

THE CAUSAL MOSAIC UNDER CONSTRUCTION: THE EXAMPLE OF EXPOSOMICS

24.1 Making mosaics

A mosaic is an assembly of tiny tiles, all different. When the tiles are carefully composed, positioned and angled, they will make a figure. They can be tiny personal pieces of craft, or vast and stunning public pieces of art, such as the mosaic on the ceiling of the Basilica di San Marco in Venice. In either case, every tile contributes to making the whole mosaic. In this book we have provided the materials to make a causal mosaic, to arrange and rearrange the different accounts of causality.

The accounts of causality. The accounts of causality developed in the literature are the tiles of any mosaic. They are described in Part 2, and there are plenty to choose from, stretching from counterfactuals and agency, to physical process, or INUS conditions. In chapter 1 we explained the five scientific problems of causality: explanation, prediction, control, inference and reasoning. We assessed all of the accounts of causality according to whether they do or do not help with one or more of the five scientific problems. It seems that no one account addresses all five problems successfully. Right from the start, we abandoned the idea that *one* account will on its own provide a full-blown causal theory, allowing us to do everything we might need to do, for all scientific domains. Instead, the various accounts—the tiles—each do *something* valuable.

The five philosophical questions. Part 3 examined different ways in which the tiles—the accounts—can be arranged, according to different perspectives or interests. In chapter 22 we explained the five philosophical questions: epistemology, metaphysics, methodology, semantics, and use. We saw the resulting philosophical theory fragments constructed as answers to these questions. We explained why it is important to clarify the scope of the question asked, to understand what purposes the different accounts of causality serve, to see how they can be used together. One tile may be better placed in epistemology, and another in metaphysics; or the same tile may need to be used in different ways to address different questions. This means that many of the Part 2 accounts can be seen as complementary, addressing different questions, rather than as competing answers to a single question. If we are clear about what question we are asking, it is more likely that we will find complementary notions, from the other fragments, to help us complete our account, and solve our problem or problems. This accords with the idea that the philosophical questions, and the theory fragments built in response to them, are distinct but not independent.

Diversity and pluralism. The idea of using the different accounts in a complementary way may sound good, but it is not simple. In chapter 23 of Part 4 we explained some

of the challenges for recomposing the accounts of causality—challenges for making a mosaic out of these materials. These were the challenges of diversity and pluralism. Worldly causes seem to be very different, we have different sources of evidence, and the methods of the sciences are pretty diverse. What form of pluralism suits a given scientific context is at least partly an empirical question. We need to study the field, its particular problems and methods, and work out what tiles and what fragments we need, so that we can start to build up a mosaic.

The brief recap on accounts (tiles), questions (fragments) and diversity serves as a reminder that there are a lot of materials available in the literature on causality. So how do we put them together? Most of the accounts address more than one of the five scientific problems, and most can address more than one of the five philosophical questions. Further, the different accounts—the tiles—also interrelate. For example, some of them form into the difference-making or dependence family, others into the production family. Alternatively, some are primarily conceptual, some primarily epistemological or methodological—or single accounts have alternative conceptual or methodological interpretations.

This is why we illustrate how we think this complex philosophical literature can be used by scientists and others trying to think about causality with the idea of making a mosaic, because a mosaic is precisely a whole made out of diverse pieces. We acknowledge a debt to Carl Craver's useful metaphor in writing of the 'mosaic unity of neuroscience' in Craver (2007). Russo (2009) also talks about the 'mosaic of causality', of which the epistemology of causal modelling is but one tile. In this chapter we extend these useful metaphors and link them to both philosophical and scientific practice. We will demonstrate how the exercise can be done, using the example for exposomics, an emerging field of science that we will explain. Note that we do not think we are giving the final answer for exposomics, nor for the philosophy of causality—we suggest just *one* useful arrangement of the tiles for exposomics.

24.2 Preparing materials for the exposomics mosaic

Exposomics, or the science of exposure

Exposomics, or the science of exposure, is an emerging area of research in the biomedical field. It draws on biology, epidemiology, environmental science, statistics, bioinformatics, and information and communication technologies (ICTs). The field is new, proposing innovative and challenging methods to study exposure to environmental factors and their effects on several diseases. Exposomics faces many challenges at the conceptual and methodological level and also in making policy recommendations. The hope is that interactions between exposomics scientists and Causality in the Sciences (CitS) philosophers will help address these challenges.³⁶

³⁶At the time of writing, there is a major FP7 project studying environmental exposure and its effects on several diseases: <http://www.exposomicsproject.eu>. The project is coordinated by Prof. Paolo Vineis at Imperial College, London. We are extremely grateful to Prof. Vineis and collaborators for the opportunity to work and discuss with them.

To understand what exposomics science tries to achieve, we need to take one step back to more traditional studies on environmental exposure and disease. This has been done by epidemiologists, who study the distribution and variation of exposure and disease in populations. Traditional epidemiology, and more precisely environmental epidemiology, managed to establish links between environmental factors and classes of diseases. They found strong and stable correlations between categories of factors (such as air pollution, chemicals, etc.) and diseases (such as cancer, allergies, etc.). *How* exactly environmental factors lead to disease is much less understood and the methods used so far will not illuminate that question. This calls for a turn in the methods used in epidemiology, which is precisely what is happening in exposomics. Scientists think that we need to study exposure and disease at the molecular level, in order to understand the molecular basis of life and disease.

To do this, the question of the connection between exposure and disease has to be translated to the molecular level: How to track changes at the micro-molecular level due to levels of chemicals in, for example, air or water? The answer to this question lies, according to exposomic scientists, in the study of *biomarkers*. ‘Biomarker’ means biological marker. A biomarker is a characteristic which is objectively measured and which indicates normal biological processes, or pathogenic processes, or a pharmacologic reaction to a therapeutic intervention. Where do we look for those markers? There are various candidates, for instance metabolites in blood, proteins, or features of gene expression. These can be detected and measured, and so tell us something about what’s going on in the body.

Carrying on the study of exposure and disease at the molecular level is, in a sense, the effect of a much bigger change—a change in the concept of *exposure*. One aspect of exposure is familiar to us: there is stuff ‘out there’ we have contact with and to which our body responds in some way. Exposomic scientists call this the ‘external exposome’. A second aspect of exposure is the novelty in the approach: there is stuff that happens *inside* the body, once we have contact with the stuff out there. But what happens inside depends not just on the stuff out there, but also on the internal environment that our body creates. Exposomics scientists call this the ‘internal exposome’. Now, to understand how, say, pollution leads to allergies, we need to understand both aspects: the internal and the external, i.e. the total, exposome (Wild, 2005; Wild, 2009; Wild, 2011; Rappaport and Smith, 2010).

In practice, how can we study the internal and the external exposome? For the external, that is relatively easy, provided that we can make accurate measurements of the environmental factors. The internal is more complicated, and this is where new statistical methods and new technologies come to the rescue.

After exposure, scientists need to collect bio-samples and analyse them, looking for relevant biomarkers of exposure. But what biomarkers? That is precisely the problem. Exposomics is seeking help from biological theory *and at the same time* is helping refine biological theory (Vineis and Chadeau-Hyam, 2011). Exposomic scientists also look for biomarkers that indicate the presence of a disease. And finally they look for the biomarkers that are in between, the biomarkers that indicate that some clinical change within the body has occurred or is occurring. Exposomics scientists are creating new

cohorts (i.e. groups of people) to carry out these studies, and are also using data available from previous studies, such as the EPIC cohort (European Prospective Investigation into Cancer and Nutrition).

Once the data are in, the issue is how to analyze them. This is where a novel ‘meeting-in-the-middle’ methodology enters the scene. Exposomic scientists try to match the biomarkers that are most correlated with the exposure, with those that are most correlated with the disease, hoping to find a sensible overlap—in the middle—and also hit on those biomarkers that may indicate early clinical changes, the very first indicators of the onset of disease. This requires highly sophisticated statistical modelling and network theory approaches.

But why are exposomics scientists bothering with all this? Well, in a sense this is a response to genomics. Studies of the genome achieved a lot, but still less than what was hoped for (Manolio *et al.*, 2009). In particular, we hoped to gain much more insight about disease mechanisms than we actually did. Exposomics scientists are trying a new venture, in quite a difficult field. If, as it seems, disease is not all in our genes, we need to study how our body interacts and reacts to stuff that may well be responsible for a number of diseases. This is the line of argument of exposomics scientists, who have been quite successful in gaining the support of European funding bodies. The hope is to get a much deeper understanding of disease mechanisms and also to advise policy makers about public health matters.

This gives a quick overview of what exposomics science is, what its aims are, and its position with respect to traditional epidemiology and to genomics. We will now see in some more detail what scientific problems are at stake. In this chapter we illustrate the CitS approach by showing how some of the different concepts and methods discussed in Part 2 play a role in a scientific project dealing with many causal issues. Exposomics enables us to touch simultaneously on epistemological, methodological, and metaphysical issues, and also on issues concerning how we use causal knowledge. It also helps us show how the distinction between difference-making and production can be useful, how probabilistic and variational issues come into play, and so on.

Our aim here is to illustrate how the resources we have laid out in the book can be marshalled to help think about the issues. We lay this out as a series of questions to address. Precisely because these questions will have different answers for different purposes, different people might find different tiles more useful. This is why we think having a guide to them *all* is helpful. We make no suggestion that the illustration we provide here is unique. It is also a timely philosophy. Exposomics is research-in-progress. In a few years, the situation will have changed and the mosaic of causal concepts useful to thinking about it may well also have changed. Here we show how to do the exercise; we don’t say that this is the one mosaic for exposomics, nor the one mosaic for science looking for causes more generally. Since this is one of the cases we both work on, this means we also give hints about what we think would be exciting research to do in the CitS spirit. We are carrying out some of these projects ourselves, but there is so much going on that we hope many other people join the venture.

Question 1: What scientific problems would it be useful to address?

Thinking through the five scientific problems of causality can help identify distinct issues for any scientific case. What are the scientists trying to do? Is there more than one aim, and do they conflict?

Inference. Is there a causal relation between X and Y ? Does X cause Y ? What are the causes of Y ? How much of X causes how much of Y ?

Basically, causal inference concerns finding out about causes. In exposomics, we want to know what the environmental causes of disease are. We want to find links between environmental factors and diseases. We also want to establish links between exposure and some biomarkers and diseases and some biomarkers, so that ultimately we can link up the whole chain from environmental exposure to disease. We need to think about how we can establish such correlations. Where can we get the data? (For exposomics, technologies help a lot, but produce big data sets, which are their own challenge.)

Explanation. How do we (causally) explain phenomena? To what extent is a phenomenon explained by statistical analyses? What level of explanation is appropriate for different audiences, such as fellow specialists, the public, or jurors in a court?

We often want to know not just what happened or will happen, but how it happened, and why. This is causal explanation. In exposomics, the explanations offered by our background knowledge are very important. We are trying to probe an area where little is known, but we depend very heavily on what we already know about the human body, and about environmental causes of disease, found in traditional epidemiology. A lot of our knowledge of the body takes the form of known mechanisms, which explain, for example, how cells make proteins, and how our immune system works. This is important to help us make sense of the many correlations we will find in big data sets. Just correlations won't do. We need a plausible biochemical story of what happens after exposure and inside the body, to think we are starting to make progress on understanding disease causation by environmental factors—at the molecular level.

Prediction. What will happen next, and how do we find out what will happen next? How accurate are our predictions about the evolution of some population characteristics (e.g. mortality, morbidity, ...)? What does a physical theory predict will happen in a given experimental setting?

Prediction is just working out what will happen. In exposomics, we want to know things like what will happen to populations if current levels of environmental contaminants remain the same, or continue to rise. We also want to know whether we can predict disease given knowledge of biomarkers. If we can better predict disease trends, thanks to biomarkers, this will help us decide what public health actions we should take.

Control. How do we control variables that may confound a relation between two other variables? In what cases is it appropriate to control for confounders? How do we control the world or an experimental setting? How do we find out about that? This extends all the way from interfering with a system to get a slightly different result, to the con-

struction of entirely new causal systems, such as smart phones, or GPS, or a political system.

Control is going beyond prediction, to alter what will happen. In exposomics, our core concern is how to prevent disease. But we also want to know how we can understand the pathways from exposure to disease, and thus better control possible confounders, and get a ‘cleaner’ link or pathway from exposure to disease.

Reasoning. What reasoning underlies the construction and evaluation of scientific models? What conceptualization of causation underpins causal methods? How do we reason about all aspects of causality? How can we sharpen up that reasoning, making it more precise and so more effective?

Causal reasoning is our broadest scientific problem, as it concerns all the ways we think about causality in science, explicitly and implicitly. In exposomics, we need to think in terms of causes that hard to find and that are fragile, i.e. whose actions or capacities are easy to disrupt. We need to adjust our rationale for our methods accordingly, because looking only for correlations isn’t enough, and we don’t have a well-developed theory of the causation of disease by interactions of genes and environmental factors within the body to help us much either. How should we model the interrelations between biomarkers?

Note that under causal reasoning is where we worry about relations between the other four problems. In exposomics, we have to think about what existing explanations are available, and what causal inference methods are available, decide how they fall short, and design the methods of the project to meet the needs of the current state of the science. It is here that it becomes useful to ask the most obviously philosophical questions, the conceptual questions. For exposomics, it is useful to wonder: how can we conceptualize the link from exposure to disease? How does the macro-environmental level connect to the micro-molecular level? How am I thinking about causality, so that these links and connections make sense? Does that thinking about causality accord with my methods?

Question 2: What do the scientists want to know or do, and what problems of diversity do they face?

As we explained in chapter 22, philosophical work on causality in the spirit of CitS can begin at any point, from science or from philosophy. So we offer these questions merely as an illustration of how one might choose to approach a scientific case study. This question is useful to ask, to get a sense of the scope of the research you are looking at, and the scale of the problem understanding it poses.

Fragility. In exposomics, scientists are looking for fragile relations, relations which disappear in some contexts. This problem is important to exposomics, but also affects every field where populations are heterogenous, where complex mechanisms exist, where there are homeostatic mechanisms that change and adapt in response to multiple stimuli, and so on. So the question is: how do we isolate such fragile relations? How can we establish that they are causal given their fragility? What are the right methods to circumscribe them? Experiments? Observations? Simulations? What?

Diverse research groups. Exposomics science is not the result of one single brain. Research groups are essential. Exposomics projects usually involve consortia of several institutions, and within each institution there are several groups (epidemiologists, statisticians, etc.). But this is not unique to exposomics. Another example is the discovery of the Higgs boson, which has been possible thanks to numerous research groups working on different experiments. Much of current research needs synergies between different groups and approaches.

Diverse fields of science. Exposomics is (molecular) epidemiology, but it is also medicine, and biomedicine, and biochemical medicine, and medicine using knowledge from nuclear physics, and from statistics, and from sociology etc. Understanding, predicting and preventing disease is a complex enterprise where each of these fields contributes something vital. Again, this is not a problem only for exposomics, but is a feature of many important collaborations nowadays.

Diverse technologies. The embeddedness of technology in modern medicine goes far beyond looking through a microscope. We don't 'see better' using omics technology; we collect data about stuff that we wouldn't see at all otherwise! Actually, we don't 'see' even using these technologies. We are able to identify signals that then need to be interpreted in order to find something in there. The technologies used to generate, collect, and analyse data in exposomics (GPS and sensors, omics technologies, network-based statistics software, etc.) change the landscape of what we can observe and study and of how we understand and conceptualize diseases. How should we conceive of the relation between science and technology in the face of techno-science? What questions does technology pose for the practice and use of medicine? Again, this does not only arise in exposomics. Think for instance of diagnostic tools, that allow us to see *possible* tumours very early and intervene. While this is generally seen as an achievement and is praised, some controversies are also arising, for instance the 'slow medicine' movements.³⁷ New technological tools have the power to enhance the production and analysis of data, and also to 'create' new phenomena, by making diseases appear too early or even when they are merely a possibility. How we should react in response to these technological tools is quite a difficult issue: think about Angelina Jolie undergoing mastectomy to avoid cancer development. Is this ethical? Is this action justified by the evidence available?

Data overload. Big data is a problem for exposomics, as for many other sciences. Having a lot of data, massive datasets, is not a solution *per se*. Big data creates a big problem of analyzing data, of using them, or re-using them for other research questions. In exposomics, for instance, they are re-using samples and data from other cohorts. Maintaining their quality is a serious challenge (Leonelli, 2009; Leonelli, 2014; Illari and Floridi, 2014).

Significant diversity in evidence and data probing methods. One idea in exposomics is that to establish how pollutants induce changes in the body at the molecular level we

³⁷See e.g. <http://blogs.bmj.com/bmj/2012/12/17/richard-smith-the-case-for-slow-medicine/>, accessed 23rd July 2013, and McCullough (2009), who initiated the "Slow Medicine" movement.

use different methods: analyse the exposure (chemistry tools), analyse samples (again chemistry *and* nuclear physics *and* omics tools, etc.), analyse data with statistics and other similar tools (e.g. network theory, calibration methods etc.). Another idea is that data should be available for re-analysis by different people with different tools. Making all this work effectively can be ferociously difficult, and also touches on political issues about accessibility of data, who funds research, who owns the data, and so on.

Clearly there are some significant challenges of diversity in understanding exposomics.

Question 3: What philosophical questions would it be useful to address?

Asking the five philosophical questions can be a useful way to push beyond the surface, to uncover assumptions that may be underlying the research methodologies or questions.

Epistemological issues. What do the Exposomics scientists want to know? What kinds of assumptions about scientific knowledge might they be making? Would the assumptions they are making have a more general bearing on causal epistemology?

Exposomics scientists want to understand how environmental exposure is linked to disease. In that sense, their epistemological question is fairly simple. But answering it is rendered complex by surrounding knowledge. They know a lot about the system, and so many mechanisms within the body are assumed, as are some known mechanisms for environmental causes of disease, such as radiation damaging DNA, and DNA methylation due to smoking. But a lot is still unknown, and in the middle of that exposomics scientists are looking for many small causes, with small effects, and large interaction effects. They expect widely different factors to be causes, including e.g. social and chemical factors. Their use of new omics technologies allows a comprehensive approach, but generates vast amounts of data that must all be stored, maintained, and processed. Naturally, this is strongly linked to the scientific problem of causal inference. But exposomics is also a fascinating lesson in how increasing knowledge happens in such a complicated way, particularly when we are taking the first tentative steps in an attempt to push back the current boundaries of knowledge.

Metaphysical issues. Are there assumptions about the nature of the domain being investigated? What are scientists assuming about the causal relations they are seeking?

In exposomics, there are certainly assumptions, based on known mechanisms described above. Nevertheless, scientists assume that there is something to track in the middle of the vast uncharted territory, whereby particular environmental factors cause disease. Even though it can be hard to get such factors to show up—hard to find the fragile correlations, in the middle of the many misleading correlations likely to exist in large datasets—they assume the causes are there, and can be found. They are interested in how to characterize them.

Methodological issues. What methodology or methodologies are being used? What do the scientists say about why they chose as they did? Are they using a novel methodology? What are the background assumptions of any models or other standard techniques?

As a reaction to the novel methodological challenges of exposomics, the scientists have articulated a new ‘meeting-in-the-middle’ methodology. Broadly, they are trying to track biomarkers of exposure, and biomarkers of early disease onset, and match up biomarkers that correlate with both in the middle, so that ultimately they can track the entire evolution of disease using biomarkers. This new methodology raises interesting issues about the relations between background knowledge and new causal discoveries. What makes us think that the meeting-in-the-middle methodology and use of omics technologies will really advance understanding of disease mechanisms? How do we—and funders—know that this is not mere hype?

Semantic issues. Do the scientists seem to be using causal language in a standard way? Are they innovating? How? What does it mean that macro-environmental factors cause micro-molecular changes in the body? What kind of worldly causation would that be?

The main reason we are personally so drawn to studying exposomics is that the scientists do seem to be innovating. They have explicitly created a new word, the ‘exposome’, by analogy with the ‘genome’ to capture the system-wide nature of the interactions they expect to see. They are then trying to think about causality from within the exposome, while simultaneously struggling with some serious methodological challenges. One thing they do, when talking about causality, is frequently talk about signalling. Do they think this has something to do with the meaning of causality? Is it a way to point to one notion that could cash out causality (see below) or is it a way of avoiding causal talk (like sometimes using the language of determinants instead of causes in biomedical sciences)?

Issues of use. What are the scientists intending to use their results for? What are other people intending to use their results for? How can results be communicated to research funders and policy makers so that actions can be taken? How can the public use the results? How is fine-grained molecular knowledge of disease mechanisms going to be used for policy purposes?

Clearly the work is ultimately intended to inform healthcare recommendations. But it is not obvious how to use exposomics results for policy. For instance, if exposomics results mainly concern relations at the *molecular* level, how can we inform governments about interventions at a *social* level, i.e. trying to change patterns of behaviour, or banning certain kinds of pollutants? How do results about ‘molecules’ translate into actions about ‘people’ and ‘behaviours’?

Again, asking all these questions about a domain of science is an interesting study, and helps us to explore aspects of the work that may otherwise not occur to us. We have argued in chapter 22 that these questions cannot be addressed wholly in isolation. Nevertheless, addressing them all simultaneously is a lot of work. We also suggested that we have to identify which question or questions are of most interest to us, in light of our goals and interests, and keep that question or questions clear in our work.

24.3 Building the exposomics mosaic

We are now in a position to begin building the mosaic itself, beginning by selecting the tiles for the mosaic.

Question 4: What accounts might be useful?

Which of the five scientific problems are most of interest to us? The summary tables of appendix A can be used to identify accounts that might be useful on that basis. Which of the theory fragments do we want to concentrate on? We can use the summary table, and the examples, ‘Core ideas’ and ‘Distinctions and warnings’ of each Part 2 chapter to track down accounts that touch on the scientific problems and philosophical questions we have identified. We selected the following tiles to help us with our thinking about exposomics:

Levels. The ‘exposome’ is a new (causal) concept put forward by epidemiologists in order to redefine the causal context in which causal relations at different levels take place, so chapter 5 on levels of causation may be useful. The issue of the levels here concerns more than the relation between the generic and the single-case level. It also concerns the integration of factors of different natures (social and biological) into the same explanatory framework. So, what understanding of ‘levels’ helps illuminate concerns in exposomics science?

Evidence. In this study the interplay between evidence of difference-making and evidence of mechanisms is crucial to establish causal relations successfully, so chapter 6 could be interesting. Biomarkers of disease are supposed to make a difference to the probability of disease, but this probability raising needs to be substantiated by a plausible underlying mechanism. How are the scientists searching for each?

Production-mechanisms. Exposomics scientists talk about mechanisms, examined in chapter 12, and clearly know a lot about mechanisms such as biochemical mechanisms in the cell and mechanisms such as the immune system in the body. How are they thinking about these things? Note that exposomics scientists also conceptualize the evolution of biomarkers as a process, so how should we think about this with respect to the distinction between difference-making and mechanisms? Is ‘mechanism’ the right way to think about causal linking in exposomics?

Production-information. As well as talking about processes, exposomic scientists talk about ‘signal detection’. Since they are looking for something underlying their data, a kind of linking, it looks like they are seeking productive causality, and they associate this idea with the idea of picking up signals. So it looks like productive causality, in the minds of exposomics scientists, is associated with the concept of information, discussed in chapter 13. So, information, as a production account, should be investigated. Would it suit exposomics? Why? Could it help solve problems elsewhere too?

Capacity. What does it mean that pollutants have the capacity to induce changes in the body at the molecular level? How can we find out exactly what this capacity is? When it is activated, what is its threshold to induce changes? Also, is the predictive power of a biomarker due to its own capacity, or to the capacity of some entity it is a proxy for? Here, the examination of capacities in chapter 14 is of interest, particularly the idea of masking of capacities and the resulting difficulty of getting evidence of capacities when they are sensitive to context.

Depending on what you want to know, and what you already know, there could be material in many other chapters that is useful to you. For example, a background grasp of probabilistic approaches to causality is given in chapter 8; more about what powerful techniques are available for probing data can be found in chapter 7; and a discussion of how diverse sources of evidence might be thought of as contributing to causal inferences is available in chapter 18.

Question 5: How can we put together these resources to help us in our thinking?

Philosophical concepts can't make science easy—*nothing can do that*. What these accounts of causality can do is to sharpen up thinking. This in turn can help the actual practice, already being done, to be done with more clarity, aiding communication among scientists themselves, and outside science and academia more generally.

Now we have accumulated the resources available, we really have to begin choosing tiles for our mosaic, according to which questions are of most interest to our project. Exposomics is such a rich case study that there are many questions of both scientific and philosophical interest. We will just have to pick some to illustrate the process of building a mosaic. Suppose we take the process of discovery as of primary interest, at least to start with. Suppose we are interested in how the conceptualisation of causes can help with the struggle to build a new methodology apt for exposomics. Then we're interested in the methodological challenges. We are interested in how semantics and metaphysics (philosophical questions) might help with constructing a novel methodology (philosophical question and scientific problem) for this case of exposomics. In turn, of course we are interested in how the construction of that novel methodology might inform semantic and metaphysical questions.

We can use Part 2 chapters, beginning with the core ideas and distinctions and warnings, as guides to place the tiles we have identified in our mosaic:

Levels resources. We can watch out for mistakes translating between the philosophical and scientific literatures. Inference between the population level and the single case level is frequently a problem, and we can be alert for mistakes. Often, we can have evidence of one level, and want to know about the other, but be unable to make the inference directly. There are also concerns about different kinds of measurements, integrating them into single model or single explanation, and so on.

Difference-making. This helps us see the difference between worrying about the relation between cause and effect variables, in isolation from worrying about what happens in between. In exposomics, as in any other case which generates a large dataset, there is a problem with finding too many correlations, and requiring some means of isolating those of interest.

Mechanisms. The idea of entities and activities organized to produce a phenomenon, initially in an attempt to explain the phenomenon, might be of interest. Can activities and entities and their organization be found? Can evidence of such things help us? Can they help answer philosophical questions? Can they help with the scientific problems? The problems with mechanisms, properly understanding their context and organization,

do seem to arise here. Context shifts make activities and entities difficult to detect in exposomics, and the organization of many mechanisms here are so complex we are not even sure how to *begin* describing them.

Information. This might offer an interesting way to conceptualize linking in exposomics. Partly, it is interesting that a causal link could be something so thin, so apparently intangible. From the problems discussed in chapter 13, we might be wary that thinking of the link as informational might not be so very informative. On the other hand, it is reasonable to think of information as something that can work with the methodology of exposomics—detect at point *x*, detect at point *y*, detect at point *z*, match up the chain, infer that there's a link right through the system.

Capacities. These help conceptualize something that is stable, at a particular locality, in a particular context. Capacities are useful if you can find them. But they are only useful if they can be found, in spite of context shifts, or if they are stable enough that they only change with known context shifts. What are the known capacities in exposomics? What are their limits?

What we end up with is a mosaic, a picture depicting the methodological problems faced by exposomics scientists, where the tiles of the mosaic have given us the language to express these problems very clearly, yielding a deep understanding of them. We can say how these factors relate, how known mechanisms and capacities structure the problem of exposomics, giving us the background knowledge against which much finer-grained causal links are sought. In this way we can clear the ground, to show the importance of the construal of some kind of linking to causal reasoning, and so to addressing the methodological problems in exposomics. In view of the innovation within the science, and the absence of existing language in the philosophical debate, it can now be no surprise that the scientists are reaching for new language to express this. Here philosophical work can feed directly back into the scientific enterprise by helping provide new concepts, designing a conceptual apparatus that can help support the science.

Note that we are not attempting to say how to write a philosophical paper, but to indicate how to use the philosophical literature on causality to understand a scientific problem and to address it philosophically. This, however, is what it is to do philosophy of science in the CitS style. Our hope is that in creating your own mosaic from the materials we have provided in this book, you end up with something adapted to your questions, your problems, allowing you to move forward with them, and communicate them to others to get any help you need. Not the least of your achievement is to have *identified and refined your questions and problems*.

How many mosaics?

There could be indefinitely many such mosaics. If we choose different questions as our starting-point then, consequently, we collect different tiles, focus on different philosophical questions, and design a different mosaic. Philosophers with different backgrounds might naturally start from a different perspective, focusing on different questions from each other. Likewise, scientists with different backgrounds may approach the

same problem differently. And, clearly, philosophers have taken a different perspective on causality, by focusing on different questions than scientists usually do.

This kind of mosaic might be useful to build not just for a single person—whether philosopher or scientist—but for a research group. For example, exposomics researchers benefit from building a common mosaic, to make their language more precise, and facilitate their communication, and their training of postdocs and doctoral researchers. This ‘mosaic of brains’ has been discussed in the literature as ‘distributed understanding’ (Leonelli, 2014), and also using resources from social epistemology, such as Andersen and Wagenknecht (2013), Fagan (2012), and Beaver (2001).

The *mosaic*, and the understanding of exposomics science that we have developed along with it, is something we have made. This does not mean that there is no reality out there. There is. Reality, to echo Floridi (2008), is a *resource* for knowledge and we interact with that reality in various ways: using evidence-gathering methods through epistemological notions. We don’t merely passively imitate, which sometimes seems to be all that is meant by ‘represent’, causal relations. It takes work to construct our knowledge about causal relations, gathering data, probing the data using advanced methods—and then of course using that knowledge to build things, like social or health policies, which in turn generate more causal knowledge. All this is very active, not at all passive. Exposomics science shows even more detail. We know there is something there, something important to us, but it is very difficult to find. We construct technology, construct research teams, control circumstances so that what we want to find is discoverable. Even after that, data processing is still necessary to find anything. There is a reality out there that we can hit and act upon, study, model and understand, using biomarkers, or conceptualizing disease causation using processes or other notions. But we have to interfere a great deal with the system and constantly design or re-design concepts in order to find out. Ultimately we get causal knowledge by a very sophisticated interaction between us and the world, using reality as a resource to construct our causal knowledge.

To conclude, in this chapter, as in the book generally, we have focused on explaining the philosophical literature on causality to people outwith that literature. From that point of view, we have used this chapter to illustrate how anyone can use the resources of the causality literature to think better about science. But, particularly in Part 3, we also examined issues about how philosophy can successfully engage with science. This chapter can also be read as the finale to chapters 20 and 22, by showing in an extended way how science can be a rich resource for the development of philosophical problems.