The metaphysics and epistemology of causal production: the prospects of variation to trace the transmission of information

Phyllis Illari and Federica Russo

Abstract.
In the sciences and in everyday talk, language is replete with cases in which we, epistemic agents, are interested in how causes produce effects, in a variety of contexts and situations. In this paper, we develop the account of causal production as information transmission which we initially presented in previous work. Specifically, in this contribution we explain the role of causal production within the 'mosaic' of causal theory, and we explain how a metaphysics for causal production as information transmission is in need of a set of epistemological strategies to trace whether information is transmitted or not; such epistemological strategies can be in terms of 'variation' as well as 'information transmission'. We explain the 'ontoepistemological' combination of information transmission and of variation with the aid of four episodes of causal production.

1. Why bother with causal production?

The sciences and everyday life are replete with cases in which we, epistemic agents, are interested in how causes *produce* effects. We want to know how SARS-COV-2 produces an inflammatory response, how particles interact at the subatomic level, how being late will make you miss your train, or how failing to properly plug in your headset will result in not hearing much of the videocall you are in.

It has already been noted in the literature that causal talk across the sciences and everyday life is very diverse; an argument made at least since Anscombe (1975), and repeated by a number of scholars in the philosophy of causality (Nancy Cartwright 2004; Weber 2007; Longworth 2010; Godfrey-Smith 2010). In fact, precisely because causal language and methods are so diverse, pluralistic strategies have been tried in many places. Just to confine the discussion to contemporary philosophy of causality, notable pluralists are Nancy Cartwright (2004), who famously pointed out that one word ('causation') in fact means many things, or Erik Weber (2007), who argued that different concepts of causality (probabilistic and process-based) are suited to different scientific contexts, and of course the whole group of scholars who, in the past two decades, developed pluralistic approaches to evidence for causal claims (see e.g. (Russo and Williamson 2007; Campaner 2011; LaCaze 2011; Clarke et al. 2013; Reiss 2015; Parkkinen et al. 2018; Pérez-González and Rocca 2021)).

In previous work, we have also defended a qualified version of causal pluralism, the *causal mosaic* approach (Illari and Russo 2014). We think causality is not to be reduced to one philosophical question or one scientific problem. We have already presented our causal mosaic and will not rehearse our arguments in full here. This is just to signal that our account of information transmission as causal production, which is the object of this chapter, is part of a much broader, pluralistic, perspective on causality and not an attempt to provide yet another monistic account. We have also presented and defended information transmission in previous work as the most general account of causal production (Illari 2011; Illari and Russo 2016; Vineis, Illari, and Russo 2017; Vineis and Russo 2018). In this chapter, we will not rehearse our argument, but provide some further qualifications of the account. In line with our pluralistic approach to causality, we hold that the account of causal production in terms of information transmission that we are developing needs to be *added* to the library of useful causal concepts. We also think we need a concept of production (information-transmission), *and* a concept of difference-making (variation), which can usefully classify some of the useful accounts of causation, although not all of them. More specifically, in this chapter, we focus on causal production from a metaphysical perspective, and explore its corresponding epistemology. Note that we think the *concept* of variation is equally interesting, but here we focus on the *epistemological strategies* of variation, and their relation to information-transmission as production.

Our distinctive approach here is practice-based on the one hand, and from an agent's perspective on the other hand (following Russo (2022)). In asking what causal production is, we are interested in a metaphysics that can be widely applicable across scientific domains. Importantly, though, this metaphysics is not a priori, but is always the product of an agent's perspective, and in this sense our approach aligns with perspectivism (Giere 2006; Massimi 2022), constructionism (Floridi 2011a; 2011b), and ontoepistemology (Barad 2007) and aligns with the lengthy discussion of Russo (2022). In Illari (2011b) and Illari and Russo (2016b) and Russo (2022) we formulated a number of desiderata, which we adopt but also further elaborate as follows.

A concept of production should:
- [scientific domains:] make sense across sciences, including physics, social sciences, life sciences, and particularly for cases of causal relations *across* these levels.
- [levels:] help us understand causal relations across micro and macro causes (and vice-versa) and across factors of different natures (sometimes called 'inhomogeneous variables').
- [technology:] be able to return a meaningful metaphysics for highly technologized contexts, in which there is arguably an important element of *construction* (so causal relations are not in any *simple* way 'out there').

The chapter is structured as follows. In section 2, we present information transmission as a thin and general metaphysics of production, with advantages and complementarity with respect to other production accounts. We explain a little further that our approach to metaphysics and ontology is in line with constructionism Floridi (2011b), and ontoepistemology approach of Barad (2007), and Russo (2022). This means we need to understand how human epistemic agents *come to establish* that there is transfer of information. In section 3, we will explore the epistemology of information-transmission. In particular, we will distinguish between what we call 'information-transmission' and ‘variation’ epistemological strategies. We will show that to know about information transmission requires both epistemological strategies that seem allied with our concept of production (such as mark transmission and process tracing) and epistemological strategies that seem allied with our concept of difference-making (variational strategies such as observational studies, studying variation across similar and different things). We will discuss how these strategies work in practice for four examples, highlighting the central role of epistemic agents. We will argue, however, that variational epistemic strategies are typically needed to establish even information-transmission, indeed most successful epistemologies use mixed strategies. In section 4, we will draw some more general conclusions from our examples. The chapter as a whole will show that we need a metaphysics of information-transmission as production, but that the *epistemology* of production that accompanies it is rather complex, and does not reduce to information-transmission. In general, good epistemology tends to use mixed strategies. Also, many of these epistemologies require a whole variety of instruments (from statistical software to lab equipment, or more simply our perceptual apparatus), but we place human epistemic agents in a central position to explain how these epistemological strategies work. In this paper, for the first time, we begin to put two sides of our work together explicitly, the causal mosaic and the information transmission account of production, and develop the epistemology of production that is at work in tracing the transmission of information.

2. Information transmission as a thin general metaphysics for causal production

2.1 Causal mosaic

In previous work, we carried out a substantial review and systematization of the literature on causation in the past 60-70 years and showed that the project of finding *The One* concept of causation that fits *any* scientific context whatsoever is highly likely to fail (Illari and Russo 2014). We characterized causal pluralism as the view according to which causality cannot or should not be reduced to one notion or kind of thing only. Causal pluralism, however, is not an entirely new enterprise either. There exist a number of pluralistic approaches, e.g. about types of causing, about inferences and evidence, about the very concept of causation, or about methods for causal inference (Illari and Russo 2014, chap. 23). The pluralistic approach we present and favour is broader in character than other, existing pluralistic approaches, and can help systematize the literature in useful ways.

We name our approach ‘causal mosaic’ because we think of specific approaches such as in terms of processes, dispositions, counterfactuals, or manipulation and invariance as *tiles* that need to be put next to one another to form an image. The image we form is the mosaic of causal theory, which is a dynamic image. We particularly like the metaphor of the mosaic, but as any metaphor in science and philosophy, it should not be taken literally, but more heuristically. Thus, for instance, the mosaic helps us in conveying the idea that tiles next to each contribute to make an image appear, but we do not think that tiles are mutually exclusive, and in fact often overlap, or with some due modification they can be used in different parts of the mosaic. Thus, for instance, depending on specific purposes, we may need a mechanistic theory of causality or a causal theory of mechanisms (for a discussion, see Gillies (2018, chap. 4)). Once we have an inventory of available approaches – the tiles –

the question is how to arrange them for particular purposes. One purpose is to account for a very specific question (e.g., how to causally explain the economic crisis of 1929), another is to give a highly general account of what causation is, and anything in between (see Illari and Russo (2014, chap. 20)). What counts as a tile, what its borders and overlaps are, and how we arrange them, depends strongly on how we epistemic agents set the purpose.

Causal mosaic has two motivations. The first, philosophical, is that no single concept is likely to fit all scientific and policy domains. In this sense, our approach is maximally liberal, in that it allows for various kinds of causing as well as various causal methods. This philosophical motivation is further substantiated by the observation that there isn't one single philosophical question of causation, but at least 5:

(i) Metaphysics (or ontology): What is causality? What are causal relata?
(ii) Epistemology: What concepts guide causal reasoning or govern causal knowledge?
(iii) Methodology: What methods to use to discover/explore/confirm causal relations?
(iv) Semantics: What is the meaning of "cause" / "causality" in natural or scientific language?
(v) Use: What can we do (or not do) in the presence/absence of causal knowledge?

We follow Cartwright (2007) who argues for prioritising use as a philosophical question. Here, 'use' does not merely refer to action, but encompasses a whole range of activities, epistemic, material, and scientific, in accord with Chang's 'epistemic practices' (Chang 2011; 2014).

The second motivation is scientific. Just as there isn't one philosophical question only, the sciences deal with different types of causal problems. We identified at least 5:
(i) Inference: Does C cause E? To what extent?
(ii) Explanation: How or why does C cause or prevent E?
(iii) Prediction: What can we expect if C does (or does not) occur?
(iv) Control: What factors should we hold fixed to understand the relation between C and E? Or to modify C so that E accordingly changes?
(v) Reasoning: What considerations enter into establishing whether, how, or to what extent C causes E, and using that knowledge?

In our view, the causal mosaic helps in making sense of the vast intellectual enterprise about understanding causality, and about finding out about causes in the sciences. It recognises that, for any given causal concept, we need to understand what it does help with, and what it does not help with. For example, take classic debates on the notion of mechanism. Mechanisms are arguably helpful in explanatory practices in e.g. biology or neuroscience, but metaphysical concepts such as capacities or dispositions might better help address ontological aspects of biological phenomena, perhaps as complementary to mechanisms.

Although our approach of causal mosaic has a strong pragmatist flavour, it does not license an "anything goes" strategy. We think of the philosophical questions and scientific problems exactly as what should guide the choice of appropriate notions for particular contexts. Making the mosaic for a particular context is primarily an exercise in conceptual clarity, and we acknowledge that, in practice, philosophical questions about and scientific problems of causality are highly intertwined and interconnected. Concepts and notions of causality are not intrinsically good or bad, or better and worse than others, but they are more or less appropriate *for a given purpose*. Thus, for instance, if probabilistic theories are criticised for not proving good enough for causal explanation, the problem is not with probabilistic theories per se, but with what we set as a target for these theories: they aim to address problems of inference, rather than explanation, and so they are to be located primarily in the realm of epistemology and methodology.

In sum, the goal of the causal mosaic approach is to offer guidance to scientists and philosophers in selecting and using notions that are appropriate for the philosophical questions and scientific problems of causality they intend to address. It is part of the debate, we think, to discuss in the open which 'tiles' best fit the intended function and why, but such debate is premised on the idea that causal theory has to be inherently pluralistic, if it has to meet and account for the practices of the sciences and of policy.

2.2 The question of production: information transmission

We locate the question of production at the level of causal metaphysics: what *is* causation? What *are* causes? We are particularly interested in the metaphysical question, understood as a question of *linking*: what is it that *links* causes and effects? The way in which we aim to answer this question is *not* to give the *one* metaphysical concept, but to offer a very general concept that can help shed light on as many contexts as possible, in combination with other concepts. Recall that according to our mosaic pluralism, we seek to add information-transmission to our library of causal concept, not to replace existing concepts. In previous work, we highlighted the complementarity of information transmission with 'mechanisms', 'processes', and 'capacities-dispositions' (Illari 2011; Illari and Russo 2016b). Here, we focus instead on its complementarity with 'variation', as is examined in Russo (2009).

Let us begin with clarifying the meaning and scope of causal production. Illari (2011b) distinguishes between giving an account of causation in its entirety as production, as, for example, Salmon-Dowe do, and giving an account of production within a pluralist perspective such as causal mosaic. Illari examines candidate accounts in the metaphysics of causality for an account of production: processes, mechanisms, and capacities/dispositions, and argues that while they are all useful, they are also insufficient. Illari sets the stage to provide an account of causal production, and of causal linking, in terms of information. This line of work has been further pursued in joint work (Illari and Russo 2016a; 2016b). Russo (2022) further expands on this, including the complementarity of mechanisms and dispositions to information-transmission. It is worth noting that ours is not the only philosophical approach in terms of information. John Collier (Collier 1999; 2011) was perhaps the first in the causality literature to develop a full-blown account of causality in terms of information, holding that causation is the transfer of a particular quantity of information from one state to another. Others appeal to information in different ways. For example, James Ladyman and Don Ross (2007, 263) speak of "information-carrying relations" that scientists study widely. Holly Andersen (2017) uses information-theoretic approaches to the notions of causal nexus and of patterns, and Billy Wheeler (2018) discusses the prospects of an information transmission account in the context of big data. Given this interest and attention to the notion of information, we hope to provide a very general framework of thinking about information, capturing a core of agreement among these approaches.

In previous publications, we presented 'information transmission' as a thin metaphysics to cash out causal production. With the term 'thin metaphysics' we mean a minimal metaphysical commitment towards what causation is or what causes what, and this minimal commitment gives us maximal flexibility to cash out causal production across micro- and macro-factors, across factors of different natures, or in highly technologized contexts.

The most general idea of 'information transmission' comes from early accounts of mark transmission, notably as developed by Reichenbach and Salmon. Specifically, the original idea of Reichenbach was to characterize causal processes in terms of worldlines, and in these worldlines we can track genidentical events (Russo 2022, chaps 11–12). Simply put, to establish the identity of some object, individual, or event, we need to establish the diachronic identity of some state of affairs that pertains to said object, individual, or event. When Salmon further developed this idea of mark transmission, he adopted a counterfactual formulation. A process is causal, following the idea of mark transmission, if, and only if, *were* we to mark it, that mark would be transmitted to later stages of the process. A stock example is the 'dented car'. Imagine we dent a car; when the car moves, the dent moves along. This mark – the dent – can be detected at later stages of the process. But this is possible because the process that is marked is causal. Instead, if we try to mark the *shadow* of the car, for instance trying to deform it somehow, we will see that the mark – the deformation – will not travel along. This is because the shadow is not a causal process, or, better said, we are not looking at the shadow in relevant causal interactions with other processes. The original account of Salmon was largely based on Reichenbach's, and both took the Special Theory of Relativity as a fundamental constraint on (physical) causation. But the main issue with the early formulation of the process account was related to its *counterfactual* formulation: a process is causal in case *were* a mark introduced, it would be transmitted at a later stage. We lack space to reconstruct this debate in detail (which we did in previous work), but the important element to retain is that, in very many cases, this counterfactual element, and the idea of introducing some kind of *material* mark, proves problematic. Put briefly, Salmon gave up his own mark-transmission account because he acknowledged that you cannot always mark physical processes without changing them, and he didn't want to save his account from this and other problems by making it *purely* counterfactual, saying that if we introduced a mark, then the process would transmit it, even in cases where that would never actually be possible. Instead, Salmon wanted his account to be in terms of what is actually being transmitted (Wesley C. Salmon 1994). Yet, we can retain the idea of mark transmission if we can generalise the notion of process beyond a physics formulation (including using physics quantities such as energy or momentum as hallmarks of causal interactions), and most importantly liberating mark transmission from the idea that marks have to be *introduced*.

This is where the concept of information and information transmission proves helpful. We can think of information as a mark in a relevant sense, because information can and is transmitted by and in causal processes. Interestingly, this is hinted at by Salmon in his earlier work, but never further developed. In his (1994, 303) paper, Salmon said:

“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, 154–57) (Salmon, 1984 p. 154-7)” However, unlike the formulation from philosophy of physics, information is not a mark we introduce, but a mark that is already there: we epistemic agents can describe processes in informational terms, and track whether information is transmitted or not. The agent’s perspective we advocate here is key, because although our claim is that causal production *is* information transmission, it is important to bear in mind that any claim about *whether* information is transmitted or not is in fact the product of us epistemic agents engaging in numerous techno-scientific practices, from observation to manipulation. Differently put, any claim about causal production is the result of an ontology that depends on epistemology, and this is the gist of ‘ontoepistemology’ and of ‘constructionism’.

This, however, does not answer yet the question of what is this information being transmitted. We primarily work with the concept of semantic information, which is a rather qualitative approach. In this way, we do not reduce information to mathematized and formalized accounts such as Shannon-Weaver’s, and at the same time we do not exclude the use of such accounts in specific situations. Likewise, uses of and reasoning about information are common in many life sciences, particularly biology and the sciences of mind and brain. Although semantic information is clearly difficult to quantify, we take this to be a virtue of the approach, as it makes it a flexible and versatile way of expressing aspects of the world into contents that are semanticized, and this from an explicit agent’s perspective. This means, to be more concrete, that we *can* use Shannon-Weaver information or biological information, but that we can also interpret marks such as the dented car as information. Or, as we shall try to show in the next section, information is what will help cash out causal production in a variety of contexts.

3. Tracing transmission: why production needs variation

3.1 The intertwining of causal metaphysics and causal epistemology

Our approach is to understand causal production as transmission of information, in the most general terms. But it is useless to have a positive account of what causal production *is*, with no account of how we human epistemic agents *know* whether some effect is produced or not. In other words, a meaningful and useful metaphysical account has to be accompanied by an epistemology.

Our arguments about information-transmission and variation begin from two more general stances. First, as we have said, we are pluralists about causation and according to our causal mosaic, questions of metaphysics and of epistemology, while distinct, are also *connected*. Second, we think that any account of causal epistemology or causal metaphysics has to consider that these are at least in some sense ‘products’ of us human epistemic agents. To repeat, we adhere to broad principles of constructionism (Floridi 2011b) and of ontoepistemology (Barad 2007); there is no ‘view-from-nowhere’ or ideal rational agents, but real epistemic agents that make inferences, whether in ordinary or scientific contexts, which is a view clearly in line with perspectivism too (Massimi 2022).

Given our views, a natural assumption might be that finding out about production requires information transmission epistemological strategies, and finding out about difference-making requires variational epistemologies. But we will show here that this is not true, at least for production. Finding out about production can be done with information transmission strategies, but it often also requires variational strategies. Broadly, we need to trace links and marks, and compare similarities and differences across cases. However, as we will get to later on, typically we need both strategies. This is consistent with the evidential pluralism about medicine that we have developed in other work (Russo and Williamson 2007; Illari 2011a; Gillies 2011; Clarke et al. 2013; 2014; Parkkinen et al. 2018). Evidential pluralism, briefly put, holds that to establish a causal claim, one typically needs evidence of difference-making (for instance of correlations) and evidence of production (for instance of mechanisms). The thesis is primarily epistemological and methodological in character, focusing on what *evidence* is needed for us epistemic agents to establish claims about causality, and it has been discussed by us and others in contexts such as medicine (see references above) and recently also in social science (Moneta and Russo 2014; Shan and Williamson 2021) . Our claim here, instead, specifically pertains to causal production, able to work across different scientific domains, addressing epistemological strategies for knowing about causal production. While the scope *is* different, evidential pluralism and our thesis about information-variation do share a pluralistic approach to evidence.

From an ontoepistemological and constructionist perspective, we therefore ask: How do we know that some C has produced some E? How do we know that information has (or has not) been transmitted? We know by observing some phenomenon and checking whether something changes or has *changed* and how it has changed. This requires making explicit a number of 'parameters', from the background knowledge used to the empirical data that are available, from the methods and techniques used for causal analysis to the situatedness of the researcher(s) – which means carrying out observations at a specific *level of abstraction*.

In this section, our main job is to address this key question: how do we *know* about information transmission? We think that while the question can be formulated in simple terms, the answer turns out to be rather complex. We can distinguish broadly two kinds of strategies that epistemic agents use. There are the classic strategies like observational studies and RCTs, which mostly seem to match the idea of causality (difference-making) as variation, but there are also strategies that seem to match more the idea of causality (production) as information-transmission, such as process tracing, or literal marking. We will illustrate how these strategies are used by epistemic agents in various contexts, but we will show that, in general, strategies used in successful causal inferences are typically mixed and intertwined. *Both* kinds of strategies are used to infer causal production.

We will use four examples to illustrate this idea. First, we examine two simple stock examples: billiard balls colliding and watering the plant; these cases capture folk intuition about causal strategies, and also begin to illustrate more fundamental, scientific issues in simple terms. Specifically, these simple cases help us to introduce the variational and information-transmission strategies that we think are at work in a vast number of techno-scientific practices. Then we move to two cases (or episodes in the sense of Chang (2012; 2014)) where we examine practices in molecular epidemiology and in astrophysics. We discuss how epistemic agents track and trace the transmission of information in the 'exposome approach'; this is quite studied as a case of production, but the second, the science of supernovae is not, so we offer novel arguments for this. We show how important it is to discuss these episodes to highlight how information and variation work together, and from the agent's perspective: we are interested in metaphysical questions, but metaphysics does not fall from the sky, it is still the product of agents, working in groups and subject to all sorts of social and epistemic dynamics.

3.2 The transfer of information in billiard balls colliding

Billiard balls colliding is a stock causality example, descending at least since Hume. It is typically taken to be the paradigm of material forms of causation, physical processes intersecting, causal interaction, and exchange of conserved quantities. This seems to be as clear a case of information transmission as is available, especially if we cash out information transfer as exchange of some physics quantities. But how do epistemic agents know that information is transmitted? We re-examine this case from the agent's perspective.

One answer that seems obvious is that moving billiard balls is a process that we can simply watch. Tracing the process is as easy as it ever gets. We can also mark the process, for example the classic idea of chalk passing from the cue, to the first ball, and onto the second ball. Detecting the mark at the end of the process confirms that information has been transmitted. However, notice that we do also use variational strategies to find out some kinds of things, even in this extremely simple case.

First, we can know that processes of two balls transmit information because we measure relevant properties. We observe that some of these relevant properties *change* after interaction. This helps us to distinguish between real processes and what Salmon calls pseudo-processes. We can compare the billiard balls with intersecting shadows of airplanes on the ground: we see the shadows on the ground, we see they intersect, but no change is detected in either trajectory of the shadows after interaction. Comparing the change (billiard balls) to no change (shadows) studies variation.

Second, epistemic agents need to identify relevant properties, and variational strategies of examining and comparing repeated experiments and trials are often vital to help drill down and identify and establish for sure the relevant properties, and whether they can be generalized. For instance, one relevant property is the weight of balls, their form, including how hard they are, and which direction they are hit, and not, for example, their colour.
All of these strategies are useful in understanding even such a simple causal process.

In this case, to repeat, the kind of information that is transmitted is exchange of some physics quantity. Our story, however, is that this simple answer needs to be complemented with an a story about *how* an epistemic agent can establish *whether* information is in fact transferred or not.

3.3 Omitting to water the plant

Omitting to water the plant is another stock example of the causality literature. Failing to water the plant causes it to die. Omission is likely the cause, but omission is metaphysically controversial because something that is *not* cannot cause something else. Thus, according to a traditional materialist perspective, there cannot be any transmission of information, unlike in the collision of billiard balls. Notice that legal reasoning is replete with cases of omissions, and this is causal reasoning. In the legal context, questions also arise about responsibility and accountability (see e.g. (Moore 2009)), but we set aside this debate. Here we only focus on the question about production and whether omissions can be properly considered as cases of causal production, re-analyzing from information transmission and the agent's perspective. This, we think, legitimatizes talking about information transmission: not because we can univocally and directly point to material, physical transfers of information, but because an epistemic agent can analyze the case and reconstruct it in terms of the transmission of (semantic) information.

An epistemic agent can observe a plant and notice that it is in a pretty good state up to a certain point. Yet at some point the plant is dead. Upon receiving information that the plant has not been watered enough, this omission becomes *informative* about the next state of the plant. There is a *change* (i.e., variation) in the plant's life process, in which *omission* of watering explains death.

The epistemic agent can study two kinds of information-transmission processes here, first the ongoing life of the plant, maintaining the plant in (living) homeostasis, and, second, the plant's constant intake of the essentials of life, including water. If one were to make experiments on information-transmission processes in the plant, one way of marking an information-transmission process would be to stop watering it. But that doesn't tell you much without variational strategies.

A variational strategy to infer information transmission would be to apply Mill's methods, and notice that, for instance, the only factor that changed was watering the plant, which then becomes informative about the production. An interesting feature of this case is that, while billiard balls colliding gives the impression that information transmission is ultimately reducible (and to be reduced) to *physical* transfer of some physical quantity, in this case it starts becoming clearer that whatever transmission *in the world* is, there is an ineliminable element of construction by the epistemic agent.

3.4 Biomarkers research and information channels

In earlier publications, we analyzed the case of molecular epidemiology, and more specifically the use made there of biomarkers in order to establish causal links (Illari and Russo 2016b; Russo and Vineis 2016; Vineis, Illari, and Russo 2017; Vineis and Russo 2018). Scientists in this field are interested in establishing links between exposure, for instance pollution or certain chemical hazards, and disease, for instance specific types of cancer, asthma, or allergies. The big challenge of molecular epidemiology is to connect macro and micro factors. How do we know that certain chemicals cause cancer? Scientists go search for biomarkers of exposure, then biomarkers of early clinical changes, then biomarkers of disease onset.

Molecular epidemiology is an interesting area to look at because statistical methods and experimental equipment, together with novel theorizing about the notion of exposure, made it possible to re-frame questions about exposure and health outcomes that classic epidemiology had previously dealt with at a rather coarse-grained level, at a very fine-grained level. For example, exposure research uses proteomics, metabolomics, genomics (sometimes called the 'omics technologies') to look at the fine-grained level of proteins, metabolites and genes in the cell, as the body reacts to environmental exposure, and disease starts to develop. While this is fascinating work, even this level of detail does not make it easy to establish links as we have discussed for instance in the previous case of the billiard ball or the unwatered plant. For one thing, the strategies to establish linking between exposure and disease are still about establishing information transmission, but we can never establish *continuous* linking, of the kind we may directly observe in two billiard balls colliding. We can only access certain points in the complex interacting thing that is the human body, and in its interactions with the environment. Notice, however, that this does not mean that the link isn't continuous, simply that *we epistemic agents* can't trace *all* the points between exposure and disease. This is why an agent's perspective is needed: *we* reconstruct the process, and how information is transmitted all along. This is why our account is constructionist or ontoepistemological: information transmission *is* causal production, but as human epistemic agents we establish *whether* information has been transmitted through a number of epistemic and ontic strategies.

This work involves quite a few different strategies. A lot of experimental analysis of biospecimens is done, at various omics levels. But once the data about omics is generated, the analysis is statistical. Researchers conceptualise their methodology as 'meeting-in-the-middle' (Vineis and Chadeau-Hyam 2011) , which is a statistical methodology that cross-references two types of correlations: correlations between exposure to a given hazard X and some biomarker (to be specified), and then correlations between that biomarker and some clinical conditions. The idea is to start to get a grip on the body reacting to environmental exposures as early as possible, and find which such biomarkers are linked to early disease onset. So in this way, we are trying to do something like tracing the moving billiard balls, but we can only access them at certain points, and using lots of technologies. The correlations found can then be understood as joint *variations* that help us track the linking. It is important to note, however, that these statistical analyses do not operate in a 'vacuum', but are based on lots of background knowledge, including knowledge of bio-chemical mechanisms that make it plausible to search for links at, say, the proteomic level rather than at the metabolomic level. It is in this sense that in Illari and Russo (2016b) we talk of mechanisms as *information channels* which impose constraints on possible and plausible linking, and so indicate where we should go to find the links. So this case shows how even research that is strongly dependent on variation strategies can still be seen as aiming to trace information.

Even more recent projects in exposure research add an important layer of complexity to the analysis, considering not only biochemical but also social mechanisms, broadly construed. The question is, for instance, how adverse childhood experiences are part of complex causal pathways that start quite early in life, and that become visible as clinical conditions much later in life. Here scientists search for links between exposure and disease that are not just biological, but also social and bio-social, and for this reason we need *socio*-markers, next to bio-markers (Ghiara and Russo 2019).

Once again, our understanding of information transmission is a thin metaphysics because our metaphysical understanding of information is at once minimal and maximally liberal. We can achieve that by working with a notion of *semantic* information, which can be cashed out as something quite material as exchange of physics quantities in the case of billiard balls and also as something more epistemically oriented as the *reconstruction* of a link by epistemic agents in the biomarkers case.

3.4 Exploring information transmission across mechanisms and processes in supernovae

SN1987A was a particularly important supernova (McCray (1997) offers a quick accessible introduction). It exploded some 168,000 years ago, and then became visible from Earth in the late 1980s. Two issues stand out. First, it extended our understanding of the cause, the particular 'mechanism' of type IIa supernovae, specifically what carried away much of the energy of the collapse (see Prialnik (2010) for a textbook presentation, Suzuki (2008) for a review, see also Walker (1987), Lattimer (1988)). Second, there was significant work immediately following SN1987A's discovery on the cause of its various peculiarities, including that it was the explosion of a blue supergiant, its light curve, and its peculiar 'squashed figure 8' appearance: there were two dimmer rings forming a 'squashed figure 8' surrounding the more usual bright center with a single bright ring around it (Joss et al. (1988), Podsiadlowski and Joss (1989), Woosley and Chevalier (1989), Podsiadlowski et al. (1991)).[1] What caused them?

These two issues above both raise questions of causal production. The first question is a general case question of what produces the kinds of bright flares that we see, and can be distinguished by their light curves from other types of supernovae. What is the *general* kind of continuous process that causes the things we detect, what the scientists call the 'mechanism' of Type IIa supernovae? In the second question, we are trying to figure out what happened in a single-case continuous physical process. So scientists detect the bright flare of SN1987A, also detect the neutrinos passing through Earth, and try to infer what happened to cause the bright flare, and in between that explosion and the neutrinos getting to Earth.

In addressing the first question, lots was already understood about supernovae. The size of the star's core is crucial, and SN1987A clearly had a core higher than the Chandrasekhar limit (about 1.4 solar masses), so would fall into Type II (although see Murdin (1993), for difficulties typecasting supernovae). This was theoretically understood to occur once nuclear burning in the star's core had turned it to iron, whereupon changing to further elements does not release energy, ceasing the nuclear burning that supports the star, and leading to its collapse (Prialnik 2010).

So the beginning of this kind of process was quite well understood, while the end was even clearer. Stars exploding makes them flare far brighter, and this is something that had been recorded for centuries before SN1987A. What was mysterious was the exact mechanism in between. Specifically, the collapsing core releases an enormous amount of energy, that physical law tells us must go somewhere. Something must carry off the energy. A contemporary theoretical model suggested that a neutrino blast could do it. SN1987A was

[1] Images are widely available. See, for example, Hubble's image http://www. spacetelescope.org/images/potw1142a/, or https://en.wikipedia.org/wiki/SN_1987A.

the first detection of neutrinos from such a source, reported by neutrino detectors at IMB and Kamiokande (so scientists were not able to *compare* to any other cases). Nevertheless, the causal process was filled in: SN1987A's core, larger than 1.4 solar masses, turned to iron, became no longer able to support itself against gravity, and collapsed, releasing an enormous neutrino blast, and creating the bright flare visible for several years (Sato, Shimizu, and Yamada 1995). This is the missing part of the mechanism of this type of supernova, in large part established on the basis of fewer than 10 neutrinos actually detected! In one sense this was relatively simple. Neutrinos were theoretically understood, known to exist, and had established detection methods. In theory they are created by the blast, and literally travel out in all directions including to Earth, and then they were indeed detected on Earth the expected time later (given their expected travel time and compared, for example, with visible light). In this case information transmission consists of the literal travelling of physical particles. And while we can only detect the information (i.e., the neutrinos) at one end of the transmission, that is ok in such cases. This reinforces the idea that information transmission does not require scientists to identify all points of the continuous process, but that instead we epistemic agents need to identify the relevant ones, just as in the case of exposure science, or even in the simple case of the plant not being watered, reconstructing the process. In another sense, this episode involves a very complex *epistemic activity of a large community* to, first, reconstruct the arrival of neutrinos on the basis of a tiny amount of data blips created in massive underground water-filled structures and, second, converge on a postulated mechanism of type IIa supernovae on the basis of a blast inferred from so few neutrinos.

Addressing the second question was important, because it was immediately obvious that something else must have happened with SN1987A. It did not merely show as a bright inner dot surrounded by a bright ring, as is usual, but also showed a fainter squashed figure eight, and it had other peculiarities too. What caused this? Work was done to modify standard models to understand what might be different about SN1987A. Stellar structure models of main sequence burning are already modified to get models of what will happen at the end of a star's life, yielding multiple possible mechanisms of supernovae. Yet SN1987A was already a nonstandard type IIa supernova. Simulations suggested that significant roiling in the star core prior to collapse could account for a violent asymmetric explosion that could lead to two 'cones' blasting far into space, the ends of which would appear from Earth as a 'squashed figure 8' (McCray (1997), Chevalier (1992)). To be sure, this is a very simplified story, cutting out a lot of enduring controversy, but for our purposes it will do. So the process for SN1987A was thought to be as the newly agreed mechanism of type II supernovae described above, but with significant roiling in the star core prior to collapse, that caused the peculiarity of the 'squashed figure 8' appearance.

Unlike exposure research, scientists theorize about parts and activities they literally think exist in relatively unproblematic ways, such as the star's core, roiling, and cones. Even neutrinos were in some ways well validated before 1987. Like the billiard balls, how those parts move and what happens to energy is a constant part of any inferences and understanding. Thus, the case of supernovae also shows how information transmission can be used across micro- and macro-levels of reality, crossing from something as large as star's core passing the Chandrasekhar limit, to something as small as a neutrino.

However, just as in exposure research, the agent's perspective is key. Unlike in the billiard ball case, where an agent may be able to get qualitative, direct observation of the transmission of information (and also get quantitative measurements of such transmission), the cases of exposure science and of supernovae show that strategies to trace information transmission need complemented, or aided by, epistemological, variational strategies to trace the transmission of information. Even in a case where significant elements of information transmission consists in the literal travelling of physical particles vast distances, a large community of epistemic agents reconstructs information-transmission, using variational strategies that got us to stellar structure models, mechanisms of supernovae, and the building of the simulations that explained the peculiarities. This all involved studying stars, building models that generate the data we see, comparing differences between different stars, and modifying models according to what theoretically *could* change, to see if we can match the differences we see, and gradually building a case for some changes – causes – over others.

This again shows how our ontoepistemological approach works. While the case of SN1987A shares some features of the simple billiard ball case, it is far from simple. Our understanding of information as at once minimal and maximally liberal works well, and here we see that our knowledge of even relatively simple transmission of physics quantities is really understood by epistemic agents reconstructing the link using significant technologies, as in the biomarkers case.

4. Discussion and conclusion

We hold a pluralistic approach to causality that we dubbed 'causal mosaic'. Within this pluralistic framework, we focused here, following previous work, on questions of causal metaphysics, and more precisely of causal production. This is the question of what *links* causes and effects, and we have been defending 'information transmission' as the most general account of causal production, able to account for cases across levels (micro-macro), across inhomogeneous factors (biological-social), and in highly technologized contexts (admittedly, we haven't discussed this extensively here, but the interested reader is referred to Russo (2022, chap. 12)). The question of causal production, however, is not exhausted within the boundaries of metaphysics, because we share views that consider the perspective of agents as vital in the reconstruction of 'what there is', such as for instance the constructionist framework of Floridi, the ontoepistemological framework of Barad, the perspectival framework of Giere and of Massimi – accounts that are collectively discussed in Russo (2022).

In introducing an explicit agent's perspective, the metaphysical question of causal production as information transmission needs to be accompanied with an epistemological question about how do we epistemic agents *know* whether information is transmitted. In the chapter, we illustrated with four examples how information transmission needs epistemological strategies, based on variation, that help us to establish the existence of information transmission. Differently put, we have sketched here the main lines of an epistemology of causality *of epistemic agents*, which in science are typically organized into epistemic communities. From these epistemic agents' perspectives we can establish facts about causal metaphysics. From the agent's perspective, we can distinguish two broad categories of strategies, although we have also shown that these are typically both used, and indeed are quite deeply intertwined.

For the sake of clarity, we keep them separate, in order to highlight their main features.

Beginning with information transmission strategies, these can be understood quite generally as directly tracking the information, or the channel it runs through. This can be done by literally marking the process, like the billiard balls, or channels, such as dropping a dye in water and detecting it later downstream, to show that chemicals could be transmitted that way. This is related to process tracing à la Steel (2008), which can involve simply watching change happen, like in the case of watching moving billiard balls, but can also involve watching stuff be maintained, like the living plant.

We showed that in exposure research conceptualizing linking as information transmission is helpful, but empirical access to that informational linking is very difficult, pushing scientists to primarily variational strategies. Note that background mechanistic knowledge does a lot to make some kinds of information transmission channels more plausible than others (discussed more extensively in Illari and Russo (2016b)). SN1987A presents a different challenge. There is a single case process we are trying to trace, in particular at the time aiming to figure out a missing step. Scientists face the challenges of studying both very small and very large entities, which are also mind-bendingly far away. They cope with severe limits on what naturally travels close to Earth, as all our empirical information on supernovae must. However, scientists are still in a sense trying to access different points on that continuous process.

In sum, information transmission strategies are particularly good for ruling in and ruling out possibilities for causal linking, probing new unknown things, and studying single or rare cases.

Turning to variation (which has been deeply studied by Russo (2009)), these strategies can be understood very generally as comparing across similar and different things, tracing back to Mill's classic methods. Standard much-discussed strategies of course include randomized controlled trials and observational studies, or analysis of population data, but we have examined other kinds of work.

Even in the simplest physical processes, when we need to establish relevant properties, we often need variational strategies like in the billiard balls. Turning to the classic example of omitting to water the plant, strictly speaking we can mark it, as stopping watering it is a mark. But we really need variational strategies here to be sure what's going on, such as comparing other plants that are regularly watered to those given no water, or holding fixed details of who was in charge of watering the plant.

Variational strategies are very widely used in exposure research, going beyond the classic established study methodologies to 'patchwork' variations across the hypothesized linking from environmental exposure to disease onset, building a different evidential picture than previously existing work such as broader observational studies. SN987A might be a single rare process, but empirical challenges mean scientists still need lots of modelling of that type of process (and comparison with empirical data for other stars) to get us going. Unsurprisingly, the need for similar variational strategies is very clear for building understanding of type IIa supernovae in general. Note that Russo (2009) develops an account of variational strategies in great detail, particularly for social science practices. But the key point there is to demonstrate how a broad variational approach can capture the rationale for both interventionist and a range of observational methodologies in social science. In sum, variational strategies are particularly good for finding fine-grained relevant properties and for establishing links when we can't directly establish links, for which counterfactual reasoning (a special kind of variational strategy) is very helpful.

As we have emphasized, though, these strategies are deeply intertwined, and typically used together. While we have drawn out and shown different kinds of methods, to distinguish the strategies, even in the simplified examples, multiple strategies are useful. Epistemic agents use strategies in concert.

To conclude, epistemic agents, organised in communities, use both information transmission and variational strategies to establish the existence of information transmission. This allows them to reconstruct continuous linking (causal production) when it is only empirically accessible at certain points, as is the case in both exposure research and the astrophysics of supernovae. Combinations of strategies yield much more than one alone – this is the key idea of the 'causal mosaic' (Illari and Russo 2014). The agent's perspective, notice, is crucial. It is important to point out that a constructionist or ontoepistemological perspective does not entail anti-realist stances, quite the contrary is the case. There is no causal metaphysics in the absolute, a priori, or detached from us, and the ontoepistemological perspective also demands that our causal metaphysics tells us something useful about causal epistemology. It is high time that we have an explicit account of causal production in which the very real role of epistemic agents is central, and this is the main message of our contribution.

References

Andersen, Holly. 2017. 'Pattens, Information, and Causation'. *The Journal of Philosophy* 114 (November): 592–622. https://doi.org/10.5840/jphil20171141142.

Anscombe, G.E.M. 1975. 'Causality and Determination'. In *E. Sosa, (Ed.) Causation and Conditionals*, 63–81. Oxford University Press.

Barad, Karen. 2007. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning.* Durham & London: Duke University Press.

Campaner, Raffaella. 2011. 'Understanding Mechanisms in the Health Sciences'. *Theoretical Medicine and Bioethics* 32: 5–17.

Cartwright, N. 2007. *Hunting Causes and Using Them: Approaches in Philosophy and Economics*. Cambridge: Cambridge University Press.

Cartwright, Nancy. 2004. 'Causation: One Word, Many Things'. *Philosophy of Science* 71: 805–19.

Chang, Hasok. 2011. 'Beyond Case-Studies: History as Philosophy'. In *Integrating History and Philosophy of Science*, edited by Seymour Mauskopf and Tad Schmaltz, 263:109–24. Boston Studies in the Philosophy of Science. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-007-1745-9_8.

———. 2012. ‘Beyond Case-Studies: History as Philosophy’. In *Integrating History and Philosophy of Science: Problems and Prospects*, edited by Seymour Mauskopf and Tad Schmaltz, 109–24. Boston Studies in the Philosophy Science 263. Dordrecht, Heidelberg, New York, London: Springer Science & Business Media.
———. 2014. ‘Epistemic Activities and Systems of Practice: Units of Analysis in Philosophy of Science After the Practice Turn’. In *Science After the Practice Turn in the Philosophy, History, and the Social Studies of Science*, edited by Léna Soler, Sjoerd Zwart, Michael Lynch, and Vincent Israel-Jost, 67–79. Routledge Studies in the Philosophy of Science. New York and London: Routledge, Taylor & Francis Group.
Chevalier, Roger A. 1992. ‘Supernova 1987A at Five Years of Age’. *Nature* 355 (6362): 691–96. https://doi.org/10.1038/355691a0.
Clarke, Brendan, Donald Gillies, Phyllis Illari, Federica Russo, and Jon Williamson. 2013. ‘The Evidence That Evidence-Based Medicine Omits’. *Preventive Medicine* 57 (6): 745–47. https://doi.org/DOI: 10.1016/j.ypmed.2012.10.020.
———. 2014. ‘Mechanisms and the Evidence Hierarchy’. *Topoi* online first. https://doi.org/10.1007/s11245-013-9220-9.
Collier, John. 1999. ‘Causation Is the Transfer of Information’. In *Causation, Natural Laws, and Explanation*, edited by Howard Sankey, 215–63. Dordrecht: Kluwer.
———. 2011. ‘Information, Causation and Computation’. In *Information and Computation: Essays on Scientific and Philosophical Understanding of Foundations of Information and Computation*, edited by Gordana Dodig Crnkovic and Mark Burgin. Singapore: World Scientific.
Floridi, Luciano. 2011a. ‘A Defence of Constructionism: Philosophy as Conceptual Engineering’. *Metaphilosophy* 42 (3): 282–304. https://doi.org/10.1111/j.1467-9973.2011.01693.x.
———. 2011b. *The Philosophy of Information*. Oxford: Oxford University Press.
Ghiara, Virginia, and Federica Russo. 2019. ‘Reconstructing the Mixed Mechanisms of Health: The Role of Bio- and Sociomarkers’. *Longitudinal and Life Course Studies* 10 (1): 7–25. https://doi.org/10.1332/175795919X15468755933353.
Giere, Ronald. 2006. *Scientific Perspectivism*. Chicago: The University of Chicago Press.
Gillies, Donald. 2011. ‘The Russo-Williamson Thesis and the Question of Whether Smoking Causes Heart Disease’. In *Causality in the Sciences*, edited by Phyllis McKay Illari, Federica Russo, and Jon Williamson, 110–25. Oxford: Oxford University Press.
———. 2018. *Causality, Probability, and Medicine*. 1 [edition]. New York: Routledge.
Godfrey-Smith, Peter. 2010. ‘Causal Pluralism’. In *Oxford Handbook of Causation*, edited by Helen Beebee, Christopher Hitchcock, and Peter Menzies, 326–37. Oxford: Oxford University Press. http://www.petergodfreysmith.com/CausalPluralismPGS-07-Final.pdf.
Illari, Phyllis McKay. 2011a. ‘Mechanistic Evidence: Disambiguating the Russo-Williamson Thesis’. *International Studies in the Philosophy of Science* 25 (2): 139–57.
———. 2011b. ‘Why Theories of Causality Need Production: An Information-Transmission Account’. *Philosophy and Technology* 24 (2): 95–114. https://doi.org/10.1007/s13347-010-003-3.
Illari, Phyllis McKay, and Federica Russo. 2014. *Causality: Philosophical Theory Meets Scientific Practice*. First edition. Oxford, United Kingdom: Oxford University Press.

Illari, Phyllis, and Federica Russo. 2016a. ‘Causality and Information’. In *The Routledge Handbook of Philosophy of Information*, edited by Luciano Floridi, 235–48. Routledge.
———. 2016b. ‘Information Channels and Biomarkers of Disease’. *Topoi* 35 (1): 175–90. https://doi.org/10.1007/s11245-013-9228-1.
Joss, P. C., Ph. Podsiadlowski, J. J. L. Hsu, and S. Rappaport. 1988. ‘Is Supernova 1987A a Stripped Asymptotic-Branch Giant in a Binary System?’ *Nature* 331 (6153): 237–40. https://doi.org/10.1038/331237a0.
LaCaze, Adam. 2011. ‘The Role of Basic Science in Evidence-Based Medicine’. *Biology and Philosophy* 26 (1): 81–98.
Ladyman, James, and Don Ross. 2007. *Every Thing Must Go*. Oxford University Press.
Lattimer, James M. 1988. ‘Supernova Theory and the Neutrinos from SN1987a’. *Nuclear Physics A* 478 (February): 199–217. https://doi.org/10.1016/0375-9474(88)90846-9.
Longworth, Francis. 2010. ‘Cartwright’s Causal Pluralism: A Critique and an Alternative’. *Analysis* 70 (2): 310–18. http://analysis.oxfordjournals.org/content/70/2/310.extract.
Massimi, Michela. 2022. *Perspectival Realism*. Oxford Studies in Philos Science Series. New York: Oxford University Press.
McCray, Richard. 1997. ‘SN1987A Enters Its Second Decade’. *Nature* 386 (6624): 438–39. https://doi.org/10.1038/386438a0.
Moneta, Alessio, and Federica Russo. 2014. ‘Causal Models and Evidential Pluralism in Econometrics’. *Journal of Economic Methodology* forthcoming.
Moore, Michael S. 2009. *Causation and Responsibility. An Essay in Law, Morals, and Metaphysics*. Oxford: Oxford University Press.
Murdin, Paul. 1993. ‘Supernovae Can’t Be Typecast’. *Nature* 363 (6431): 668–69. https://doi.org/10.1038/363668a0.
Parkkinen, V-P, C. Wallmann, M. Wilde, B. Clarke, P. Illari, M. P. Kelly, C. Norell, F. Russo, B. Shaw, and J. Williamson. 2018. *Evaluating Evidence of Mechanisms in Medicine. Principles and Procedures*. New York, NY: Springer Berlin Heidelberg.
Pérez-González, S., and Elena Rocca. 2021. ‘Evidence of Biological Mechanisms and Health Predictions: An Insight into Clinical Reasoning’. Preprint. PhilSci Archive. http://philsci-archive.pitt.edu/id/eprint/19210.
Podsiadlowski, Ph., A. C. Fabian, and I. R. Stevens. 1991. ‘Origin of the Napoleon’s Hat Nebula around SN1987A and Implications for the Progenitor’. *Nature* 354 (6348): 43–46. https://doi.org/10.1038/354043a0.
Podsiadlowski, Ph., and P. C. Joss. 1989. ‘An Alternative Binary Model for SN1987A’. *Nature* 338 (6214): 401–3. https://doi.org/10.1038/338401a0.
Prialnik, Dina. 2010. *An Introduction to the Theory of Stellar Structure and Evolution*. 2nd ed. Cambridge ; New York: Cambridge University Press.
Reiss, Julian. 2015. *Causation, Evidence, and Inference*. Routledge.
Russo, Federica. 2009. *Causality and Causal Modelling in the Social Sciences. Measuring Variations*. Methodos Series. New York: Springer.
———. 2022. *TECHNO-SCIENTIFIC PRACTICES: An Informational Approach.* S.l.: ROWMAN & LITTLEFIELD.
Russo, Federica, and Paolo Vineis. 2016. ‘Opportunities and Challenges of Molecular Epidemiology’. In *Philosophy of Molecular Medicine*, ch.12. Routledge.

Russo, Federica, and Jon Williamson. 2007. ‘Interpreting Causality in the Health Sciences’. *International Studies in Philosophy of Science* 21 (2): 157–70.
Salmon, W.C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
Salmon, Wesley C. 1994. ‘Causality without Counterfactuals’. *Philosophy of Science* 61: 297–312.
Sato, K., T. Shimizu, and S. Yamada. 1995. ‘Explosion Mechanism of Collapse-Driven Supernovae’. *Nuclear Physics A* 588 (1): c345–56. https://doi.org/10.1016/0375-9474(95)00160-3.
Shan, Yafeng, and Jon Williamson. 2021. ‘Applying Evidential Pluralism to the Social Sciences’. *European Journal for Philosophy of Science* 11 (4): 96. https://doi.org/10.1007/s13194-021-00415-z.
Steel, Daniel. 2008. *Across the Boundaries. Extrapolation in Biology and Social Science*. Oxford: Oxford University Press.
Suzuki, A. 2008. ‘The 20th Anniversary of SN1987A’. *Journal of Physics: Conference Series* 120 (7): 072001. https://doi.org/10.1088/1742-6596/120/7/072001.
Vineis, Paolo, and Marc Chadeau-Hyam. 2011. ‘Integrating Biomarkers into Molecular Epidemiological Studies’. *Current Opinion in Oncology* 23 (1): 100–105.
Vineis, Paolo, Phyllis Illari, and Federica Russo. 2017. ‘Causality in Cancer Research: A Journey through Models in Molecular Epidemiology and Their Philosophical Interpretation’. *Emerging Themes in Epidemiology* 14 (1): 7. https://doi.org/10.1186/s12982-017-0061-7.
Vineis, Paolo, and Federica Russo. 2018. ‘Epigenetics and the Exposome: Environmental Exposure in Disease Etiology’. In *Oxford Research Encyclopedia of Environmental Science*, by Paolo Vineis and Federica Russo. Oxford University Press. https://doi.org/10.1093/acrefore/9780199389414.013.325.
Walker, Terry P. 1987. ‘Making the Most of SN1987A’. *Nature* 330 (6149): 609–10. https://doi.org/10.1038/330609a0.
Weber, Erik. 2007. ‘Conceptual Tools for Causal Analysis in the Social Sciences’. In *Causality and Probability in the Sciences*, edited by Federica Russo and Jon Williamson, 197–213. College Publications.
Wheeler, Billy. 2018. ‘How to Find Productive Causes in Big Data: An Information Transmission Account’. *Filozofia Nauki* 26 (4): 5–28. https://doi.org/10.14394/filnau.2018.0021.
Woosley, S. E., and R. A. Chevalier. 1989. ‘Was the Millisecond Pulsar in SN1987A Spun up or Born Spinning Fast?’ *Nature* 338 (6213): 321–22. https://doi.org/10.1038/338321a0.