

# Abstract versus Causal Explanations?

Alexander Reutlinger<sup>1</sup> and Holly Andersen<sup>2</sup>

Forthcoming in *International Studies in the Philosophy of Science*

## Abstract

In the recent literature on causal and non-causal scientific explanations, there is an intuitive assumption (which we call the ‘abstractness assumption’) according to which an explanation is non-causal by virtue of being abstract. In this context, to be “abstract” means that the explanans in question leaves out many or almost all causal microphysical details of the target system. After motivating this assumption, we argue that the abstractness assumption, in placing the abstract and the causal character of an explanation in tension, is misguided in ways that are independent of which view of causation or causal explanation one takes to be most accurate. On major accounts of causation, as well as on major accounts of causal explanation, the abstractness of an explanation is not sufficient for it being non-causal. That is, explanations are not non-causal by dint of being abstract.

**Keywords:** explanation; abstraction; non-causal explanation; causal explanation; causation

---

<sup>1</sup> Alexander Reutlinger, Ludwig-Maximilians-Universität München, Munich Center for Mathematical Philosophy, Ludwigstr. 31, 80539 München, Email: Alexander.Reutlinger@lmu.de.

<sup>2</sup> Holly Andersen, Simon Fraser University, Department of Philosophy, 8888 University Drive, Burnaby, B.C., Canada, V5A 1S6, Email: holly\_andersen@sfu.ca.

## 1. Introduction

In the recent literature on scientific explanations, a consensus is emerging that one finds two significantly different kinds of explanation: there are both causal and non-causal explanations in the sciences. This emerging consensus raises the question of how one should distinguish between causal and non-causal explanations, what precisely renders an explanation causal or, respectively, non-causal, and how causal and non-causal explanations relate to one another. Particularly in recent discussions of mathematical explanations in science, one finds reliance on an assumption that the more abstract an explanation is, the less causal it is. While this equation of abstractness with non-causality is not always made explicitly, the underlying assumption is roughly that explanation is non-causal by virtue of being *abstract*, and, respectively, what makes an explanation causal is its not being abstract. We will call this the ‘abstractness assumption’. In this paper, we will fill out what is involved in this assumption, partly by characterizing the assumption more clearly, and partly by demonstrating how it is relied on in recent work.

How is the abstractness assumption motivated? There is an intuitive conceptual connection between abstractness and the distinction between causal and non-causal explanation. This intuitive connection can be motivated by considering the following observations: *on the one hand*, causal explanations are ubiquitous in the sciences and many causal explanations are rich in detail, involving reference to ornate nesting and interlocking mechanisms and sub-mechanisms responsible for their explananda (such as mechanistic explanations in neuroscience, see Craver 2007). *On the other hand*, there clearly are some explanations that are not rich in detail, where some may even derive their explanatory character from relationships that do not carry specific physical meaning,

such as mathematical relationships. Their explanantia (and often though not necessarily their explananda) are ‘abstract’. Abstractness as a quality of an explanation that can vary along a continuous parameter is thus taken to be inversely related to the causal content of that explanation. Reliance on the abstractness assumption occurs, we are arguing, when noting the abstract character of an explanation is part of arguing for why that explanation is not causal. This reliance often takes the form of characterizing an abstract explanation as one that leaves out many or almost all microphysical details about the target phenomenon of the explanation. We will use the notion of multiply realizability to characterize explanantia leaving out many or almost all microphysical details about the target phenomenon (see Section 2.1 for a detailed argument).

Interestingly, many paradigm cases of non-causal explanations of empirical phenomena have this abstract character, such as mathematical explanations (including geometric, topological, and statistical explanations), renormalization group explanations, dimensional explanations, kinematical explanations, and explanations based on symmetry principles and inter-theoretic relations (see, for instance, Batterman 2002, 2010; Batterman and Rice 2014; Bokulich 2008; Chirimuuta 2014; Huneman 2010; Lange 2009, 2011, 2013a, 2013b; Pincock 2012; Saatsi forthcoming a; Weatherall 2011; Reutlinger 2014, forthcoming a, b). These examples of abstract explanations can be contrasted with scientific causal explanations relying on the description of ‘non-abstract’ and detailed biological-physical mechanisms. One classic example is the Krebs’s citric acid cycle involving a great deal of physical machinery organized in specific ways so that

coordinated processes can run through stable physical entities in order to produce concrete outputs.<sup>3</sup>

These observations have led some philosophers to invoke abstractness in arguing that some particular type of explanation is non-causal. Most prominently in the recent discussion, Batterman (2000: 28, 128, 138-139; 2002: 40-42; 2010: 2, 21) and Pincock (2012: 53, 203-210, 221-229) rely on this intuitive assumption connecting abstractness and non-causal explanations. Although the abstractness assumption is sometimes adopted implicitly, it is a common argument pattern in recent work on scientific explanation.

The core idea of the abstractness assumption is that explanations are non-causal because of being abstract, and explanations are causal because of being *non-abstract*, or concrete and detailed. However, it is unclear to what phrases such as ‘because of’ might amount. In characterizing the relevant notion of abstractness, we use as guidance the role that this assumption is supposed to play when used in connection with arguments for the non-causal character of a specific type of explanation. One plausible, simple and helpful interpretation of connectives such as ‘because of’ in the context of the abstractness assumption is the following: an explanation is *non-causal* to the extent that its explanans is *abstract* (that is, the explanans leaves out almost all microphysical causal details of the target system); an explanation is *causal* to the extent that its explanans is *not abstract*.

However, is the abstractness assumption reliable? Does the feature of abstractness *make* an explanation non-causal? Are abstract explanations non-causal *because of being*

---

<sup>3</sup> Even though the output of the Krebs’ citric acid cycle can be characterized abstractly as ‘energy’, it is usually characterized in non-abstract terms, e.g. two extra units of adenosine tri phosphate.

*abstract*? Since this assumption figures in discussions of how to characterize causal and non-causal explanations in particular, it is also key to ask: does the answer to these questions depend on the specific view of causation or causal explanation one takes?

We think that abstractness is not a defining feature of non-causal explanation, and indeed, that once this apparent tension between causation and abstraction is made explicit it loses much of its intuitive appeal. Importantly for how this bears on the issue of explanation, however, we will focus our discussion on showing that abstractness and causation are not in tension for every major account of causation and causal explanation currently under discussion. Rejection of the tension between causation and abstraction is not dependent on one's other philosophical commitments with respect to causation or to explanation. Increasing the abstractness of an explanation does not necessarily thereby decrease the causal character of the explanation, or so we will argue.

We hold that, first, there can be and there in fact are *abstract causal* explanations in the sciences, and second, that recognition of this is compatible with every major theory of causation and causal explanation. We show that, on distinct major approaches to causation and to causal explanation, genuinely causal explanations can be, and in fact frequently are, abstract. To emphasize, abstract causal explanations are *not only possible but actually quite numerous* in the social and life sciences. Moreover, many philosophers of science want to capture this abstract aspect of causal explanation in science and of causal modeling practices in general. While abstractness assumption may be intuitively plausible, it is not, on reflection, to be endorsed.

The structure of the paper is as follows. In section 2, we reconstruct how advocates of the abstractness assumption use the notion of abstractness. In section 3, our strategy is

to present two arguments against the abstractness assumption: first, abstract explanations can and often do represent causal relations according to major accounts of causation; second, abstract explanations can be and often are causally explanatory according to major causal accounts of explanation. Our main conclusion is twofold: first, that, contrary to the abstractness assumption, abstractness is not necessarily in tension with the causal character of an explanation; consequently and second, that relying on the feature of abstractness is not a reliable marker for distinguishing between causal and non-causal explanations, regardless of one's other commitments concerning causation and causal explanation.

Finally, we wish to stress that we have diverging views about how to positively explicate the distinction between causal and non-causal explanations (as articulated in Andersen 2016; Reutlinger 2014, forthcoming a, b). However, the purpose of this paper is not to discuss the merits of our own positive views about this subject – this will be left for another paper. Instead we will focus on a point of agreement between us: the claim that the abstractness assumption is not a viable strategy for characterizing the distinction between causal and non-causal explanations. That is, the central focus of the paper consists in articulating a criticism.

## **2. What Abstractness Is (Not)**

In the introductory section, we have sketched the abstractness assumption, according to which the more abstract an explanation is, the less causal it thereby is. But what does being *abstract* amount to? To address this question, we will first rule out a number of

concepts of abstractness that we think are clearly inappropriate for articulating the abstractness assumption (Section 2.1). Then, we will turn to the interpretation that (to our mind) best fits with the role of the abstractness assumption in arguments regarding non-causal explanation (Section 2.2). In articulating the relevant notion of abstractness, we are already arguing against abstractness as a useful way to distinguish causal and non-causal explanations. These other construals of ‘abstract’ also have their own intuitive pull, even though they are relied on less often or not at all in contemporary discussions, and our elaboration of why each is not the right construal of abstractness reveals also why they could not themselves stand as what makes explanations non-causal.

## **2.1 What Abstractness Is Not**

We believe that the following concepts of abstractness are not suited for articulating the abstractness assumption we are isolating here, although they are genuine uses of the term for other purposes.

- *Abstract explanations = mathematized explanations.*

Suppose one takes explanatory statements to be abstract by virtue of having a mathematized form. Mathematized law statements (for instance in physics) would thereby be abstract and many non-mathematized descriptions of biological-physical mechanisms, especially those including diagrams, would not be. However, this interpretation of abstractness is a non-starter as a way of arguing for abstract explanations being thereby non-causal, since it is uncontroversial that causal relationships can be represented by mathematized statements. For instance, causal Bayes nets and structural

equations are paradigmatic mathematical representations of causal relationships (as has been emphasized, for instance, by Cartwright 1989; Spirtes, Glymour and Scheines 2000; Pearl 2000; Hitchcock 2001, 2007; Woodward 2003). Further, many explanations may be mathematized without thereby being non-causal, distinctively mathematical explanations, at least on many characterizations (Lange 2013; Andersen 2016; Reutlinger forthcoming b). Hence, being in a mathematical form is not sufficient to be abstract in the relevant regard.

- *Abstract explanations = explanations referring to abstract entities.*

Another possible way to understand the notion of an *abstract* explanation is to take it to refer to abstract mathematical entities, i.e. to entities that do not exist in space and time. But, again, this reading of ‘abstract’ cannot be what those relying on the abstractness view have in mind for the following two reasons.

First, if an explanation is abstract by referring to abstract entities (such as certain kinds of graphs or geometric structures) and not also to actual instantiations (or realizations) of these abstract entities (that is, actual objects instantiating a certain kind of graph or geometric structure), then it is unclear as to how such an explanation can account for the occurrence of *contingent* phenomena at all, e.g. physical or biological phenomena (Lange 2013; Humphreys 2014; Reutlinger forthcoming b). However, the abstractness assumption is supposed to refer to explanations of clearly contingent phenomena (as opposed to explanatory proofs in pure mathematics). Hence, this reading of ‘abstract’ is not suited for their concerns.



Second, if we take abstract explanations to refer not only to abstract entities but also to their actual *instantiations* (or realizations), then abstract explanations may or may not be causal. Consider a simple example. Suppose we take our abstract entity to be a graph. If our explanatory assumptions rests on the assumption that a certain kind of graph is realized, then the explanation can be a causal explanation if the graph has a causal interpretation (Pearl 2000; Spirtes et al. 2000). Or it may be a non-causal explanation if the graph has a non-causal interpretation (as in the case of using the actual instantiation of an Eulerian graph for explanatory purposes, see Van Fraassen 1989; Pincock 2012; Lange 2013; Reutlinger forthcoming b). In other words, by stating that some objects instantiate a mathematical structure we have not answered the question whether this structure has a causal or a non-causal interpretation. Thus, this notion of abstractness is also not useful for proponents of the abstractness assumption.

- ***Abstract explanations = explanations involving Aristotelian idealizations (aka abstractions).***

An explanation may be abstract in yet another sense that is commonly used in the literature on idealizations. It is common to distinguish between Aristotelian idealizations and Galilean idealizations (Frigg and Hartmann 2012: Sect. 1.1). Aristotelian idealizations are also discussed in terms of “abstraction” (Cartwright 1989) and “isolation” (Mäki 1992; Hüttemann 2004). Aristotelian idealizations “strip away” some feature(s) that the target system (whose behavior is to be explained) in fact possesses (for instance, a model of a pendulum strips away the color of the pendulum), or Aristotelian idealizations assume that some causal factor actually influencing the target system is

absent. By contrast, Galilean idealizations deliberately distort the target system, for instance, by making the assumption that agents are perfectly rational, that the number the number of molecules in a gas goes to infinity, and so on (McMullin 1985; Strevens 2008; Weisberg 2013).

Now, suppose one takes an explanation to be abstract if its explanans involves an Aristotelian idealization (or “abstraction”). However, the fact that an explanans is abstract in this sense is compatible with it being the explanans of a causal explanation, because we can assert that X causes Y, under the ‘abstract’ assumption that some other causal factor Z actually influencing Y is assumed to be absent or neutralized. As a matter of fact, most of Cartwright’s, Mäki’s, and Hüttemann’s examples of models or explanations involving Aristotelian idealizations have a causal interpretation. Hence, to say the very least, involving Aristotelian idealizations cannot be a sufficient condition for being non-causal.

- *Abstract explanations = explanations referring to “non-mathematical-abstract-objects”.*

Psillos (2010) draws attention to explanations referring to (what he calls) “non-mathematical-abstract-objects” such as “the linear harmonic oscillator, or the two-body Newtonian system, or the frictionless inclined plane” (Psillos 2010: 950), and, perhaps also the center of mass of a physical object. These “objects” may be said to be abstract in the sense that they do not actually exist in space in time (as abstract mathematical objects) – they are fictional objects. As a referee remarked, this sort of abstractness does not imply that “non-mathematical-abstract-objects” cannot be causes – or, at least, be part

of a causal explanation (Pincock 2012: 66-68). The crucial philosophical question regarding such objects is this a question of realism and truth, as Psillos points out: what is the justification for referring to fictional entities in the context of an explanation, if the explanans is required to be (approximately) true – as it is assumed in many theories of explanation? This is a difficult question that we do not aim to answer in this paper. But, more importantly, this question is logically independent from the causal/non-causal distinction, because it also arises if a non-causal explanation refers “non-mathematical-abstract-objects” and if we require non-causal explanantia to be true. Thus, the kind of abstractness Psillos highlights is connected to the issue of truth and falsity, rather than the distinction between causal and non-causal explanation.

## **2.2 What Abstractness Is**

Having ruled out a number of concepts of abstractness, the question remains: what is a sensible interpretation of the concept of abstractness underlying the abstractness assumption such that it appears to fill a certain role in arguing for non-causal explanation? In this section, we examine how Batterman and Pincock – sometimes rather explicitly, sometimes merely implicitly – rely on the abstractness assumption in the recent literature as illustrations of the role this assumption plays. We hold that, when articulating the abstractness assumption, Pincock and Batterman rely on the following notion of abstractness:

- *Abstract explanations = explanations with an explanans that has multiple microphysical realizers.* An explanation is *abstract* insofar as its explanans leaves out many or almost all causal microphysical details about the target phenomenon

of the explanation. We will use the notion of multiple realizability<sup>4</sup> to refer to such an explanans, since the explanans in question can have different microphysical realizations. We feel partly justified to choose this terminology as, for instance, Batterman adopts it to describe abstractness (see below). Put in a slogan, the more multiply realizable an explanans is, the more abstract it is.

Now let us turn to Pincock's and Batterman's strategies for spelling out the abstractness assumption based on the notion of 'abstract' just presented.

(i) **Pincock.** In his recent book *Mathematics and Scientific Representation* (2012), Pincock introduces a distinction between abstract acausal and concrete causal representations. Although Pincock (2012) is primarily concerned with distinguishing causal and acausal *representations*, his distinction applies to explanations utilizing those representations, because (a) he often describes the explanatory use of scientific representations (Pincock 2012: 53, 203-210, 221-229), and (b) Pincock himself explicitly applies the distinction "abstract acausal" versus "concrete causal" to scientific explanations in later work (see Pincock 2015).

Moreover, Pincock distinguishes between *abstract acausal* and *abstract varying* representations (Pincock 2012: 66-68). According to his framework, only the former kind of abstract representation is crucial for understanding what makes a representation (and, derivatively, an explanation) non-causal. The "varying" kind of

---

<sup>4</sup> We use the notion of multiple realizability to refer to (a) the fact that some *types* of higher-level behaviors can be realized by multiple types of microphysical behaviors, and (b) the fact that some *tokens* of higher-level behaviors can be instantiated by multiple tokens of microphysical behaviors. However, the differences between (a) and (b) are not central to our concern in this paper.

abstractness, Pincock argues, need not “deprive the representation of its causal content” (Pincock 2012: 67). Pincock’s chief example of an *abstract varying* representation is the linear harmonic oscillator that may be a *causal* representation of physical systems that differ significantly in microphysical causal details (such as metal springs and pendulums). From our perspective, there is nothing to disagree with. The main challenge we pose in this paper is that the same moral holds for the other kind of abstract representation: the abstract acausal character – as defined by Pincock – also needs not (and often does not) deprive a representation of its causal content.

Pincock illustrates his notion of a concrete causal representation by invoking the representation of “the causal structure of the development of a traffic jam” (Pincock 2012: 5):

If everything worked properly, these transformations [i.e. the causal laws about the propagation of cars along the road] would mimic the causal effects of the cars on one another, for example, how slowing down a car would change the velocity of the cars behind it. (ibid.)

He contrasts the concrete causal representation with an abstract acausal representation of the development of a traffic jam:

There are mathematical representations where causes are absent. [...] I call such cases *abstract acausal* because they abstract away from the causal relationships found in the concrete causal representation. In our traffic case, we might imagine a representation of a given system of cars that simply

represented what the system would look like after all the cars had achieved some fixed velocity for a given density of cars. (Pincock 2012: 6)

Pincock's point regarding the traffic jam example seems to be that the second representation is non-causal (or 'acausal' in his terminology) because it abstracts away from many of the "concrete" details of "the causal effects of the cars on one another" that are built into the concrete causal representation. The abstract acausal representation replaces those causal details of "how slowing down a car would change the velocity of the cars behind it" by "some fixed velocity for a given density of cars". We take this example to indicate that Pincock's notion of abstractness rests on the intuition that causal representation is (broadly) mechanistic in character: causal representations and explanations must involve explicit reference to detailed mechanisms, i.e. reference to details of how the components of the target system of the explanation (or representation) engage in causal interactions over time. According to Pincock's notion of abstractness, the explanans of a non-causal explanation abstracts away from specific details of the causal interactions between the components of a system.

Moreover, Pincock argues that abstracting away from details is best understood in terms of counterfactual stability. He claims that the counterfactuals needed for defining counterfactual stability "are quite special and should not be mixed in with the sort of counterfactuals that are that central to causal representations" (Pincock 2012: 54). In order to define counterfactual stability, he distinguishes two kinds of counterfactuals about the system whose behavior is supposed to be explained: difference-making counterfactuals and stability counterfactuals (this is our terminology). Pincock holds that *difference-making counterfactuals* assert that, roughly, had the (purported) cause been

different then the effect would have been different as well. By contrast, *stability counterfactuals* reveal that many of the microphysical details of a system could have been different but the explanatorily relevant factors would have remained the same.<sup>5</sup> In other words, the explanans at issue is multiply realizable. We believe that Pincock's discussion of abstractness in terms of stability counterfactuals supports our interpretation that the notion of abstractness underlying is this: *explanations are abstract iff their explanantia are multiply realizable by microphysical states in the sense defined above.*

(ii) **Batterman.** We now turn to a second illustration of the abstractness assumption. Batterman argues that causal explanations do not and in fact cannot ignore causal "micro details" because causal explanations provide "detailed causal-mechanical accounts of the workings of the mechanisms leading to the occurrence of the explanandum phenomenon" (Batterman 2000: 28). A causal explanation "tells us all of the gory details" (ibid.) about why a particular effect occurs (Batterman 2010: 2, 21). Moreover, Batterman holds that ignoring certain details about the causal interactions of the components of a physical system is essential for non-causal explanations. One of Batterman's central examples supporting this claim is the central role of abstractness in

---

<sup>5</sup> Pincock uses the following example as an illustration. Suppose a traveler attempts to traverse all of the bridges of Königsberg without crossing one bridge twice. The traveler fails. Why did the traveler fail? The non-causal explanation of this failure consists in realizing that the bridge system of Königsberg instantiates the structure of a non-Eulerian graph and due to the formal properties of such a graph it impossible to traverse all of the bridges exactly once (Pincock 2012: 51-53). Pincock claims that this explanation is abstract non-causal because the explanatorily relevant factor (i.e. Königsberg instantiating the structure of a non-Eulerian graph) is counterfactually stable under changes in the micro-constitution of the bridges: "Supposing that the bridges are made of stone, we can see clearly that changing the bridges to gold, while preserving the isomorphism to the relevant graph, would keep the impossibility of the circuit intact." (Pincock 2012: 54).

renormalization group explanations. Explanations of this kind are used to explain why microscopically different physical systems exhibit universal macro behavior near the point of critical temperature. Batterman argues that one crucial element of this sort of explanation is to abstract away from the microphysical causal details of condensed matter physical systems. In the specific case of a renormalization group explanation, the abstraction away from the microphysical causal details takes the form of reiterating renormalization group transformations (and asymptotic reasoning, such as taking the thermodynamic limit). Batterman *verbatim*:

In effect, the renormalization group transformation eliminates degrees of freedom (microscopic details) which are inessential or irrelevant for characterizing the system's dominant behavior at criticality. (Batterman 2000: 127)

R[enormalization] G[roup] type analysis illuminates those physical features that *are* relevant for the upper level universal behavior, *and at the same time demonstrates that all of the other details which distinguish the systems from one another are irrelevant*. (Batterman 2000: 128; see also Batterman 2000: 138-139, 2002: 40-42, 2010: 7)

Batterman seems to suggest that, while a causal explanation “tells us all of the gory details”, the non-causal character of renormalization group explanations consists in the abstraction away from or the “systematic throwing away” (Batterman 2010: 3) of microphysical causal details.

What is more, Batterman makes the very same point in terms multiple



realizability. In various publications, Batterman equates universality and multiple realizability in stating that the explanandum of RG explanations is why a certain macroscopic property is multiply realized (Batterman 2000, 117; 2002, 72, 2015, 8; see also Reutlinger 2014, 2016). That is, he interprets the universality of macro-behavior in terms of the multiple realizability of a kind of macro-behavior. To take one prominent example, Batterman interprets renormalization group explanations as meeting the “challenge of multiple realizability” expressed in the question “how can systems that are heterogeneous at some (typically) micro-scale exhibit the same pattern of behavior at the macro-scale?” (Batterman 2015: 8).

Having reviewed Batterman’s and Pincock’s versions of the abstractness assumption, we would like to point out that there is a substantial issue of disagreement between Batterman and Pincock that is, however, immaterial to our concern. Although Batterman seems to be drawn to the intuition that the non-causal character of an explanation is closely connected to abstractness via “the systematic throwing away of various causal and physical details” (Batterman 2010: 3), he does not agree with all aspects of Pincock’s argument concerning abstractness. In particular, Batterman wants to resist indispensability arguments for mathematical realism. Batterman (2010: section 2) does not want to accept that the central role of mathematics in scientific (non-causal) explanations necessarily leads to a metaphysical commitment to the existence of mathematical entities and their properties (such as numbers, graphs, sets, and their properties). Rather he contrasts mathematical entities and mathematical operations (such as taking the thermodynamic limit and renormalization group transformations; see Batterman 2010: section 3) and argues that the workings of non-causal explanations such

as renormalization group explanations “appeal to (or better ‘involve’) mathematical operations” (Batterman 2010: 4). Improving our understanding of how explanations involving mathematical operations work does not imply any commitment to the existence of mathematical entities, or so Batterman argues. This disagreement about ontological commitment is irrelevant for our present concerns, because Batterman and Pincock agree on a notion of abstractness and take it to be useful for distinguishing between causal and non-causal explanations.

In sum, we take Pincock and Batterman to use a notion of abstract explanation, according to which an explanation is *abstract* if and only if its explanans leaves out many or almost all causal microphysical details about the target phenomenon of the explanation. We describe an explanans that is abstract in this way as being multiply realizable. Plugging this notion of abstractness into the *abstractness assumption*, we obtain the following statement: an explanation is non-causal to the extent that its explanans leaves out many or almost all causal microphysical details about the target phenomenon of the explanation). In the next section, we will raise two objections to the abstractness assumption.

### **3. Two Problems for the Abstractness Assumption**

Our main concern with the abstractness assumption is that it is in conflict with many *standard accounts of causation* and also with *major causal theories of scientific explanation*. More precisely, we will now raise two problems for the abstractness assumption:

1. Abstract explanations can, and in fact often do, identify causes according to standard accounts of causation (Section 1.1).
2. Abstract explanations can be, and in fact often are, causally explanatory according to the standard causal accounts of scientific explanation (Section 1.2).

By considering a range of theories of causation and of causal explanation, we aim to show how none of them supports the opposition between abstractness and causation on which the abstractness assumption is based. Regardless of one's philosophical commitments regarding causation and causal explanation, abstractness will not serve as a distinguishing marker for non-causal versus causal explanations.

### **3.1 First Problem: Abstractness and Theories of Causation**

Causal explanations can be, and sometimes actually are, abstract. We back up this claim, in this section, by going through a number of established theories of causation, many or arguably all of which are mutually incompatible, and showing how each nevertheless allows for abstract causal explanations, with concrete examples of such explanations. This might strike some as excessively thorough; our goal in doing so is twofold. First, it highlights how the intuitive plausibility of the opposition between abstractness and causation fades when considered in light of specific theories of causation. Second, it clearly shows how the failure of the abstractness assumption to track the distinction between causal versus non-causal explanation is independent of one's other philosophical commitments regarding causation. In order to establish these points, we will go through the following accounts of causation: (i) counterfactual accounts, (ii) mechanistic accounts,

(iii) process accounts, and we will conclude with (iv) a brief outlook on regularity, dispositionalist, and primitivist accounts of causation.

**(i) *Counterfactual accounts of causation.***

To begin with, consider counterfactual accounts of causation, including but not limited to Woodward's (2003) popular interventionist account, which are the most widely accepted accounts of causation in current philosophy of science.<sup>6</sup> The common core of such (broadly construed) counterfactual accounts is that X causes Y iff Y counterfactually depends on X, in an appropriate way. Distinct versions of the counterfactual account mostly differ with respect to how they specify the 'appropriate' kind of counterfactual dependence. For instance, in contrast to David Lewis (1973), interventionists claim that the 'appropriate' counterfactuals are those referring to possible interventions on the (purported) cause variable. Counterfactual accounts are extremely powerful as they capture (a) causal relationships holding between types<sup>7</sup> and tokens of events (both are usually represented by random variables), (b) deterministic and probabilistic causation, and (c) counterfactual accounts successfully capture different degrees of causal strength (typically expressed by the coefficients figuring in structural equations).

---

<sup>6</sup> We take counterfactual theories, broadly construed, to include Lewis' (1973) classic counterfactual account; recent interventionist developments such as Hausman (1998), Hitchcock (2001), and Halpern and Pearl (2005); and also some agency theories of causation such as von Wright (1974), and Menzies and Price (1993).

<sup>7</sup> It is a common place in today's philosophy of causation that there is token-level and type-level causation. As a matter of fact, type-level causes play a far more prominent role in the sciences than token-level causation. However, Batterman seems to hold that there is only token-level causation (Batterman 2002: 24-25; 2010: 21). Not acknowledging the existence of type-level causation is another point regarding which Batterman is in tension with the causation literature and, strikingly, science.

The crucial question is whether counterfactual accounts allow for ‘abstract causation’. Using our earlier characterization of the abstractness assumption, the answer is: yes, they do. According to the counterfactual account there may be causal relations that remain stable under (actual and counterfactual) changes in the physical micro-details, and also under changes in initial conditions and background conditions (see Hitchcock and Woodward 2003: sect. 3; Woodward 2003: ch. 6; Woodward 2010). This point can also be rephrased in semantic terms: abstract causal *statements* are compatible with the counterfactual account, since the variables (representing cause and effect) can be indefinitely abstract, if they still have some empirical content, if they involve a well-defined notion of change, and if counterfactual dependence holds (that is, if it is true that if variable X were to take the value of x rather than x\*, then variable Y would take the value of y rather than y\*). As a result, a statement can be a causal statement, even though each of the variables figuring in this causal statement (representing cause and effect) abstracts to an enormous degree from the physical micro details of any concrete situation in which the statement is true. Consider two classic examples of causal statements: “regular smoking is a cause of lung cancer” and “HIV is a cause of AIDS”. Both causal statements are abstract (since the description of the cause and effect leave out many microphysical causal details). Proponents of counterfactual accounts can straightforwardly consider ‘abstract’ causal relata such as HIV/AIDS and regular smoking/lung cancer as well as, for instance, democratic governance, natural resources, civil unrest, inflation, aggression, temperature, and so on.<sup>8</sup>

---

<sup>8</sup> A similar claim seems to hold for the epistemology of causation, if one adopts a counterfactual account. The existence of a causal relationship can be clearly established using evidence that does not ever have to involve a detailed mechanism, or any detailed

To sum up our discussion of counterfactual accounts, abstract causal relationships can be stated and captured in the framework of counterfactual accounts. Finally, we would like to stress that it is not merely a logical point to insist that counterfactual accounts are *compatible* with abstract causal claims. Instead it is the *primary goal* of many advocates of counterfactual theories to capture and understand the abstract causal claims of the life sciences and the social sciences, such as “regular smoking is a cause of lung cancer” and “HIV is a cause of AIDS” (Hausman 1998; Hoover 2001; Woodward 2003 clearly pursue this goal).

However, one need not accept a commitment to counterfactual accounts of causation for the abstractness assumption to be problematic. The same point can be made if one draws on alternative accounts of causation.

***(ii) Mechanistic accounts of causation.***

Let us begin with mechanistic accounts of causation. Even within mechanist accounts of causation, there is room for abstractness. (Note that we distinguish between mechanistic accounts of *causation* and mechanistic accounts of *explanation*, as the latter do not imply the former; see Andersen 2014a, 2014b). Mechanistic accounts of causation – such as Glennan’s (1996, 2010) – portray the world as nested layers of mechanisms, each of which is constituted by the entities and interactions at a (relatively) lower level until one reaches the fundamental physical level of reality. Taking the notion of a mechanism as primitive, the core idea of Glennan’s mechanistic account of causation is that, roughly

---

or concrete connection between the cause and effect in question. For instance, a randomized controlled trial may provide suitable evidence for the effective workings of a new drug without providing evidence about the detailed mechanisms of how the drug influences the test subjects (see Woodward 2003: 95-98).

put, causal statements (if true) describe a mechanism in this layered hierarchy of mechanism. The mechanistic account of causation allows for abstractness in at least two ways: (a) the representation of a mechanism need not “bottom out” (i.e. need not refer to the most fundamental microphysical mechanisms) and still be a causal representation. (b) Each of these layers of mechanisms (including the fundamental layer) can be represented in a way that leaves out certain details without compromising the causal character of representation (for instance, the representation may omit details regarding the dynamical interactions between components). As such, mechanists about causation readily acknowledge that every kind of mechanism allows for abstractness (see Craver 2013; Andersen 2014a, 2014b). For instance, a mechanist about causation may acknowledge that the model of a mechanism may be abstract because the model emphasizes certain features of the spatial organization of the causally interacting components of a mechanism (while idealizing away from the dynamical details of the interaction among components). However, an emphasis of this kind does not render the resulting explanation non-causal.

***(iii) Process accounts of causation.***

Even the process accounts of causation developed by Salmon allow for abstract causal claims. Salmon’s earlier mark transmission account (Salmon 1984: 148) as well as his later conserved quantity account (Salmon 1998: 257) are explicitly taken to apply to a wide range of higher-level or macroscopic causes, which he takes to be unproblematically causal even though they leave out many microphysical details and can be microphysically multiply realized. For instance, Salmon counts a moving car (Salmon

1984: 143) and “recovery from neurotic symptoms” with psychotherapy as causal processes (Salmon 1998: 110) – the description of both causal processes is articulated in macroscopic terms that do not (and need not) involve any references to the microphysical causal details of cars and patients.

Similarly, Dowe’s conserved quantity theory of causation allows for abstractness. Dowe defines a causal process as “the worldline of an object that possesses a conserved quantity” (Dowe 2000: 90). He explicitly permits that tracking such a process may happen at a coarse-grained or abstract level of description. The object involved in a causal process may be “anything found in the ontology of science (such as particles, waves and fields), or common sense (such as chairs, buildings and people).” (Dowe 2000: 91). Surely, describing causal processes involving chairs, buildings and people is abstract in leaving out various microphysical causal details. Hence, there are many items in the ontology of science that can be represented in very abstract ways, yet still qualify as causal processes on Dowe’s account.<sup>9</sup>

*(iv) A brief outlook on other accounts of causation.*

Although we cannot go through all prominent accounts of causation in the same detailed manner, we believe it is at least plausible that, for instance, regularity theories (Mackie 1980, Baumgartner 2008) and dispositionalist accounts of causation (Cartwright 1989, 1999; Hüttemann 2004; Bird 2005) are often intended to capture abstract causal claims. According to regularity and dispositionalist theories, scientists often describe regularities

---

<sup>9</sup> Note that Dowe rejects many cases often considered causal, putting them in the category of “causation\*” instead, making his conserved quantity theory arguably one of the most restrictive accounts of causation. Yet still, even this account allows for the kind of abstractness taken by the abstractness assumption to be indicative of non-causality.



and, respectively, dispositions in an abstract vocabulary. The main motivation for capturing the abstract regularities and abstract dispositions is to account for causal claims of the life sciences and social sciences. Moreover, *primitivist accounts* of causation take causation to be a conceptually and metaphysically fundamental ‘productive’ relation. Primitivists either explicitly claim, or suggest by the biological or common sense examples of causal claims they use, that claims about productive relations need not refer to microphysical causal details (Anscombe 1971; Cartwright 1989; Machamer et al. 2000; Machamer 2004; Bogen 2005).

To sum up this section, influential standard accounts of causation either allow or explicitly aim for capturing abstract causal claims (that may and often do figure in causal explanations). Thus, abstractness does not track the causal versus non-causal status of an explanation, independently of one’s commitment to a particular theory of causation.

### **3.2 Second Problem: Abstractness and Causal Theories of Explanation**

The abstractness assumption is also unsupported by the literature on causal explanations, and again, this is independent of particular theories of causal explanation. This is for two reasons.

*(1) There are abstract causal explanations according to standard causal accounts of explanation.* Relying on the intuitive pull of the abstractness assumption, Batterman, for instance, argues that causal explanations cannot ignore causal “micro details” because causal explanations provide “detailed causal-mechanical accounts of the workings of the mechanisms leading to the occurrence of the explanandum phenomenon”

(Batterman 2000: 28). A causal explanation “tells us all of the gory details” (ibid.) about why a particular effect occurs (Batterman 2010: 2, 21).

While we agree that there *are* causal explanations that include “all the gory details,” there are also many that do *not* include all those details, yet remain causal and explanatory. Consider an example of an abstract causal explanation from non-fundamental physics. Woodward and Hitchcock present this example as their prime illustration of causal explanations in science. The gist of the explanation can be outlined as follows. The explanandum in question is the “magnitude of the electric field created by a long, straight wire with a positive charge uniformly distributed along its length” (Woodward and Hitchcock 2003: 4; also Woodward 2003: 187-189). Coulomb’s law is the central piece of the explanans, as this law

“can be used to tell us how the electric field would differ if the charge density of the wire were increased, or if the wire twisted into a circle or a solenoid. In this way, [Coulomb’s law] shows us that certain factors, such as charge density and geometrical configuration of the conductor, make a systematic difference to intensity and direction of the field. In short, Coulombs law is explanatory because it tells us what the electric field *depends on*.” (ibid.; original emphasis)

Woodward and Hitchcock take this explanation to be causal – an assumption they justify by drawing on their interventionist account of causation. Moreover, they argue that their account of causal explanations, a counterfactual account of explanation, adequately

captures the explanation based on Coulomb's law. However and most importantly for our concerns, the explanation is not only causal but also clearly abstract, because it rests on a non-fundamental law (Coulomb's law) that leaves out many microphysical causal details.

Moreover, Batterman's characterization of causal explanation as *non-abstract* applies only to a very specific version of the causal account of explanation – such as, perhaps, Railton's account of an ideal explanatory text, to which Batterman (2002: 28) refers. The majority of causal theories of explanation are explicitly designed to account for the fact that many excellent causal explanations “ignore details” (that is, are abstract in the sense in question). For instance, Strevens (2008) and Franklin-Hall (forthcoming) explicitly hold that merely identifying the causes of the explanandum phenomenon is necessary but not sufficient for providing a causal explanation. Besides identifying causes, an “optimizing procedure” (Strevens) and a “biggest bang-for-your-buck” procedure (Franklin-Hall) have to be applied for causally explaining a phenomenon in order to filter out the explanatorily relevant causes. These procedures are recipes for omitting irrelevant causal information, namely, the mechanistic details assumed by the abstractness assumption to be necessary and sufficient for causal explanation.

Similarly, Woodward (2010) uses the notion of causal specificity to determine the accurate level of explanatory abstraction from causal details (see also Hitchcock and Woodward 2003: sect. 3). Good causal explanations are, on Woodward's view, often (but not necessarily always) those that are coarse-grained with respect to microphysical causal details. In the same spirit, many proponents of causal accounts of explanation hold that it is a virtue of the explanatory generalization (figuring in the explanans of a causal explanation) if it is stable under (actual and counterfactual) changes in the micro details,

the initial conditions and background conditions (see Hitchcock and Woodward 2003; Strevens 2008; Woodward 2010). Therefore, the abstractness assumption is in contradiction with several major accounts of causal explanation.

*(2) Abstractness is compatible with mechanistic accounts of explanation.* In analogy with mechanist accounts of *causation*, mechanist accounts of *explanation* have considerable room for abstractness without compromising the causal character of mechanistic explanations. For instance, in the Machamer, Darden, and Craver's (2000) version of mechanistic explanation, models of mechanisms need not bottom out at any privileged level of description. In the overwhelming majority of cases, any model of a mechanism will abstract away from some of the microphysical causal details in order to provide a more salient explanation (see Craver 2013; Andersen 2014a, 2014b). One reason for this is that mechanistic explanations typically are based on the model of a mechanism *type*: that is, the explanans of a mechanistic explanation usually refers to a type of mechanism, rather than merely a single token instance, and the explanandum of such an explanation is a certain type, rather than a token, of a phenomenon. A model of a type of mechanism is paradigmatically abstract, as it leaves out many irrelevant microphysical causal details of any particular instantiation of this type.<sup>10</sup>

To sum up, according to several major causal theories of explanation, there can be and there in fact are abstract causal explanations. This result provides another reason to doubt that the abstractness assumption is an adequate strategy for distinguishing causal and non-causal explanations, regardless of one's other commitments concerning causal explanation.

---

<sup>10</sup> Kuhlmann's (2014) recent work on structural mechanistic explanations in econophysics provides another interesting example.

## 4. Conclusion

We started out with the question of whether an explanation is non-causal by virtue of being abstract, or, more cautiously, to what extent increasing the abstractness of an explanation reduces the causal character of that explanation. According to the abstractness assumption, the answer is that the abstract character of an explanation makes it non-causal; abstractness is in tension with the causal character of an explanation where more of one means less of the other. We have undermined this apparent tension in arguing that abstractness is not a sufficient condition for being a non-causal explanation, nor does increasing abstractness by leaving out many microphysical causal details thereby render an explanation non-causal. Key to our arguments is that this rejection of the abstractness assumption is independent of one's other commitments regarding theories of causation and causal explanation. We conclude from our discussion that the abstractness assumption fails to adequately capture what makes a scientific explanation causal or, respectively, non-causal, and that the *prima facie* intuitive pull of this assumption should be resisted.

Looking forward from this conclusion, what changes if one resists the intuitive plausibility of the abstractness assumption? First and foremost, although this paper is written in an exclusively critical mode, we believe that there are (more) convincing ways to draw a distinction between causal and non-causal explanations (as elaborated in Andersen 2016; Reutlinger forthcoming a, b). However, this paper is not the place to present our own views. We believe that the way forward in the debate is to elaborate the alternatives to the abstractness assumption.

Another useful take-away point of our paper is that the issue of distinguishing causal and non-causal explanations should not be conflated with other closely related but distinct topics in the philosophical analysis of scientific explanations. These closely related discussions need extra care to avoid accidentally invoking abstractness in arguing for the non-causal character of a particular type of explanation. For instance, Batterman (2010) and Batterman and Rice (2014) seem to be mainly interested in understanding how it is possible that highly idealized models in physics and biology can be used for explanatory purposes. In particular, they are concerned with the question whether the time-honored requirement that the explanans of an explanation be true – a requirement built into many theories of explanation, as they correctly stress – can be maintained for highly idealized explanations (especially if the idealizations are “indispensable”, as they assume). However, answering this genuinely important question does in no way require reliance on the abstractness assumption (see, for instance, Saatsi forthcoming b). Similarly, Pincock’s (2012) main goal is to articulate ways in which mathematics contributes epistemically to the success of scientific representations. While perhaps some way of differentiating causal and non-causal representations is needed, the abstractness assumption does not seem to be required for Pincock’s own goal of understanding the epistemic role mathematics in scientific representation. Realizing that these issues are logically independent from the causal/non-causal distinction also constitutes progress in the debate on scientific explanation.

## Acknowledgements

We are grateful to Maria Kronfeldner, Insa Lawler, Chris Pincock, and Juha Saatsi, as well as to audiences in Helsinki, Leeds, and Munich for their constructive comments and suggestions. Particularly, we would like to thank four anonymous referees and the editor of this journal for their critical remarks that greatly helped to sharpen the argument of the paper. Andersen is grateful for the opportunity to live and work on unceded Coast Salish territory.

## References

- Andersen, H. [2014a]: ‘A field guide to mechanisms: part I’, *Philosophy Compass* 9(4), pp. 274-283.
- Andersen, H. [2014b]: ‘A field guide to mechanisms: part II’, *Philosophy Compass* 9(4), pp. 284-293.
- Andersen, H. [2016]: ‘Complements, Not Competitors: Causal and Mathematical Explanations,’ *British Journal for Philosophy of Science* Advanced Access August 18, 2016, doi:10.1093/bjps/axw023.
- Batterman, R. [2000]: ‘Multiple Realizability and Universality’, *British Journal for Philosophy of Science* 51, pp. 115-145.
- Batterman, R. [2002]: *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction and Emergence*. New York: Oxford University Press.

- Batterman, R. [2010]: 'On the Explanatory Role of Mathematics in Empirical Science', *British Journal for Philosophy of Science* 61, pp. 1-25.
- Batterman, R. and C. Rice [2014]: 'Minimal Model Explanation', *Philosophy of Science* 81, pp. 349-376.
- Baumgartner, M. [2008]: 'Regularity Theories Reassessed', *Philosophia* 36, pp. 327-354.
- Bechtel, W. and R. Richardson [1993]: *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*, Princeton, New Jersey: Princeton University Press.
- Bird, A., [2005] 'The Dispositionalist Conception of Laws', *Foundations of Science* 10: 353-370.
- Bogen, J. [2005]: 'Regularities and causality; generalizations and causal explanations', *Studies in History and Philosophy of Biological and Biomedical Science* 36, pp. 397-420.
- Bokulich, A. [2008]: 'Can Classical Structures Explain Quantum Phenomena?' *British Journal for the Philosophy of Science* 59(2), pp. 217-235.
- Cartwright, N. [1983]: *How the Laws of Physics Lie*, Oxford: Oxford University Press.
- Cartwright, N. [1989]: *Nature's Capacities and Their Measurement*, Oxford: Clarendon Press.
- Cartwright, N. [1999]: *The Dappled World: A Study of the Boundaries of Science*, Cambridge: Cambridge University Press.



- Chirimuuta, M. [2014]: ‘Minimal models and canonical neural computations: the distinctness of computational explanation in neuroscience’ *Synthese* 191(2), pp 127-153.
- Craver, C. [2007]: *Explaining the Brain*, New York: Oxford University Press.
- Craver, C. [2013]: ‘The Ontic Account of Scientific Explanation’, In *Explanation in the Special Sciences. The Case of Biology and History*, *Synthese Library*, A. Hüttemann and M. Kaiser (eds.), pp. 27-54, Dordrecht: Springer.
- Dowe, P. [2000]: *Physical Causation*, New York: Cambridge University Press.
- Franklin-Hall, Laura [forthcoming]: ‘The causal economy account of scientific explanation’, In *Minnesota Studies in the Philosophy of Science*, K. Waters and J. Woodward (eds.).
- Glennan, S. [1996] ‘Mechanisms and the Nature of Causation’, *Erkenntnis* 44: 49-71.
- Glennan, S. [2010]: ‘Mechanisms, causes, and the layered model of the world’, *Philosophy and Phenomenological Research* 81(2), pp. 362–81.
- Halpern, J.Y., and Pearl, J. [2005]: ‘Causes and Explanations: A structural-model approach’, *British Journal for the Philosophy of Science* 56(4), pp. 843-847 and 889-911.
- Hausman, D. [1998]: *Causal Asymmetries*, Cambridge, UK: Cambridge University Press.
- Hitchcock, C. [2001]: ‘The Intransitivity of causation revealed in equations and graphs’, *The Journal of Philosophy* 98(6), pp. 273-299.
- Hitchcock, C. [2007]: ‘What Russell Got Right’, In H. Price and R. Corry (eds.)

- Causation, Physics, and the Constitution of Reality. Russell's Republic Revisited*,  
Oxford: Clarendon Press, pp. 45-65.
- Hitchcock, C. and Woodward, J. [2003]: 'Exploratory Generalizations Part II: Plumbing  
Exploratory Depth', *Nôus* 37: 181-199.
- Hoover, K. [2001]: *Causality in Macroeconomics*, Cambridge, UK: Cambridge  
University Press.
- Humphreys, P. [2014]: 'Explanation as Condition Satisfaction', *Philosophy of Science*  
81: 1103-1116.
- Huneman, P. [2010]: 'Topological Explanations and Robustness in Biological Sciences',  
*Synthese* 177: 213-245.
- Hüttemann, A. [2004]: *What Wrong with Microphysicalism?* London: Routledge Press.
- Kuhlmann, M. [2014]: 'Explaining Financial Markets in Terms of Complex Systems',  
*Philosophy of Science* 81(5): pp. 1117-1130.
- Lange, M. [2009]: "Dimensional Explanations." *Noûs* 43: 742-775.
- Lange, M. [2011]: 'Conservation Laws in Scientific Explanations: Constraints or  
Coincidences?', *Philosophy of Science* 78: pp. 333-352.
- Lange, M. [2013a]: 'What Makes an Explanation Distinctively Mathematical?' *British  
Journal of the Philosophy of Science* 64(3), pp. 485-511.
- Lange, M. [2013b]: 'Really Statistical Explanations and Genetic Drift', *Philosophy of  
Science* 80, pp. 169-188.
- Lange, M. [2015]: "On 'Minimal Model Explanations': A Reply to Batterman and Rice",  
*Philosophy of Science* 82: 292-305

- Lewis, D. [1973]: 'Causation', *The Journal of Philosophy* 70(17), pp. 556-567.
- Machamer, P., L. Darden, and C. F. Craver [2000]: 'Thinking about mechanisms', *Philosophy of Science* 67(1), pp. 1–25.
- Machamer, P. [2004]: 'Activities and Causation: The Metaphysics and Epistemology of Causation', *International Studies in the Philosophy of Science* 18(1), pp. 27-39.
- Mackie, J.L. [1980]: *The Cement of the Universe*, Oxford: Oxford University Press.
- Mäki, U. [1992]: 'On the Method of Isolation in Economics', *Poznan Studies in the Philosophy of the Sciences and the Humanities* 26: 319-354.
- McMullin, E. [1985]: 'Galilean Idealizations', *Studies in the History and Philosophy of Science* 16: 247-273.
- Menzies, P. and H. Price [1993]: 'Causation as a secondary quality', *British Journal for the Philosophy of Science* 44, pp. 187-203.
- Pearl, J. [2000]: *Causality: Models, Reasoning, and Inference*, Cambridge University Press.
- Pincock, C. [2012]: *Mathematical and Scientific Representation*, New York: Oxford University Press.
- Pincock, C. [2015]: 'Abstract Explanations in Science', *British Journal for the Philosophy of Science* 66.4: 857-882.
- Psillos, S. [2010]: 'Scientific Realism: Between Platonism and Nominalism', *Philosophy of Science* 77: 947-958.
- Reichenbach, H. [1956]: *The Direction of Time*, Berkeley and Los Angeles, California: University of California Press.

- Reutlinger, A. (forthcoming a) “Does the Counterfactual Theory of Explanation Apply to Non-Causal Explanations in Metaphysics?”, *European Journal for Philosophy of Science*.
- Reutlinger, A. (forthcoming b) “Is There A Monist Theory of Causal and Non-Causal Explanations? The Counterfactual Theory of Scientific Explanation”, *Philosophy of Science*.
- Reutlinger, A. (2016) “Do Renormalization Group Explanations Conform to the Commonality Strategy?”, *Journal for General Philosophy of Science*, Online First, DOI: 10.1007/s10838-016-9339-7.
- Reutlinger, A. (2014) “Why Is There Universal Macro-Behavior? Renormalization Group Explanation As Non-causal Explanation.” *Philosophy of Science* 81: 1157-1170.
- Saatsi, J. [forthcoming a]: “On Explanations from ‘Geometry of Motion’” *The British Journal for the Philosophy of Science*.
- Saatsi, J. [forthcoming b]: ‘On the “Indispensable Explanatory Role” of Mathematics’, *Mind*.
- Salmon, W. [1998]: *Causality and Explanation*, Oxford: Oxford University Press, 1998.
- Spirtes, P., C. Glymour, and R. Scheines [2000]: *Causation, Prediction, and Search*, Cambridge, MA: The MIT Press.
- Strevens, M. [2008]: *Depth*, Cambridge, MA: Harvard University Press.
- von Wright, G.H. [1974]: *Causality and Determinism*, New York: Columbia University Press.
- Weatherall, J. [2011]: ‘On (Some) Explanations in Physics’, *Philosophy of Science* 78: pp.

421-447.

Weisberg, M. [2013]: *Simulation and Similarity*, New York: Oxford University Press.

Woodward, J. [2003]: *Making Things Happen*, New York: Oxford University Press.

Woodward, J. and Hitchcock, C. [2003]: 'Explanatory Generalizations, Part I: A Counterfactual Account', *Noûs* 37(1): 1-24.