

Information and causality

Phyllis Illari and Federica Russo

1 Philosophy of causality meets information

Philosophical theorizing has been concerned at least since ancient Greek thinkers with the problem of connecting events as causes and effects. For Aristotle causes are first principles that explain the ‘why of things’, but they are also ‘efficient’ in that they are the ‘source of change or rest’. In this sense Aristotelian efficient causation is very close to the attempts made by contemporary philosophy of science to give an account of how something gives rise to something else.

Recent debates in philosophy of causality have highlighted that it is one thing to establish *that* C causes E and another thing to establish *how* C causes E. This derives from the work of Hall (2004), who distinguishes two concepts of causation – dependence (that) and production (how) – and is followed up by philosophers interested in analyzing the different evidential components (dependence or association (that) and production or mechanisms (how)) which enter into causal assessment (Russo and Williamson, 2007, Illari, 2011a, Clarke et al., 2014). Recent philosophical literature exploring how C causes E has focused on examining the ways in which mechanisms explain such connections. Here, we will focus on understanding production, which is broader in scope, as will become clear.

Concerning *how* C and E are connected, so far we have two dominant accounts. One is in terms of physical processes, characterized using concepts from physics such as conserved quantities. For instance, there is a physical process explaining how hitting a billiard ball makes it move on the table, involving conservation of momentum. Another account is in terms of mechanisms such as: there is a complex bio-chemical mechanism that explains how

proteins are synthesized, or there are complex socio-economic mechanisms explaining how education affects wealth and vice-versa. In brief, mechanistic explanation of the link between C and E involves finding the parts and their activities by which C brings about E.ⁱ While these approaches certainly have merits, there are many situations in which we would look for a *productive relation* between cause and effect and yet we wouldn't characterize it as either a physical process using quantities from physics, or a mechanism in the sense just sketched.

Suppose you just installed your new smart TV, together with the bluray and the home theatre system. You then try out a DVD, and the image appears, but there is no sound. This *absence* suggests that something went wrong with plugging in the cables between the bluray player and the loudspeakers. But it is not clear how a physical process or a mechanism can connect this cause to the *absence* of sound.

Consider a different case. Doctors fighting an epidemic might reason in a similar way to decide whether they have two separate outbreaks, or a single virus or bacterium that has spread to a distinct population. Epidemiologist John Snow famously stopped the cholera epidemic in London in 1854, arguably by figuring out the 'channels' through which the disease was spreading. To stop an epidemic it is important to understand the mode of communication of the disease. This means understanding how a bacterium (or other agent) spreads, and also how the disease is transmitted from person to person. Snow's innovation was to realize that cholera was being transmitted by water, at a time when the dominant medical theories suggested only two transmission mechanisms, one by touch (contagion) and one by transmission through the air (miasmas). Snow hypothesized poor hygiene in behaviour and living conditions were the main channels for the spread of the disease. He managed to plot cholera deaths and contaminated water by comparing cholera deaths in

different parts of London; it turned out that different water suppliers were active in these neighbourhoods. He managed to convince the authorities to block a suspected water pump and the epidemic gradually stopped (Paneth, 2004). In other words, Snow managed to block what was *linking* different cholera deaths. But this link is not clearly either a physical process or a mechanism.

Snow's question, and the question about the bluray player, are questions about what can cause what; more precisely, these are questions about how C and E are connected, i.e. *causal linking*. As we will show in this chapter, this is reasoning about linking, it is about how cause and effect can – or cannot – be connected, and it seems to be distinct from reasoning about difference-making, which is broadly about plotting variations in one variable against variations in another variable, in abstraction from the explanation for that variation.ⁱⁱ This reasoning is important in daily life, and in science. We will show how current work has turned to giving an account of this in terms of informational linking.

2 Towards an informational account of causality

Hans Reichenbach (1956) and Wesley Salmon (1984) were the first to try to express the idea of tracing linking for causality, giving an account of causality as mark-transmission. A main goal, at least of Salmon's approach, was to distinguish causal processes from pseudo-processes, in the context of physics. Reichenbach and Salmon's core claim can be expressed in simple terms: a process is causal when, if you mark it at an earlier point, the mark is transmitted to later points in the process. So, for example, a moving car is a causal process because, if you mark or dent the side of the car at an early point in the process, the dent will be carried along with the moving car, and will be detectable later on. On the other hand, the car's shadow is a pseudo-process because, if you mark or interrupt the shadow, that kind of mark will not be transmitted, and will not be detectable later.

However, the problem with this approach is that some causal processes cannot be marked without changing the process itself, such as those involving fundamental particles, as any change would profoundly alter the process. Some less delicate processes, such as transmission of bacteria, might be altered by introducing a dye or other marker. So causal processes are not those which actually transmit marks, but those which *would* transmit a mark, if only a mark could be introduced.

The counterfactual characterization of mark transmission, presented in detail in Salmon (1984) was criticized by Dowe (1992) , which led Salmon to reformulate his theory. In Salmon's revised theory, processes are world lines of objects, and causal processes are those that transmit conserved quantities when they interact. These are any quantities that are universally conserved, as described by physical theory (e.g., mass-energy, linear momentum, or charge). Causal interactions happen when causal processes intersect, exchanging conserved quantities, so changing each other. When pseudo-processes meet, such as car shadows falling on each other, no quantity is transmitted, and nothing is changed by such apparent 'interactions'. (See Salmon (1994) and Illari and Russo (2014a, Ch. 11).)

This change solves the original problem, but the new account, often called the 'Salmon-Dowe' account, now lacks the very general applicability of the idea of mark transmission. On the mark transmission view, causal linking is beautifully general, because we can think of so many different kinds of processes as being marked. We can try to alter the signal we think might be interfering with the loudspeakers, and see if the sound they emit changes. We could put floats, or a dye, into a river, and watch to see where the currents take them, to see if the route matches the outbreaks of cholera. The idea of mark transmission applies across many different scientific fields. Indeed, the idea also matches some of the ways we might reason about linking, and try to establish routes of linking. In contrast, the Salmon-Dowe view is set

up using the terms of physical theory, and using examples from physics, but such physical quantities do not seem to be relevant to understanding causality in other sciences or everyday cases.

Nonetheless, it might be possible to redeploy the Salmon-Dowe view of process tracing, making the notion of process more general, applying also outside physics, while still avoiding the key problem for the mark transmission account. In order to reclaim that generality, we need to introduce *information*. Some little-noticed remarks of Salmon actually give us this hint. For example, in his 1994 paper, Salmon (1994 p. 303) comments on his *own* earlier work:

It has always been clear that a process is causal if it is capable of transmitting a mark, whether or not it is actually transmitting one. The fact that it has the capacity to transmit a mark is merely a symptom of the fact that it is actually transmitting something else. That other something I described as information, structure, and causal influence (Salmon, 1984 p. 154-7).

In trying to give an account of causal linking, a major problem is that there is an enormous number of links that we might want to trace, that are of very different types. The examples of causal links that we used above lead us to formulate the question: what do bacteria and signals in cables have in common? The diversity of worldly causal links is recognized by Elizabeth Anscombe (1975), who draws our attention to the richness of the causal language we use to describe different kinds of linking, such as pulling, pushing, breaking, binding, and so on.

It is a real problem to understand what features are shared by cases of causal linking, given how diverse they are. But information theory gives us a very general formal framework that can be used to represent and assess *any* kind of process. Anything can be described

informationally, from a person to a supernova to a tsunami. The formal framework of information theory ensures that the description, in spite of its unprecedented generality, is not empty or vacuous. Information theory itself is part of mathematics (see Chapter 4), but the math gives us new ideas, new ways of thinking we did not have before. The views of this chapter all, in one way or another, hold that the idea of information helps us understand linking. The crude idea is that all these diverse kinds of causal links, energy, radio waves, electrons, bacteria, and bits, are all forms of information. Put this way, all these scientists are asking a version of the same very general question: Can information be transmitted between *C* and *E*? And how? We will also examine how thinking about information *alongside* thinking about mechanisms can help us understand causal linking.

John Collier was probably the first philosopher who explicitly gave an informational account of causality: “The basic idea is that causation is the transfer of a particular token of a quantity of information from one state of a system to another.” (Collier, 1999 p. 215.)

Collier fills this out by offering an account of what information is and an account of information transfer. The account of information is given using algorithmic information theory (AIT), deriving from the work of Kolmogorov (see Chapter 5), to define formally the information in anything, and formalizing ideas of complexity and compressibility (Kolmogorov, 1965, Kolmogorov, 1983). The idea is that something, say a car, is more ‘complex’ than something else, such as a rock, the longer its description *needs* to be: a complete description of a car that cannot be shortened – compressed – without loss of information will be longer than an incompressible complete description of a rock.

The next step for Collier is to give an account of information transfer, to describe a flow of information, which happens over time, such as a moving car. Collier initially describes this in terms of identity of at least some part of the information at the beginning and at the end of the

process (Collier, 1999 p. 222). This is refined in more recent work, where Collier says that an information channel is a family of infomorphisms (Collier, 2011). The idea of an ‘infomorphism’ derives from work by Barwise and Seligman (1997), subsequently refined by Dretske (1999) and Floridi (2010). The covariance model of an infomorphism states that if two systems a and b are coupled in such a way that a ’s being (of type, or in state) F is correlated to b being (of type, or in state) G , then such a correlation carries for the observer of a the information that b is G . For example, the dishwasher’s yellow light (a) flashing (F) is triggered by, and hence is informative about, the dishwasher (b) running out of salt (G) for an observer O , like Alice, informed about the correlation. Collier’s use of infomorphism can be understood in a very similar way, by supposing you have two systems, each consisting of a set of objects, where each object has a set of attributes. For example, a switch has possible attributes on or off, and a bulb also has attributes on or off. If knowing the attributes of the switch tells you about the attributes of the bulb, there is an infomorphism. So in a torch, with the main working components being bulb, battery, switch and case, the information channel is a series of infomorphisms, connecting switch to bulb via battery and case. Of course, knowing the attributes of the switch might not tell you everything about the state of the bulb, as information might be lost.

Collier’s final view is:

P is a causal connection in a system from time t_0 to t_1 if and only if there is a channel between s_0 and s_1 from t_0 to t_1 that preserves some part of the information in the first state. (Collier, 2011 pp. 10-11.)

On this view, information flow is characterized in terms of the identity of information at various stages in the information channel (Collier, 2011 pp. 11-12). Consider Salmon’s example of the dented car. The car is a real causal process, and that is why it transmits marks,

like dents. Collier, though, doesn't have to think in terms of marks that are introduced, like the dent. For Collier, the car itself is an informational structure, and as it moves, that identical item of information exists at each moment of the process. Information, however, can be lost in an information channel, and this is important to thinking about the transmission of cholera by water. We don't need to introduce a mark, as we can think of the bacteria itself in the sewage system as informational. In this kind of case there will be information loss inherent to the system, as not all of the bacteria will be transmitted from the source to a particular downstream town. Some will die, perhaps be eaten, or be diverted; others will reach different towns. Nevertheless, some part will be transmitted, and so we can construe the sewage system as an information channel. Note that when engaged in causal inference, we will usually think in terms of being able to *detect* the relevant informational structure – the bacterium or the car – only at various points in the route of transmission. However, this is about how we gather evidence of transmission. Collier's idea is that there is an informational structure at every point in the process, and part of the information will exist at least at multiple points in the process. This has a great deal in common with Reichenbach's 'at-at' theory of mark transmission, which was also developed by Wesley Salmon (Salmon, 1977, Salmon, 1984, Reichenbach, 1956). According to the 'at-at' theory, a mark is transmitted from *A* to *B* if the mark appears at each point between *A* and *B*. When two processes intersect and undergo modifications that persist after the interaction, that interaction is causal and the processes are also causal, rather than pseudo-processes.

Collier says that a major virtue of his theory is its generality. He has given a view that “applies to all forms of causation, but requires a specific interpretation of information for each category of substance (assuming there is more than one)” (Collier, 1999 pp. 215-6). Collier also claims that his view subsumes other theories of causality, most notably the

Salmon-Dowe conserved quantities view, simply by interpreting the conserved quantities view as limiting the kind of informational connection we find in its domain of application.

2 What problems can an informational account of causality solve?

Recall that the purpose of an account of production is to help us conceptualize causal linking, and understand how it functions in our causal reasoning. This means this chapter focuses on production accounts of causality, which can be seen as complementary to difference-making or variation accounts of causality. The philosophical literature pointed to two problems that beset production accounts: applicability and absences (Schaffer, 2000, Dowe, 2008). Below, we briefly present each and explain how an informational account can help address each of these problems, so deepening our understanding of causal linking.

Applicability is the prime virtue of the informational account, as might be expected as this is what it has been designed to achieve. Previous accounts that bear on causal linking have been the Salmon-Dowe theory, focusing on the exchange of conserved quantities, Reichenbach-Salmon mark-transmission, and the idea of Glennan (1996) that there are causes where there are mechanisms. The informational account is more widely applicable than all three. It does not require the presence of conserved quantities, or the introduction of a mark. It can merge usefully with the mechanistic approach, deepening that account, as we will see shortly. The informational account conceives of the causal linking in a way that can be formally defined in terms of computational information theory. But we do not always have to *specify* the information theoretic structure of a phenomenon. Much of our causal language provides an informal, but meaningful, account for an informational description. This description gives the ‘bones’ of the causal linking, in a way that is applicable to phenomena studied in physics, as well as psychology, or economics. So information is a general enough concept to express what diverse kinds of causal links in the sciences have in common.

The second problem, the problem of causation by absences, has undermined several production accounts. Everyday language, as well as scientific language, allows absences to be causes or effects. Someone apologizing for missing a meeting might say ‘I’m so sorry I wasn’t there, my bus didn’t turn up.’ This intends to claim that the *absence* of the bus caused the person to miss the meeting. Similarly, cerebral hypoxia – i.e., lack of oxygen in the brain – causes brain damage and even death. But how can absences, like missing buses or lack of oxygen, be connected to something else by conserved quantities, or mark transmission, or anything? Absences seem to introduce *gaps* in any causal connection, gaps that traditional production concepts were unable to account for. Schaffer (2004), for instance, argues that causation by absences shows that causation does not always involve a persisting line, or a physical connection. The problem of absences bothered scholars to the point that Dowe (2001) conceptualized them not as cases of genuine cases, but of *quasi*-causation.

The solution to this problem that informational accounts offer is entirely novel. Notice, first, that whether or not you think a gap exists depends on what you think the gap is in. There seem to be no gaps in a table, but if you are considering it at an atomic level, well, then there are gaps. This is what our most advanced physical theories tell us. If you happen to visit CERN in Geneva, stop by the shop; one thing you can buy is a bracelet with the following printed sentence: “The silicon in this bracelet contains 99.9% of empty space”. That we always need to take care concerning what features of the world we are prioritizing, and for what purposes, is a lesson of the Method of Levels of Abstraction (see Chapter 7.) For the purposes of dining, we consider the table at the level of abstraction where it exhibits properties such as solidity and stability, and we do not think there are any gaps in it. For the purposes of physical theorizing, we consider the world at a very different level of abstraction, paying attention to much smaller constituents of the world, and so to many features of atoms, including their non-continuous nature, which then does imply that there are gaps in the table.

Now, information can be transmitted across what, from a purely physical point of view, might be considered gaps. Suppose the person missing the meeting leaves a message for her boss: ‘If I’m not there, it’s because my bus didn’t turn up.’ Then her boss knows about the absence of the bus from her absence at the meeting. Information channels can also involve absences. Recall that a binary string is just a series of 1s and 0s, such as 11010010000001, which can be conveyed as a series of positive signals, and *absences* of a positive signal. Gaps in information-transmission will not be the same as gaps in continuous spacetime. Floridi (2011 p. 31) argues that a peculiar aspect of information is that absence may also be *informative*.

However, it is worth noting that this potential is not fulfilled either by the ‘at-at’ theory of causal transmission of Salmon (1977), nor yet by the closely allied persistence of the identical item of information through multiple places in a process view of Collier (1999). Since they both rely on something persisting at least at some points in a process, merely physical gaps may still interrupt the process, and so seem to break the causal linking, as it is difficult to see how either a mark or an item of information can be continuously transmitted between, say, an absent bus and being late for a meeting. This is in need of future work.

From considering absences, we can see both that information-transmission offers a possible novel account of causal connection, causal linking, and also that a novel account is needed. The persistence of the problem of absences indicates that we have not yet fully understood causal linking. An informational account allows greater flexibility, offering the possibility that the kinds of connections that exist in different domains is an empirical discovery, that can be understood as further constraints on kinds of information transmission discovered there.

The final problem that then arises for the informational account is the problem of vacuity.

There are so many different ways to describe information. The field of mathematical information theory has flourished since Shannon, so there are even multiple formal measures of information. This is important because it yields the applicability that has eluded previous accounts of causal linking. But it might be a weakness if the account is vacuous, if it does not seem to say anything. This might be thought to be the case if there is no one concept of information that is always applied, that can be understood as meaning something substantive.

Alternatively, the rich variety of informational concepts available can be seen as a huge advantage of the informational approach. There are two points worth noting. First, the formal measures of information available, whatever they apply to, however general, are not vacuous. They are also increasingly connected to information-theoretic methods for causal inference. Second, what is *needed* to make any account of causal linking work is something like a light-touch generality. To illuminate our reasoning about linking, we need to be able to see causal linking, in a way that does not obscure the important differences between kinds of causal linking. The informational account offers this, the opportunity to describe – perhaps formally describe – patterns that cannot be described in other ways. Ultimately, the problem of saying something general enough to be widely applicable, while still saying something substantive enough to be meaningful, is going to be a problem for *any* account of production that aims for generality. The challenge that has to be met is precisely to find a concept that covers the many diverse kinds of causal linking in the world, one that nevertheless says something substantive about causality.

In sum, we seem to reason about possible causal linking, and attempt to trace causal links, in many important causal inference tasks in the sciences. Informational approaches to causal

production offer a novel approach to conceptualizing causal linking in a way that assists in this task.

3 How to integrate an informational account into a mechanistic approach

An informational account of causality can be useful to help us reconstruct how science builds up understanding of the causal structure of the world, assisting with the questions of linking we have described. We have seen that traditional accounts of production such as the Salmon-Dowe account do not focus their attention explicitly on *linking*. The core of the mechanisms literature focuses on causal explanation, examining how we (causally) explain natural and social phenomena by identifying the mechanisms underlying them, i.e. identifying their key entities, activities, and organization. The question arises whether such mechanistic approaches, which have been very fruitful in understanding mechanistic explanation, are complementary, or in opposition, to an informational account of linking.

Illari (2011b) and Illari and Russo (2014b) do not attempt to give an account of causality *tout court*. Instead, they seek to give an account only of a part of causality – of production, or causal linking. Ultimately, their guiding idea is that this account will be complementary to difference-making, in that evidence of linking provides further support to evidence such as joint variation between variables. They also argue that an informational account is complementary to mechanistic accounts, helping illuminate the scientific practice and conceptualization of causal linking in the emerging field of ‘exposomics’ research, for example (see below). Broadly, we find mechanisms that help us grasp causal linking in a coarse-grained way. Then we can think in terms of causal linking in a more fine-grained way by thinking informationally. An informational account of causality may also give us the prospect of saying what causality *is*, in a way that is not tailored to the description of reality provided by a given discipline. And it carries the advantage, over other causal metaphysics,

that it fares well with the applicability problem for other accounts of production (processes and mechanism).

Illari (2011b) is interested in how an informational account of causality can be combined with our recent better understanding of mechanisms to solve two problems. The first problem is that the informational account has undeniable generality due to its formal properties. Yet, how can a formal informational account give us understanding of the richness of causal links like ‘binding’, ‘growing’, ‘preying’ or ‘repressing’ in specific domains like biology or psychology (Machamer et al., 2000), or the social sciences (Russo, 2009)? Describing these links informationally allows a very general account, but at the cost of losing rich details that are far too useful to discard.

The second problem is this: when scientists look for a causal link, they often speak of looking for a ‘mechanism’ for the effect. For example, finding mechanisms of disease transmission, which spell out how diseases spread, has been very important. But this raises the question of how we understand mechanisms as causal links. It is widely agreed that mechanisms are activities and entities organized to produce some phenomenon (Illari and Williamson, 2012, Glennan, 2008). But this looks like taking a whole, the mechanism, and breaking it up into parts, rather than linking anything. How should we understand such arrangements of parts as linking cause and effect? Harold Kincaid explains the problem using the terminology of ‘vertical’ and ‘horizontal’ mechanisms (Kincaid, 2011 p. 73). Vertical or constitutive explanations consider a system and explain it by invoking the properties that constitute it and their organization. An etiological or horizontal explanation, instead, considers a system and explains it by invoking the intervening causes (entities and activities) that lead up to some phenomenon. So it is not clear how finding a ‘vertical’ mechanism helps us with causal linking that happens in the ‘horizontal’ mechanism.

The problem of how to understand information substantively enough for it to become meaningful in the special sciences, and the opposite problem of how to understand causal linking in mechanisms are entangled, and it can be difficult to see any solution to both. Illari (2011b) argues that mechanisms (as characterized by the mechanists discussed in Illari and Russo (2014a, Ch. 12)) are the channels through which the information flows. On the one hand, this allows us to integrate causality as information flow in the style of Collier with the rich detail of causal relationships we understand from mechanisms. The functional organization of mechanisms structures, or channels, where information can and cannot flow in many sciences. On the other hand, connecting informational causality to mechanisms can allow us to trace the ‘horizontal link’ – information – across the more familiar ‘vertical’ or constitutive mechanism. This allows us to ally the resources of our understanding of mechanisms to an information-transmission approach to causality. Note that this is in accord with Collier’s view (Collier, 2011 p. 8) .

Broadly, mechanisms are what connect C and E. We can find, study and describe them in science. But we study them so assiduously because they hold together the conditions for certain kinds of information transmission. So building up our understanding of mechanisms builds up understanding of information channels – possible, impossible, probable and improbable causal links. This is what we know of the causal structure of the world. We have come to understand many different specific kinds of linking, from radio waves, to hormone signaling in the human brain, to protein receptors on the surface of cancer cells that can be used to signal to the damaged cell to kill itself. We can think of all these very generally, as forms of informational linking, but we can also categorize the different kinds of information transmission we find. In some cases we can even measure them, although much of the time they will be described more informally, as are the many activities in mechanisms.

Illari and Russo (2014b) try to develop these ideas and pull together other strands of the causality literature, using exposomics (the science of exposure) as an example. Exposomics is an emerging field of research within the health sciences, aiming to push back the frontiers of what we know about the causal role of environmental factors for a number of diseases, for instance cancer or allergies. While traditional epidemiology (notably environmental epidemiology) managed to find stable correlations (or joint variations) between categories of determinants (e.g. certain health conditions, socio-economic status, dietary and various life habits) and categories of disease, molecular epidemiology seeks to find correlations at the *molecular* level. The goal is then to measure levels of chemicals and hazards in water or air *and then* changes in our bodies at different ‘omics’ levels (proteomics, genomics, metabolomics, etc). This way, scientists try to reconstruct *linking* between exposure and disease, reconstructing how disease evolves, from exposure to early clinical changes to proper disease manifestation. But such linking has to be reconstructed from the biological and statistical interpretation of very complex data analyses. In addition, exposomics provides useful insights about how reasoning about mechanisms, processes, and difference-making complement each other. This has been examined by Russo and Williamson (2012), and Illari and Russo (2014b) build on this work. Illari and Russo examine how ideas of causal linking are used in cutting-edge science, particularly when the science is exploring an area with great uncertainty, due to the existence of both known unknowns, and unknown unknowns. Illari and Russo argue that, in this case, while known mechanisms are used in study design, too little is known for the possible causal links to be sufficiently illuminated using known mechanisms. Mechanisms can give some coarse-grained connections, but what is sought is considerably more fine-grained linking. Instead of reasoning about mechanisms, the scientists reach for the language of chasing signals in a vast, highly interactive search space. Here, the level of unknowns means that linking mechanisms are generally unavailable. In the

discovery phase, and possibly beyond it, scientists also need to conceptualize the linking they are attempting to discover in terms of something that can link profoundly inhomogenous causal factors.

Finally, understanding the relationship between mechanisms and information helps us see why one mechanism supports multiple causes, in both the discovery phase and when much more is known. A single mechanism may have more than one function, producing a certain cause effectively, and if the mechanism malfunctions, it may produce one or a few alternative causes reliably, or cease to produce anything reliably at all.

4 Connected debates

A great deal of the history of theorizing about causality is structured by Hume's work. Hume famously denied that we see any 'secret connexion' between causes and their effects – we can only observe effects regularly following their causes in time. Much work in philosophy of science is still in the broadly Humean tradition (Psillos, 2002), although others have sought to find what Mackie (1974) dubbed the 'cement of the universe'. An attempt to give an account of causal linking in terms of information could very well be construed as an attempt to describe the cement of the universe informationally. If construed in this way, it would appear to be a poor attempt.

However, within the Philosophy of Information (PI), giving an account of causality should not be construed in Humean terms, as a search for some elusive causal link. Instead, an informational account of causality, possibly combined with a theory of mechanisms, is very much a *post*-Humean project. Indeed, it can be seen as an attempt to give an account of causality in the spirit of the timely philosophy advocated by Floridi (2011), and apply that account to particular scientific cases such as exposomics science above. Understood as a post-Humean project, an informational account of causality has three aims:

1. Metaphysical: say what causality itself is, starting from interesting cases in science;
2. Epistemological: provide a concept of productive causality that can answer needs that have been recognized in the causality literature;
3. Methodological: provide a concept of productive causality that can answer the needs of scientific cases that present interesting challenges, such as exposomics science.

The boundaries between these aims are permeable and the choice of the labels themselves is also idiosyncratic, as it depends on one's objective (for a discussion, see Illari and Russo (2014a, Ch. 22.) In particular, there is a very thin line between the epistemological and methodological aims; in this context, the emphasis is on the contribution to philosophical theorizing (whence the label 'epistemological') and to scientific method (whence the label 'methodological'). Any of these aims is individually worth achieving. For instance, Illari and Russo (2014b) lay out in detail how they take themselves to meet aim three with regard to exposomics science. Here, we explain the project more broadly within the context of PI.

One important aspect of that context is the consideration of philosophy of technology alongside philosophy of science. To begin with, there are two – somewhat artificial – distinctions that are worth considering as they illuminate current thinking. One distinction places science on one side and technology on the other side; the other distinguishes between the epistemic agent (or knowing subject) and the object of study.

This distinction between science and technology has a lengthy pedigree, with the famous view that science (*epistêmê*, i.e. pure theory) is epistemically superior, while technology (*technê*, i.e. art) is merely a means to 'make crafts' – a view that, by and large, we inherit from Greek philosophy. However, today the situation is quite different. Arguably, without science it would be impossible to build any complex experimental apparatus to examine bio-

specimens in exposomics or to accelerate particles at CERN. *At the same time*, without technology science would not progress at all (Russo, 2012). So the interesting question is not ‘what comes first’ or ‘what is more essential’, but how *techno-science* deeply changes epistemological, metaphysical, and methodological questions, as well as our relation with the world, with ourselves, and among ourselves. In other words, the interesting questions on techno-science are asked from a PI perspective, notably one that takes the fourth revolution as a starting point (Floridi, 2011).

This perspective makes the second distinction – between epistemic agent and object – crumble away. The reason is that scientists are no longer (if they ever were) ‘just’ passive observers of a Nature that stands in front of them. At least since the scientific revolution the scientist, now a *techno*-scientist, is increasingly a *maker*. The techno-scientist makes artifacts, such as computers, software for the analysis of data, particle accelerators, and of course experiments under specific and controlled conditions, etc., but the techno-scientist also makes *knowledge* – i.e. the techno-scientist is a *homo poieticus* (Floridi and Sanders, 2003). We are not passive observers but active learners and creators. This does not necessarily lead to a constructivist position à la Bruno Latour (Latour and Woolgar, 1986) or Isabelle Stengers (Stengers, 1993), but instead leads to a *constructionist* position, according to which we ‘shape’ the objects of inquiring by studying them, *and* the objects of inquiry constrain knowledge construction (Floridi, 2011). So this is not a traditional realist position, but it is not an antirealist position either, as it does *not* deny reality. What it undermines is the view that reality is totally other, detached from us, and in this sense the position is neo-Kantian in spirit (Floridi, 2011).

The relation between science and technology, and between the epistemic agent and reality, have a bearing on questions about causality. In fact, this very active process of construction

and reconstruction is an accurate representation of exposomics science as it is practiced. Exposomics scientists go to great lengths to construct the links between exposure and disease. On the one hand, they need to find the right ‘intermediate’ biomarkers, the ones that are linked to exposure and to disease. On the other hand, they need to place this reconstructed link into a plausible network of relations. Whether scientists hit upon the right intermediate biomarker will be theoretically justified to the extent that the complex (internal) biochemical mechanisms also include that biomarker. This means that linking cannot be seen with the naked eye, nor using experimental set-ups, and not even with found correlations. Instead, linking is *reconstructed* by putting together the many pieces of the evidential puzzle. And it is scientists who carry out this work of reconstruction. This requires much empirical evidence and a great deal of interpretation of the evidence using the right concepts. The thought is that information is precisely one concept needed to do that. It is worth noting that this problem is not specific to exposomics science. It is shared by experimental and observational methods alike. In fact, any scientific conclusion is the result of a reconstruction and interpretation of evidence.

So, more generally, any causal claim derived from techno-scientific research will be the scientists’ interpretation of very many pieces of the ‘evidential’ puzzle. It will be a reconstruction of information coming from experimental analyses, plus statistics, plus biological or physical theories, for example. It will be an a posteriori reconstruction of data- or technology-driven research. In this context, informational thinking helps with conceptualizing production (the linking) as the evolution of biomarkers, from exposure to early clinical changes, to disease.

Against this background, giving an account of causal linking in terms of information is not to give an account of the cement of the universe, as Mackie originally construed it, nor is it an

attempt to present the hidden nature of causality. Indeed, the project that Reichenbach and Salmon engaged in was already different from this, as in various ways they attempted to spell out what we were learning from science about causality. But these earlier production accounts did not explicitly include any epistemic agent in the process of finding, conceptualizing, or using linking. They still saw Nature as separate from the techno-scientist or any other person. However, from a PI perspective, the epistemic agent is an integral part of the process of finding and conceptualizing causal linking, and the ‘poietic’ practices of the scientific community craft our knowledge, builds the technology we need to test it against reality, and then crafts the artifacts our enhanced knowledge allows us to make, testing that knowledge again in changing our lives.

It is in this context that Illari and Russo (2014b) argue that information is the most general possible characterization of causal production or linking. It provides a very general concept of causal linking, and a “lite” metaphysics of causal production which can be widely applicable. If informational linking helps in these complex poietic practices, then informational linking is as real as it needs to be. The informational structural realist approach (see chapter 18), in so far as it is also embarked on a project of understanding the world informationally, is in the same spirit. Informational structural realists share the wish to identify generalities in the post Stanford School age of pluralism (Cartwright, 1999, Dupré, 1995), where general concepts are unfashionable.

Structural realism is a view in the scientific realism debate that says that what is real, what science ultimately tracks through time, is the fundamental structure of the world. It is this structure that is described, for example, in the mathematical expressions that are so important to physical theory. In their theory, Ladyman and Ross (2007) set out an extended attempt to explain how structural realist ideas, originally developed in the philosophy of physics, can

actually be extended into the – prima facie very different – special sciences. They are explicit about their reasons for using informational language, and about their influences:

As we noted at the top of the chapter [chapter 4], special sciences are incorrigibly committed to dynamic propagation of temporally asymmetric influences – or, a stronger version of this idea endorsed by many philosophers, to real causal processes. Reference to transfer of some (in principle) quantitatively measurable information is a highly general way of describing any process. More specifically, it is more general than describing something as a causal process or as an instantiation of a lawlike one: if there are causal processes, then each such process must involve the transfer of information between cause and effect (Reichenbach, 1956, Salmon, 1984, Collier, 1999); and if there are lawlike processes, then each such process must involve the transfer of information between instantiations of the types of processes governed by the law. (Ladyman and Ross, 2007 pp. 210-11.)

So they are clear that generality is an important reason for using informational language. For Ladyman and Ross, as for Collier, the idea of compressibility is important to their theory, which they call ‘information-theoretic structural realism’.

As of today, the only other major informational structural realist is Luciano Floridi (2011), who of course situates his work explicitly within the philosophy of information. Floridi’s motivations are in some ways quite different from those of Ladyman and Ross. He uses informational language in a neo-Kantian effort to describe what we know of the world, with the minimal metaphysical commitments possible. Again, though, it is the generality of informational language, in this case allied to its minimal commitments, that is so attractive.

Neither Floridi, nor Ladyman and Ross, are trying to address the issue of causal linking. Nevertheless, they are trying to argue for a view about the nature of the world, and in that

sense they are offering a metaphysics, as well as a conceptualization, using informational concepts. But the driving aim is generality, to describe different things in a way that illuminates what they have in common, in a minimal way. This is not to deny the differences, but to describe things at a level of abstraction (see chapter 7) that is appropriate for some purposes. The description will only be useful if it does capture some features of the world – what Ladyman and Ross call ‘real patterns’. So this informational metaphysics, and the informational account of causality, is minimally realist. Thinking of causality informationally captures useful generalities, generalities that can illuminate our causal reasoning. It does not describe the hidden nature of causality, or the ‘cement of the universe’. Rather, it makes the process of knowledge construction explicit, showing how general concepts such as a concept of informational linking function in this process. If, in cases like exposomics, thinking of the link informationally is the best way to describe what is sought – and found – then we have the best possible reason to think that link is real, and is informational.

Further reading

Salmon (1994) provides Salmon’s own reassessment of his earlier mark-transmission theory, and his shift to the conserved quantities view. Collier (1999) is a good introduction to his approach to informational causality. Illari (2011b) explores the aims of an informational account of productive causality, while Illari and Russo (2014b) apply such an approach in detail to the emerging scientific field of exposomics.

Related topics

Related chapters include:

16. Bayesianism and information, Jon Williamson & Michael Wilde

18. Informational Metaphysics (the informational nature of reality), Terry Bynum

21. Philosophy of Science and Information, Ioannis Votsis

23. The Philosophy of Biological Information, Barton Moffatt

References

- ANSCOMBE, G. E. M. 1975. Causality and Determination. In: SOSA, E. (ed.) *Causation and Conditionals*. Oxford: Oxford University Press.
- BARWISE, J. & SELIGMAN, J. 1997. *Information Flow: The Logic of Distributed Systems*, Cambridge, Cambridge University Press.
- CARTWRIGHT, N. 1999. *The Dappled World: A Study of the Boundaries of Science*, Cambridge, Cambridge University Press.
- CLARKE, B., GILLIES, D., ILLARI, P., RUSSO, F. & WILLIAMSON, J. 2014. Mechanisms and the evidence hierarchy. *Topoi*, online first.
- COLLIER, J. 1999. Causation is the transfer of information. In: SANKEY, H. (ed.) *Causation, natural laws, and explanation*. Kluwer.
- COLLIER, J. 2011. Information, Causation and Computation. In: CRNKOVIC, G. D. & BURGIN, M. (eds.) *Information and Computation: Essays on Scientific and Philosophical Understanding of Foundations of Information and Computation*. World Scientific.
- DOWE, P. 1992. Wesley Salmon's process theory of causality and the conserved quantity theory. *Philosophy of Science*, 59, 195-216.
- DOWE, P. 2001. A counterfactual theory of prevention and 'causation' by omission. *Australasian Journal of Philosophy*, 79, 216-226.
- DOWE, P. 2008. Causal Processes. In: ZALTA, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*.
- DRETSKE, F. 1999. *Knowledge and the Flow of Information*, Cambridge, Massachusetts, Cambridge University Press.
- DUPRÉ, J. (ed.) 1995. *The Disorder of Things*, Harvard: Harvard University Press.
- FLORIDI, L. 2010. Semantic Information and The Correctness Theory of Truth. *Erkenntnis* 74, 147-175.
- FLORIDI, L. 2011. *The philosophy of information*, Oxford, Oxford University Press.
- FLORIDI, L. & SANDERS, J. W. 2003. Internet Ethics: the Constructionist Values of *Homo Poieticus*. In: CAVALIER, R. (ed.) *The Impact of the Internet on Our Moral Lives*. New York: SUNY.
- GLENNAN, S. 1996. Mechanisms and the nature of causation. *Erkenntnis*, 44, 49-71.
- GLENNAN, S. 2008. Mechanisms. In: PSILLOS, S. & CURD, M. (eds.) *Routledge Companion to the Philosophy of Science*. Routledge.
- HALL, N. 2004. Two concepts of causation. In: PAUL, L. A., HALL, E. J. & COLLINS, J. (eds.) *Causation and Counterfactuals*. MIT Press.
- ILLARI, P. 2011a. Mechanistic evidence: Disambiguating the Russo-Williamson Thesis. *International Studies in the Philosophy of Science*, 25, 139-157.
- ILLARI, P. 2011b. Why theories of causality need production: An information-transmission account. *Philosophy and Technology*, 24, 95-114.
- ILLARI, P. & RUSSO, F. 2014a. *Causality: Philosophical theory meets scientific practice*. Oxford: Oxford University Press.
- ILLARI, P. & RUSSO, F. 2014b. Information Channels and Biomarkers of Disease. *Topoi*, Online first.
- ILLARI, P. & WILLIAMSON, J. 2012. What is a Mechanism?: Thinking about mechanisms across the sciences. *European Journal of the Philosophy of Science*, 2, 119-135.

- KINCAID, H. 2011. Causal Modeling, Mechanism, and Probability in Epidemiology. In: ILLARI, P. M., RUSSO, F. & WILLIAMSON, J. (eds.) *Causality in the Sciences*. OUP.
- KOLMOGOROV, A. N. 1965. Three approaches to the quantitative definition of information. *Problems of Information and Transmission*, 1, 1-7.
- KOLMOGOROV, A. N. 1983. Combinatorial foundations of information theory and the calculus of probabilities. *Russian Mathematical Surveys*, 38, 27-36.
- LADYMAN, J. & ROSS, D. 2007. *Everything must go*, Oxford, Oxford University Press.
- LATOUR, B. & WOOLGAR, S. 1986. *Laboratory Life: The Construction of Scientific Facts*, Princeton, Princeton University Press.
- MACHAMER, P., DARDEN, L. & CRAVER, C. 2000. Thinking about Mechanisms. *Philosophy of Science*, 67, 1-25.
- MACKIE, J. L. 1974. *The cement of the universe. A study on causation*, Oxford, Oxford University Press.
- PANETH, N. 2004. Assessing the Contributions of John Snow to Epidemiology. 150 Years After Removal of the Broad Street Pump Handle. *Epidemiology*, 15, 514-516.
- PSILLOS, S. 2002. *Causation and Explanation*, Acumen Publishing.
- REICHENBACH, H. 1956. *The Direction of Time*, University of California Press.
- RUSSO, F. 2009. *Causality and causal modelling in the social sciences. Measuring variations*, Springer.
- RUSSO, F. 2012. The *Homo Poieticus* and the Bridge between *Physis* and *Techne*. In: DEMIR, H. (ed.) *Luciano Floridi's Philosophy of Technology: Critical reflections*. Dordrecht: Springer.
- RUSSO, F. & WILLIAMSON, J. 2007. Interpreting causality in the health sciences. *International Studies in Philosophy of Science*, 21, 157-170.
- RUSSO, F. & WILLIAMSON, J. 2012. EnviroGenomarkers. The interplay between difference-making and mechanisms. *Medicine Studies*, 3, 249-262.
- SALMON, W. C. 1977. An "At-At" Theory of Causal Influence. *Philosophy of Science*, 44, 215-224.
- SALMON, W. C. 1984. *Scientific explanation and the causal structure of the world*, Princeton, Princeton University Press.
- SALMON, W. C. 1994. Causality without counterfactuals. *Philosophy of Science*, 61, 297-312.
- SCHAFFER, J. 2000. Causation by Disconnection. *Philosophy of Science*, 67, 285-300.
- SCHAFFER, J. 2004. Causes need not be physically connected to their effects: The case for negative causation. In: HITCHCOCK, C. (ed.) *Contemporary debates in Philosophy of Science*. Oxford: Blackwell.
- STENGERS, I. 1993. *L'invention des sciences modernes*, Paris, La Découverte.

ⁱ In philosophy of science there is a lively debate on the concept of mechanism. Here, we adopt the proposed consensus definition of Illari and Williamson (2012) "A mechanism for a phenomenon is composed of entities and activities organized so that they are responsible for the phenomenon". The debate on mechanisms rapidly expanded from biology to many other disciplines, including psychology and neuroscience and the social sciences. (Illari and Russo, 2014a, Ch. 12.)

ⁱⁱ For a detailed discussion of variational reasoning in causal methodology see Russo (2009) and Illari and Russo (2014a, Ch. 16).