new problems for it. However, while, in theory, there are many reasons to prefer a socialized approach, in practice, the responsibility for assessing and managing inductive risks will still often fall on individual scientists (or teams of scientists). Moreover, in many cases, a broadly individualistic approach might actually be the most suitable approach, as the costs of a more socialized approach might often far outweigh its benefits. It is for this reason that, I believe, we need a better understanding of how the responsibility for the mitigation of inductive risk is currently distributed and how we can improve on this by developing better rules and procedures to balance individual and collective responsibilities especially in cases in which the two problems I discussed in this paper are particularly likely to arise.¹²

For example, at the time of this writing, many countries are under more or less stringent lockdown restrictions aimed at reducing the rate of transmission of the SARS-CoV-2 virus. The advice of epidemiologists and public health experts played a crucial role in the adoption of these restrictions. Given the significant non-epistemic consequences of some of the epistemic decisions of epidemiologists and public health experts in this context, this is clearly a case in which inductive risk looms large. While it is still too early to determine whether the relevant inductive risks were properly assessed and managed, this seems to be a case in which the two problems introduced in this paper are particularly likely to arise. In fact, this might be a case in which the two problems compound each other. On the one hand, epidemiologists and public health experts seem to have consistently opted to err on the side of caution—i.e., of overestimating (rather than underestimating) the threat of the virus.¹³ While this abundance of caution might have been understandable in the initial stages of the pandemic, this is also the sort of scenario that is likely to result in a precautionary cascade if one leaves the mitigation of inductive risk to individual researchers and advisors. On the other hand, epidemiologists and public health experts might not be in the best position to fully appreciate the broad range of serious consequences of lockdown restrictions, as many of those consequences do not fall within their specific domain of expertise.¹⁴ This seems to be an instance of what I have called the problem of exogenous inductive risk. For example, epidemiologists and public health experts seem to have underestimated how, in richer countries, lockdown restrictions might exacerbate the phenomenon that some economists call “deaths of despair” (Case and Deaton 2020) or how, in poorer countries, they might contribute to an increase in food insecurity and mortality (see, e.g., (Oxfam 2020)). If this is the case, then this seems to be an instance of what I have called asymmetric salience—the deaths directly caused by the virus are more salient to the relevant experts than those indirectly caused by the lockdown restrictions. Moreover, the two problems seem to compound each other. Asymmetric salience might lead to uniform value judgments, which give rise to a particularly

¹² I would like to thank one of the reviewers for this journal for suggesting this formulation of the issue.

¹³ For a preliminary critical assessment of the two most influential epidemiological models of the spread of SARS-CoV-2, see (Avery et al. 2020).

¹⁴ For similar criticisms of the current response to the pandemic, see, e.g., (Winsberg, Schliesser, and Levy 2020).

problematic kind of precautionary cascade, one which relies on a lopsided assessment of the relevant risks.

Given the high stakes and the likelihood of running into the two problems discussed in this paper, this would have been a case in which it would have been appropriate to adopt a more socialized approach to the mitigation of inductive risk—one in which the advisory committees include not only epidemiologists and public health experts but also experts from a range of relevant scientific fields (from virologists to aerosol physicists and from psychologists to economists) as well as variety stakeholders. However, this is not what seems to have happened in most jurisdictions. While it is too early to determine whether inductive risks have been mismanaged as a result or whether a more socialized approach to the mitigation of inductive risk might have led to better outcomes, the criteria I outlined in this paper suggest that this is a case in which a more socialized approach would have been more effective in mitigating the relevant inductive risks.

ACKNOWLEDGEMENTS

[...]

REFERENCES

- Avery, Christopher, William Bossert, Adam Clark, Glenn Ellison, and Sara Fisher Ellison. 2020. “An Economist’s Guide to Epidemiology Models of Infectious Disease.” *Journal of Economic Perspectives* 34 (4): 79–104. <https://doi.org/10.1257/jep.34.4.79>.
- Biddle, Justin B., and Rebecca Kukla. 2017. “Inductive Risk, Deferred Decisions, and Climate Science Advising.” In *Exploring Inductive Risk: Case Studies of Values in Science*, edited by Kevin C. Elliott and Ted Richards, 215–38. Oxford: Oxford University Press.
- Case, Anne, and Angus Deaton. 2020. *Deaths of Despair and the Future of Capitalism*. Princeton, NJ: Princeton University Press. <https://press.princeton.edu/books/hardcover/9780691190785/deaths-of-despair-and-the-future-of-capitalism>.
- Douglas, Heather. 2000. “Inductive Risk and Values in Science.” *Philosophy of Science* 67 (4): 559–79.
- . 2009. *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh Press.
- . 2018. “From Tapestry to Loom: Broadening the Perspective on Values in Science.” *Philosophy, Theory, and Practice in Biology* 10 (8). <https://doi.org/10.3998/ptpbio.16039257.0010.008>.
- EFSA. 2015. “Conclusion on the Peer Review of the Pesticide Risk Assessment of the Active Substance Glyphosate.” *EFSA Journal* 13 (11): 4302. <https://doi.org/10.2903/j.efsa.2015.4302>.
- Elliott, Kevin C. 2017. *A Tapestry of Values: An Introduction to Values in Science. A Tapestry of Values*. Oxford: University Press. <https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780190260804.001.0001/acprof-9780190260804>.
- Elliott, Kevin C., and Ted Richards, eds. 2017. *Exploring Inductive Risk: Case Studies of Values in Science*. New York: Oxford University Press.
- Elliott, Kevin C., and Daniel Steel. 2016. *Current Controversies in Values and Science*. Routledge.
- Hempel, Carl G. 1965. “Science and Human Values.” In *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York, NY: The Free Press.

- IARC. 2015. *Some Organophosphate Insecticides and Herbicides*. IARC Monographs on the Evaluation of Carcinogenic Risks to Humans 112. Lyon: International Agency for Research on Cancer. <https://publications.iarc.fr/Book-And-Report-Series/Iarc-Monographs-On-The-Identification-Of-Carcinogenic-Hazards-To-Humans/Some-Organophosphate-Insecticides-And-Herbicides-2017>.
- Jeffrey, Richard C. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science* 23 (3): 237–46.
- JMPR. 2016. "Summary Report of the Joint FAO/WHO Meeting on Pesticide Residues." Food and Agriculture Organization and World Health Organization.
- Klotz, Katrin, Wobbeke Weistenhöfer, Frauke Neff, Andrea Hartwig, Christoph van Thriel, and Hans Drexler. 2017. "The Health Effects of Aluminum Exposure." *Deutsches Ärzteblatt International* 114 (39): 653–59. <https://doi.org/10.3238/arztebl.2017.0653>.
- Landemore, Hélène. 2017. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, NJ: Princeton University Press.
- Lidsky, Theodore I. 2014. "Is the Aluminum Hypothesis Dead?" *Journal of Occupational and Environmental Medicine* 56 (5 Suppl): S73–79. <https://doi.org/10.1097/JOM.0000000000000063>.
- Longino, Helen E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.
- Oxfam. 2020. "The Hunger Virus: How COVID-19 Is Fuelling Hunger in a Hungry World." <https://www.oxfam.org/en/research/hunger-virus-how-covid-19-fuelling-hunger-hungry-world>.
- Plutynski, Anya. 2017. "Safe or Sorry? Cancer Screening and Inductive Risk." In *Exploring Inductive Risk: Case Studies of Values in Science*, edited by Kevin C. Elliott and Ted Richards, 149–70. Oxford: Oxford University Press.
- Rudner, Richard. 1953. "The Scientist Qua Scientist Makes Value Judgments." *Philosophy of Science* 20 (1): 1–6.
- Schroeder, S. Andrew. forthcoming. "Democratic Values: A Better Foundation for Public Trust in Science." *British Journal for the Philosophy of Science*, axz023. <https://doi.org/10.1093/bjps/axz023>.
- Stegenga, Jacob. 2017. "Drug Regulation and the Inductive Risk Calculus." In *Exploring Inductive Risk: Case Studies of Values in Science*, edited by Kevin C. Elliott and Ted Richards, 17–36. Oxford: Oxford University Press.
- Tarazona, Jose V., Daniele Court-Marques, Manuela Tiramani, Hermine Reich, Rudolf Pfeil, Frederique Istace, and Federica Crivellente. 2017. "Glyphosate Toxicity and Carcinogenicity: A Review of the Scientific Basis of the European Union Assessment and Its Differences with IARC." *Archives of Toxicology* 91 (8): 2723–43. <https://doi.org/10.1007/s00204-017-1962-5>.
- Winsberg, Eric, Eric Schliesser, and Neil Levy. 2020. "Coronavirus: There Is No Reliable Science Yet." *RealClear Science*, June 19, 2020. https://www.realclearscience.com/articles/2020/06/19/coronavirus_there_is_no_reliable_science_yet_111434.html.