From Ideal to Real Risk. Philosophy of Causation Meets Risk Analysis.

Rani Lill Anjum[1] and Elena Rocca[1*]

[1] Norwegian University of Life Sciences, Ås, Norway.

[*] Address correspondence to Elena Rocca, Norwegian University of Life Sciences, Centre of Applied Philosophy of Science, Pb 5003, 1432 Ås Norway; tel: +47 67231155; elena.rocca@nmbu.no

ABSTRACT

A question has been raised in recent years as to whether the risk field, including analysis, assessment and management, ought to be considered a discipline on its own. As suggested by Aven (2012), unification of the risk field would require a common understanding of basic concepts, such as risk and probability; hence, more discussion is needed of what he calls 'foundational issues'. In this paper, we show that causation is a foundational issue of risk, and

that a proper understanding of it is crucial. We propose that some old ideas about the nature of causation must be abandoned in order to overcome certain persisting challenges facing risk experts over the last decade. In particular, we discuss the challenge of including causally relevant knowledge from the local context when studying risk. Although it is uncontroversial that the receptor plays an important role for risk evaluations, we show how the implementation of receptor-based frameworks is hindered by methodological shortcomings that can be traced back to Humean orthodoxies about causation. We argue that the first step toward the development of frameworks better suited to make realistic risk predictions is to re-conceptualize causation, by examining a philosophical alternative to the Humean understanding. In this paper, we show how our preferred account, causal dispositionalism, offers a different perspective in how risk is evaluated and understood.

KEY WORDS: Local knowledge; causation; risk assessment; norms; practice

1. INTRODUCTION

Risk is traditionally treated as a discipline-specific matter, requiring different tools and approaches depending on the area of application. In recent years, however, it has been suggested that risk fields, including analysis, assessment and management, ought to be considered a discipline on its own (Aven & Kristensen, 2005). Aven urges that in order to develop the risk field as a discipline, we should aim for a common scientific platform and common concepts (Aven, 2012). This, however, requires more discussion about what he calls foundational issues, for instance about the concepts of risk or probability.

A reason why such discussions are important when dealing with risk is offered in a recent paper by Rocca and Andersen, where they show how scientific disagreements about risk evaluations can arise also in cases where there is empirical agreement over the available data (Rocca & Andersen, 2017). When scientists disagree, not over facts, but over their implications, this suggests that their conclusions are based on premises that are not stated explicitly. The example discussed by Rocca and Andersen shows how different understandings of complexity led to different risk evaluations over stacked genetically modified plants in cases when the individual genetically modified plants had already been assessed and found safe. Once these implicit assumptions are made explicit and subject to open and critical discussion, the debate can move forward.

In this paper, we argue that causation is a foundational issue of risk. A proper understanding of causation is thus crucial for risk analysis, risk assessment and risk management. When seeking to identify which exposures contributed causally to a harm, or predict potential harms from an exposure, we are dealing with causal matters. How we understand causation will therefore influence how risk is understood and predicted. This is a conceptual or even ontological discussion, and ought to be philosophically informed (Anjum, 2016; Eriksen, Kerry, Mumford, Lie, & Anjum, 2013).

Recently, it has been proposed that the risk field can only advance if we abandon some old ideas and adopt some new ones (Sexton, 2015). Here we propose that something that should be abandoned are some standard ideas about the nature of causation, and that this can help us overcome certain persistent challenges facing risk experts in the last decade. In particular, we discuss the challenge of including causally relevant knowledge from the local context when studying risk. Although it is uncontroversial that the receptor plays a crucial role for risk, we show that the implementation of receptor-based risk assessment frameworks is hindered by methodological shortcomings that can be traced back to an orthodox view of causation. As we will see, although this orthodoxy has been criticized by a number of contemporary philosophers, standard scientific thinking has been resilient to revising its most fundamental assumptions.

2. WHAT IS LOCAL KNOWLEDGE?

Traditionally, scientific risk assessment is described as consisting of a first phase in which the potential hazards are identified, followed by an analysis of the likelihood that each hazard actually provokes harm. The aim of such assessment is to gain some general understanding about the causal potential of a certain stressor. One might for instance assess whether a pollutant *can* cause a harmful effect in a variety of experimental and observational contexts, and investigate the dose-effect relationship. The next step is to evaluate the significance of this potential in a particular population or context. For this evaluation, it is critical to consider levels of exposure, vulnerability of the exposed system, relevant conditions, and so on, of the local context.

We should note, then, that local context is crucial for risk assessment in two ways. The first is that causal inferences concerning hazards are drawn from multiple evidence, collected under specific contexts: experimental animals (e.g. toxicology), exposed individuals (e.g. pharmacovigilance), communities (e.g. epidemiology), population (e.g. experimental trials) or ecosystems (e.g. ecology). Secondly, these data are then used to make evaluations and predictions for different contexts, for instance about the impact of a certain hazard on a specific individual, economy, workplace, area, community or ecosystem. As such, local-level evidence and knowledge is ubiquitous in the process of risk assessment, both in generating causal knowledge and in applying it. We thus define *local knowledge* as all the available information about (i) the specific settings from which the causal inferences are drawn, and (ii) the contexts to which the causal inferences are applied.

Consider an example. Suppose we want to assess whether the intensive use of neonicotinoids in agricultural practices can have a causal role in the decline of honeybee colonies. This could be investigated using lab toxicological tests, in which single honeybees are exposed to field-realistic dosages of insecticide. We could also perform chronic feeding tests using entire colonies in the field. In the first case, local knowledge would include dosage of the insecticide, lab conditions and variables linked to the tested bees, such as their genetics and health conditions. In the second case, we would need to include additional knowledge concerning the tested bee community, such as local fauna, landscape, nutritional availability, feeding habits, pathogens or exposure to other chemicals. If the experimental data show a slight impact of neonicotinoids exposure on honeybee health, we might want to see how this potential hazard applies to a specific intensive agricultural area. We could then perform a long-term observational study on the beehives surrounding the area. In this case, we should extend *local knowledge* to include also the patterns of exposure, interactions between the insecticide and local factors, interactions within the bee community and with the human communities, bee-keeping practices, et cetera.

What counts as relevant local knowledge will therefore depend on the overall aim of the risk assessment process. Specifically, it depends on the type of evidence we analyze, the level of complexity and the question addressed.

3. LOCAL KNOWLEDGE AND RISK ASSESSMENT: PERSISTENT CHALLENGES

The importance of local knowledge for risk assessment has been stressed increasingly over the last decades, especially for environmental risk (Gallagher et al., 2015; Løkke, 2010; Williams, Dotson, & Maier, 2012). One has moved beyond the assessment of single stressors, outcomes, sources and pathways, acknowledging the crucial role of interactions between stressors and environmental factors for the insurgence of harmful effects in complex systems (Sexton, 2015). Improved risk assessment frameworks are developed explicitly to meet the challenge of evaluating the combined harms from co-exposure to multiple stressors, including chemicals, biological, physical and psychosocial entities. The Cumulative Risk Assessment (CRA) framework, for instance, was proposed by US EPA with the purpose of making risk assessment 'more reliable, realistic and relevant'(U.S. EPA, 2003). One significant novelty in CRA and similar frameworks was that they promoted community-level, local-based risk assessment (Callahan & Sexton, 2007). Official agencies and scientists recognized the importance of focusing on the receptors, and not only the chemical, in order to uncover local processes of harm and time-dependent pathways of toxicity. Overall, more emphasis was placed on high-level, social interactions that can amplify or hinder the effect of a particular stressor, opening up to a multi-disciplinary process of risk assessment. This would include collaborations with social scientists and anthropologists (Checker, 2007; Dendena & Corsi, 2015).

Researchers from environmental toxicology, immune-toxicity, social science and anthropology have urged that genuinely complex phenomena cannot be studied through isolation of causal factors (Abolins et al., 2017; Mesnage et al., 2013; Moretto et al., 2017; Peterson et al., 2003; Sawyer & Loja, 2015; Seok et al., 2013; Suryanarayanan, 2013).

US EPA stated in 2003 that conventional risk assessment needs to advance to include qualitative data, evidence about combined effects, non-chemical stressors and realistic analyses of exposed scenarios (U.S. EPA, 2003). Although some progress has been made in this direction, a definite approach has not yet been developed. Not only are risk assessment processes still judged as inefficient, but there are also more alarming concerns that 'a single unified and comprehensive approach might not be practical, given this complexity'(Moretto et al., 2017). We now present briefly some of the practical challenges that remain unsolved.

How to study causal complexity and interactions through separation, isolation and addition. Current approaches for assessing potential risks from exposure to multiple chemicals prefer

'whole mixture methods'. Indeed, the most accurate way to account for complex interactions, such as synergism and inhibition, as well as for unidentified components, is to assess the toxicity of the whole mixture (Boobis et al., 2011; Callahan & Sexton, 2007). However, this approach is often unfeasible. In real world conditions, the variety of exposure to different types of mixtures is enormous (Løkke, 2010). This forces us to simplify the assessment of toxicity by studying separate chemicals and stressors, and then to combine them following an approximate model. Effectively, this involves an adoption of the principle of additivity as the default assumption. That is, the toxicity data of different chemicals are assumed to combine in an additive way, thereby excluding the possibility of toxicodynamic interactions, such as synergistic effects. We see this in the recent framework of RISK21, which applies the additivity assumption to chemicals with common mechanism of toxicity and low levels of exposure, compatible with real-world situations, in order to identify a 'threshold of toxicological concern' for every specific mixture (Moretto et al., 2017; Solomon et al., 2016). The problem is that this perpetuates the danger of overlooking potential synergistic effects, thus under-estimating the total toxicity of low exposure to multiple chemicals, which was exactly the type of problem that these risk frameworks aimed to overcome in the first place. To this date, the effects of interactions also at low concentrations remains a subject of uncertainty (Boobis et al., 2011).[1]

How to study all relevant causal knowledge through quantitative approaches and measurements. Receptor-based approaches should start from a thorough characterization of the local context for which the risk is evaluated. Such extensive analyses can produce multiple lines of evidence, including qualitative data. These are usually descriptive, and aim to uncover the unique characteristics of certain cases by finding the stories behind a particular person, household or community. For the purpose of local-level risk assessment, qualitative methods might have an epistemological advantage because they allow us to study a local community over time and in detail, including information about local population, their lifestyle, health, diet, occupation, culture, as well as environmental factors. Moreover, qualitative approaches often focus on identifying how the analysis can be affected by the context in which it is carried out. Typically, these approaches emphasize the deep understanding of a single case, rather than the comparison of many cases. A story reported by Checker from a low-income African-American neighborhood might illustrate the dangers of excluding such knowledge: "…in collecting surface soil samples, testers had actually sampled new dirt that residents had imported and put over their old, contaminated dirt… 'They sent out some people to do that testing out here and they scooped a little bit of dirt with spoons on the ground. Hey, I done put dirt on top of dirt trying to get rid of the floods and things we been having out here for years'"(Checker, 2007). This might

[1] A large meta-study from 2011 found studies indicating that synergistic interactions at low concentrations are unlikely to occur. However, the studies focusing on low concentrations were very few and short-term, performed on animal models. Because of the scarcity of evidence, the authors of the study concluded that it is 'too early to draw firm conclusions, particularly for cumulative and low level chronic exposures' (Boobis et al. 2011, p. 381).

seem like a trivial example of bad practice that could easily be avoided. However, no framework has so far been able to genuinely integrate socio-anthropological narrative data with quantitative indicators of environmental risk. Factors linked to host, life-style and environment are acknowledged in the RISK21 framework to modulate the toxic effect of chemical stressors; yet, such modulators should only be included 'if deemed necessary… and the number of variables to consider should be kept to a minimum'(Solomon et al., 2016). While it is recognized that such a study requires 'additional types of scientific expertise', because of practical constraints, it also demands a 'rigorous methodology for a clear delimitation of the analysis so that it is manageable and provides meaningful and applicable results'(Solomon et al., 2016). By demanding such restrictions in order to enable parameter quantification, the challenge remains how not to lose crucial information provided by the case studies.

How to predict propensities using statistics and probability theory. Predictions about risk in specific cases are typically probabilistic and informed by statistical data. How to understand probability, however, is a controversial issue in philosophy, as well as in the risk field (Aven & Reniers, 2013; Flage, Aven, Zio, & Baraldi, 2014). We can illustrate the conceptual tension within probability theory with an example from medicine. Suppose a doctor is evaluating whether to give a certain medicine to a child. This would involve an estimate also of the child's risks of undesired effects of the medicine. In doing this, the doctor will consider the statistical evidence, both for targeted and untargeted effects. But she will also have to assume that the dispositions of the medicine to produce benefits and harms are not purely statistical matters, but caused by some intrinsic properties of the molecule. For instance, the medicine can bind a specific receptor, interact with certain hormones or be metabolized by the liver. Similarly, the susceptibility of the patient to be harmed depends on some intrinsic characteristics, if he carries a rare mutation in that specific receptor, or a liver impairment. In other words, the risk of harm (and benefit) depends on intrinsic propensities of the specific combination drug - patient. Probabilistic and statistical data, on the other hand, are simply mapping the outcomes of drug - patient interactions in a population. Such data might be indicative of the propensities of the specific patient, but this cannot be the default assumption. For instance, a patient with a point mutation in the gene POLG, which codes for the mitochondrial protein DNA polymerase γ, will have a high propensity to develop fatal liver failure as a secondary effect of the anti-epileptic drug valproic acid (Stewart et al., 2010). However, this mutation is rare and therefore not easily detectable using statistical data from the general population, or even specific sub-populations. In order to arrive at the hypothesis of a link between the genetic mutation and the secondary effect of the drug, one could therefore not rely on statistical studies on large populations, but would have to make detailed observations of certain specific cases of harm (Rocca, Anjum, & Mumford, 2018). However, even if we accept that individual or local propensities are not fully described by statistically derived probabilities, we are left at the moment without a valid alternative in the

field of risk assessment, at least not one that is generally acknowledged by the scientific community.

4. HUMEAN INFLUENCE ON SCIENTIFIC METHODOLOGY

The methodological problems described above might seem of a purely practical nature: how to deal with complexity, causation and probability given the available scientific tools. The way in which a scientific enquiry is approached, however, is not only conditional upon technology or knowledge advance, or on the specific evidence available. It is also restricted by a set of norms for "the correct, systematic acquisition of empirical knowledge", which we here refer to as *norms of science* (Anjum & Mumford, 2018a). Note that there are different types of norms in science, for instance concerning ethical practice (e.g., transparency, disinterestedness, communality)(Hansson & Aven, 2014). The norms that we discuss in this paper, however, are restricted to epistemic ones, specifying how scientists should best collect and process empirical knowledge (Anjum & Mumford, 2018a). For instance, evidence based medicine explicitly gives epistemological priority to statistical evidence over mechanistic evidence (Clarke, Gillies, Illari, Russo, & Williamson, 2013). Highest priority is given to randomised controlled trials (RCTs), at least for establishing a causal relationship between treatment and outcome. That medical decisions ought to be 'evidence based' thus means that they should be informed by evidence of a particular type. Similar non-empirical, foundational assumptions can be found in all scientific disciplines, although these are usually much more implicit than they are in evidence based medicine.

Contemporary discussions in philosophy and science have focused on the relationship between theory of causation and scientific practice (Beebee, Hitchcock, & Menzies, 2009; Illari & Russo, 2014; Mumford & Tugby, 2013; Phyllis, Russo, & Williamson, 2011) . The reason why it is important to identify foundational assumptions in science, including causation, is that they shape and restrict every aspect of the scientific process: type of research question, choice of method and study design, interpretation of data and application of the results. The point of making these foundational assumptions in science explicit, is to motivate a rethink of the basic norms of science when practice falls short for our purposes. Only by changing the norms of science can we change scientific practice (Anjum & Mumford, 2018a). The following discussion is meant to prepare the ground for a bottom-up change in risk approaches (from the most fundamental level to the practical one).

ONTOLOGY → EPISTEMIC NORMS → METHODS → PRACTICE

We will now show that the persistent challenges for risk assessment methodology can be traced back to shortcomings of a certain scientific framework. This framework emphasizes use of

correlation data as the starting point for causal analysis, use of quantitative methods and large data sets, use of statistical (Bayesian) tools for prediction and probability understood as frequentism. By uncovering some of the basic ontological assumptions underlying this epistemic-normative framework, we wish to offer an alternative starting-point for discussing the practical limitations that this places on scientific risk assessment. In particular, we want to show how the understanding of causation, first suggested by David Hume (Hume, 1739), has deeply influenced the established scientific norms and practice. Hume stipulated an empiricist criterion of knowledge, that one can only trust as true knowledge what can be traced back to observation.

Causation is linked to regularity. Scientific methodology often starts from correlation data when seeking to establish causation. To be empirically based, causation must then either be treated as something observable, or met with skepticism. Hume did both (Hume, 1739). First, he analyzed causation into three observable features: i) regularity, or the constant conjunction of cause and effect, ii) temporal priority, meaning that the cause precedes the effect and iii) contiguity, the contact between cause and effect in time and space. His skepticism toward making causal claims follows from his problem of induction. Due to the lack of empirical proof of a necessary connection between cause and effect, one cannot predict from the observation of (i-iii) that next time the cause occurs, the effect will follow. Any generality assumed in the causal claim that A causes B is an inductive, hence logically invalid, inference. As a scientific claim, it is therefore empirically unwarranted. The practical consequence of this is that causal theories and predictions are a source of scientific uncertainty. On this view, therefore, correlation data are seen as more dependable for causal conclusions and predictions than theoretical knowledge of causal mechanisms.

Causation must be studied by comparison and in isolation. In risk assessment, one might be interested in predicting the possible harms of a certain stressor, which means we start with the cause and investigate its potential effects. Another way to study causation is to start with an outcome, such as a harm, and try to pin it to a causal factor that produced it (Sexton, 2015). For an empiricist, none of these are easy tasks. Most causal processes include a range of factors, thus one might be mistaken about which factor (or factors) was responsible for the effect. Lewis, a modern-day Humean, suggested that we think of causation as a *sine qua non*, or necessary condition, for the effect (Lewis, 1973). Such necessary conditions might be observed by comparing two causal set ups: one in which the factor is present, and one in which it is absent. In the lab, one will try to isolate the causal factor as much as possible from its usual complexity to see what it does on its own, but this might not always be practically possible. RCTs are one way to do this, the idea being that if there is a difference in outcome between the two situations, it is because of the controlled factor. Only if the cause can be isolated or controlled in this way, can we say to have 'observed' its causal contribution. Without such isolation, we cannot know if some other factor was present that produced the effect instead. This is called the problem of

overdetermination and is a problem for studying causation in its natural complexity (Collins, Hall, & Paul, 2004). From an empiricist point of view, we thus see that causal complexity becomes a challenge for establishing what exactly the causal contribution of each factor is, which means that separation and isolation are necessary steps of the scientific procedure.

Quantitative data gives better science. Scientific studies are typically judged by the quality of their data, which often means the quantity. Even qualitative studies have to report how large the sample is, and how representative it is of its targeted population. Experiments should be repeated, preferably under the same or similar conditions. These criteria for scientific quality are rational from an empiricist perspective. Hume denied that there could be such a thing as a unique causal event, such as the creation of the universe (Hume, 1739, 1748). According to him, for epistemological reasons, causation requires repetition. But because of the problem of induction, we cannot assume that we have enough repetitions to make a correct prediction. For this, we would need a complete set of past, present and future data, which is practically impossible. What we must do instead is to have the largest data set that our resources allow. For the empiricist, we thus see that quantitative approaches, including lots of data and repetitions, will provide better evidence of causation than qualitative or in-depth studies of a few cases.

Causal predictions must be probabilistic. Because we never have complete data, causal predictions are rarely made with 100 percent certainty in science. Many will assume, however, that if one had complete data and scientific knowledge, infallible predictions would be possible. For instance, one might think that the outcome of a coin toss is at least in principle predictable from complete knowledge of the initial conditions. This assumption can be traced back to the philosophical idea of determinism. But even if we assume that determinism is true, causal predictions remain a practical problem. From the empiricist criterion that we can only know what can be observed, the future becomes unknowable to us. This type of reasoning would compel us to only accept scientific predictions as probabilistic claims. Motivated by a commitment to determinism, which states that no events are genuinely probabilistic, such less-than-certain predictions are typically generated within a Bayesian framework of subjective belief (Pearl, 2000; Swinburne, 2002). So, although causal predictions would be infallible under the assumption of determinism, they still have to be probabilistic in order to be scientifically (i.e. empirically) grounded.

Probabilities are statistically derived. Regardless of one's commitment to Bayesian frameworks, statistics plays a central role in science. Statistical tools are used to generate, present and interpret data, and to make predictions. In epidemiology and evidence based decision-making, predictions about a particular case is explicitly derived from statistical data concerning the relevant population. This way of calculating probability can be rationally justified by the philosophical theory of frequentism. Frequentism defines probability as the proportion of the outcome over a sequence of trials. The best method for predicting probabilities for an individual

would then be exactly by observing what happens on population level. Every individual case will, on this view, be a statistical average of its relevant sub-population or 'twin model'. We might even define 'relevant population' as one that makes the individual case an average. Frequentism is an empiricist theory of probability, with probabilities generated by quantitative data.

Returning to the persistent challenges discussed above, we can now see how Humean ontology, epistemic norms and scientific practices are closely related. Starting with the practice, we indicate the underlying norm and show which Humean assumption would motivate it.

Practice: Causal complexity and interactions are studied through separation, isolation and addition. *Norm*: Individual factors must be distinguished from confounders in order to establish their causal contribution. *Ontological assumption*: Causation is a difference-maker, or necessary condition, for its effect. Such difference-makers are observable, but not in cases of overdetermination.

Practice: Causal knowledge is generated through quantitative approaches and measurements. *Norm*: Causal knowledge ought to be confirmed quantitatively, through correlation data, including large-scale comparative studies. *Ontological assumption*: Causation is a type of regularity, only established through repeated observations of two types of events following each other.

Practice: Predictions about individual cases are uncertain, but can be calculated using statistically generated probabilities. *Norm*: Estimates of the likelihood of an effect for an individual case should be informed by what happens in other similar cases. *Ontological assumption*: Probabilistic predictions reflect lack of knowledge, plus estimated degree of subjective belief in the outcome. Alternatively, probability is objective and given by the frequency of an outcome over a sequence of trials.

5. CHALLENGING THE HUMEAN NORMS AND PRACTICES

We said that in order to change scientific practice, we must also change their underlying epistemic norms (Anjum & Mumford, 2018a). How can we do this? To change such norms of science, we must start from different ontological assumptions than the Humean ones. In particular, we must start from a different understanding of causation.

A number of anti-Humean approaches to causation have already been suggested(Armstrong, 1983; Bird, 2007; Ellis, 2001; Heil, 2017; Marmodoro, 2016). Typical for these is that they supplement the regularity theory with some extra modal strength of necessity (Anjum & Mumford, 2018b; Mumford, 2004). Although Hume denied any form of necessity in the world, the Humean and anti-Humean traditions generally agree on most of the observable features

relevant for scientific methodology: law-like regularities, causes as difference-makers, probability understood in terms of frequentism, and so on. It would therefore not be necessary to replace any of the epistemic norms or practices discussed above if we were to replace one of these frameworks with the other. In response to this, a theory of causation has recently been developed by Anjum and Mumford, called causal dispositionalism (Mumford & Anjum, 2011b), based on a neo-Aristotelian ontology of causal powers. This reconceptualization of causation has sparked philosophical debate and the theory has already been applied to other disciplines (Edwards, 2018; Kerry, 2017; Trivino & Nuno de la Rosa, 2016). Application also includes harm detection from medicines and industrial contamination (Rocca & Anjum, 2018; Rocca et al., 2018), and development of a clinical framework for person-centered practice (Evans, Lucas, & Kerry, 2017; Low, 2017). We will now briefly present some features of causal dispositionalism and then go on to show how this ontological approach would challenge some standard scientific norms and practices informed by Humeanism.

Causes are intrinsic dispositions. On the Anjum-Mumford theory, causes are dispositions. A disposition is a certain type of property that can exist unmanifested (Mumford, 1998). Reproduction is the manifestation of the disposition of fertility, but one can still be fertile without reproducing. Dispositional properties are a problematic issue for the empiricist, and Hume denied any knowledge of dispositions, forces or powers (Hume, 1739). For an empiricist, the only reason for believing that something is fertile, fragile or explosive is because we have repeatedly observed similar things reproducing, breaking or exploding. For a dispositionalist, however, causation cannot simply be an extrinsic relation between two types of occurrences. Something counts as a cause only insofar as there is an intrinsic disposition that contributes to the effect. For instance, a medical treatment should have an intrinsic disposition to cure. In the case of placebo, the outcome is the same, but it is produced by a disposition in the patient, not the treatment. Establishing causation between A and B must thus involve establishing that there is such an intrinsic disposition of A to bring about B. This requires theoretical, not only empirical, justification.

Causation requires interaction of dispositions. We often focus on single causes and effects at the time for practical purposes, and take all other factors as background conditions. A dispositionalist would say that any contextual factor that might dispose toward or away from the effect should be considered a causally relevant factor. From this starting point, an effect is never produced by only one factor, or disposition, but by the (often nonlinear) interactions of many. Any effect is thus produced by the interaction of multiple dispositions, or what Martin calls 'mutual manifestation partners'(Martin, 2008). Someone can be fertile, for instance, but this disposition cannot manifest itself without the interaction with an appropriate manifestation partner. When studying the effect of an intervention, therefore, one cannot focus only on what is added to the situation. Instead, one must consider how the intervention interacts with what is already there

as part of the context, that is, which manifestation partners are involved. Ibuprofen has a disposition to cure pain, but it can also cause pain in some people, depending on their intrinsic properties (e.g. genetics). The same type of intervention in two different contexts would then give two different causal processes.

Causation is singular and context-specific. Dispositionalism is a singularist theory about causation, which means that causation happens at a particular time and place that need not be repeated. This contrasts with the regularity theory, which explicitly denies single, unique instances of causation. A regularity theorist must be able to observe that the same type of cause is repeatedly followed by the same type of effect, under some set of ideal, normal or sufficiently similar conditions. For a dispositionalist, a full set of causal conditions are rarely, if ever, repeated. We saw that, from a Humean notion of causation, one cannot claim causation for unique cases. The creation of the universe could for Hume not count as a cause, insofar as it was a unique event (Hume, 1739). Repetition is thus a Humean requirement for establishing causation, but for the dispositionalist this is a purely epistemic matter. Ontologically, there is nothing problematic about a unique case of causation if one is a causal singularist. Rather than trying to repeat a causal process in exactly the same way many times, a dispositionalist would look at what is unique to a specific causal situation and how the various dispositions provided by the context could potentially influence the outcome.

Causal predictions are tendential. Causal tendencies come in degrees, from very weak to very strong. But essential for causal tendencies is that no matter how strongly they dispose towards an effect, the effect can in principle be counteracted by other tendencies. A cause will tend to produce its effect, but it cannot guarantee or necessitate it. A heater that has an intrinsic disposition to heat a room to 25 degrees might still fail to do so if a window is open. Because of the tendential nature of causation, dispositionalists acknowledge that causal predictions are fallible. The primary reason for this fallibility is not that we are unable to predict what type of contribution a disposition will make (e.g. a heater will tend to warm a room rather than cooling it), but that any such causal contribution is irreducibly Tendential (Mumford & Anjum, 2011a).

Causal tendencies are given by individual propensities. Since causal tendencies are dispositional, intrinsic and singular, they are also qualitative. Causal tendencies are thus not the same as statistical tendencies; rather, they are propensities(Mellor, 1971; Popper, 1990). They cannot be revealed simply by observing the incidence of outcome in a population. This has to do with the very nature of such tendencies: they do not always result in a detectable outcome. A tendency can be very common, but too weak to ever manifest itself (e.g. the toxic disposition of cyanide in apple seeds), or it can be very strong, but too rare to be detected statistically (e.g. the anaphylactic shock of vaccines) (Rocca et al., 2018). Propensities must be assigned uniquely to each individual situation, rather than derived from the frequency of outcomes. For instance, the 0.5 propensity of a fair coin to land heads is given by its properties (shape, weight distribution,

etc.), not by the proportion of outcomes in a series of previous trials. Similarly, the propensity of a woman to develop breast cancer is given by her properties (genetics, age, lifestyle, etc.), not by the frequency of breast cancer in her relevant sub-group. On a singularist view, such frequency might indicate the propensity, but it does not generate it (Anjum & Mumford, 2018a).

To sum up, a dispositionalist sees complexity, context-sensitivity, uniqueness, interaction and fallibility as essential to causation(Mumford & Anjum, 2011b). This points to the importance of local, contextual knowledge. To focus only on a single causal factor, intervention or stressor when making predictions about the outcome, will be to disregard most of the causally relevant factors of the causal interaction. From this perspective, frameworks such as CRA, as proposed by US EPA in 2003(U.S. EPA, 2003), seem dispositionalist in nature, at least if we consider the original intentions and motivations. The importance of studying causal complexity, such as combined risk, co-exposure, multiple stressors and sources, was explicitly acknowledged. The possibility of interaction (synergism and antagonism) was used to show the shortcoming of assuming the principle of additivity as default. Instead, nonlinear interaction of various stressors were emphasized (Callahan & Sexton, 2007; Williams et al., 2012). Another shortcoming of the standard statistical approach that CRA was intended to overcome, is the distinction between population risk and the risk of individuals or sub-populations(U.S. EPA, 2003; Williams et al., 2012). Instead of focusing on single stressors, isolated from their complex interactions, and population level risk, CRA would have a focus on locally based risk assessment and individual risk factors, emphasizing the role of community and stakeholder engagement as a valuable source of knowledge (Gallagher et al., 2015a).

Frameworks such as RISK21, promoted as a new way forward for risk assessment, develop the CRA outlines, but without being able to overcome persistent challenges, as we have shown. This limitation comes from an attempt to adapt the original 'dispositionalist intentions' to standard Humean methodology for dealing with causation, complexity and probabilities. If a new way forward is what we are looking for, it ought to start with an honest revolution at the foundations of our intuitions. From a causal dispositionalist starting point, the scientist should prefer a framework that analyses risk by starting from case studies of real-world interactions. And if the renewal is genuine, when faced with methodological challenges, she should not look for compromises with previous practices, for instance by assuming additivity, downgrading results from qualitative studies or consider as few modulating factors as possible in order to maintain the analysis feasible. On the contrary, she should dare to extend the boundaries of scientific thinking and invest resources in a search for genuine alternatives.

6. FROM HUMEAN TO DISPOSITIONALIST METHODOLOGY: TRANSFORMING SCIENCE?

So far, we have proposed that a discussion on the basic assumptions about causation is needed in the risk field, and that this should be done in dialogue with philosophers. We have offered a dispositionalist reconceptualization of causation. We have not tried to defend this theory against other contemporary theories of causation, since this is done elsewhere (Anjum & Mumford, 2018b; Mumford & Anjum, 2011b). Instead we address another question: if we *do* accept the dispositionalist reconceptualization of causation, how could this contribute to a better yet feasible methodological approach?

One might get the impression that dispositionalism makes causal predictions practically impossible, because of the uniqueness of each causal context. We should explain why this is not the case. Although such uniqueness remains a practical challenge for science in general, we think that dispositionalism can give us a better way to make causal predictions of harm for the context of application. In practice, this involves adjusting the prediction in light of detailed knowledge about the dispositions of a certain stressor and the way in which these interact with dispositions and propensities provided by the local context (individual, community, ecosystem etc.). The more we know about these mechanisms of interactions, and which dispositions are involved in the interaction, the better our causal understanding, hence predictions.

Let's illustrate this with an example. While MRI-scanning is generally a safe and effective diagnostic tool, one would still use it with caution in patients with an artificial pacemaker. First of all, this is because of the mechanistic understanding of the interaction between the magnetic field created by the MRI-scanner and the metal implant attached to the heart. But it also requires that one has relevant knowledge about that individual patient. "Relevant knowledge" here refers to those causal mechanisms that will be involved in this particular case. This means that, although there are a number of dispositions and propensities involved, not all of them will be relevant for making a good causal prediction. Each causal set-up is thus in some sense unique, because it will have a unique combination of dispositions, but it is still possible to make causal predictions if we understand how various dispositions interact in different combinations. This, however, requires theoretical and explanatory knowledge that goes beyond regularities and frequencies of occurrence.

Dispositionalism emphasizes that causation is singular and context-specific, and that a causal process might never be repeated in its entirety. Causal knowledge, even of the general type, must then be derived from what can be found in single cases. How would this work? Especially, this seems difficult if a case is very rare or even unique. However, we suggest that detailed analyses of specific cases can contribute to general causal knowledge. Most importantly, the cases that we think are more useful for this purpose are those in which the expected causal outcome fails to happen: deviations or outlier cases, rather than confirmatory ones (Rocca et al., 2018). While confirmatory studies work well to indicate that there is a causal link between an intervention and an outcome, they usually reveal little about how the causal process happens or what might

interfere with it. In contrast, when the expected outcome fails to happen, we can investigate the case in detail to identify causally relevant factors, or interferers.

This reconceptualization of causation would have certain implications for how risk ought to be studied and practiced. For instance, one dispositionalist assumption is that a cause tends to produce different effects, depending on the dispositions with which it interacts. This would motivate a norm stating that contextual variation and heterogeneity should be part of the causal investigation of detecting the dispositions of a stressor. This norm suggests a practice in which causal dispositions of a stressor are studied by observing it in different types of interactions or contexts.

Another dispositionalist assumption is that effects are produced by the interactions of multiple dispositions, or mutual manifestation partners. A norm following this could be that the causal investigation ought to start from an in-depth, qualitative study of a real-life setting, in its full complexity. This norm motivates the practice that assessment of risk starts from a detailed, qualitative study of the contexts of harm.

A third dispositionalist assumption is that causal predictions are about what *tends to* happen, depending on the intrinsic propensities and dispositions that are involved and how they interact. From this, we could postulate that causal predictions of harm should consider which propensities and dispositions are involved in a particular, local context and must be sensitive to epistemic limitation. This norm favors the practice in which causal predictions are based on theoretical understanding of the dispositions and propensities involved, while also identifying which other factors could affect the predicted outcome and how.

How feasible is this approach? Although they are not yet mainstream, scientific approaches already exist that seem dispositionalist in nature. A characteristic of these is that the causal insights come from detailed studies of post-harm cases. For instance, in a study designed to investigate the role of pesticides in the decrease of bee colonies, Chakrabarti and colleagues sampled Indian honeybees (Apis cerana) from populations in the vicinity of an intensively cultivated area and from populations in areas of low or no pesticide intensity (Chakrabarti et al., 2015). Individual bees from the two populations were compared for several morphological factors, and significant differences were found. These results were then corroborated by controlled experiments, where bees from the low cultivation area were fed increasing doses of agro-toxic. Without the detailed analyses of individual bees from a particular setting, it would be impossible to generate the study's working hypotheses. In other words, a single-case study gave a better opportunity to develop deep understanding of the causal process, which again contributed to an understanding of general causal mechanisms. So although the starting-point of the study was singular and local, the causal insights that it generated were general.

Such use of contextual, in-depth studies as a starting point for causal hypotheses, which would otherwise be overlooked using more orthodox methodologies, is not unique. The field of medical anthropology offers several examples. One emblematic example is the study of the causal factors underlying the epidemiological synergy of drug abuse, HIV and other blood-borne diseases (Singer, 1994). Without the help of contextualized case studies of specific communities, the sole epidemiological analysis of AIDS/HIV could not provide a causal understanding of the complexity underlying risk behaviors related to the disease transmission (Loomis Marshall, Singer, & Clatts, 1999).

We see, therefore, that crucial causal insights come from qualitative approaches that give access to rich and varied data about the local context for which harm is assessed. This shows the importance of integrating social science methodology in the scientific risk assessment, from start to finish. Anthropologists have stressed the methodological shortcomings of assessing risk without considering hidden arenas of experience and social interaction when measuring the interaction between chemicals and human (and even animal) health (Checker, 2007; Sawyer & Loja, 2015). In particular, progress in the field of environmental justice increasingly reveals that large epidemiological studies can be deceptive in their estimate of real exposure to chemicals. For a more realistic estimate, one should begin the risk assessment process by in-depth interviews or ethnographical studies of the contaminated context (Arcury et al., 2005). Such post-harm contexts can be found in all risk related areas and should be used more systematically in developing new risk assessment frameworks.

7. CONCLUSION

The aim of this paper has been to encourage an interdisciplinary dialogue between philosophers and analysts, in order to create awareness about some tacit basic assumptions that guide scientific enquiry. We have argued that the first step toward the development of more realistic frameworks for risk assessment is to re-conceptualize causation, by acknowledging philosophical alternatives to the Humean orthodoxy. Of these, we presented our preferred account, causal dispositionalism. Following this conceptual shift, it becomes possible to look for new methodologies that start from a detailed study of the local context of harm and the causal mechanisms involved, and let the insights from this study guide the research hypotheses and study design. In practice, this means that stressor focused and effect focused approaches to risk assessment can inform each other, rather than being separated or restricted by disciplinary traditions(Moretto et al., 2017). Risk assessment would then become a genuinely multidisciplinary effort, where insights and approaches from various disciplines are treated as integrated and mutually informative.

ACKNOWLEDGMENTS

This research was funded by the Norwegian Research Council. We thank two anonymous reviewers for helpful feedback on previous drafts of this article.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

REFERENCES

Abolins, S., King, E. C., Lazarou, L., Weldon, L., Hughes, L., Drescher, P., … Riley, E. M. (2017). The comparative immunology of wild and laboratory mice, Mus musculus domesticus. *Nature Communications*, 8: *14811*, DOI 10.1038/ncomms14811.

Anjum, R. L. (2016). Evidence-based or person-centered ? An ontological debate. *European Journal for Person Centered Health Care*, *4*(2), 221–229.

Anjum, R. L., & Mumford, S. (2018a). *Causation in Science and the Methods of Scientific Discovery*. Oxford: Oxford University Press.

Anjum, R. L., & Mumford, S. (2018b). *What Tends to Be. The Philosophy of Dispositional Modality.* London: Routledge.

Arcury, T. A., Quandt, S. A., Rao, P., Doran, A. M., Snively, B. M., BArr, D. B., … Davis, S. (2005). Organophosphate Pesticide Exposure in Farmworker Family Members in Western North Carolina and Virginia: Case Comparisons. *Human Organization*, *64*(1), 40–51.

Armstrong, D. M. (1983). *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Aven, T. (2012). Foundational Issues in Risk Assessment and Risk Management. *Risk Analysis*, *32*(10), 1647–1656.

Aven, T., & Kristensen, V. (2005). Perspectives on risk: Review and discussion of the basis for establishing a unified and holistic approach. *Reliability Engineering and System Safety*, *90*(1), 1–14.

Aven, T., & Reniers, G. (2013). How to define and interpret a probability in a risk and safety setting. *Safety Science*, *51*(1), 223–231.

Beebee, H., Hitchcock, C., & Menzies, P. (Eds.). (2009). *The Oxford Handbook of Causation.* Oxford: Oxford University Press.

Bird, A. (2007). *Nature's Metaphysics: Laws and Properties.* Oxford: Oxford University Press.

Boobis, A., Budinsky, R., Collie, S., Crofton, K., Embry, M., Felter, S., … Zaleski, R. (2011). Critical analysis of literature on low-dose synergy for use in screening chemical mixtures for risk assessment. *Critical Reviews in Toxicology*, *41*(5), 369–83.

Callahan, M. A., & Sexton, K. (2007). If cumulative risk assessment is the answer, what is the question? *Environmental Health Perspectives*, *115*(5), 799–806.

Chakrabarti, P., Rana, S., Bandopadhyay, S., Naik, D. G., Sarkar, S., & Basu, P. (2015). Field

populations of native Indian honey bees from pesticide intensive agricultural landscape show signs of impaired olfaction. *Scientific Reports*, *5*(July), 12504, DOI 10.1038/srep12504

Checker, M. (2007). "But I Know It's True": Environmental Risk Assessment, Justice, and Anthropology. *Human Organization*, *66*(2), 112–124.

Clarke, B., Gillies, D., Illari, P., Russo, F., & Williamson, J. (2013). The evidence that evidence-based medicine omits. *Preventive Medicine*, *57*(6), 745–747.

Collins, J., Hall, E., & Paul, L. (Eds.). (2004). *Causation and Counterfactuals*. Cambridge MA: MIT Press.

Dendena, B., & Corsi, S. (2015). The Environmental and Social Impact Assessment: A further step towards an integrated assessment process. *Journal of Cleaner Production*, *108*, 965–977.

Edwards, I. R. (2018). Living with complexity and big data. *Uppsala Reports*, *78*, 28–29.

Ellis, B. (2001). *Scientific Essentialism*. Cambridge: Cambridge University Press.

Eriksen, T. E., Kerry, R., Mumford, S., Lie, S. A. N., & Anjum, R. L. (2013). At the borders of medical reasoning : aetiological and ontological challenges of medically unexplained symptoms. *Philosophy, Ethics, and Humanities in Medicine*, *8*(1), 1–11.

Evans, D., Lucas, N., & Kerry, R. (2017). The form of causation in health, disease and intervention: biopsychosocial dispositionalism, conserved quantity transfers and dualist mechanistic chains. *Medicine, Health Care and Philosophy.*, *20*(3), 353–363.

Flage, R., Aven, T., Zio, E., & Baraldi, P. (2014). Concerns , Challenges , and Directions of Development for the Issue of Representing Uncertainty in Risk Assessment. *Risk Analysis*, *34*(7), 1196–1207.

Gallagher, S. S., Rice, G. E., Scarano, L. J., Teuschler, L. K., Bollweg, G., & Martin, L. (2015). Cumulative risk assessment lessons learned : A review of case studies and issue papers. *Chemosphere*, *120*, 697–705.

Hansson, S. O., & Aven, T. (2014). Is Risk Analysis Scientific? *Risk Analysis*, *34*(7), 1173-1183

Heil, J. (2017). Real Modalities. In J. Jacobs (Ed.), *Causal Powers*. Oxford: Oxford University Press.

Hume, D. (1739). *A Treatise of Human Nature*. (L. A. Selby-Bigge, Ed.) (1888th ed.). Oxford: Clarendon Press.

Hume, D. (1748). *An Enquiry Concerning Human Understanding*. (L. Selby-Bigge, Ed.). Oxford: Clarendon Press.

Illari, H., & Russo, F. (Eds.). (2014). *Causality*. Oxford: Oxford University Press.

Kerry, R. (2017). *Causation in Evidence Based Medicine*. Retrieved from http://eprints.nottingham.ac.uk/40494/

Lewis, D. (1973). Causation. In *Philosophical Papers v.2* (1986th ed., pp. 159–213). Oxford: Oxford University Press.

Loomis Marshall, P., Singer, M., & Clatts, M. C. (1999). *Integrating cultural, observational, and epidemiological approaches in the prevention of drug abuse and HIV/AIDS*. Bethesda, MD: U.S. Dept. of Health and Human Services, National Institute on Drug Abuse, Division of Epidemiology and Prevention Research.

Low, M. (2017). A novel clinical framework: the use of dispositions in clinical practice. A person centred approach. *Journal of Evaluation in Clinical Practice*, *23*(5), 1062–1070.

Løkke, H. (2010). Novel methods for integrated risk assessment of cumulative stressors — Results from the NoMiracle project. *Science of the Total Environment*, *408*(18), 3719–3724.

Marmodoro, A. (2016). Dispositional modality vis-a-vis conditional necessity. *Philosophical Investigations*, *39*, 205–214.

Martin, C. B. (2008). *The Mind in Nature*. Oxford: Oxford University Press.

Mellor, D. (1971). *The matter of Chance*. Cambridge: Cambridge University Press.

Mesnage, R., Clair, E., Gress, S., Then, C., Székács, A., & Séralini, G. E. (2013). Cytotoxicity on human cells of Cry1Ab and Cry1Ac Bt insecticidal toxins alone or with a glyphosate-based herbicide. *Journal of Applied Toxicology*, *33*(7), 695–699.

Moretto, A., Bachman, A., Boobis, A., Solomon, K. R., Pastoor, T. P., Wilks, M. F., & Embry, M. R. (2017). A framework for cumulative risk assessment in the 21st century. *Critical Reviews in Toxicology*, *47*(2), 85–97.

Mumford, S. (1998). *Dispositions*. Oxford: Clarendon Press.

Mumford, S. (2004). *Laws in Nature*. London: Routledge.

Mumford, S., & Tugby, M. (Eds.). (2013). *Metaphysics in Science*. Oxford: Oxford University Press.

Mumford, S., & Anjum, R. L. (2011a). Dispositional Modality. In *Deutsches Jahrbuch Philosophie 3: XXI Deutscher Kongress für Philosophie*. Hamburg: Meiner Veralg.

Mumford, S., & Anjum, R. L. (2011b). *Getting Causes from Powers*. Oxford: Oxford University Press.

Pearl, J. (2000). *Causality. Models, Reasoning, and Inference* (2nd ed.). New York: Cambridge University Press.

Peterson, C. H., Rice, S. D., Short, J. W., Esler, D., Bodkin, J. L., Ballachey, B. E., & Irons, D. B. (2003). Long-Term Ecosystem Response to the Exxon Valdez Oil Spill. *Science*, *302*(5653), 2082–2086.

Phyllis, I., Russo, F., & Williamson, J. (2011). *Causality in the Sciences*. Oxford: Oxford University Press.

Popper, K. (1990). *A World of Propensities*. Bristol: Thoemes.

Rocca, E., & Andersen, F. (2017). How biological background assumptions influence scientific risk evaluation of stacked genetically modified plants : an analysis of research hypotheses and argumentations. *Life Sciences, Society and Policy*, DOI 10.1186/s40504-017-0057-7

Rocca, E., & Anjum, R. L. (2018). Why causal evidencing of risk fails. An example from oil contamination. *Ethics, Policy & Environment*, forthcoming.

Rocca, E., Anjum, R. L., & Mumford, S. (2018). Causal Insights for Failure. Post-Marketing Risk Assessment of Drugs as a Way to Uncover Causal Mechanisms. In *Uncertainty in Pharmacology: Epistemology, Methods and Decisions* (Springer). Boston Series for the History and Philosophy of Science, forthcoming.

Sawyer, S., & Loja, N. (2015). Crude contamination: law, science, and indeterminacy in Ecuador and beyond. In H. Appel, A. Mason, & M. Watts (Eds.), *Subterannean Estates: Life Worlds of Oil and Gas* (pp. 126–146). New York: Cornell University Press.

Seok, J., Warren, H. S., Cuenca, A. G., Mindrinos, M. N., Baker, H. V, Xu, W., … Tompkins, R. G. (2013). Genomic responses in mouse models poorly mimic human inflammatory diseases. *Proceedings of the National Academy of Sciences*, *110*(9), 3507–3512.

Sexton, K. (2015). Cumulative Health Risk Assessment : Finding New Ideas and Escaping from the Old Ones. *Human and Ecological Risk Assessment*, *21*(4), 934–951.

Singer, M. (1994). AIDS and the health crisis of the U.S. urban poor; the perspective of critical medical anthropology. *Social Science & Medicine*, *39*(7), 931–948.

Solomon, K. R., Wilks, M. F., Bachman, A., Boobis, A., Moretto, A., Pastoor, T. P., … Embry, M. R. (2016). Problem formulation for risk assessment of combined exposures to chemicals and other stressors in humans. *Critical Reviews in Toxicology*, *46*(10), 835–844.

Stewart, J. D., Horvarth, R., Baruffini, E., Ferrero, I., Bulst, S., Watkins, P. B., … Chinnery, P. F. (2010). POLG determines the risk of sodium valproate induced liver toxicity. *Hepatology*, *52*(5), 1791-1796.

Suryanarayanan, S. (2013). Balancing control and complexity in field studies of neonicotinoids and honey bee health. *Insects*, *4*(1), 153–167.

Swinburne, R. (Ed.). (2002). *Bayes's Theorem*. Oxford: Oxford University Press.

Trivino, V., & Nuno de la Rosa, L. (2016). A causal dispositional account of fitness. *History and Philosophy of the Life Sciences*, DOI 10.1007/s40656-016-0102-5

U.S. EPA. (2003). Framework for Cumulative Risk Assessment. PA/630/P-02/001A. Washington, DC: US EPA, Risk Assessment Forum, Office of Research and Development.

Williams, P. R. D., Dotson, G. S., & Maier, A. (2012). Cumulative Risk Assessment (CRA): Transforming the Way We Assess Health Risks. *Environmental Science and Technology*, *46*(20), 10868–10874.