

Causal Models and the Asymmetry of State Preparation

Mathias Frisch

University of Maryland, College Park

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 mature physical theorizing. This view has proponents even among those who believe that causal notions have an important place in our folk conception of the world. In this paper I critically examine Bas van Fraassen's articulation of this view in a debate with Nancy Cartwright. I argue that there is an asymmetry of state preparation characterizing our experimental interactions with physical systems that can provide a justification for asymmetric causal assumptions even within the context of a physics with time-reversal invariant dynamical equations.

1. Introduction¹

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 will first critically examine considerations in favor of this causal skepticism and then discuss a strategy for defending a role for causal notions in physics, focusing on the asymmetry of the causal relation.

Many recent arguments questioning the legitimacy of causal notions in physics are descendents of Bertrand Russell's famous attack on the notion of cause (Russell 1918). Russell's paper has received a fair amount of attention in the recent literature (see, e.g., Price and Corry 2007), but there is another precursor to the contemporary debate that prefigures many of today's arguments—an exchange between Bas van Fraassen and Nancy Cartwright—and I want to focus on that exchange here.

In the next section I will outline what I take to be van Fraassen's main arguments for the claim that the distinction between causes and non-causes can only be drawn extra-scientifically. In section three I will sketch a general framework for incorporating asymmetric causal relations into a theory's models. In section four I will discuss one pervasive asymmetry—the asymmetry of state preparation—that

¹ This material is based upon work supported by the National Science Foundation under Grant No. 0646677.

suggests that asymmetric causal relations can play a legitimate role even in physics. I will end with a brief conclusion.

2. Van Fraassen's challenge

In a critical review of (van Fraassen 1989) Cartwright asks why “van Fraassen does not want to allow causality anywhere inside [the] models,” which, according to his version of the semantic view, comprise the content of a physical theory. (Cartwright 1993, 424) Van Fraassen responds to this question as follows:

To me the question is moot. The reason is that, as far as I can see, the models which scientists offer us contain no structure which we can describe as putatively representing causings, or as distinguishing between causings and similar events which are not causings. Cartwright says that if models contain [parts representing] ordinary objects around us (such as cats, and cats lapping milk) then they contain [parts which represent] causes. The question will still be moot if the causes/non-causes distinction is not recoverable from the model. 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. If Cartwright herself draws, extra-scientifically, a distinction between causes and non-causes, she can describe models furnished by science in terms of that distinction. But it may be a ‘hidden variable’ description. She may be thinking of the structures scientists use to model data as themselves parts of larger, more articulated structures that carry the distinctions she makes. (van Fraassen 1993, 437-8)

Van Fraassen's answer to Cartwright's question appears to be not only that the models of physical theories as a matter of fact do not draw a distinction between causes and non-causes but, what is more, that it is impossible to interpret these models causally. A theory's models, he says, contain no structures which we “*can* describe as representing causings” (my emphasis). But what is it about the structure of our theories or models that might lead us to such a causally austere view?

Van Fraassen's remarks echo a claim by Russell, who maintained in his discussion of Newtonian cosmology as paradigmatic physical theory, 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.*” (Russell 1918, 141, my emphasis) While Russell took a theory to be identified with a set of formulas, van Fraassen argues that a theory consists of a set of state-space models. But even though they disagree on whether theories ought to be understood syntactically or semantically, Russell and van Fraassen agree that there is no place for causal notions in physical theorizing. One might think that their argument against causes in physics is simply this:

1. The content of a physical theory is exhausted by a set of state-space models or a set of formulas.
2. Causal relations are not part of the formulas or models of a theory.

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

As it stands, however, premise (1) 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. But once we acknowledge that the radically 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 why that framework does not allow us to “describe [certain structures] as putatively representing causings.”

This point is also stressed by Cartwright, who says that van Fraassen’s state space models

are models of the equations, not models of the physical systems the equations are supposed to treat. When science constructs a picture of bit of the world, the image is far richer. [...] The scientific image of nature is no more devoid of cause and causings than is our everyday experience. The appearance to the contrary arises from looking only at science's abstract statements of law, and not how those are used to describe the world. (Cartwright 1993, 426)

The question, thus, is how rich an image of the world a theory can present and how rich the interpretive frameworks are within which our theories’ state space models are embedded. While Cartwright maintains that causal descriptions can be part of a theory’s interpretive framework, and, hence, of the scientific image of the world, van Fraassen thinks that causal notions are an extra-scientific addition to that image.

One extremely quick argument for the claim that causal notions constitute an extra-scientific addition to the interpretation of scientific theories is suggested by Russell’s famous remark that the word ‘cause’ is not used in the advanced sciences. Thus, one might try to argue that causal notions cannot be part of our theories’ interpretive framework, since scientists do not use causal discourse in describing the world. However, the premise of this argument can easily be shown to be false. As has been pointed out repeatedly—for example by Patrick Suppes (Suppes 1970) and more recently by Chris Hitchcock (Hitchcock 2007)—the words ‘cause’, ‘causal’, and related words are still widely used in contemporary physics. But what is the status of causal discourse in which physicists engage? On the one hand, one might think that when physicists use causal language, they are merely offering an informal commentary on the science using everyday language. This seems to be van Fraassen’s view, who agrees with Russell that mature *physics* only provides us with non-causal models embodying functional dependencies but nevertheless acknowledges that *physicists* often describe the world in causal terms. On the other hand, some of the examples of causal discourse in physics certainly seem to suggest a more central role to causal notions than that of an informal gloss. For example, a widely used textbook on classical

electrodynamics singles out a principle of causality as “the most sacred tenet in all of physics” (Griffiths 1989, 399).

Van Fraassen invokes two different kind of considerations in support of this view, appealing to the asymmetry of the causal relations and to its modal character, respectively. Cartwright’s disagreement with van Fraassen focuses largely on the latter issue. According to van Fraassen, scientific theories present us with a Humean picture of the world, free from modal properties. Allowing causings into a theory’s models, he maintains, takes us outside of the realm of science proper and into that of “woolly metaphysics,” as Jim Woodward has put it in characterizing van Fraassen’s worry (in Woodward 2003). By contrast, Cartwright argues that a demodalized Humean picture of the world is incoherent. Many events, such as milk lappings and photon scatterings, are intrinsically causal, according to her, and therefore the idea of a world exactly like ours but stripped of causings is “ridiculous.” (Cartwright 1993, 427) Thus, far from it being the case that the models of a theory “contain no structure which we can describe as putatively representing causings”, it is not clear that, in a world causal through and through, we can coherently think of the models as not representing causings.

The other considerations to which van Fraassen appeals concern the asymmetry of the causal relation. One reason for why the distinction between causes and non-causes is not recoverable from our theories’ models, he suggests, is that these models do not allow us to draw an asymmetric distinction between cause and effect. The causal relation is asymmetric—there is a difference between shoving and being shoved—but the group-theoretic models of our theories do not reflect this asymmetry. Since the causal asymmetry is generally taken to line up with the temporal asymmetry in the sense that effects do not precede their causes—at least in the kind of circumstances with which we are familiar—this point is often expressed by appealing to the time-reversal invariance of the laws of (most of) our mature physical theories. As Russell (1913) put it, “the laws make no difference between past and future.” That is, the problem for causal notions is not that we cannot, as van Fraassen says, describe structures as representing causings but rather that our models do not allow us to represent causes *as* causes, since the causal relation is asymmetric and our models do not allow us to represent this asymmetry.

The underlying argument is most directly made in terms of our theories’ fundamental equations:

4. The fundamental equations of all mature physical theories are time-reversal invariant.
5. There is no place for an asymmetric notion of cause within a physical theory with time-reversal invariant laws.
6. Therefore, there is no place for an asymmetric notion of cause in mature physical theories.

Of course, it does not follow from the fact that a theory’s dynamical equations are time-reversal invariant that all of the theory’s state space models will be time-symmetric. In fact ‘most’ models, in some

intuitive sense, will not be time symmetric, reflecting asymmetries between the initial and final conditions characterizing a given model. But the entire class of a theory's models will be time-symmetric in the sense that for each time-asymmetric model M there will be a time-reversed model M^* that also satisfies the theory's dynamical equations. From this symmetry van Fraassen concludes that the distinction between causes and non-causes must be a distinction made outside the theory.

Despite his belief that causal notions play no role in how physics represents the world, van Fraassen agrees with Cartwright that our overall conception of the world ineliminably involves causal notions and that causal discourse is irreducible to non-causal notions. The way to avoid a causal metaphysics and nevertheless acknowledge an important place for an asymmetric notion of cause in our conception of the world, according to him, is to locate causal discourse within psychological and intentional discourse. The ultimate source of causal notions, van Fraassen suggests, is our conception of ourselves as agents and the use of causal notions in physics is a metaphorical and analogical extension of language that has its basic meaning only in the context of folk psychology.

Van Fraassen does not make this point explicitly, but his discussion may give the impression that the two types of consideration are closely linked and that the only way in which we can grant an important role to an irreducible asymmetric notion of cause in our conception of the world without embracing a rich causal metaphysics is to locate the source of such notions outside of the image of the world presented to us by physics. One of my aims in this paper is to challenge this impression. I think one can afford a scientifically legitimate role to asymmetric causal notions in physics without embracing a rich causal metaphysics. In the next section I want to suggest how causal notions can be incorporated into a theory's class of models in a way that is metaphysically neutral or non-committal.

3. Causal models

Let us assume van Fraassen's framework of theories as represented by a class of state space models. We can then think of interpreting a theory causally as the equivalent of embedding the theory's state space models into larger model-theoretic structures by introducing asymmetric relations between state space variables. The state of a system $S(t)$ is given by the values of a set of variables $s_1(t), s_2(t), \dots, s_n(t), \dots$, 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 then define an asymmetric, transitive, and non-circular relation $C = \langle S(t_i), S(t_j) \rangle$ over the set of states S , which defines a partial ordering over the set of states in a model. C is 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 I want to call *potential causal models* of a

theory. Depending on the theory in question, we can also introduce more fine-grained causal relations $\langle s(t_i), s(t_j) \rangle$ defined over individual state variables s_i .

One might object that physicists do not explicitly represent their theories in terms of causal structures of this kind. But the reason for this might be that the causal structures at issue will usually be quite ‘boring’ and hence not in need of an explicit representation: often their content is exhausted by the claim that the state of a system is a cause of future states of the system.² By making the formal framework explicit, however, we see that asymmetric causal notions need not constitute a vague or imprecise addition to a physical theory, as Patrick Suppes (1970) and Chris Hitchcock (2007) have suggested, but can be introduced into a theory in a mathematically well-defined manner. Moreover, embedding a theory’s state space models into richer causal structures does not in itself carry any weighty metaphysical commitments with it. As van Fraassen himself has argued, accepting a theory that embeds the theory’s observational substructures into a richer set of models does not yet settle the question as to what our metaphysical commitments to the entities posited by our models ought to be. We might be realists, who believe in the unobservable substructures postulated by our theories, or we might be constructive empiricists, who do not think that our grounds for accepting a theory are not also good reasons for believing in its unobservable substructures. Similarly, taking the theory’s state space models to be embedded into richer, causal structures in itself is metaphysically non-committal. In particular, adopting causal models carries with it no metaphysical commitment to the existence of anything like a ‘causal glue’ between causally related events. Thus, there appears to be room for a stance on the issue of causation that “involves no metaphysics” (van Fraassen, 1993, 439) yet does not follow van Fraassen in taking causal discourse to be ineliminably psychological.

Causal structures do, however, introduce asymmetric relations that are not part of the non-causal state space models, which van Fraassen identifies with the content of a physical theory. Thus, van Fraassen worries that adding such structures amounts to providing a “‘hidden variable’ description” and that there can be no empirical or (more generally) scientifically legitimate reason for accepting the richer causal structures instead of non-causal state space models. Yet it seems to me that a closer look at theorizing in physics reveals a far more prominent role for causal notions than van Fraassen wants to allow. First, there are theoretical contexts in which physicists appeal to time-asymmetric causal constraints to restrict the range of physically possible models. In terms of the terminology introduced above, causal constraints are invoked to restrict the class of *potential causal models* to a proper subclass of models that are *causally possible*. For example classical dispersion relations are derived from an explicitly time-asymmetric causal constraint (see Frisch forthcoming-a; forthcoming-b).

² There is, however, at least one field in fundamental physics in which causal relations of this kind are introduced explicitly—the causal set approach to quantum gravity (see, e.g. Rideout and Sorkin 2000).

Second, there are phenomena that exhibit asymmetries in prevailing initial or final conditions. If actual systems in the domain of a time-reversal invariant theory are best represented by models most of which exhibit the same kind of temporal asymmetry—that is, if there is an asymmetry between the initial and final conditions characterizing models of typical actual systems—then this might be evidence for causal relations among the physical quantities involved. A paradigmatic example of this is the temporal asymmetry characteristic of waves in the presence of wave sources, as I have argued elsewhere.³

Both these cases involve an asymmetry in the class of state-space models representing actually occurring phenomena. But, as I will argue in the next section, there can even be scientifically legitimate evidence for time-asymmetric causal structures in cases where the class of state-space models representing the phenomena is time-reversal invariant.

4. The asymmetry of state preparation

In this section I will argue that our experimental interactions with physical systems exhibit a temporal asymmetry even in the case of systems that are best modeled with the help of theories with time-reversal invariant laws and that this asymmetry is best thought of as a causal asymmetry. In explicit premise-conclusion form, my argument is this:

1. There is a temporal asymmetry characterizing experimental interventions into otherwise closed systems.
2. If there is such an asymmetry, it is best explained by appealing to a causal asymmetry.
3. If a concept plays a role in the best explanation of a phenomenon, the concept plays a legitimate role in science.
4. Therefore, asymmetric causal notions play a legitimate role in science.

Why should we accept premise (1)? It is a striking fact about experimental interactions that we can only intervene into a system ‘from the past,’ as it were. Consider a system S that is governed by both past and future deterministic laws. That is, let us assume that the final state $S_f(t_f)$ of the system is uniquely determined by the initial state $S_i(t_i)$, where $t_i < t_f$, together with the dynamical laws and the boundary conditions; and that the initial state $S_i(t_i)$ is similarly determined by the final state $S_f(t_f)$. Thus, if S is closed between t_i and t_f , then the initial and final states are both dependent on each other. Nevertheless there is an asymmetry of state preparation in the following sense. We can prepare the system in its initial state $S_i(t_i)$ without making use of any knowledge we might have of the system’s dynamical evolution between t_i and t_f ; and we can subsequently calculate the system’s future evolution for times $t > t_i$ from the initial state, the dynamical laws, and the boundary conditions. But we could not similarly first prepare the system’s final state at t_f without using our knowledge of the dynamics and then take the final state

³ See (Frisch 2005), and also (Frisch 2000; 2006; 2008).

together with the laws to calculate the system's past evolution for $t < t_f$. (Of course we cannot first prepare the system in S_f and then let it *evolve* into S_i . That is not what the asymmetry consists in. Rather the asymmetry consists in the fact that we cannot first prepare the system in S_f without *making use* of facts about the dynamical evolution and *then calculate* what the system's past evolution from S_i to S_f must have been, given the dynamical laws and the boundary conditions.)

There are two ways in which we can prepare the system in a specific final state at t_f . First, we can make use of our knowledge of the dynamical laws to determine the initial state in which the system has to start out at t_i in order to evolve into the final state in question and prepare the system in the appropriate initial state. In that case our ability to prepare the system in its final state relies crucially on our knowledge of the system's evolution between t_i and t_f . If the system is closed between t_i and t_f , we need to know which state S_i will evolve into the state S_f we are trying to set up, because we can only 'prepare' the system in the state S_f by preparing it in S_i . By contrast, we can prepare the system in an initial state S_i without any knowledge of the dynamical evolution of the system between t_i and t_f , even though the dynamical laws and boundary condition determine what that evolution is. While, given the dynamical laws and boundary conditions, the initial and final states determine each other, we do not need to make use of that fact if our goal is to set up the system in some specified *initial* state. But we do need to make explicit use of the dynamical evolution between t_i and t_f if our goal is to set up the system in a specified *final* state.

We can imagine, for example, that one experimenter is responsible for preparing a system S in an initial state S_i or a final state S_f and that a different experimenter is responsible for setting up the boundary conditions. If the first experimenter wants to prepare the system in a certain initial state, he can do that without knowing what boundary conditions the second experimenter chooses to set up. But if the first experimenter wants to make sure that the system ends up in a certain final state, he needs to know what the boundary conditions will be in order to make sure he prepares the system in the appropriate initial state.

A second way of preparing the system in a final state S_f is to prepare the system in that state directly by intervening into the system between t_i and t_f . In this case we do not need to make use of our knowledge of the evolution between initial and final times. But then the system will not be closed between t_i and t_f and we cannot use the dynamical laws and boundary conditions governing the closed system to retrodict the initial at t_i . Thus, there is a way for first experimenter to prepare the system in S_f without knowledge of the boundary conditions which the second experimenter tries to set up, but at the cost of having to violate these conditions through his intervention and, thus, by losing any ability to retrodict the evolution of the system with the help of the dynamical equations.

That is, even systems that are governed by both past and future deterministic dynamical equations exhibit an asymmetry of state preparation. In the case of a system that is closed between t_i and t_f and initial and final states mutually determine each other, given the dynamical equations and boundary conditions, we can only prepare a system in a given final state S_f by making explicit use of the dynamics and prepare the system in the corresponding initial state S_i . If instead we directly intervene on the state of the system at t_f , we can no longer use the state at t_f to retrodict the state at t_i (while we can, of course, directly intervene on the state of the system at the earlier t_i and then predict the system's evolution until some later time t_f). We can only intervene into a system from its past.

My reason for claiming that the asymmetry of state preparation is best understood as a causal asymmetry is that the asymmetry is a paradigm case of the asymmetry characterizing interventions, as understood by interventionist accounts of causation (see, e.g., Woodward 2003). In particular, if S_f is an effect of S_i , then according to an interventionist account of causation there are two ways by which one can intervene on the system to set S_f to a particular value: first, we can intervene on S_i , which in turn will affect the value of S_f ; or, second, we can intervene directly on S_f , which 'breaks the causal arrow' from S_i to S_f and, therefore, makes it impossible to retrodict the value of S_i on the basis of the value of S_f . Thus, interventionist accounts of causation predict that experimental systems will exhibit an asymmetry of state preparation, if earlier states of the system are causes of later states. Positing an asymmetric causal relation between the states of a system at different times clearly provides *an* explanation of the asymmetry of state preparation. Moreover, it seems to me that there is no other fully worked out and equally as successful non-causal alternative explanation of the asymmetry. Thus, I take it that this asymmetry provides us with empirically justified reasons for embedding non-causal state space models of a system into richer causal models.

Now, van Fraassen recognizes that causal notions such as manipulation and control are an integral part of science but maintains that these notions play a role only in applied science and that their use there can be accounted for entirely as an analogical and metaphorical extension of psychological discourse. Thus, van Fraassen might argue that the asymmetry of state preparation does not point to a genuinely scientific asymmetry, since it reveals itself only in experimental interactions with physical systems and in virtue of being an asymmetry of manipulation and control falls under the domain of our folk-scientific conception of ourselves as agents.

Yet there clearly is no direct argument from the claim that our use of causal discourse in science arises from our experimental interactions with physical systems to the conclusion that causal distinctions are drawn extra-scientifically. There are many scientifically legitimate accounts of phenomena for which our only evidence comes from experimental interactions. Indeed, the asymmetry of state preparation seems to show that the causal asymmetry is a more robustly objective feature of such interactions than

van Fraassen's account of causal talk as metaphorical or analogical extension of everyday life allows. If describing experimental interactions with physical systems in causal terms is *only* metaphorical or analogical, then a causal account cannot provide a scientific explanation of the asymmetry of state preparation. As a useful contrast case consider the following. We might metaphorically describe a ball rolling up an inclined plane as 'struggling to reach the top of the plane.' In this case there are no facts about the (macroscopic) physical situation that are not fully accounted for in terms of the initial and boundary conditions and the dynamical equations and the metaphorical description truly cannot serve any genuine scientifically explanatory purpose. By contrast, in the case of the asymmetry of state preparation we do not merely choose to describe in asymmetric causal terms facts about physical systems that also can exhaustively be captured in a non-causal description of that system. Rather, the asymmetry characteristic of our interactions with such systems goes beyond what can be captured in an acausal and time-symmetric description of the system.

In general, two kinds of account of the asymmetry of state preparation seem possible. Either one can appeal to asymmetric causal relations between the states of a system at different times, along the lines I have suggested here. According to this account, the asymmetry is due to an intrinsic asymmetry characterizing the system. Or one might try to argue that the asymmetry is due to what are ultimately non-causal physical features of the kinds environment into which the systems with which we interact are embedded. This second kind of explanation either might appeal to general asymmetries of the physical environment, such as thermodynamic features, or might more narrowly focus on asymmetries characterizing physical agents intervening in experimental systems. Variants of the latter account have been defended by Richard Healey (1983) and Huw Price (2007). Now, I do not know of a principled argument that can establish that no non-causal account of the asymmetry of state preparation can succeed. But I also do not know of any existing account that offers more details than might be contained in a promissory note. The task for any such account is to show that the asymmetry of state preparation varies with hypothetical changes to the environment. Thus, Healey's and Price's arguments rely crucially on intuitions concerning what the direction of causation would be in hypothetical anti-thermodynamic environments, but it is not clear why we should share their intuitions. That is, to my mind neither has shown convincingly that in a different thermodynamic environment state-preparation would be from the future, as it were. But in the absence of a fully developed convincing alternative, the causal account remains the best explanation of the asymmetry of state preparation.

4. Conclusion

I argued that a theory's non-causal state space models can be embedded into richer structures containing asymmetric causal relations in a way that does not carry with it a commitment to a particular causal

metaphysics. The use of causal models does not merely constitute the introduction of a ‘hidden variable description’ but can be empirically justified in various ways. In this paper I presented one such justification that appeals to a pervasive asymmetry characterizing all our interactions with experimental systems: the asymmetry of state preparation.

References

- 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. 1989. *Laws and Symmetry*. 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**, pp. 381-410.
- . 2005. Inconsistency, Asymmetry and Non-locality: A Philosophical Investigation of Classical Electrodynamics. New York: Oxford University Press.
- . 2006. A tale of two arrows. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics* 37, no. 3 (September): 542-558.
- . 2008. Philosophical Issues in Classical Electrodynamics. *Philosophy Compass*. Blackwell Publishing.
- . forthcoming-a. 'The most sacred tenet?' Causal Reasoning in Physics. *British Journal for the Philosophy of Science*.
- . forthcoming-b. Causality and dispersion: a reply to John Norton. *British Journal for the Philosophy of Science*.
- 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: 79-104.
- 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.
- Newton-Smith, W. H. 1983. Temporal and Causal Asymmetry. In *Space, Time, and Causality*, edited by Richard Swinburne. Dordrecht: D. Reidel Publishing Co: 105-124.
- 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.
- Rideout, D. P. and R. D. Sorkin. 2000. A Classical Sequential Growth Dynamic for Causal Sets. *Physical Review D* **61**.
- Russell, Bertrand. 1918. On the Notion of Cause. In *Mysticism and Logic and other Essays*. New York: Longmans, Green and Co.
- 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.

———. 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.