No Place for Causes?
Assessing Causal Skepticism in Physics

Mathias Frisch
University of Maryland, College Park

[Draft. Please do not quote.]

Abstract

According to a widespread view, causal notions have no legitimate role to play in mature physical theorizing. This view, which can be traced back to Russell's famous attack on the notion of cause, has proponents even among those who believe that causal notions have an important place in the special sciences and in our folk conception of the world. In this paper I critically examine a range of general arguments for the view, argue that none of them succeed and propose two routes by which asymmetric causal assumptions can be legitimated even within the context of a physics with time-symmetric laws.

1. Introduction[1]

It appears to be both natural and intuitive to think of the world as causally evolving. We conceive of events in the present as being caused by events in the past and, in turn, as acting as causes for what happens in the future. But it is also a widespread view—at least among philosophers of physics—that this conception is not part of how mature physics represents the world. According to this view, the notion of cause survives—if at all—as part of a 'folk' scientific conception of the world but has no place in our mature theories of physics. In this paper I want to examine critically a cluster of arguments in favor of this causal skepticism, focusing in particular on the asymmetry of the causal relation. Many of these anti-causal arguments are descendents of Bertrand Russell's famous attack on the notion of cause (Russell 1918), but they also have found more recent defenders (see, for example, van Fraassen 1993, Norton 2003, 2007; Hitchcock 2007).[2]

In the next section I will distinguish three dimensions of our intuitive, pre-philosophical notion of cause that prima facie might appear to play—or have played—a particularly prominent role in theorizing in physics: causation as determinism, causality as imposing locality constraints, and the asymmetry of the causal relation. In section three I will critically examine a range of arguments aimed at showing that there

[1] Ancestors and earlier versions of this paper have been presented at the Boston Colloquium in 2006, a workshop at the University of North Carolina, Chapel Hill in 2006, the conference "Causation and Mechanism" at the University of Maryland in 2007, and the 1st conference of the European Philosophy of Science Association in Madrid in 2007. I want to thank the audiences there for extremely useful comments and criticisms.

[2] As I learned only recently, (Steiner 1986) argues for a conclusion quite similar to mine, also using Russell's discussion as his main target. But Steiner's argument centrally relies on certain interpretive assumptions about quantum mechanics, while my defense of causal notions makes no such assumptions.

is no room for an asymmetric causal relation in mature physics. My conclusion will be negative: I will argue that none of these general arguments succeed in showing that an asymmetric notion of cause cannot play a legitimate role in mature physical theories. In section four I will present two routes by which asymmetric causal assumptions may be scientifically justified even within the context of a theory with time-reversal invariant laws.

One of my aims in this paper is to urge a shift in the focus of the debate: if we want to make genuine progress in assessing the role of causal assumptions in physics, we need to examine concrete examples of the use of putatively causal notions in physics rather than try to rely on general anti-causal (or pro-causal) arguments. John Norton's discussion of the alleged failure of causal determinism in Newtonian physics (Norton 2003) is an example of this kind of detailed engagement with actual physics and my (Frisch manuscript), aims to provide a similarly detailed discussion—albeit with different conclusions—of the use of asymmetric causal notions in modeling scattering phenomena.

It may be helpful at the outset to contrast the view in which I am interested with the views that Norton criticizes in (Norton 2003; 2007). While some of Norton's arguments are meant to show that the notion of cause is merely part of 'folk science' and hence are directed at the thesis I will be examining here, his main target is a stronger claim: the claim he calls "causal fundamentalism" and which maintains that "nature is governed by cause and effect and the burden of individual sciences is to find the particular expression of the general notion in the realm of their specialized subject matter" (Norton 2003). The thesis on which I will focus is not committed to either of the two conjuncts of causal fundamentalism, as defined by Norton. It is not committed to the first conjunct, since it allows for the possibility that on the most fundamental level—whatever that may be—nature is not governed by cause and effect, even though there are certain mature sciences which involve causal representations of nature. And it is not committed to the second conjunct, because the thesis allows that even if nature were fundamentally governed by cause and effect, it need not be the job of each and every individual science to find expressions of this notion. There may be many fruitful ways of representing various parts of nature that do not involve causal notions. Thus, I will not here be concerned with Norton's weaker thesis—the claim that causal fundamentalism is false—but only will critically examine arguments for the stronger claim that the notion of cause is at best part of a folk science and plays no role in physics.

A believer in 'metaphysically robust' causal relations in the world would certainly welcome the claims of this paper, but I myself want to resist drawing any conclusions about the metaphysics of causation. I am interested in whether asymmetric causal constraints play a role in scientific theorizing on a par with other physical constraints, such as the principle of energy conservation or other nomic constraints. How the postulates of a theory, including causal constraints, are to be interpreted metaphysically is a separate question. The debate between van Fraassen's constructive empiricism and scientific realism can serve as a useful comparison here (see van Fraassen 1980). Both van Fraassen's empiricist and the scientific realist agree that there are scientific theories that properly understood posit unobservable entities, such as electrons or quarks, yet they disagree on whether acceptance of these theories entails a commitment to the reality of the entities in question. By contrast, a traditional instrumentalist would deny that the theories at issue, properly understood, posit unobservable entities. Like the instrumentalist, the causal skeptic against whom I will argue in this paper denies that asymmetric causal notions play a role in scientific theorizing proper, even though the causal skeptics arguments might be quite different from that of a traditional instrumentalist.[3] I am here interested in the claim that corresponds to that shared by the constructive empiricist and the scientific realist—the claim that asymmetric causal notions, literally understood, play a role in theorizing in physics. But I do not here want to take sides in the debate as to what metaphysical conclusions we should draw from this fact.

2. Dimensions of 'cause'

Russell famously claimed that the word 'cause' is not used in the advanced sciences. Yet, as has been pointed out repeatedly—for example by Patrick Suppes (Suppes 1970) and more recently by Chris Hitchcock (Hitchcock 2007)—this claim is simply false. Contrary to what Russell maintained, the words 'cause', 'causal', and related words are still widely used in contemporary physics. A popular textbook on classical electrodynamics even maintains that a principle of causality is "the most sacred tenet in all of physics" (Griffiths 1989, 399). But what do physicists mean when they use causal language and what are their reasons for invoking causal notions? Unfortunately this question does not permit of a simple answer, since the use of causal notions is not much more well-regimented in physics than it is in everyday life and causal terms are used to express a variety of different claims in a variety of different contexts.

One of the few physicists who is careful explicitly to distinguish different aspects of the notion of cause in physics is Fritz Rohrlich. According to Rohrlich, there are three different meanings of causality in classical physics: "(a) predictability or Newtonian causality, (b) restriction of signal velocities to those not exceeding the velocity of light, and (c) the absence of 'advanced' effects of fields with finite propagation velocity." (Rohrlich 2007, 50) The third sense refers to the requirement that disturbances of a field associated with a field-source propagate into the future and not into the past. A cursory survey of contexts in which causal talk is used in physics seems to confirm that these three dimensions of the notion of cause indeed play a particularly important role in physical theorizing: first, that causes determine their effects; second, that causes act locally; and third, that the causal relation is asymmetric and that this asymmetry is closely related to the temporal asymmetry. All three aspects are, for example, part of Erwin

---

[3] Note that with respect to causal relations, van Fraassen himself is of course not a constructive empiricist, but more of a traditional empiricist.

Schrödinger's "principle of causality," which is the requirement that "the exact situation at *any* point P at a given moment is unambiguously determined by the exact physical situation within a certain surrounding of P at any previous time, say $t-\tau$." (Schrödinger 1951, 28). And determinism together with a temporal asymmetry constitute what Niels Bohr called a "causal description" and which rests on the "assumption that the knowledge of the state of a material subsystem at a given time permits the prediction of its state at any subsequent time" (Bohr 1948, 312).[4] My main focus in this paper will be on the causal asymmetry, but in this section I want to briefly discuss the other two causal dimensions as well.

In the first several decades after the quantum revolution it seems to have been not uncommon to maintain that quantum mechanics forces us to abandon a principle of causality. Bohr, for examples, contrasted causal descriptions with the non-deterministic descriptions of the new quantum physics and he took the latter to pose a threat to causality precisely because it is an indeterministic theory. It seems, however, that we have learned to live with genuine indeterminism and, correspondingly, the notion of cause has broadened to allow for the possibility of probabilistic causation. Thus, it appears to be less common among physicists today than it was perhaps in the first half of the twentieth century to refer to a condition of determinism as *the* 'principle' or 'law' of causality. Rohrlich, as we have seen, says that predictability is the hallmark of "Newtonian" or "classical" causality, which allows for other, indeterministic notions of causality as well. Nevertheless, many of the most general anti-causal arguments advanced by Russell and his more recent disciples assume deterministic micro-laws and I will do the same in much of my discussion below.

The causal locality constraints that physicists invoke fall into two broad classes: first, there are constraints against 'gappy' causation, which take the form of prohibitions against causes acting across spatial, temporal, or spatiotemporal gaps; and, second, there are constraints on the speed of causal propagation. Relativistic field theories satisfy both kinds of constraint: the presence of the field ensures that causes do not act across gaps, while relativity theory posits a finite upper limit on the speed of causal propagation. In fact, we can distinguish two constraints that relativistic theories satisfy: the condition that there is a finite, invariant velocity—the velocity of light; this condition is often expressed as demanding that spacetime has a lightcone structure. And the condition that there is no propagation in matter faster than the speed of light. Both constraints and the spacetime structures satisfying them are usually characterized in causal terms in the literature. For example, two points in spacetime that can be connected by a signal traveling at most at the speed of light—that is, points that are either timelike or lightlike related to each other—are called "causally connectable" (see, for example, Hawking and Penrose 1970). And curves in spacetime representing points moving at less than or equal to the speed of light are

[4] Hermann Weyl also refers to a principle of determinism as the "the law of causality" (Weyl 1989, 40).

called “causal curves.” In general relativity, models of Einstein’s field equations are said to satisfy a “causality condition” if they do not contain closed causal curves. In quantum field theories relativistic constraints are implemented in the form of a condition called “micro-causality,” which demands that the commutator between fields at spacelike separated spacetime points vanishes. Micro-causality is meant to capture the intuitive condition that the value of the field at one spacetime point can make no difference to the value of the field at another point, if the spacetime points are spacelike separated—that is, the two spacetime points could not be connected by a light signal or by any object moving slower than the speed of light.

What role does the causal language used to characterize relativistic locality constraints play in physical theorizing? Norton has argued that these constraints cannot plausibly be understood to be part of a general causal principle, since this would imply that

> any theory not complying with the causal principles of modern physics is causally deficient. The immediate consequence is that older theories, notably Newton’s mechanics, were causally defective in not admitting a finite upper bound to speeds of propagation. And that has the odd consequence that we were mistaken for hundreds of years in extolling the causal perfections of Newtonian mechanics.” (Norton 2007, 223)

We have to be careful, however, to distinguish the general framework given by Newton’s *laws of motion*, which survive in amended form even in relativistic theories and which, as Norton points out, are indeed often taken to be paradigmatically causal, from the particular *force law of gravitational attraction*. As Norton himself emphasizes, worries about the latter, which is an action-at-a-distance law and violates both kinds of locality constraint that I distinguished above, go back even to Newton himself. While it is true that physicists came to accept the law of gravitational attraction in light of its astounding empirical success, it may be that conceptual worries about the law were pushed into the background rather than being successfully resolved. Thus, rather than pointing to Newton’s law of gravity as a counterexample to a general condition of causal locality, one might instead take the development of relativistic theories as a vindication of the ‘ancient’ demand that causes act locally. And this does not undermine the intuition that Newton’s second law, in either its classical or relativistic incarnation, is the paradigm of a causal law.

One might think that the use of causal principles in the context of relativistic theories amounts to mere ‘labeling’. In general relativity, as Norton explains, subsets of the solutions to Einstein’s field equations are classified as causal in various senses. That is, causal notions are used to pick out proper subclasses of the theory’s models and causality conditions might, thus, be understood as devices for cataloging different solutions to Einstein’s field equations. Thus, Norton argues that none of the various causal conditions can be understood as universal constraint on physically possible solutions, since it is “routine to consider solutions to the Einstein equations that do not conform to them.” (Norton 2007, 228) Implicit in such an argument is a premise identifying models of the field equations with the range of what is physically possible. Yet someone who wished to defend the view that a relativistic causality condition provides a factual constraint could deny that all models of Einstein's equations represent physically possible universes or situations and could insist that causality conditions present additional potential constraints on what is physically possible. It is no argument against this view to point out that physicists also consider models that violate the condition. If a causality condition has the status of a hypothesis that is not yet sufficiently well-confirmed, then investigating what the world would be like if the condition failed may be an important component of testing the condition. And even if we took it to provide a well-confirmed constraint, exploring situations which violate the constraint—and hence are taken to be unphysical—can be a useful and fruitful exercise, since it can help up understand the theory better.

Finally one might think that there could be no causal constraint in general relativity prohibiting closed causal curves—so-called 'causal loops'—since there are strong arguments suggesting that the possibility of causal loops need not result in inconsistencies. But any argument along these lines is in danger of confusing physical possibility with conceptual possibility. It may be the case that causal loops are *conceptually* possible yet *physically* impossible and causal constraints are proposed as constraints on what is physically possible. A defender of the facticity of causal constraints, thus, would have to insist that we carefully distinguish between conceptual possibility, causal possibility, and what is possible according to a well-confirmed theory's basic equations.

One important usage of causal language in the context of relativistic theories is to express the time-*symmetric* locality constraint that spacetime has a lightcone structure and that there is no propagation outside the lightcone. But there is a second aspect to causal talk in relativity: causal notions are also used to mark a time-*asymmetric* distinction between the future lightcone, which is called "the causal future" of an event, and the past lightcone, which is the "causal past" of an event. A condition of local causality in physics is often introduced as the time-asymmetric condition that the fields at a point are fixed by the fields in the causal past of the point. This is in accord with the fact that the causal relation is asymmetric: if c is a cause of e, then it is not the case that e is a cause of c. Both Bohr's and Schrödinger's principles reflect this fact: According to Bohr's assumption, a causal description of a deterministic system is one that characterizes the system's evolution in terms of an initial value problem. An 'anti-causal' description, by contrast, would then be one that describes the evolution of a system in terms of final value problem. And Schrödinger's condition of local determination is likewise a principle of past-to-future determination.

This third aspect of the notion of cause—its asymmetry—is arguably the most central of the three dimensions and will be the main focus of this paper. The asymmetry is clearly an integral part of our intuitive idea that causes 'bring about' or 'produce' their effects, but it is also an integral part of less metaphysically 'weighty' notions of cause. As the identification of the causal future with the future lightcone in relativity theory attests, the causal asymmetry is intimately related to a temporal asymmetry, even though what precisely the relation is, is somewhat of a delicate issue. On some accounts of causation, such as Humean regularity accounts, it is a conceptual truth that effects do not precede their causes, but even those accounts that allow for the conceptual possibility of backward causation would presumably maintain that causation in our world (or at least in the spatiotemporal region of the universe accessible to us) is forward directed, and hence, causal constraints are often taken to imply time-asymmetric constraints.

Fortunately, for our present concerns we do not need to resolve the issue of what ultimately the relation between the asymmetry *of* time and the causal asymmetry is. The arguments I will consider in the next section concern properties of scientific theories that govern the behavior of physical system *in* time. My concern in the present paper, thus, is what—if anything—we can conclude from how the fundamental equations of physical theories represent the evolution of physical systems in time for the prospects of interpreting these equations causally. And investigating the relation between symmetries or asymmetries of physical phenomena *in* time and the causal asymmetry can be done independently of having an account of the asymmetry *of* time.

3. No place for asymmetric causes?

I want to begin this section, too, with Russell's famous paper (Russell 1918). Russell argues that imprecise common sense causal regularities are replaced in physics by precise laws that have the form of functional dependencies. The argument appears to be roughly 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. 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. That is, instead of vague regularities of the form "When an object at rest experiences a sufficiently strong force it tends to begin to move" or a precise but perhaps only singly instantiated regularity of the form "Whenever a ball of mass m_0 in circumstances c_0 is struck with a force f_0 it accelerates at rate a_0" physics presents us with logically much stronger functional dependencies, such as that all massive bodies obey Newton's law $\boldsymbol{F}=m\boldsymbol{a}$. Russell claims that such functional dependencies have replaced putatively causal regularities in physics, but of course it does not follow from the fact that physical theories present us with functional dependencies that these dependencies themselves cannot be understood causally. How, then, might we try to establish the claim that there is no legitimate place for an asymmetric notion of cause in fundamental physics? I want to examine several arguments for this claim, some of which are meant to spell out some of Russell's own suggestions while others are suggested by the more recent literature.

Let us first consider what may be the most ambitious argument—an argument that does not explicitly invoke the asymmetry of the notion of cause and is aimed at showing that any notion of cause that does not take causal relations to be reducible to non-causal facts is meaningless or incoherent. This argument is suggested, for example, by some of van Fraassen's discussions of causation and begins by asking us to contrast a putatively causal world with a 'Hume world' replica of that world—that is, a world, that is identical to the first world as far as its Humean matters of fact are concerned but that does not include any causal relations between events (see van Fraassen 1993). If we further assume that the Hume world is empirically indistinguishable from its putatively causal twin, and also assume a weak verification principle according to which for a concept to be meaningful, there must be some empirical differences between those situations when the concept applies and those when it does not, then it follows that the notion of causation is meaningless. Thus, in explicit premise-conclusion-form the argument is:

1. If there are asymmetric causal relations between events in a world (a c-world), then there is a qualitatively distinct possible world that is a replica of the c-world except that there are no causal relations between any events in that world (a Hume-world replica).
2. A c-world and its Hume-world replica are qualitatively distinct, only if the notion of cause is meaningful.
3. For the notion of cause to be meaningful, a c-world and its Hume-world replica must be empirically distinguishable, at least in principle. (Verification principle)
4. The c-world and its Hume-world replica are empirically indistinguishable.
5. Therefore, the c-world and its Hume-world replica are not qualitatively distinct. (2,3,4)
6. Therefore, the c-world contains no asymmetric causal relations. (1,5)

Note that this argument does not, of course, undermine any account of causation that takes causal relations to be reducible to Humean facts, since such accounts would deny (1). Any world containing the same Humean matters of fact would contain the same causal facts, if causal facts were reducible to non-causal facts. Yet advocates of a Humean account of causation are not the only ones who would reject (1). A defender of a notion of cause richer than that allowed by a Humean can question whether the notion of a Hume world replica is coherent. Thus, Nancy Cartwright has argued that many events, such as shovings or milk lappings, are intrinsically causal (Cartwright, 1993, 427). If this is correct, then a Hume-world replica of a causal world, in which events have all their non-causal properties but none of their causal properties, does not constitute a coherent possibility. Imagining a world that is like ours except for its causal content—that is, among other things, a world in which the cat's tongue moves as it does in the actual world and the milk disappears in its mouth—but in which there are no intrinsically causal lappings, Cartwright contends, "is ridiculous." (427) Thus, a defender of causation could, with Cartwright, reject (1).

It is not easy to assess the relative merits of the two competing positions here in ways that do not simply beg the question against one of the two views. Both the defender of causation and the causal critic want to conclude that a c-world is not qualitatively distinct from its Hume-world replica. According to the argument of the causal critic, a putatively causal world already is a Hume world, while the defender of causation insists that any replica of a causal world will still contain causal relations, if that world is to present a coherent possibility at all. Is it not possible to coherently imagine qualitatively distinct pairs of c-worlds and Hume worlds, because the notion of cause itself is incoherent or because causal relations are woven into the very fabric of causal worlds in ways that makes it impossible coherently to imagine their absence? I worry that at this point we have simply reached an impasse of clashing intuitions.

Our first argument does not yet involve any appeal to physical theorizing. Perhaps we can sidestep the issue concerning the coherence of a Hume world and do so in a way that also explicitly includes an appeal to mature theories of physics. For a concept to be meaningful we might assume, there have to be conditions of its employment that are ultimately underwritten by the results of science—science, that is, is the ultimate arbiter of meaningfulness. But, the argument contends, causal notions are not part of mature physics:

7. For the notion of cause to be meaningful, it has be part of the mature theories of physics. (Verification principle*)
8. The notion of cause is not part of any mature physical theory.
9. Therefore, the notion of cause is not meaningful. (7,8)

A serious problem with this argument is that—like the previous argument—it relies on a version of a verification principle of meaning and the history of twentieth century philosophy has taught us how dubious such principles are. And (7) is even more questionable than its cousin (3), since there obviously are many meaningful concepts that are not part of mature physics. Moreover, the argument's second premise, (8), is itself in need of an argument. Thus, it seems more promising to grant the meaningfulness of an asymmetric notion of cause and argue against the claim that such a notion has a legitimate role to play in physical theorizing. In what follows, then, I want to examine possible arguments for (8), restricting myself to arguments that grant, at least for the purposes of the arguments under consideration, that causal notions are meaningful.

Russell suggests one argument for (8) during his discussion of Newton's law of gravity. He 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 remark is echoed by Bas van Fraassen, who answers Cartwright's question "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 causing, 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 wherewithall to distinguish those from shovings of big sisters by kid brothers. The distinction is made outside the theory. (van Fraassen, 1993, 437-8)

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 physical theorizing. We might try to reconstruct at least part of their remarks in terms of the following explicit argument:

10. The content of a physical theory is exhausted by a set of state-space models or a set of formulas.
11. Causal relations are not part of the formulas or models of a theory.
12. Therefore, causal relations are not part of the content of physical theories.

As it stands, however, (10) is false. Mathematical physics provides us with mathematical models or representations of the world, yet on their own mathematical models do not represent anything. How a given model or class of models represents the world depends on how the model is interpreted. Thus, no theory of physics can be strictly identified with a set of formulas or state-space models, since, minimally, a 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.[5] But once we see that the austere view of theories as consisting solely in a mathematical formalism or set of models is untenable and that an interpretive framework needs to be part of a theory, it is no longer obvious that this framework cannot be rich enough to include causal assumptions as well. That is, we cannot conclude from the mere fact that an *uninterpreted* formula does not on its own mark, say, **F**'s as causes and **a**'s as effect, that one cannot interpret the formalism causally. That causal distinctions are not part of the uninterpreted mathematical formalism does not imply that they are made outside the theory.

One might reply that while a theory clearly cannot be identified with a bunch of uninterpreted mathematical squiggles, its content is exhausted by a formalism together with a minimal interpretation fixing the theory's ontology. But instead of answering the question as to what is allowed to be part of the

---

[5] One might think that on van Fraassen's view (or at least according to a view van Fraassen once held) 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.

theory's interpretive framework, this reply merely postpones it. What can properly be part of the minimal interpretation and why can causal interpretations of certain mathematical relations not be part of a 'minimal interpretation'? One 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 any 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 as "physically well-founded assumptions" (Jackson 1975, 312), as "fundamental assumption" (Nussenzveig 1972, 4) or as "general physical property" (Nussenzveig 1972, 7), or even, as we have seen, as "most sacred tenet in all of physics" (Griffiths 1989, 399).

One common suggestion is that the specific form of the fundamental equations of physics partly determines the range of their admissible interpretations. These equations, it is argued, are time-symmetric and hence do not permit of an asymmetric causal interpretation. Such an argument, which may also be suggested by van Fraassen in the quote above, can again be traced to a remark by Russell, who said that "the laws make no difference between past and future: the future 'determines' the past in exactly the same sense in which the past 'determines' the future." One way to interpret Russell here is as claiming that the purported fact that physics provides us with time-symmetric functional dependencies undermines the claim that asymmetric causal relations can play a role in physics. On another reading, the premise of the argument is that the fundamental equations are both future- and past-deterministic—that is, define both a well-posed initial and a well-posed final value problem—irrespective of whether the laws are time-symmetric or not. These two readings result in two distinct arguments. First, the argument from determinism:

13. The fundamental equations of classical physics are both past- and future-deterministic.
14. There is no place for an asymmetric notion of cause in the context of a theory with fundamental equations that are both past- and future-deterministic.
15. Therefore, there is no place for an asymmetric notion of cause in mature physical theories.

Of course, we no longer believe that the fundamental laws of nature are deterministic, but there are independent reasons for rejecting the conclusion. The argument relies on the assumption that in situations where causes determine their effects, the set of effects of an event cannot in turn determine its causes, 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, 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.[6] 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.

Once we move to the context of genuinely probabilistic theories, the case of the defender of causal relations may be even stronger, since, as (Callender 2000) shows, any non-trivial theory that specifies transition probabilities possesses a time-asymmetry. Any theory that specifies both non-trivial forward and on-trivial backward transition probabilities for a system has the consequence that the expected state of the system cannot change with time. Hence, any interesting physical theory could specify forward transition probabilities or backward transition probabilities, but not both. In particular, this argument shows that a quantum theory with transition probabilities cannot be time-reversal invariant. A proponent of a causal interpretation of a probabilistic theory might appeal to this asymmetry and argue that the direction of causation ought to be identified with the direction of the theory's transition probabilities and that this asymmetry is precisely the kind of asymmetry of the formalism that Russell was looking for.

The second argument that one might extract from Russell's remark appeals to the time-reversal invariance of a theory's fundamental equations:

16. The fundamental equations of all mature physical theories are time-reversal invariant.
17. There is no place for an asymmetric notion of cause in the context of a theory with time-reversal invariant laws.
18. Therefore, there is no place for an asymmetric notion of cause in mature physical theories.

Both premises of this argument, too, are open to challenge. Quantum mechanics may seem to contradict (16). Since the Schrödinger equation is first-order in time, it is not time-reversal invariant (see Callender 2000). The Schrödinger equation is invariant under the joint operation of time-reversal and complex conjugation (in the coordinate representation), but since the equation is not invariant under time-reversal alone, some philosophers wish to conclude that, as Callender puts it, "time in a quantum world is handed." Curiously, then—and perhaps somewhat surprisingly—certain anti-causal arguments are even less compelling in the context of quantum mechanics than in the context of a Newtonian and putatively deterministic physics.

[6] 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.

But even if were to grant (16), the move from the time-reversal invariance of the dynamical laws to the claim that there is no room for causes in a theory with such laws is far from compelling. First, (17) papers over the distinction between the claim that the relation between cause and effect is asymmetric and the claim that it is *temporally* asymmetric, yet there are cases where we may want to take a cause to operate simultaneous with its effect. Newton's Second Law, for example, is often interpreted causally in that the external force acting on a body is taken to be the cause of the body's simultaneous acceleration. And it is not immediately obvious how the time-symmetry of Newton's law could have any bearing on the question whether the force at a time is properly thought of as asymmetric cause of the simultaneously occurring acceleration.

Yet someone who is inclined to think of Newton's Laws causally presumably would also hold that causally affecting the acceleration of an object has an effect on the object's future evolution, but not on its past evolution, and it is this idea, presumably, that is supposed to be undermined by the observation that Newton's equations are time-reversal invariant. The claim, to repeat, is that time-symmetric dynamical laws are incompatible with time-asymmetric causal relationships. The intuitive idea that an external force on an object is a contributing cause to the object's later states but not its earlier states is undermined, according to this argument, by the fact that the dynamical equations also have a solution for which initial and final state are exchanged (or more, precisely where initial and final states are replaced by the time-reversed final and initial state, respectively).

I see no reason, however, why a defender of causation in Newtonian physics should be moved by this argument. One might try to argue that the notion of causation is incoherent. But this, it is important to recall, is not the claim we are currently considering. The current argument grants the prima facie legitimacy or meaningfulness of causal notions but maintains that it is impossible for a world governed by time-symmetric laws to be time-asymmetrically causal. That is, the present argument grants that we can conceive of a world that evolves causally but insists that there could be no causally evolving world that is governed by time-symmetric dynamical laws. It is unclear to me what an argument for this latter claim might be. It does not follow from the fact that a formalism is time-*asymmetric* that that formalism has to be interpreted causally. (Hence I do not want to suggest that the claim that the Schrödinger's equation is not time-reversal invariant implies that it has to be interpreted causally.) And similarly, it does not follow from the fact that a formalism is time-*symmetric* that it cannot be interpreted as representing states of a causally evolving system. Independently of whether theories are construed syntactically or semantically, it seems to me that there can be no general argument of the kind suggested by Russell or van Fraassen to show that causal assumptions cannot form an integral part of theories in physics.

A related argument tries to derive a contradiction from the conjunction of time-symmetric dynamical laws and an asymmetric causal assumption that picks out a subset of the theory's models as causally—and, hence, physically—possible. Let us postulate, as a constraint on what is physically possible, a causal principle according to which future states of a system causally depend on past states. And let us assume that only the members of a proper subset of the theory's laws, comprised of certain time-asymmetric models, satisfy this principle. Then it follows from the time-reversal invariance of the laws that for each causal model there is a time-reversed model in which relations of dependence isomorphic to that of the causal model obtain, except that all dependence relations will be time-reversed. Hence the time-reversals of the causal models will violate the causal principle. But since the time-reversed models are also models of the theory's dynamical equations, they, too, are physically possible. Thus, it follows that the causal principle provide does not provide a constraint on what is physically possible after all.

But this argument conflates the notion of what is physically possible with that of being a model of the theory's dynamical equations. A defender of the idea that causal principles can function as constraints on what is physically possible would deny that all the models of a theory's basic equations do represent genuine physical possibilities. The time-reversed models are dynamically possible, but if they violate the causal condition they are causally, and hence physically impossible.

There is one final kind of argument I want to consider, which offers perhaps the most promising strategy for the causal skeptic. This argument does not question the possibility or conceivability of a causal world with time-symmetric dynamical laws, but maintains that we could never have good reasons to postulate a time-asymmetric causal dependence of future states on past states in such a world. If there cannot be scientifically justified reasons for adopting a certain causal interpretation, then picking one amounts to a mere convention. Indeed, Bohr's notion of a *causal description* appears to rely on a purely conventional asymmetric distinction. The basic equations of a theory that is both future- and past-deterministic theory define both an initial and a final value problem and, hence, the very same time-evolution of a system can be characterized in two different ways—as solution to an initial value problem and as solution to 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. Bohr suggest that the former approach is causal while the latter is not, but since both approaches yield one and the same temporal evolution, calling the former "causal" seems to be only a question of labeling. This view is endorsed explicitly by Fritz Rohrlich, who maintains that the "identification of *causality* with *prediction* rather than *retrodiction* in a time-symmetric system of equations is completely arbitrary." (Rohrlich 2007, 51 italics in original)[7] Thus, while there may be no good arguments that disallow interpreting a theory causally,

[7] 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)

there might also be no scientifically legitimate reasons for an asymmetric causal interpretation—or so one might claim. On the most charitable interpretation it may in fact be this kind of argument to which van Fraassen is alluding in the passage quoted above: there is no place for causal assumptions in mature physics, because we could never have scientifically legitimate reasons for interpreting time-reversal invariant laws causally. In premise-conclusion form the argument is:

19. The fundamental equations of all mature physical theories are time-reversal invariant.
20. There are no scientifically justified reasons to include causally asymmetric notions in the interpretive framework of a theory with time-reversal invariant laws.
21. Therefore, there can be no scientifically justified reasons to include causally asymmetric notions of cause in mature physical theories.

Although this argument strikes me as offering the most promising anti-causal strategy, its premises are not self-evidently true. (20) implies that any legitimate justification for interpreting a theory causally would have to appeal to the character of the theory's dynamical laws. But there are at least two other avenues through which one might support a causal interpretation. First, one might argue that our experimental interactions with physical systems provide us with scientifically legitimate reasons for adopting an asymmetric causal interpretation even of a theory with time-symmetric dynamical laws; and, second, it may be the case that asymmetries in prevailing initial or boundary condition are best explained by appealing to causal assumptions. In the next section I want to examine the first of these avenues in more detail. For a detailed discussion of the second strategy see (Frisch 2005).

4. Justifying Causal Interpretations

While some of our physical theories are meant to have in principle universal scope, many experimental tests even of putatively universal theories concern applications to finite systems and our interactions with such systems exhibit a striking asymmetry even in the case of systems governed by both future and past deterministic laws. Consider how we prepare a system that is closed during a time-interval $t_1<t<t_2$ and is in a specified initial state S_i at t_1 and a final state S_f at t_2, subject to the constraint that S_f is the final state into which S_i evolves dynamically (or, equivalently, that S_i is the state resulting from evolving S_f backward in time). While the dynamical laws do not single out a temporal direction, possible state-preparation procedures are time-asymmetric. We can prepare the initial state of the system at t_1 and then calculate the system's future evolution for times $t>t_1$ from the initial state and the dynamical laws, but what we cannot do is first prepare the system's final state at t_2 and then use this state together with the laws to calculate the system's past evolution for $t<t_2$. The only way for us to 'prepare' the system in a specific final state at t_2 is by making use of our knowledge of the dynamical laws to determine the initial state in which the system has to start out at t_1 in order to evolve into the final state in question and then

prepare the system in the appropriate initial state. Thus, our ability to prepare the system in its final state relies crucially on our knowledge of the system's evolution between t_1 and t_2. By contrast, we can prepare a system in an initial state without knowledge of its evolution between t_1 and t_2. This experimental asymmetry also manifests itself in isolated interventions into an otherwise closed system. Two experimenters can, without knowledge of the other's actions, prepare a system's *initial* state and engage in a later isolated intervention into the system; but they cannot similarly set up the system's *final* state and perform an earlier intervention independently of each other. Despite the fact that a system may be governed by time-reversal invariant dynamical laws and is both past- and future-deterministic, our experimental interactions with the system exhibit a temporal asymmetry—we interact with the system from the past, as it were.

Given the close connections between the notions of intervention and causation, the asymmetry characterizing interventions can be taken as evidence for a causal asymmetry. These connections have been exploited in recent interventionist or manipulability accounts of causation, such as (Woodward 2003). Woodward treats causation as a relation between variables. Roughly, and ignoring a number of qualifications that are important in other contexts but need not concern us here, according to Woodward's account, a variable C is the cause of a variable E, if intervening on the value of C results in changes to the value of E. Since our experience with physical systems suggests that interventions into a physical system affect its future evolution but not its past, an interventionist account of causation supports interpreting the relation between the state of the system at different times causally. Conversely, the causal asymmetry provides an explanation for the asymmetry of intervention: it is precisely because earlier states of a system *cause* later states that our experimental interactions with otherwise closed systems exhibit the temporal asymmetries just discussed. It is because of the causal asymmetry that interventions into an otherwise closed system percolate through the system forward in time.

One might worry that even an appeal to the asymmetry of intervention and manipulation cannot establish the scientific legitimacy of a causal interpretation, since the notion of cause is too vague and imprecise and, therefore, a causal interpretation can never have be other than an alien element that is externally grafted onto a theory. For example, Christopher Hitchcock has recently argued that causal notions play a role in the more fundamental sciences only as provisional and imprecise guides to future research (Hitchcock 2007). According to Hitchcock, the use of causal notions in a field of physics is a sign that physics has "regressed" (p. 56) in that we have encountered new phenomena that cannot yet be given a precise mathematical treatment. In support of his view Hitchcock cites an earlier discussion by Patrick Suppes, who maintains that the notion of cause can play a legitimate role in science, but only as long as the phenomena at issue "have not yet succumbed to comprehensive mathematical analysis" that provide a "more precise characterization" (Suppes 1970, 6). Thus, both Hitchcock and Suppes contrast causal talk with precise mathematical relationships, such as differential equations and maintain that it "becomes appropriate to eliminate causal talk" (Hitchcock 2007, 56) from physics proper once the phenomena at issue have been characterized with the help of precise mathematical relations.

However, the contrast between precise mathematical relations and causal notions is a false dichotomy, as I have argued above. A theory's mathematical formalism can be interpreted causally and the resulting causally interpreted formalism inherits the precision from the underlying formalism. A bit more explicitly the notion of a causal relation between events or between the states of a physical system can be introduced as follows. Let the state of a system $S(t)$ be given by the values of a set of variables $s_1(t), s_2(t), \ldots, s_n(t), \ldots$, 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 a state-space model. We can then define an asymmetric, transitive, and non-circular 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. If $S(t_2)$ bears C to $S(t_1)$, then $S(t_2)$ is an effect of $S(t_1)$. If two states do not stand in relation C then they are not causally related. The result is a class of what we might call *potential causal models*. Depending on the theory in question, there may also be more fine-grained causal relations defined over individual state variables s_i.

In the case of relativistic theories, the relation C holds between the local values of the variables at spacetime points. The state-variables at a spacetime point p are a cause of the state at q just in case either there is a future directed timelike curve from p to q or q lies on the future lightcone of p. If there are dynamical models over which C cannot be defined, then there are dynamical models for which there is no corresponding potential causal model. Models of Einstein's field equations with closed timelike curves are an example of this. Thus, the demand that the dynamical models must permit a causal interpretation can lead to a constraint on the class of permissible models. (See, for example (Hawking and Penrose 1970) for how causal conditions can be introduced formally in this context.)

In some instances we might want to impose additional causal constraints restricting the class of potential causal models to a proper subset comprised of those models that are *causally possible*. I argue in (Frisch manuscript) that causality conditions in the derivation of dispersion relations function impose such an additional constraint. Curiously, Suppes himself refers to the example of dispersion relation as an instance of causal talk that is merely preliminary. Yet causality conditions function in this context precisely in the manner that Suppes himself contrasts with the putatively preliminary role of causal assumptions—that is, "as simple fundamental law to derive exact relationships such as those expressed in differential equations." (Suppes 1970, 6)[8]

---

[8] Hitchcock's choice of examples is equally as puzzling as Suppes's: One of Hitchcock examples of causal talk as preliminary is the use of causal notions in causal set cosmology. Causal sets are sets with a transitive, asymmetric,

The causal relation at issue is in some sense quite minimal and might, in fact, seem to be too minimal to qualify as properly causal. Thus, (Woodward 2007) maintains that the notion of cause may not be applicable to theorizing in physics, since, as he argues, the distinction between genuine causes and background conditions cannot be drawn in this context. Woodward points out that common sense causal claims, such as 'striking a match causes the match to light' relate relatively small numbers of events (or, in his framework, variables) to one another. Of course such claims hold only if certain additional conditions apply, but these conditions have the status of background conditions and are not part of the causal regularity itself. Within the context of a complete physical theory of the match and its surroundings, by contrast, we can no longer draw a distinction between causes and background conditions and the physical state of the room next door will be a cause of the precise lighting of the match just as the micro-state of the match will be. But this, it may seem, is an absurd conclusion—clearly, the striking of the match is a cause of the lighting in a way in which the state of the room next door is not—and in order to resist it we may want to conclude that the concept of cause is essentially a concept meant to express the kind of imprecise, limited, and course-grained dependencies between small numbers of variables, with which the special sciences or common sense are concerned.

However, we can grant that the distinction between salient causal factors and background conditions plays an important role both in ordinary discourse and in the special sciences and yet resist the claim that there can be no notion of cause that does not recognize this distinction. The importance of the distinction, one can argue, reflects the close connections between the concepts of cause and of explanation is an important *pragmatic* component of our concept of cause. Causes, intuitively, are 'difference-makers,' but not everything that in principle can make a difference to the occurrence of a particular event is salient in every context. Yet in addition to its pragmatic dimensions we can also identify a non-pragmatic, minimal core of the concept of cause which is given by the causal asymmetry—and it is the applicability of such a minimal concept of cause that can be justified by the asymmetry of our experimental interactions with physical systems.

Corresponding to the asymmetry of causation and intervention, there is a counterfactual asymmetry: a certain kind of counterfactual that we might call 'causal counterfactual' or 'intervention counterfactual' is time-asymmetric. This is not to say that there is a *general* asymmetry of counterfactuals in 'standard contexts', as, for example, David Lewis has claimed. Indeed, it seems to me that any attempt to derive such a general asymmetry for counterfactuals—for example along the lines of Lewis's account—is bound to fail, for there are perfectly standard contexts in which backtracking counterfactuals are true. One kind of counterfactual that is time-symmetric is underwritten by time-

---

and non-cyclical relation that defines a partial ordering over the set. Defenders of causal set cosmology propose causal sets as fundamental structures at the heart of a theory of quantum gravity.

reversal invariant dynamical laws that define an initial or final value problem. These laws provide us with assertability conditions for counterfactuals of the form "If a closed system were in state S_i at t_1, then it would be in state S_f at the later time t_2". We evaluate the counterfactuals by determining whether the dynamical laws evolve the state S_i into the state S_f. But the laws also underwrite corresponding backtracking counterfactuals of the form "If a closed system were in the state S_f at t_f, then it would have to have been in the state S_i at the earlier time t_i", which we evaluate analogously by evolving S_f backward in time in accord with the laws. Intervention counterfactuals, by contrast, are asymmetric and are evaluated by holding the actual causal history of the counterfactual antecedent fixed and determining how the changes posited in the antecedent affect events causally 'downstream.'

Now, one might worry that if this counterfactual asymmetry cannot be underwritten by considerations along Lewis's lines, then treating intervention counterfactuals asymmetrically is a matter of choice or perhaps merely a reflection of our particular perspective on the world (see, for example, Price 2007). We cannot simply observe that counterfactual interventions affect the future and not the past, for the simple reason that we can never observe counterfactual cases, as Price argues (Price 2007, fn. 22). Still, the experimental asymmetry supports our asymmetric treatment of intervention counterfactuals. Despite the obvious truth of Price's claim that we cannot actually observe counterfactual cases, in many cases we can experimentally probe counterfactual claims about a system by setting up replicas of the system that instantiate the counterfactual states in question. Of course we need to be careful that we do not import an asymmetry from the start by treating as relevant replicas only systems that are prepared in identical initial states. Rather we have to consider *both* replicas that agree with the original system in their initial states *and* replicas that agree in their final states. But it turns out to be a fact about our experimental interactions with such systems that it is generally much easier to set up multiple systems with identical initial states but different subsequent interventions than it is to set up systems that evolve into identical final states despite different prior interventions. The reason for this asymmetry is precisely the fact that the only way for us to set up systems that evolve into identical final states despite different earlier isolated interventions is to prepare the systems in a carefully differently chosen initial states that compensate for the later interventions. Thus, there is a sense in which systems with initial states identical to that of the actual system and non-actual final states are 'experimentally closer' to the actual system than systems with final states identical to the actual final state but non-actual initial states.

As a concrete example consider a system consisting of a charged object in an external electromagnetic field. We are given the actual initial and final states of the field and the charge and are asking what the effects of interventions on the motion of the charge would be. It is, at least in principle, relatively easy to set up different systems in the same initial state characterized, say, by zero incoming fields and then intervene differently on the motion of the charge by accelerating it by different amounts.

It is much harder to make sure that the final fields for several different systems are the same, even though the systems' charges have earlier been accelerated by different amounts. It is a mathematical fact that we can represent the total electromagnetic field at any point *either* as a sum of a source-free incoming field—the initial field—and diverging fields associated with the charge *or* as a sum of a source-free outgoing field—the final field—and converging fields associated with the charge. That is, as far as the possible mathematical representations of the total field are concerned, there is no asymmetry. Yet our experimental interactions with systems of charges and fields suggest that, at least for points sufficiently distant from the charge, we can independently manipulate the source free-incoming field and the field diverging from the source. By contrast, we can only manipulate the formally 'source-free' outgoing field through manipulations of the incoming field together with carefully calibrated interventions on the motion of the charge.

I have argued that there exist formally precise ways in which causal relations can be introduced into a theory and that a causal interpretation of the relation between the states of a system at different times can provide an explanation of a pervasive asymmetry characterizing our experimental interactions with physical systems. Does this justify interpreting mature physical theories causally? There are two remaining worries that need to be addressed. First, my discussion relies crucially on the fact that we use our theories to model *finite* systems that we can manipulate experimentally. But many of the mature theories of physics are intended to have universal scope and are taken, at least in principle, to govern entire 'possible worlds' and not just finite systems. How can the interventionist considerations be brought to bear in this case? Second, some have suggested a different source for the asymmetry in our experimental interactions—the entropic or thermodynamic behavior of the world. One might argue that while a causal interpretation provides *a* potential explanation for the experimental asymmetry it is not *the best* explanation and that a better explanation is provided by broadly thermodynamic considerations.

Even many of those theories with potentially universal scope, however, are used to model finite systems. And it seems to be part of our inductive practice in science that—to the extent that we take our theories in principle to provide us with models of 'possible worlds'—we take these worlds to be characterized by the same kind of relations which govern the finite systems with which we interact experimentally. If we think of a theory as having universal scope, then we take models of finite systems to be embeddable into models of the theory representing richer possible worlds, which are characterized by the same type of relations as the finite systems. The fact that our evidence for causal relations comes from models of 'small worlds' does not imply that we cannot thereby also be justified in interpreting 'large worlds' causally. It may be the case that our justification for positing causal relations relies crucially on our interactions with finite systems; but, first, we can consistently embed causal models of such systems into causally possible physical worlds, and second, our warrant for positing causal relations in the case of large worlds arguably transfers from our interactions with finite systems.

What are the prospects of a rival explanation of the experimental asymmetry or of asymmetries in initial conditions? Rival explanations that have been proposed all appeal, in some form or other, to the foundations of thermodynamics. David Albert (Albert 2000) and Barry Loewer (Loewer 2007) have argued that the assumption that the universe began in a state of extremely low entropy—what Albert calls the 'past hypothesis'—underwrites the causal asymmetry. Price (Price 2007) has argued that the causal asymmetry is a perspectival asymmetry characteristic of beings like us whose epistemic access to the world, due to the entropy gradient in our spatiotemporal region of the universe, is time-asymmetric. While offering fewer details of a reductive account, Norton, too, maintains that any putatively causal relations can be generated from more fundamental non-causal laws through reduction relations.

Now, it is important to distinguish between different theses that a reductive account of the causal asymmetry might advance. One possible thesis is that the notion of cause is not part of our most fundamental physics but may nevertheless play a scientifically legitimate role in less fundamental yet mature theories. As an analogy consider the notion of a classical field in classical electrodynamics. We do not think that classical fields are part of our most fundamental theory of the world, which presumably is a quantum theory, but this does not mean that the concept of a classical field cannot have explanatory power and is not an integral and immensely fruitful part of contemporary physics. Classical fields theories are a legitimate part of mature physics, have explanatory value, and arguably even provide us with the best explanations for certain phenomena, and this is so even though we think that classical field theories are empirically adequate only within a certain domain and may ultimately be reducible to quantum field theory. Thus, if the claim of the reductionist is that the causal asymmetry is not fundamental in the same sense in which classical fields and the Maxwell equations are not fundamental, then what I have argued here is compatible with the reductionist's view. I have not argued for the thesis that causal notions are part of our most fundamental science—whatever that may be—but only that there can be good reasons for interpreting the models of certain mature theories causally.

Alternatively, a reductionist might argue that the role of causal notions is in crucial respects analogous to that of the notion of 'caloric'—that is, it is a notion which even though it might have some explanatory value can be 'reduced away' and be shown not to have a place in mature theories of physics. According to the first thesis, a reductive account are not proposed as a rival to causal explanations in all contexts but merely helps to establish the domain of validity of such explanations—just as quantum field theory is not a rival theory to classical electrodynamics challenging its explanatory adequacy in the domain of classical physics but merely establishes its domain of validity. According to the second thesis, there exist alternative and scientifically superior accounts of the asymmetries that I have discussed, which show that causal relations have no legitimate role to play in mature physics.

Which of these two theses do existing reductive accounts defend? It seems to me that by and large their aim is to establish the second claim. Price, for example, calls the causal asymmetry a piece of "naïve physics" (Price 2007) and argues that it merely reflects our particular epistemic perspective on the world, while Norton explicitly compares the notion of cause to that of caloric and argues that causal relation are part of 'folk science' (Norton 2003).[9] But the situation is not entirely clear since Norton also compares the notion of cause to that of gravitational force—that is to a notion that retains even today a legitimate role in theorizing in mature classical physics. To the extent that reductive accounts aim to show that causal notions can at best be part of a folk science but not of mature physics, these accounts present a challenge to the view I defended her in addition to the critical arguments examined in this paper. I have here argued that a range of general anti-causal arguments fail and I have pointed to a range of phenomena that can be—and in fact often are—explained causally. The fact that there are asymmetries that do not manifest themselves in asymmetries of the dynamical laws and that *can* be given a causal explanation undermines any anti-causal argument purely from the time-reversal invariance of the dynamical laws. This does not yet settle the question whether an appeal to causal considerations provides the *best* explanation of these asymmetries, but it provides a challenge to the causal skeptic to provide an explanation for these asymmetries that are superior to a causal account. I have argued elsewhere that one particular reductive strategy—Albert and Loewer's attempts to ground the causal or counterfactual asymmetry in a Boltzmannian account of thermodynamics—is unsuccessful (Frisch 2006b; 2007; forthcoming). I am skeptical of other reductive accounts as well, but there is no space to argue this point here.

A second observation that might lead us to posit time-asymmetric causal notions in a world with time-symmetric laws is that it does not follow from the fact that the laws of a theory are time-reversal invariant that all the models of the theory are time-symmetric. To the contrary, 'most' models of a theory with time-reversal invariant laws will not be time-symmetric. But if actual systems governed by a certain theory are best represented by models that all (or at least overwhelmingly) exhibit the same kind of asymmetry—that is, if there is an asymmetry between the initial and final conditions characterizing models of typical actual systems—then this, too, might be evidence for causal relations among the physical quantities involved. Again, the temporal asymmetry characteristic of electromagnetic waves in the presence of wave sources is one example of this. This asymmetry intuitively consists in the fact that there are coherently *diverging* waves in nature but not their time-reverse—coherently *converging* waves—even though the dynamical laws governing waves are time-symmetric. One can show that this

[9] Yet Norton also compares the notion of cause to that of a gravitational force, which suggests the first thesis.

asymmetry in fact consists in an asymmetry between prevailing initial and final conditions—roughly, the fact that incoming waves are often approximately equal to zero, while outgoing waves in the presence of wave sources are generally non-zero. One explanation for this asymmetry, which I defend in more detail in (Frisch 2000; 2005; and 2006a), is that wave sources asymmetrically cause disturbances propagating into the future. In fact, Griffith's appeal to causation as the most sacred tenet of physics, quoted above, appears in this very context and this is also the view defended by Rohrlich (2006), who argues that only retarded (i.e. diverging) fields exist physically, even though the Maxwell equations have solutions that describe converging fields associated with a source (i.e., fields that, intuitively, radiate toward the past).

5. Conclusion

In this paper I distinguished three uses of causal notions in physics, expressing a principle of determinism, imposing locality conditions, and expressing time-asymmetric constraints. I focused on the third usage and examined a range of arguments that have been advanced against the claim that an asymmetric notion of cause can play a role in mature physical theorizing. None of these anti-causal arguments, I argued, succeed. I then argued that there are two routes by which the use of causal models can be justified even in the context of a theory with time-reversal invariant laws. The first route appeals to the pervasive experimental asymmetry that our interventions into physical systems are future directed; the second route proposes time-asymmetric causal relations as explanation of asymmetries in prevailing initial or boundary conditions.

References

Albert, David Z. 2000. *Time and Chance*. Cambridge, Mass.: Harvard University Press.

Bohr, Niels. 1948. On the Notions of Causality and Complementarity. *Dialectica* 2: 312-319.

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.

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. 2000. “(Dis-)Solving the Puzzle of the Arrow of Radiation,” *British Journal for the Philosophy of Science*, 51, 381-410.

———. 2002. “Non-Locality in Classical Electrodynamics,” *British Journal for the Philosophy of Science*, March, 1-19

———. 2005. *Inconsistency, Asymmetry and Non-locality: A Philosophical Investigation of Classical Electrodynamics*. New York: Oxford University Press.

———. 2006a. Causal Asymmetry, Counterfactual Decisions and Entropy. *Philosophy of Science* 72 (5):739-750.

———. 2006b. A tale of two arrows. *Studies in the History and Philosophy of Modern Physics* 37 (3):542-558.

———. 2007. Causation, Counterfactuals and the Past-Hypothesis. In *Russell's Republic: The Place of Causation in the Constitution of Reality*, edited by H. Price and R. Corry. Oxford: Oxford University Press.

———. manuscript. "The most Sacred Tenet"? Causal Reasoning in Physics.

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

Hawking, S. W. and R. Penrose. 1970. The Singularities of Gravitational Collapse and Cosmology. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, Vol. 314, No. 1519: 529-548.

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.

Lange, Marc. 2002. *An Introduction to the Philosophy of Physics: Locality, Fields, Energy, and Mass*. Oxford: Blackwell Publishers.

Loewer, Barry. 2007. Counterfactuals and the Second Law. In *Causality, Physics, and the Constitution of Reality: Russell's Republic Revisited*, edited by H. Price and R. Corry. Oxford: Oxford University Press.

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

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.

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.

Rohrlich, Fritz. 1990. *Classical Charged Particle*s. 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.

Schrödinger, Erwin. 1951. *Science and Humanism*. Cambridge: Cambridge University Press.

Steiner, Mark. 1986. Events and Causality. *The Journal of Philosophy* 83 (5): 259-264.

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

Weyl, Hermann. 1989. *The Open World*. Woodbridge: Ox Bow Press. Original edition, 1932.

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

———. 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.