No Place for Causes?
Causal Skepticism in Physics
Mathias Frisch

Abstract

According to a widespread view, which can be traced back to Russell's famous attack on the notion of cause, causal notions have no legitimate role to play in how mature physical theories represent the world. In this paper I critically examine a number of arguments for this view that center on the asymmetry of the causal relation and argue that none of them succeed. More positively, I argue that embedding the dynamical models of a theory into richer causal structures can allow us to decide between models in cases where our observational data severely underdetermine our choice of dynamical models.

1. Introduction

Is there a place for the concept of cause in physics? There is a venerable tradition, dating back at least to Bertrand Russell's famous "On the Notion of Cause" (1918), arguing that causal notions can play no legitimate role in how physics represents the world. Russell himself, of course, argued not only that causal notions have no place in modern physics, but also that this implies that the notion of cause ought to be expunged entirely from our conception of the world. By contrast, recent defenders of broadly neo-Russellian arguments concede that some concept of cause may play a useful role in our folk understanding of the world or even in the special sciences—causal concepts, they argue, are an important and perhaps ineliminable component of our conception of the world from our perspective as intentional agents embedded into that world—but they maintain, in agreement with Russell, that such concepts have no place in the conception of the world presented to us by our well-established theories of physics.

The list of philosophers who have expressed skepticism about causation in physics is both extensive and distinguished, including Patrick Suppes (1970), Richard Healey (1983), Bas van Fraassen (1992), Hartry Field (2003), Huw Price (2007), John Norton (2003, 2007, 2009), and John Earman (2011). Others, including Jim Woodward (2007) and Chris Hitchcock (2007) have argued that at least certain aspects of our common sense causal notion become increasingly strained when applied to physics. The general argumentative strategy consists in pointing to one or several supposed contrasts between causal relations and the kind of structures presented to us by the

theories of physics and to argue that the existence of these contrasts undermines the applicability of causal reasoning to physics. Perhaps the most telling contrast is supposed to be that between the asymmetry of the causal relation and the putative time-reversal invariance of the dynamical laws of our mature physical theories. In this paper I want to examine whether this particular contrast can be forged into a successful argument to show that there is no place for causal representations in physics and will argue that anti-causal arguments appealing to the asymmetry of causation fail.

Even though I will defend the thesis that causal notions can play a role in physics, I am not interested in arguing for a range of views that Norton has grouped together as "causal fundamentalism" or "causal foundationalism": I am not defending the view that causal notions play a role in all our well-established theories of physics, or that our most fundamental theories of physics provide us with a conception of the world as causal, or that it is the job of all theories of physics to describe causal relations in their domain of investigation. Rather I am here interested in the much more modest thesis that asymmetric causal notions play an important role in the way in which at least some of our well-established physical theories represent the world. That is, I am not advocating what Huw Price and Brad Weslake have called a "hyperrealism" about causation and which holds that causation is "something 'over and above' physics." According to Price and Weslake's hyperrealist, "physics itself may be time-symmetric [yet] there is a further, causal, aspect of reality which is asymmetric." (Price and Weslake 2009) By contrast, I wish to argue that causal representations are themselves an integral part of physics. While the *dynamical equations* of many of our mature theories of physics are time-reversal invariant, this does not imply that the *physical theories* as a whole are time-symmetric.

Yet I do not here wish to defend a causal metaphysics. I want to maintain that causal relations are more fully engrained in how physical theories represent the world than Russellian and neo-Russellian arguments seem to allow. But my thesis that asymmetric causal structures play a role in physics only aims to put causal relations on a par with other physical properties and relations that are represented in our physical theories and models. As such the thesis is metaphysically neutral and, in particular, does not entail a commitment to a 'weighty' non-Humean causal metaphysics.

My overall defense of my thesis consists of two prongs: a negative project aimed at showing that arguments intended to establish that asymmetric causal relations cannot be an integral part of theorizing in physics fail—this will be the focus of sections 2 and 3 of this paper—and a second, positive, part arguing that there are certain domains of physics in which time-asymmetric causal notions do in fact play a genuinely explanatory role. In section 4 I will discuss one specific example

of such causal reasoning, concerning our inference to the existence of a field source based on local observations of disturbances in the radiation field.

The arguments I examine in this paper are broad in scope and appeal to what are supposed to be generic features of physical representations. Their aim is to establish that there is no room for causal notions in any of what one may call "our mature theories of physics". Mature theories are not merely our most fundamental theories, such as general relativity, quantum field theory or a putative 'final theory' of quantum gravity. Rather, a mature theory of physics is any theory that is taught to physics students today and that present-day physicists use to represent physical phenomena—that is, a mature theory is any theory currently 'on the books,' including classical theories such as Newtonian mechanics or classical electrodynamics.

In the next section I discuss the suggestion that the fact that physical theories centrally involve abstract mathematical structures shows that these theories do not, or perhaps even cannot, also involve causal relations. Section 3 criticizes arguments that aim to show that asymmetric causal relations are incompatible with the time-symmetric laws of physical theories. While the discussion in these two sections is at a general level without making reference to any specific scientific theory, the argument I will consider in section 4—an argument for the claim that causal principles can play no explanatory role in physics—is best addressed in a concrete setting. In a previous paper (Frisch 2010) I argued that causal representations play an important role in accounting for our experimental interventions into physical systems. In the present paper I want to focus on a purely observational case—that of stellar observations. Against recent claims by John Norton (2009) and John Earman (2011) that causal notions are both vague and dispensable, I will that causal representations play an essential inferential and explanatory role in radiation theory.

2. Formulas and State-Space Models

I want to begin my discussion with an argument by Russell (Russell 1918), in which he argues that imprecise common sense causal regularities are replaced in physics by precise laws that have the form of functional dependencies. His argument appears to be this. Putatively causal claims need to be underwritten by universal causal regularities of the form "All events of type A are followed by events of type B." But in trying to find such regularities, we are faced with the following dilemma. Either the events in question are specified only vaguely and imprecisely. In this case the resulting regularities might be multiply instantiated, but they are formulated too imprecisely to be properly scientific. Or the events in question are specified precisely, but then the resulting regularities are instantiated at most once. Physics avoids this dilemma by providing us with precise functional

dependencies between variables representing properties of event types. Russell claims that such functional dependencies have replaced putatively causal regularities in physics, but it does not follow from the fact that physical theories present us with functional dependencies that these dependencies themselves cannot be interpreted causally. How, then, might we try to establish the claim that there is no legitimate place for an asymmetric notion of cause in mature theories of physics?

During his discussion of Newton's law of gravity Russell says that "in the motion of mutually gravitating bodies, there is nothing that can be called a cause and nothing that can be called an effect; *there is merely a formula*." (141, my emphasis) This claim is echoed decades later by Bas van Fraassen, who answers Nancy Cartwright's question (Cartwright 1993) "why not allow causings in the models?" as follows:

> To me the question is moot. The reason is that, as far as I can seen, the models which scientists offer us contain no structure which we can describe as putatively representing causings, or as distinguishing causings and similar events which are not causings. [...] Some models of group theory contain parts representing shovings of kid brothers by big sisters, but group theory does not provide the wherewithal to distinguish those from shovings of big sisters by kid brothers. The distinction is made outside the theory. (van Fraassen, 1993, 437-8)[1]

While Russell's remark suggests that a theory ought to be strictly identified with a set of formulas, van Fraassen argues that a theory consists of a set of state-space models. But even though the two disagree on whether theories ought to be understood syntactically or semantically, they agree that there is no place for causal notions in physics. According to van Fraassen's view of scientific representation more generally, a scientific theory presents us with a class of abstract mathematical structures that we use to represent the phenomena. These structures, van Fraassen suggests, cannot be used to represent causal relations or do not contain a distinction between causal and non-causal relations. Russell's and van Fraassen's general idea seems to be that it is the abstract mathematical nature of physical theories that is inhospitable to causal notions.

As a first attempt, we might try to reconstruct Russell's and van Fraassen's suggestions in terms of the following explicit argument:

---

[1] This remark occurs within the context of a discussion that also involves modal aspects of the causal relation. It is important, however, to keep the two aspects—the putative modality and the asymmetry of the causal relation—distinct. In the present quote van Fraassen unambiguously is concerned with the asymmetry: the question he asks is the mathematical formalism of the theory allows us to distinguish asymmetric shovings of brothers by sisters from their inverse, and not whether the shovings in some sense are necessitated.

2.1 The content of a physical theory is exhausted by a set of formulas or state-space models.

2.2 Causal relations are not part of the formulas or state-space models of a theory.

2.3 Therefore, causal relations are not part of the content of physical theories.

One might think that (2.2) is false (or at least not obviously true) and that causal relations can be part of a model. After all, in many disciplines scientists speak of 'causal models' or 'causal equations'.[2] But according to van Fraassen's view, the models presented by a scientific theory are, in the first instance, uninterpreted abstract mathematical structures—structures, which only acquire a representational role when they are used or taken by us as representing certain phenomena. A causal model or causal equation, on this view, could only be causal insofar as the model or equation is used to represent what are taken to be causal relations.[3] The point might be obscured by van Fraassen's use of the term 'model', which has several different meanings in the philosophy of science, but (2.2) is simply a consequence of the claim that the core of a physical theory consists of abstract mathematical structures, which *on their own* are uninterpreted and do not represent anything. Put in terms of Russell's syntactic framework, at the core of the theory there is 'merely a formula'.

If understood as referring to uninterpreted formulas or mathematical structures, (2.2) appears to be true, but under this disambiguation (2.1) is obviously false. No theory of physics can be strictly identified with a set of formulas or uninterpreted state-space models, since in order to make any claims about the world, the theory has to contain an interpretation which tells us which bits of the formalism are hooked up with which bits of the world. Minimally, a theory's interpretation has to specify the theory's ontology—it has to specify which parts of the world the different components of the mathematical structures are intended to represent.[4] But once we see that the austere view of theories as consisting solely of a mathematical formalism or set of abstract mathematical structures is untenable and that an interpretive framework needs to be part of a theory the question arises why this framework could not be rich enough to include causal assumptions as well. For example, an interpretive framework for Newton's laws might not merely

---

[2] This point was raised by an anonymous referee.

[3] Just as models of $\mathbf{F}=m\mathbf{a}$ are not intrinsically models of Newtonian systems but only if we use $\mathbf{F}$, m, and $\mathbf{a}$ to represent force, mass, and acceleration, respectively.

[4] As Richard Healey has reminded me in conversation, one might think that on van Fraassen's view (or at least according to the view van Fraassen appears to have defended in *The Scientific Image* (van Fraassen 1980)) a theory does not require an interpretational framework. On that view a theory is true, if it has models that are isomorphic to the phenomena. The problem with that view, however, is that there may be too many isomorphisms and hence that almost all theories come out as (almost) trivially true.

specify that *m* represents mass, **F** force, and **a** acceleration but might include the causal assumption that forces are causes of accelerations. That is, we cannot conclude from the fact that an *uninterpreted* formula **F**=*m***a** does not on its own mark **F** as cause and **a** as effect, that the causal "distinction is made outside the theory."

Thus, our first attempt at distilling a successful argument out of Russell's and van Fraassen's remarks failed. A second suggestion is that there might be constraints on what can be part of the formalism's interpretation that exclude causal notions. Thus, John Earman has proposed that a theory's content is exhausted by a formalism together with what he calls a "minimal interpretation" (Earman 2011). Thus, we should replace (2.1) with the following claim:

> 2.1' the content of a physical theory is exhausted by a set of state-space models or a set of
> formulas *together with a minimal interpretation*.

But instead of answering the question as to what is allowed to be part of the theory's interpretive framework, this proposal merely postpones the question. What can properly be part of the minimal interpretation and on what grounds can causal interpretations of certain mathematical relations be excluded? Earman's suggestion, echoing Russell's view, is that a minimal interpretation is one that is free from "philosophy-speak" and, thus, cannot involve the notion of cause. But this still does not provide an argument in support of the causal skeptic, for what is lacking is an account of what distinguishes 'philosophy-speak' from legitimate 'physics-speak'. The criterion cannot be to exclude notions that are employed by philosophers but not by physicists, since the physics literature is replete with appeals to causality—for example, as "physically well-founded assumption" (Jackson 1975, 312), as "fundamental assumption" (Nussenzveig 1972, 4) or as "general physical property" (Nussenzveig 1972, 7), or even as the "most sacred tenet in all of physics" (Griffiths 1989, 399).

Van Fraassen's remarks quoted above suggest an alternative way of spelling out the idea of a minimal interpretation: only those terms are a legitimate part of the interpretive framework that correspond to a part of the formalism: for each physical correlate of the mathematical models posited in the interpretation we have to be able to identify the structure in the models that represents that part of the world. The anti-causal claim then is that there are no substructures in the mathematical models presented to us in physics that can be described as representing causings. For example, one might ask where in the state space models defined by Newton's laws we find anything that could be taken to represent causal relations. What is more, van Fraassen even appears to suggest that causal relations *cannot* be represented structurally in a mathematical model, for he says that "group theory does not provide the wherewithal to distinguish" asymmetric

causal relations from their inverses or, in his example, shovings of kid brothers by big sisters from shovings of big sisters by kid brothers. Asymmetric causal distinctions, therefore, would have to be drawn outside a theory.

To take this last worry first, causal assumptions can be represented structurally. Interpreting a theory causally may be thought of as embedding the theory's state space models into larger model-theoretic structures that contain asymmetric relations between state space variables. The state of a system $S(t)$ is given by the values of a set of variables $s_1(t)$, $s_2(t)$, ..., $s_n(t)$, ..., which may be finite or infinite. The dynamical laws of a theory define a class of dynamical models specifying dynamically possible sequences of states, which can be represented in terms of state space models. We can then define an asymmetric and transitive relation $C=<S(t_i),S(t_j)>$ over the set of states S, which defines a partial ordering over the set of states in a model. C can be interpreted as the causal relation: $S(t_2)$ bears C to $S(t_1)$ exactly if $S(t_1)$ is a cause of $S(t_2)$. If two states do not stand in relation C then they are not causally related. The result is a class of what one might call *potential causal models* of a theory. Depending on the theory in question, we can also introduce more fine-grained causal relations $<s_m(t_i), s_n(t_j)>$ defined over individual state variables s_i.

Applying this to van Fraassen's toy example, we can see that group theory does provide the wherewithal to distinguish shovings of sisters by brothers from shovings of brothers by sisters: we can define an asymmetric relation R over the domain of objects consisting of all sisters and brothers, which we interpret as the 'a shoves b' relation, and there will be models in which some a that are sisters stand in relation R to some b that are brothers, but in which no brothers stand in relation R to any sisters. In these models it will be true that some sisters shove their brothers but it will not be true that any brothers shove their sisters.

One might reply that all that group theory allows us to do is to define an asymmetric relation, but the formalism itself does not allow us to distinguish between the relation 'x is a *cause* of y' and the relation 'x is an *effect* of y': the formalism alone does not distinguish between the 'shoves' and 'is shoved by' relation—and perhaps this is van Fraassen's point. But just as the formalism on its own cannot determine which objects its different variables represent and represents certain objects only in virtue of it being *used* to represent these objects, our use can equally determine that a given asymmetric relation represents the 'cause' rather than the 'effect' relation.

What remains is van Fraassen's first worry, for while there might be structures which we "*can* putatively describe as representing causings," van Fraassen could maintain that these structures are not part of the models scientists do*, as a matter of fact,* use. To present a theory, he

would insist, is simply to present a class of (suitably interpreted) state space models, defined by a theory's basic equations. But it is not obvious to me that this last claim is correct. As the above quotes suggest, physicists often do invoke causal assumptions. And such informal appeals to causal principles could be understood as implicitly defining a causal structure into which models of the dynamical equations are thought to be embedded. Since the causal structures in question often are very simple there may be little or nothing gained from adding a formal representation of this structure to the theory's equations. Nevertheless, if we want to offer a formal reconstruction of a theory, say, along the lines of van Fraassen's semantic view, we would have to include a representation of any causal assumptions, even if physicists themselves never represent these assumptions formally in terms of a partial ordering relation. Of course, to establish whether physicists' explicit appeal to causal assumptions in any particular theory ought to be indeed understood as a commitment to a causal structure would require a more detailed case-by-case investigation. My present point is merely that we cannot conclude simply from the fact that the models *of a set of equations* do not contain structures representing asymmetric causal relations that *scientific theories* contain no asymmetric causal assumptions and that any causal "distinction is made outside the theory."

If we want to establish that causal relations cannot be part of how physics represents the world, we need a more substantive argument than merely an appeal to the formal character of functional dependencies or state-space models. What is it exactly about the mathematical machinery of physical theories that may appear to render it incompatible, or at least make it sit ill, with causal notions? The suggestion I want to examine now is that the time-reversal invariance of the dynamical equations is incompatible with asymmetric causal relations.

3. Time-reversal invariance

The contrast between the time-reversal invariance of the dynamical laws and the time-asymmetry of the causal relation can be fashioned into an explicit anti-causal argument as follows:

3.1 Causal relations are temporally asymmetric.

3.2 The physical laws of our well-established theories have the same character in both the forward and backward temporal directions.

3.3 Therefore, there is no place for time-asymmetric causal relations in a theory with time-symmetric laws.

3.4. Therefore, there is no place for the causal relations in our well-established theories of physics.

Both premises (3.1) and (3.2) would deserve further comment: (3.1) appears to deny the possibility of instantaneous causation, while (3.2) might strike one as obviously false: there are many well-established but non-fundamental theories that are not time-symmetric, and there are even arguably fundamental theories that are not time-reversal invariant. But we can restrict our attention to well-established theories that are not explicitly phenomenological, like thermodynamics is, and follow the perhaps unjustified practice of ignoring failures of time-reversal invariance in more fundamental physics. Thus, here I want to focus on the inference from premises (3.1) and (3.2) to (3.3).

One might read this inference as simply relying on the same assumptions about the content of a physical theory as the argument in the preceding section—the assumption that the content of a theory is exhausted by a set of state space models with a minimal interpretation that associates the "mathematical squiggles" of an equation with physical quantities. This is how Alyssa Ney apparently reads her reconstruction of what seems to be essentially the same argument (Ney 2009, 748). Alternatively, the appeal to the time-reversal invariance of the laws might be taken to add something to the argument above. The claim then is that while in principle causal notions could be part of the interpretive framework of a theory, the fact that a theory has time-reversal invariant laws prohibits that it be interpreted causally. But it is unclear why we should accept this claim.

Assume, for instance, that the theory in question is deterministic (or near-deterministic). What the argument denies, then, is that if the full set of causes of an event determine the occurrence of an effect—that is, if the theory is causally forward deterministic, a complete set of effects cannot similarly determine the occurrence of their joint cause, and this premise does not appear to be defensible. Mackie's INUS condition account, for example, has the consequence that under some very weak additional assumptions, causes are not just INUS conditions of their effects, but effects are also INUS conditions of their causes. But it does not follow from this fact, that it is impossible to supplement Mackie's account with some condition that allows us asymmetrically do distinguish causes from effects.[5] More generally, it is hard to see why the notion of an event asymmetrically causing certain effects should be incompatible with the effects determining the occurrence of their causes. The claim that causes in some sense bring about their effects does not seem to preclude the possibility that the occurrence of certain events can be used to infer the occurrences of their causes.

Usually the claim that time-reversal invariant laws are incompatible with time-asymmetric causal relations is made without offering any further argument. For example, Erhard Scheibe

---

[5] See (Newton-Smith 1983) for a proposal of how one might introduce the asymmetry of the causal relations into a Mackie-style account of causation.

maintains, after pointing to the contrast between time-symmetric laws and time-asymmetric causal relations, that "this suffices to seal the fate of event-causality." (Scheibe 2006, 213)[6] One of the few exceptions to this rule is an argument by Norton (in his reply to my (2009) discussion of the role of causal assumptions in the derivation of dispersion relations) that aims to show that one can derive a contradiction from the conjunction of time-symmetric dynamical laws with a time-asymmetric asymmetric causal assumption. I want to quote Norton's argument in full:

> Now imagine a universe completely empty excepting two processes that we will call 'A' and 'B'. Process A has an incident wave, a dielectric, and a scattered wave. Process B is the time reverse of A. The two processes are completely isomorphic in all properties. Any property of one will have its isomorphic correlate in the other. Any fact about one will have a correlate fact obtaining for the other. One might be tempted to imagine that one of the two processes is 'really' the ordinary one, progressing normally in time; while the other is a theoretician's fantasy, a possibility in principle, but in practice unrealizable. The essential point of the example is that no property of the A and B systems distinguish which is which. Every property of one has a perfect correlate in the other. Let us assume that Frisch's principle of causality applies to one of these processes, the A process, for example. That will be expressed as a condition that
>
> the present state of the process depends only on its past states. Exactly what 'depends' may amount to is to be decided by the principle. All that matters for our purposes is that an exactly isomorphic condition of dependence will be obtained in the B process, except that it will be time reversed. Indeed, using the time order natural to process A, we would have to say that the principle of
>
> causality requires the present states of process B to depend upon its future states. In short, if the principle applies to process A, it fails for process B; and conversely. This is a *reductio ad absurdum* of the applicability of Frisch's principle of causality to scattering in classical electrodynamics. (Norton 2009, 481-2)

The argument appears to be this. Let us begin by postulating time-symmetric dynamical laws that allow a certain process *A* to occur, which is itself time-asymmetric. Since the laws are time-symmetric they also allow the time-reverse of *A*, the process *B*, to occur. If we then posit in addition a general causal principle, according to which future states causally depend on past states (but not

---

[6] "Schon [mit diesem Kontrast] scheint mir das Schicksal der Ereigniskausalität als fundamentaler Gesetzlichkeit besiegelt zu sein." (the translation into English is my own)

past states on future states) and which we assume that *A* satisfies, we can derive a contradiction: on the one hand, since *A* satisfies the causal principle but the dynamical laws are time-symmetric, *B*, in virtue of being the time reverse of *A*, will satisfy an *inverse* causal principle according to which a *past* state of the process *B* causally depends on its *future* states (but not vice versa). But on the other hand, since the causal principle is assumed to be general, *B* will also satisfy the original principle and *future* states of the process should depend on the *past* state (but not vice versa). This concludes the *reductio ad absurdum.*

Schematically, the argument may be presented as follows:

3.5   There is a time-asymmetric dynamical process *A* governed by time-symmetric dynamical laws.

3.6   *B*, the temporal inverse of *A*, is dynamically possible. (5)

3.7   *A* and its temporal inverse *B* have exactly the same physical properties. (5, 6)

3.8   For all processes, future states causally depend on past states (but not vice versa). (Causal Principle)

3.9   Future states of *A* causally depend on its past states. (8)

3.10   Past states of *B* causally depend on its future states (but not vice versa). (7, 9)

3.11   Future states of *B* causally depend on its past states (but not vice versa). (8)

That is, the conjunction of time-symmetric dynamical laws with a time-asymmetric causal principle results in a contradiction.

Premise (3.5) cannot be assailed, since even though we assume the laws to be time-symmetric, many—and in some intuitive sense, most—models of the laws will be time-asymmetric. But a defender of a causal principle should resist the steps of the argument leading to (3.10), and in particular the inference from (3.5) to (3.7) and (3.10): It does not follow from the assumptions that *B* is the dynamical time-reversal of process *A* and that *A* satisfies a time-asymmetric causal principles, that *B* will satisfy an inverse causal principle. Since Norton's inference might strike one as initially plausible I want to be belabor this point a bit. Let us assume that purely dynamical models of *A* and *B* can be represented by non-directed graphs:

 (*A*)   and   (*B*)
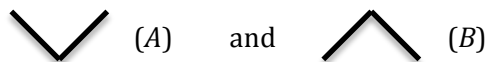
According to the causal principle, both models can be embedded into richer structures that include an asymmetric relation, which can be represented by adding a direction to the graphs:

 (A_{causal}) (B_{causal})

Norton points out that there is nothing in the purely dynamical models (that is the mathematical structures satisfying the dynamical laws) that tells us which model is which—there is no intrinsic difference between the two dynamical models—but the symmetry is broken in the directed causal models. And since the principle of causality is a general principle, once we 'add the arrowheads' to *one* graph, as it were, this fixes the direction of the arrows in the other graphs. Whatever models are used to model two physical processes A and B, respectively, we know that their orientations are opposite to each another.

According to Norton "the principle of causality requires" also that we represent the putatively causal process B by the inverse graph:

($B_{anti\text{-}causal}$)

Norton's reason is that it follows from the fact that the two processes A and B are time-reverses of each other that there is no property that distinguishes them. Since there is no physical difference between the two processes A and B, whatever reasons we might have for adding arrows to the graph representing A from the vertex of degree two—that is, the vertex with two edges at the bottom of the graph above—to the two vertices of degree one, the very same reasons would imply that we have to draw arrows in the graph representing B from the vertex with degree two at the top of the graph to the two arrows with degree one at the bottom.

But this step of the argument begs the question against a defender of a causal principle, who maintains that it is precisely the causal properties of the two processes that distinguish them: In the causal process A the event represented by the vertex of degree two causes the events represented by the two vertices of degree one, whereas in the causal process B the events represented by the two vertices of degree one cause the event represented by the vertex of degree two. That is, (3.7) does not follow from (3.5) alone, but requires as additional assumption the claim that the physical properties of a physical system are exhausted by those captured in the dynamical equations governing the system and this is precisely the assumption that a defender of a principle of causality wishes to deny, who would want to insist that the arrows in the causal model also represent physical properties of the system—properties that cannot be derived from the dynamical

equations alone.[7] While there is no difference between the two purely *dynamical* models represented by the two non-directed graphs—indeed they arguably are one and the same—there is a difference between the two physical processes represented and that difference consists in the difference in *causal* structure represented in the two directed graphs A_{causal} and B_{causal}.

Of course the causalist's assumption that there are properties that distinguish the two processes but are not represented in the dynamical equations is open to challenge and one can try to argue that any appeal to asymmetric causal structures in addition to a theory's purely dynamical models is unfounded or scientifically unjustified. Indeed, immediately after presenting his *reductio* argument, Norton suggests considerations to this effect and I will discuss these considerations below. Yet once we add as additional premise the claim that positing causal structures is unjustified, it becomes unclear what the overall structure of Norton's argument is meant to be. The additional premise would itself have to be supported by an argument, but such an additional argument would, if it could be made successfully, render the *reductio* proposed by Norton superfluous. If one were able to show that there are no legitimate reasons for positing a physical difference between processes *A* and *B*, the causalist would be defeated and there would be no work left to be done for the *reductio* argument. Thus, without an additional argument for the implicit premise in the argument from 2.5 to 2.7 the *reductio* begs the question against the causalist, but if we had such an additional argument, that alone would suffice to make the case against the causalist.

4. Broadcast waves and starlight

Up until now I have considered appeals to time-reversal invariance as aiming to show that it is impossible to interpret a theory with time-reversal invariant laws time-asymmetrically causally.[8] But I have just suggested that there may be another (and I take it arguably more plausible) position according to which it would be a mistake to accept a causal principle not because it is strictly incompatible with time-symmetric laws but because *there are no good reasons* for positing causal properties and relations in addition to the non-causal properties represented in the dynamical laws. The basic equations of a theory that is future- as well as past-deterministic define both an

---

[7] Even though I am following Norton here in expressing the argument in terms of real physical properties of a process, the point I wish to make here is independent of the debate about scientific realism. A defender of a principle of causality can also be an instrumentalist and argue that the causal relations in our models no more represent real properties than other properties or relations in our models. My claim here is that there is no principled reason for treating causal properties differently from other kinds properties and relations of our models (and of the real world systems we are modeling).

[8] This is also what (Field 2003) appears to argue.

initial and a final value problem. If we begin with the system's initial state, then the dynamical equations determine the system's subsequent evolution; if we take the system's final state to be given, then the dynamical equations determine the system's earlier evolution. Thus, one might be tempted to agree with Fritz Rohrlich, who once maintained that the "identification of *causality* with *prediction* rather than *retrodiction* in a time-symmetric system of equations is completely arbitrary." (Rohrlich 1990, 51 italics in original)[9] Thus, while there may be no good arguments that strictly *disallow* interpreting a theory causally, it might nevertheless be the case that there could be no scientifically legitimate reasons for supporting an asymmetric causal interpretation of a theory.

My discussion has been rather general so far, but the question whether positing causal relations is justified can only be addressed in a concrete setting and by examining specific examples of physical theories in which causal notions might be thought to be play a role. The example Norton himself discusses is the putatively causal asymmetry characterizing radiation fields associated with an oscillating source and I want to focus on that asymmetry here as well.[10] Consider an antenna that broadcasts into empty space.[11] The radiation emitted by the antenna is a coherently diverging field traveling to spatial infinity. Call this 'process A' in accordance with the terminology introduced above. The temporal inverse of this process—one that is also allowed by the wave equation in the presence of wave-sources—is a field that is coherently collapsing into the antenna and is absorbed. Call this 'process B'. Norton in effect argues that since both processes are physically possible, there can be no scientifically legitimate reason for representing A and B in terms of the causal structures A_{causal} and B_{causal}, respectively. Appealing to a causal principle would be justified, Norton suggests, only if the principle allowed us correctly to exclude certain *dynamically* possible processes as *causally*, and hence physically, impossible. But in the case at issue both processes—the diverging

---

[9] Despite what Rohrlich says here, he seems to believe now that there can be good reasons for interpreting a theory with time-reversal invariant laws causally asymmetrically. (See Rohrlich 2006)

[10] Norton's paper is a response to my (2009a), where I argue that *dispersion theory* invokes causal assumptions. Norton apparently assumes in his reply that the relevant asymmetry in this case just is an instance of the electrodynamic *wave-asymmetry* that he discusses. I am less sure than he is, however, that the two cases are as closely related as he suggests. I have two concerns: first, dispersion theory is just one particular application of the general framework of linear response theory, which also has many applications outside of electrodynamics; and, second, there appears to be a formal disanalogy between the Green's functions used in the two cases: while the Green's function for the wave equation has a temporal inverse, the inverse Green function in the case of linear response theory is not mathematically well-defined. That is, while one can solve the wave-equation both as an initial value problem and a final value problem, one cannot, it seems, represent a linear response system through a final value problem.

[11] In order to bring out the asymmetry as simply and clearly as possible I am adjusting Norton's specific example slightly. The example of the antenna is discussed in (Earman 2011).

and the converging radiation fields—*are* physically possible. To be sure, we generally do not observe radiation coherently converging on a source, but it is in principle possible to set up such processes, for example with the help of perfect mirrors or the carefully correlated action of multiple sources of radiation. Hence adding a causal principle does no real work in the present case—it amounts to adding a distinction that makes no difference—and therefore is not scientifically justified.

Yet, as I want to argue now, causal notions play a scientifically legitimate role in radiation theory, even though they do not provide an additional constraint on what is dynamically possible. Understanding the relation between sources and fields causally plays both an important *inferential* and an important *explanatory* role in the theory.

I want to preface my argument with a clarification of the argument's target. There is a sizeable literature defending the thesis that the asymmetry between diverging and converging waves ultimately is of thermodynamic origin. This literature suggests one possible kind of anti-causal argument—an argument that can concede that a causal account could in principle provide an explanation of the asymmetry, but argues that the thermodynamic account offers a superior explanation. That is, the argument compares two prima facie scientifically legitimate explanations and argues that one of them is superior to the other or, perhaps, that one of them is more fundamental than the other.

Now, it is not *prima facie* obvious how a thermodynamic explanation of the absence of converging radiation would proceed in the case of an antenna broadcasting into empty space. For in this case we are not dealing with a closed system and the asymmetry between coherently diverging and coherently converging radiation seems to be entirely an asymmetry in initial conditions: coherent radiation coming in from spatial infinity in one case and the absence of any coherent incoming radiation in the other. Yet since there is not enough room to more fully assess this argumentative strategy here,[12] my present aim is more modest: Instead of showing that causal representations are an irreducible part of radiation theory, I merely want to show that introducing causal relations can play both an explanatory and an inferential role. Thus, my target here is the stronger thesis that a causal principle fails dramatically even when considered on its own and not only by comparison to some other putative explanation of the asymmetry of radiation phenomena.

Imagine you are looking up at the night sky and observing the light emitted by a particular star. What licenses your inference that the observed radiation was indeed emitted by a star as its

---

[12] But see (Frisch 2005; 2006; forthcoming) for criticisms of a version of the thermodynamic argument.

source rather than it being source-free radiation coming in from spatial infinity? It appears to be almost religious dogma among certain philosophers of physics that the content of a physical theory is exhausted by the models of its dynamical equations and, hence, that the *only* way to use a theory to make empirical predictions is to solve an appropriate initial (or final) value problem. In our case the initial value problem requires a specification of the state of the field and of all sources on a Cauchy surface—a spacelike cross section of a lightcone centered on the putative source—that is then fed into the Maxwell-Lorentz equations. But in fact the only data we have at our disposal for inferring the existence of the star are our highly localized observations of the radiation fields. Neither do we know the fields on anything close to a complete initial value surface nor do we have independent access to the trajectories of the source—the star emitting the radiation. Thus, if the only tools at our disposal for making inferences about the putative sources of stellar radiation were the dynamical laws applied to an initial value-problem, it would be a complete mystery as to how we can ever come to know of the existence of a star.

But we do seem to be able to make justified inferences about the state of the cosmos at earlier times. What then is the structure of these inferences? The answer is that our inference to the existence of a star as the source of the observed radiation is a paradigmatically causal inference: the radiation fields observed at different spacetime points are highly correlated with one another and we infer from these correlations to the existence of a common cause. In fact, the locally observed fields are correlated in several different ways. What we observe are relatively strong disturbances in a very weak background field. There is an (almost) perfect coincidence in the luminosities and spectral distributions of the radiation observed at different locations. And even more strikingly, the shapes of the field disturbances received at different spatial locations match so closely, that they can be made to interfere with one another—a fact that is exploited in stellar interferometry. The degree of partial coherence in this last sense can be expressed in terms of so-called 'coherence functions' associated with the waves (see, e.g., Born and Wolf, 1999, ch. 10).

Since the directions from which the correlated field disturbances at different times are observed—their celestial latitudes and longitudes—are such that the fields can be associated with the trajectory of a single localized source in relative motion to us, we infer the existence of a star as common cause of our observations as providing the best explanation for the observed correlations.

Thus, the inference to the existence of a single star as source of the radiation is a standard example of an inference to a common cause—an inference pattern is widely used, both in the sciences and in common sense reasoning. Indeed, in discussions of correlation functions in the physics literature it seems to be simply presupposed, without need for a justification, that strong

correlations among field disturbances generally allow us to infer to a common source and, conversely, that the absence of a common source would result in uncorrelated fields. Thus, Born and Wolf elucidate the idea of measuring correlations between locally received disturbances in terms of possible interferences by contrasting the case of strong correlations, when light comes from a single "very small source of a narrow spectral range," with the case of zero correlations, when the two observation points "each receive light from a different physical source" (Born 1999, 555).

Here is an example of the same kind of inference from outside physics: in a string of bank robberies the police discover similar kinds of little plastic toys left behind at each crime scene—a fact that is not made public by the investigators. The police infer from these discoveries that the robberies were committed by one and the same gang rather than by different groups that operate in complete independence from one another.

There is a large literature on how properly to reconstruct such inferences to a common cause. One formulation of a principle of the common cause (PCC) underwriting such inferences states that if two quantities X and Y are correlated and X and Y are not related as cause and effect, then X and Y are the joint effects of a common cause Z. In Reichenbach's original formulation the principle also contains a screening-off condition according to which the common cause Z screens off any probabilistic dependence between X and Y. As has been much discussed in the literature, however, there are several counterexamples to the principle of the common cause (see Arntzenius 2010 for a survey). These counterexamples can be avoided, if we think of the principle as a defeasible epistemological guide rather than as a metaphysical principle and if, following Eliot Sober (Sober 1984; 2001), we construe common cause inferences comparatively as involving a comparison of the likelihood conferred on the evidence by a common cause explanation with the likelihoods of competing separate cause explanations. Thus, as I want understand it, the principle of the common cause asserts that it is reasonable to posit the existence of a localized common cause to explain distant correlations, unless there is a separate cause explanation that (i) confers a higher likelihood on the explanandum and (ii) does not merely engage in "explanatory buck passing" (Sober 1984, 221) by proposing to explain distant correlations at one time by even earlier distant correlations.

How should we cash out the claim that the disturbances in the radition field are correlated? One option would be to try to express the degree of correlation non-probabilistically entirely in terms of the correlation function. The problem with this suggestion is that the correlation function depends on properties of the source that seem not to track very well the degree of confidence with

which we are willing to infer the existence of a common source. For example, the absolute value of the correlation function is much smaller for sunlight than for starlight (since the distant stars appear almost as point sources), nevertheless we just as readily infer the existence of the sun from the direct sunlight we observe as we infer the existence of a star from the observed stellar radiation.

Alternatively, we could try to express the idea of correlation in terms of perfect coincidences: Within a generally very weak background field, the luminosities and spectral decompositions of stronger field disturbances at different observation points match (almost) perfectly and (due to the spatio-temporal distribution of the disturbances) can therefore be associated with the trajectory of a single source.

But one could also express the correlations probabilistically: There is a certain small but non-zero probability of observing a relatively strong non-zero field disturbance coming in from an arbitrary direction in the sky. But the conditional probability of detecting radiation from one direction, given that we detected radiation from the same (or rather appropriately related) direction earlier is muc higher. That is, we are allowed to infer the existence of a star as common cause form the fact that the field disturbances at two different locations are strongly correlated, since the correlations are far more probable, given the hypothesis that the disturbances were produced by one and the same source, than if we assume separate causes, be it different sources acting independently of each other or source-free fields coming in from spatial infinity. We can express this more formally, where the relata of the causal relation are taken to be the values of variables:

$$\Pr(F(t_1,x_1) \,\&\, F(t_2,x_2)\,/\, S(t_{ret1}, x_{ret1}) \,\&\, S(t_{ret2}, x_{ret2})) \gg \Pr(F(t_1,x_1) \,\&\, F(t_2,x_2)\,/\, S(t_{ret1}, x_{ret1}) \,\&\, S^*(t_{ret2}, x_{ret2}))$$

$$(1)$$

Here $F(t_1,x_1)$ and $F(t_2,x_2)$ are the values of the electromagnetic fields at the observation points (t_1,x_1) and (t_2,x_2), respectively, and $S(t_{ret1}, x_{ret1})$ represents the state of the source s at what is known as the "retarded" point—the point on the backward lightcone of (t_1,x_1) at which the source s would have to have been located to give rise to the field $F(t_1,x_1)$. S^* represents a separate source s^*, while the state of the source s at (t_{ret2}, x_{ret2}), $S(t_{ret2}, x_{ret2})$, follows (quasi-) deterministically from its earlier state $S(t_{ret1}, x_{ret1})$. That is, $S(t_{ret1}, x_{ret1})$ acts as a common cause of both field observations $F(t_1,x_1)$ and $F(t_2,x_2)$.

The overall process is, of course, deterministic (at least as long as we treat it purely within classical electrodynamics): initial fields in the remote past together with the field associated with the stars as sources nomologically determine what the observed fields will be. This, one might

think, leads to two problems for my account: First, demanding that the common cause screen off its effects from one another might seem problematic. Not only is it trivially true that there will be a screening-off event in the case of deterministic theories, but there also will be screening-off events in the future of the distant correlations to be explained. What, then, licenses our inference to the existence of a past common cause? The deterministic setting, however, does not guarantee that the events in question will be localized. That is, a principle positing the existence of a *localized* common cause is not trivial. Future screening-off events, in the case of our example, will generally be highly non-local: disturbances in the radiation field diverge from their common source and will be more and more dispersed in the future.

A second worry is that all the probabilities in question, and in particular the likelihoods, will be either zero or one, since the process is deterministic, and that, hence, will not be able to apply likelihood reasoning. Yet, since we have direct observational access neither to the free fields prior to the putative emission events nor to the present fields on a complete initial value surface, we cannot set up a full-fledged initial- or final-value problem. Rather we somehow have to infer, based on our localized observations, what *both* initial fields *and* the fields associated with any sources might have been. That is, our problem is somehow to carve up the total locally observed field into a component associated with past sources and a source-free incoming field. Without any additional assumptions this would be impossible: our localized observations simply do not provide us with enough information to infer both the state of any sources and the initial field and are compatible with multiple different combinations of initial fields and radiation fields.

What we observe are relatively focused packets of radiation (which we interpret as light emitted by stars) within a background field that is approximately equal to zero, or at least is very weak. We infer from this that the field disturbances coming from a single direction are due to a star as their common cause. But it is also consistent with our evidence that there existed strong correlations among source-free initial fields in spatially distant regions at some remote time in the past that resulted in macroscopic fields converging onto the putative trajectory of the star, passing over it, and then rediverging—mimicking (for later observers) the presence of a star. The further back in time we followed such source free fields, the weaker these fields would be as they become more and more dispersed toward the past, originating in what ultimately would have been extremely delicately coordinated microscopic correlations among very distant field regions.

If such fields are equally as compatible with our observational evidence, why do we infer the existence of localized common cause rather than the existence of delicately set up correlations among source-free incoming fields? The latter hypothesis violates our prohibition against

explanatory buck passing, since it accounts for the spatially distant correlations in terms of even more spatially dispersed correlations in the remote past. We can exclude this explanation by positing that the initial fields are effectively random—or rather as random as possible, given our observational evidence.[13] Making this assumption explicit we get:

$$\Pr(F(t_1,x_1) \& F(t_2,x_2)/ S(t_{ret1}, x_{ret1}) \& \textit{random initial fields}) \gg \Pr(F(t_1,x_1) \& F(t_2,x_2)/ \textit{random initial fields}) \quad (2)$$

The initial randomness assumption can itself be motivated by a causal representation of the phenomena at issue. The principle of the common cause states that it is reasonable to explain distant correlations among quantities by positing an earlier localized common cause. The contrapositive of that principle maintains that in the absence of a common cause it is reasonable to assume that the quantities in question are uncorrelated. Thus, if we do not expect there to be a localized common cause of distant field values at whatever time we choose as the initial time, we expect the randomness assumption to be satisfied.

What is the nature of the probabilities in (1) and (2)? I think for our present purposes we can be ecomenical and allow the probabilities to be whatever probabilities one thinks are needed to underwrite explanatory inferences in the sciences. If one believes that epistemic probabilities are sufficient to do this job, then one can read the randomness assumption as reflecting our ignorance of the precise initial conditions. Alternatively, the probabilities may be construed more objectively, as some form of frequencies or as propensities.

I have argued that common cause reasoning plays an important inferential role. Analogously, the principle of the common cause and its converse (which states that distant quantities are uncorrelated in the absence of a common cause) also play an explanatory role and, for example, allow us to explain why we observe diverging but not coherently converging waves in nature. Let us return to the example which we began our discussion in this section. The highly correlated behavior of the radio signals in processes A and B above are precisely the kind of phenomena that would call for an explanation in terms of localized common causes. In the case of process A—a radio antenna broadcasting into empty space—such an explanation is readily available: the action of the antenna acting as common cause of the field disturbances can explain the strong correlations among them. Contrast this with the process B of an anti-broadcast wave collapsing into the antenna. In this case there is an antenna located at the point on which the radio waves are centered, just as in process A. Yet by hypothesis the correlations in this case possess no

---

[13] See (Arntzenius 2010) for the sketch of a general argument that all inferences to a common cause presuppose an initial randomness assumption.

earlier common cause. Thus, by the converse of the principle of the common cause, we would not expect such a process to occur.

The field coherently converging onto the antenna presents a solution to the dynamical equations just as much as the diverging field. In fact, purely dynamically, and considered 'atemporally,' the two processes are completely equivalent: in both cases there are coordinated fields that are correlated with the localized action of an antenna. The symmetry between the two cases is broken, however, once we embed purely dynamical models of the two processes into richer, causal structures. Process A then appears to be normal and entirely to be expected, since the correlations among the disturbances can be explained by their common cause. Process B, by contrast, seems 'contrived', 'mysterious' or 'improbable', since the correlations do not have a common cause.

Note that the preceding discussion does not depend on the fact that an antenna or a star is a macroscopic object. The inferential and explanatory structure would be exactly the same in the case of a microscopic charge, such as a single oscillating electron, at least as long as we model the charge classically. The only difference is that the disturbances in the radiation field diverging from the electron as their source would be less easy to detect empirically.

In invoking a causal explanation of why we normally do not find coherently converging waves I am not denying the possibility of carefully setting up such waves. The way to do this is to arrange a large number of radiating objects and set them into coherent motion such that the waves diverging from each individual source combine to an overall converging wave. But notice that in this case the coordinated behavior of the wave can once again be explained by appealing to a local common cause in its past, namely the mechanism we used to set the collection of distant sources into coherent motion. Thus the explanatory role of the causal principle in the present case is not to prohibit certain dynamically permissible processes; rather the principle explains why certain dynamically allowed processes are radically improbable, while their temporal inverses are utterly familiar to us.

Our local observations of packets of stellar radiation severely underdetermine which dynamical model of the wave equations correctly represents the phenomena, but, as I have argued, embedding the dynamical models into richer causal structures allows us to decide between possible models by inferring the existence of a localized common cause of observed correlations. That is, we infer that the locally observed radiation was emitted by a star not merely by an appeal to the dynamical laws governing electromagnetic radiation but rather by using a combination of dynamical and causal reasoning. Moreover, while I have only discussed one type of example here, it

is easy to see that the discussion generalizes to a wide class of observations in physics, since only very rarely—if ever—are we in a position to know the state of a system on a full initial (or final) value surface.

But have I really shown that the notion of *cause* plays an important role in reasoning in physics? Someone might argue that even if one were to grant that our inference to the existence of a star are justified by the kind of reasoning I have described, this still does not show that this reasoning needs to invoke the notion of *cause*. It would be enough, one might claim, simply to postulate the initial randomness assumption and then use the likelihood principle or Reichenbach's principle to infer the existence of the star without ever using the term 'cause'. The explanatory work is done by the randomness assumption and the probabilistic principles. Calling the reasoning 'causal' seems to be no more than add an empty honorific, or so one might argue. I want to make the following four points in reply to this objection.

First, as I have argued, the initial randomness assumption can itself be motivated by a causal picture of the world and fits well with our explanatory enterprise of aiming to explain distant correlations in terms of *earlier* common causes. Let us assume a set of variables F that has no determinants outside of F. That is, for each variable it is the case that it has a complete set of causes in F and conditional on that set what value the variable takes is independent of the values of any variables outside of F. Then it follows from the converse of the principle of the common cause that it is very likely that the values of any variables that are not related as cause and effect and do not have a common cause will be uncorrelated.

Second, recall that my aim here is not to defend a causal metaphysics but merely to argue that certain causal representations play an important role in physics. I want to leave it open what metaphysical conclusions, if any, one might want to draw from this fact. The explanatory and inferential patterns that I identified are commonly taken to be an instance of paradigmatically causal reasoning. That strikes as a good enough reason to conclude that there is a place for causal reasoning in physics. Yet ultimately it is not crucial what label we attach to this kind of reasoning. What is important is that it is a time-asymmetric reasoning that goes beyond what is implied by the time-reversal invariant dynamical laws and our local observations but nevertheless plays an important explanatory and inferential role in physics.

Third, at least some skeptics of causal reasoning in physics are willing to grant that causal reasoning plays a role in common sense and other sciences. My arguments here put pressure on this distinction. To the extent that the use of common cause reasoning is evidence for the role of causal reasoning elsewhere, it also points to a role of causal reasoning in physics. And conversely: if

this kind of reasoning is not taken as evidence for the presence of causal reasoning in physics, then neither can it be evidence for causal reasoning elsewhere.

Fourth, postulating the initial randomness assumption merely as *de facto* constraint on initial conditions does not allow us to make explanatory distinctions of a kind even causal skeptics like Earman seems to want to make. In his discussion of the asymmetry between a diverging broadcast wave and a coherently converging anti-broadcast wave—that is, between processes *A* and *B* above—Earman says this:

> It would seem nearly miraculous if the time reverse of [the broadcast wave] were realized in the form of anti-broadcast waves coming in from spatial infinity and collapsing on the antenna. The absence of such near miracles might be explained by an improbability in the coordinated behavior of incoming source free radiation from different directions of space.

Yet in what sense exactly is the converging wave miraculous or improbable? It follows from the inhomogenous wave equation that a source will be associated with the coordinated behavior among distant disturbances of the field. *Purely dynamically* coherently converging waves are no more improbable than diverging waves. All that the dynamics tells us is that there will be distant correlations in the field *somewhere*—be it in the past of the action of the source in the form of a converging wave, or in the future of the source as diverging wave, or perhaps as a linear combination of the two.

Moreover, dynamically there is nothing especially odd about coordinated behavior of incoming *source-free* radiation. Dynamically, both diverging and converging waves can be represented both in terms of an initial value problem and a final value problem. An initial value problem represents the total wave F as a sum of a source-free incoming field and a diverging field associated with the antenna: $F=F_{in}+F_{div}$. A final value problem represents the total wave as a sum of a source-free outgoing field and a converging field: $F=F_{out}+F_{con}$. Both representations are representations of one and the same field. Now, if we represent a *converging* wave in terms of an initial value problem, the wave will appear as incoming source-free radiation. But if we represent the same wave in terms of a final value problem, the converging wave appears as being associated with the source. By the same token, a *diverging* wave can be represented either as outgoing source-free radiation (in a final value problem) or as associated with the source (in an initial value problem).

Positing the initial randomness assumption breaks the symmetry between the different representations. And it may be that when Earman says that the absence of converging waves might

be explained by an improbability of coordinated behavior in the incoming radiation he is appealing to nothing more than the inintial randomness as de facto constraint. But the quote above suggests otherwise. Earman seems to suggest that there is something apparently 'near-miraculous' about one class of solutions to the dynamical equations and that there is a need to explain the absence of such 'near-miracles'. But it is unclear what, from a strictly non-causal perspective, this miraculousness might consist in. Coherently converging anti-broadcast waves represent dynamically perfectly possible situation, which happen to be rendered unlikely by the *de facto* initial conditions.

Consider the following analogy: a ball is released at the top of a wedge- or roof-shaped inclined plane and rolls down on the wedge's left hand side. As an explanation for the ball's trajectory we appeal to the ball's initial conditions according to which the ball was released slightly to the left of the peak of the wedge. It also is dynamically possible for the ball to roll down the ride hand side. Even though the ball does not follow this second trajectory, since it is incompatible with the ball's initial conditions, there is nothing 'near-miraculous' about this second dynamically possible solution. Similarly, absent additional causal considerations, there is nothing miraculous or 'contrived' about the non-actual solutions to the inhomogenous wave-equation.

Moreover, it is natural to assume that it converging radiation is improbable precisely because it would require coordinated behavior in the "incoming source free radiation from different directions of space," as Earman says. But again a strictly non-causal construal of radiation phenomena cannot support this intuition. On a strictly non-causal picture, there is no more reason to expect that incoming radiation from different directions should be uncoordinated as there is to expect that source free outgoing radiation will be uncoordinated.

4. Conclusion

I have argued that a widely held view—the view that asymmetric causal relations cannot play a legitimate role in physical theories—is not underwritten by convincing arguments. Neither the fact that physical theories present us with classes of abstract mathematical structures nor the fact that (at least many of) the dynamical laws of our mature physical theories are time-symmetric provides a compelling argument against the possibility of causal notions playing a role in physics. Finally I have argued, by examining the asymmetry of wave and radition phenomena, and in particular the case of stellar observations as a concrete example where this asymmetry manifests itself, that causal assumptions can play an important inferential and explanatory role in physics.

References

Arntzenius, Frank, "Reichenbach's Common Cause Principle", *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2010/entries/physics-Rpcc/>

Born, Max and Emil Wolf. 1999. *Principles of Optics, seventh edition*. Cambridge: Cambridge University Press.

Callender, Craig. 2000. Is Time 'Handed' in a Quantum World? *Proceedings of the Aristotelian Society,* June, 247-269.

Cartwright, Nancy. 1993. Defence of `This Worldly' Causality: Comments on van Fraassen's Laws and Symmetry. *Philosophy and Phenomenological Research* 53 (2): 423-429.

Earman, John. 2011. "Sharpening the Electromagnetic Arrow(s) of Time". in *The Oxford Handbook on Time*. Edited by Craig Callender. Oxford: Oxford University Press.

Field, Hartry. 2003. "Causation in a Physical World". In *The Oxford Handbook of Metaphysics*, edited by M. Loux and D. Zimmerman. Oxford: Oxford University Press.

Fraassen, Bas van. 1980. *The Scientific Image*. Oxford: Oxford University Press.

Fraassen, Bas van. 1993. Armstrong, Cartwright, and Earman on Laws and Symmetry. *Philosophy and Phenomenological Research* 53 (2): 431-444.

Frisch, Mathias. 2005. *Inconsistency, Asymmetry, and Non-Locality: A Philosophical Investigation of Classical Electrodynamics*, Oxford University Press.

_____. 2006. "A Tale of Two Arrows," *Studies in the History and Philosophy of Modern Physics* 37, 542-558.

_____. 2009a. "'The most Sacred Tenet'? Causal Reasoning in Physics," *British Journal for the Philosophy of Science* 60 (September 2009) 459-474.

_____. 2009b. "Philosophical Issues in Electromagnetism," *Philosophy Compass* 4/1: 255–270.

_____. 2010. "Causal Models and the Asymmetry of State Preparation" in *EPSA Philosophical Issues in the Sciences. Launch of the European Philosophy of Science Association*, eds*. M. Suárez, M. Dorato and M. Rédei*. Springer Verlag.

_____. forthcoming. "Statistical mechanics and the asymmetry of causal influence," in *David Albert's Time and Chance*, eds. Barry Loewer, Eric Winsberg, and Brad Weslake, Harvard University Press.

Griffiths, David. 1989 *Introduction to Electrodynamics*. 2nd ed. Prentice-Hall.

Healey, Richard. 1983. Temporal and Causal Asymmetry. In *Space, Time, and Causality*, edited by Richard Swinburne. Dordrecht: D. Reidel Publishing Co.

Hitchcock, Cristopher. 2007. What Russell Got Right. In *Causality, Physics, and the Constitution of Reality: Russell's Republic Revisited*, edited by H. Price and R. Corry. Oxford: Oxford University Press.

Jackson, John David. 1975. *Classical Electrodynamics*. 2nd ed. New York: John Wiley & Sons.

Newton-Smith, W. H. 1983. Temporal and Causal Asymmetry. In *Space, Time, and Causality*, edited by Richard Swinburne. Dordrecht: D. Reidel Publishing Co.

Ney, Alyssa. 2009. "Physical Causation and Difference Making." *British Journal for the Philosophy of Science* 60 (4):737-764.

Norton, John D. 2003. Causation as Folk Science. *Philosophers' Imprint* 3 (4), http://www.philosophersimprint.org/003004/; reprinted as pp. 11-44 in H. Price and R. Corry, *Causation and the Constitution of Reality: Russell's Republic Revisited*. Oxford: Clarendon Press.

Norton, John D. 2007. Do the Causal Principles of Modern Physics Contradict Causal Anti-Fundamentalism? In *Thinking about Causes: From Greek Philosophy to Modern Physics*, edited by P. K. Machamer and G. Wolters. Pittsburgh: University of Pittsburgh Press.

Norton, John D. 2009. "Is There an Independent Principle of Causality in Physics?" *British Journal for the Philosophy of Science*. pp. 475-86.

Nussenzveig, H. M. 1972. *Causality and Dispersion Relations*. New York: Academic Press.

Price, Huw. 2007. Causal Perspectivalism. In *Causality, Physics, and the Constitution of Reality: Russell's Republic Revisited*, edited by H. Price and R. Corry. Oxford: Oxford University Press.

Price, Huw & Weslake, Brad (2009). The time-asymmetry of causation. In Helen Beebee, Peter Menzies & Christopher Hitchcock (eds.), *The Oxford Handbook of Causation*. Oxford University Press.

Rohrlich, Fritz. 1990. *Classical Charged Particles*. Reading, MA: Perseus Books.

Rohrlich, Fritz. 2006. Time in Classical Electrodynamics. *American Journal of Physics* **74** (4): 313-315.

Russell, Bertrand. 1918. On the Notion of Cause. In *Mysticism and Logic and other Essays*. New York: Longmans, Green and Co.

Scheibe, Erhard. 2006.*Die Philosophie der Physiker.* Munich: Verlag C.H. Beck.

Sober, Eliot. 1984. "Common Cause Explanation" *Philosophy of Science* 51, 212-241.

Sober, Eliot. 2001. "Venetian Sea Levels, British Bread Prices and the Principle of the Common Cause," *British Journal for the Philosophy of Science* 52, 331-346.

Suppes, Patrick. 1970. *A Probabilistic Theory of Causality*. Amsterdam: North-Holland.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Woodward, James. 2007. Causation with a Human Face. In *Causality, Physics, and the Constitution of Reality: Russell's Republic Revisited*, edited by H. Price and R. Corry. Oxford: Oxford University Press.