November 3, 11, 14, 2021. February 12, 2022.
File opened, September 1, 2021

The Metaphysics of Causation: An Empiricist Critique

John D. Norton[1]
Department of History and Philosophy of Science
University of Pittsburgh
http://sites.pitt.edu/~jdnorton/jdnorton.html

Contrary to Hume, science has found many ways in which things connect with other things in the world. Causal metaphysics, however, has failed to add anything factual to the relations discovered by science. It is at best an exercise in labeling that may have practical uses.

1. Overview

This paper presents an empirically based, skeptical critique of the metaphysics of causation. The mere mention of empiricism is apt to lead to a misreading of the approach. It is not a Humean critique. Hume's celebrated critique of causation was a mixture of inductive skepticism and causal skepticism. He doubted that we could know of causal connections because he doubted the inductive inferences that offered that knowledge. I do not share Hume's inductive skepticism. It was based on an impoverished and unsustainable conception of inductive inference. In its place, I am an inductive optimist and have mounted a far-reaching defense of inductive inference that includes what I believe is a quite serviceable dissolution of Hume's skeptical attack on inductive inference. (See Norton, 2014; manuscript, Ch. 6).

My inductive optimism extends to the connectedness of things in science. I have no dispute with the idea that we have learned inductively of many ways that things connect in the world. We have good reason to accept that electric currents are set in motion by electric fields; that chemical reactions are driven forward by gradients of free energy; that natural selection

[1] I thank Jim Woodward for helpful comments, but completely absolve him of any responsibility for or agreement with the content of this paper. I also thank Brian McLoone and Yafeng Shan for helpful comments on a first draft.

leads to a diversity of species adapted to their ecological niches; and so on and on over the full range of our sciences.

Call these last connections "causal" if you like. In my view, attaching the label "causal" has only pragmatic uses such as will be developed below. It may be little different from the practice of florists when they declare some plants flowers and others weeds. There is no deeper basis for the division in botany. However, it is a definite fiscal benefit to the florists to offer their customers what the customers will identify as desirable flowers. Or it may be akin to the way chemists divide substances into organic or inorganic. There is no prior necessity that the world must provide us substances of either type. However, once the division is found empirically, it is useful since chemists treat each type of substance by different methods.

The danger with causal labeling is that it invites confusion. Calling a process[2] "causal" is too often understood as identifying it as a manifestation of a factual, causal order in the world that has been identified by prior metaphysical analysis. The present study of the *metaphysics* of causation, I shall argue here, has nothing useful to add factually to the catalog of the connections discovered by science. Rather the metaphysical investigations of causation are confused in their foundation and purpose. They proceed under the illusion that diligent reflections on the meaning of terms and drawn-out investigations of how we ordinary folk use them can somehow provide a deeper illumination of the causal nature of the world. These efforts are, in my view, a failed attempt to pursue a priori science.

The principal means for developing the argument of this paper is in a dilemma, posed in Section 2: either conforming a science to cause and effect places a restriction on the factual content of a science; or it does not. The first horn is explored in Section 3. It shows that the history of efforts in the metaphysics of causation to place factual restrictions on processes in the world is one of sustained failures. Earlier, failed maxims of causation were replaced in the nineteenth century by the single idea that causation is just determinism. When that dogma failed with the discovery of the indeterminism of quantum processes, it was replaced by a new dogma: that causation is probabilistic. This new dogma is now widely accepted, but it is ill-fated. It rests

---

[2] Here and below, the term "process" is used informally merely to designate a part of the world of interest as it undergoes changes.

on the false assumption that all indefinitenesses in the world are probabilistic. In sum, this first horn is rejected; and the second horn accepted.

The second horn of the dilemma is developed in Sections 4 and 5. If causal analysis cannot provide novel factual restrictions on processes, then it is reduced to an exercise in labeling. We are free to choose how we might like to assign these labels. Done poorly, the assignment is unsystematic. If it attempts to mimic ordinary causal talk, the result will be inconsistent, since the ordinary use of causal language is inconsistent.[3] Done well, the exercise can have definite utility. Once we conceive a process in causal terms, we connect it with others we find analogous. The result is commonly an enhanced comprehension of the processes. The choice of causal labeling can be especially helpful if it tracks useful factual distinctions in processes. In Section 6, my version of the interventionist account of causation is identified as such a case. Its causal labels distinguish those processes that admit interventions. We thereby learn of processes that allow us to manipulate the world. This is useful knowledge. In Section 7, I show how this reduced conception of interventionist causation protects it from standard objections: that the account does not apply where interventions are impossible; and that it is circular since it takes intervention itself to be a causal notion.

Modern writing in the metaphysics of causation has become noticeably reluctant to propose causal maxims with factual content. An exception is recent work that seeks to impugn the notion of downward causation. This type causation is defined as arising between different levels of description of the world. Common examples are proposals of causal processes from the mental to the physical in theories of mind; and from the thermodynamic to the statistical mechanical levels in thermal physics. Section 8 introduces the notion of downward causation and agrees with Woodward that downward causation is supported by an interventionist account of causation. Section 9 reviews the causal maxims that have been proposed as refuting the notion of downward causation. The maxims include the ideas of causal exclusion and closure; the prohibition of whole to part causation and synchronic causation; and the requirement that causal relations must be asymmetric and proceed forwards in time. In a continuation of the project of

[3] Causation by omission occurs when the absence of a factor means that some process is not blocked. For some, the absence causes the process. For others, it is merely blamed.

examining the first horn of the dilemma above, the section shows how each maxim fails and thus fails to impugn an interventionist account of downward causation.

2. The Dilemma

What sort of activity is the pursuit of the metaphysics of causation? More narrowly, what sort of knowledge does it aspire to provide us? The title[4] of Mackie's influential treatise, *The Cement of the Universe: A Study of Causation*, indicates a goal of illuminating at its foundations the connectedness of things in the world. If this illumination is factual, the simplest vehicle for communicating its factual content is through the assertion of a principle of causality. This rather distinctively nineteenth century notion[5] found multiple, largely similar expressions. Here is one example from Flemings' "*Vocabulary"* (1860, p. 78, Fleming's emphasis):

> The belief that every exchange implies a cause, or that every change is produced by the operation of some power, is regarded by some as a primitive belief, and has been denominated by the phrase, the *principle of causality*.

That the factual illumination should come in the form of a principle need not be assumed. *The Cambridge Dictionary of Philosophy* (Audi, 1999) has many entries on causation but, as far as I can see, the terms "principle of causality" or "principle of causation" appear nowhere in it. We cannot even assume that factual illumination is sought. Twentieth and twenty-first century writing on causation is, at least in my reading, quite evasive on the question.

This obscurity of purpose cannot be left unchallenged. Whether factual illumination is provided by the present literature is something that should be decided. To this end, in my (2003) "Causation as Folk Science," I tried to capture in the broadest form what would underlie efforts to provide factual illumination of things in the world. I proposed that such efforts presume (p.3):

---

[4] We learn from the front matter of Mackie's volume that the expression "Cement of the Universe" is derived from an eighteenth century abstract of Hume's *Treatise*.

[5] A search on Google's ngram viewer (https://books.google.com/ngrams/) indicates that the terms "principle of causation" or "principle of causality" seem to have burst with force into the literature in the nineteenth century and enjoyed a strong presence through the century. That presence waned in the twentieth century.

> *Causal fundamentalism:* Nature is governed by cause and effect; and the burden of individual sciences is to find the particular expressions of the general notion in the realm of their specialized subject matter.

A common element in this strain of thinking is the idea that the issues addressed in causal metaphysics come prior to the mundane empirical work of the sciences. That this is the wrong way to proceed was pointed out forcefully by Ronald Fisher, after the coming of quantum theory shook the nineteenth century confidence in determinism. In the first issue of the new journal *Philosophy of Science*, Fisher described what he saw as the failed approach taken to causation in prior centuries (1934, p. 100):[6]

> In the last few centuries a certain uneasy compromise may be observed, between inventing the world *a priori*, and looking to see what it is like. The framework of our thoughts, our preliminary concepts, or basic ideas, have been supposed to be given by philosophy; it was the business of observational science to fit into this framework its missing details, in all confidence that whatever could really be found, had already its place prepared for it in our conceptual framework of ideas about what the world is like.

One might expect this ill-fated approach to have faded away with the new century and its unexpected scientific discoveries. It did not. For example, Hugh Mellor in his treatise, *The Facts of Causation*, wrote (1995, p. 5)

> So while I will accommodate the relevant results of modern physics I will not for example leave it to quantum physics to tell me whether causation can act immediately across spacelike intervals. On the contrary: only when our metaphysics has told us what causation is can we see if physics could reveal unmediated action at a distance (it couldn't).

This insistence sharpens the problem. Metaphysical analysis is providing us something prior to the empirical work of a science—in this case quantum physics. What is the analysis providing? Is it providing us universal facts of causation? Or something else? And, from an empiricist perspective, how can it inform us about quantum processes antecedent to the empirical

[6] I thank Brian McLoone for drawing my attention to Fisher's paper.

investigation of quantum processes. To bring the problem into clearer focus, in my 2003 paper, I posed a dilemma for causal fundamentalism (pp.3-4):

> Causal fundamentalist's dilemma: EITHER conforming a science to cause and effect places a restriction on the factual content of a science; OR it does not. In either case, we face problems that defeat the notion of cause as fundamental to science. In the first horn, we must find some restriction on factual content that can be properly applied to all sciences; but no appropriate restriction is forthcoming. In the second horn, since the imposition of the causal framework makes no difference to the factual content of the sciences, it is revealed as an empty honorific.

This early formulation of the dilemma, I now see, oversimplified the second horn. In it, causal analysis is reduced to labeling and does not restrict a science to conform with a prior factual principle of causation. However, unlike the suggestion of the 2003 wording, the labeling need not be empty. It may provide pragmatic benefits to our understanding[7] and practical benefits if it tracks practically useful factual distinctions already in present in processes, prior to the labeling.

In the following two sections, I review the prospects for each horn and explain why efforts at accepting the first horn have failed.

3. First Horn: The Failed Quest for Factual Content

The prospects of meeting the requirements of the first horn are poor. To be clear, it is insufficient in this horn for causal metaphysics merely to recite some fact of a connection in the world. It is the case that electric fields produce electric currents if free charges are present. This result and a myriad like it are just the routine results of science. Causal metaphysics pretends to knowledge of causes antecedent to the specifics of such results. It supposes a grasp of causation factually at a quite general level. What is this general grasp? What in *addition to the individual results of the sciences* has causal metaphysics provided factually?

The short answer is: nothing. That is, the centuries-old history of efforts by causal metaphysicians to deliver factual claims of foundational import has been one of sustained failure.

---

[7] The term "understanding" is used informally here and below, in a way that is antecedent to the foundational accounts of the term in the recent philosophy of science literature.

It is a story of proposal and refutation, of confident assertion and apologetic retraction. Here is a sampling of these episodes.

3.1 Early Failures

According to Aristotle we explain the fall of heavy bodies and the rise of light bodies in part through their final causes: they move in order to occupy their natural places. A standard, early example of the repudiation of a causal notion was the seventeenth century elimination of final causes from accounts of the mechanics of bodies. Those accounts were to be based on efficient causes only: the motions of the bodies, their collisions and, after Newton, the forces acting. Part 1, Article 28 of René Descartes' *Principles of Philosophy* (1982, p. 14) was headed with the proposition:

> That we must not examine the final causes of created things, but rather their efficient causes.

As he makes clear in the text of Article 28, Descartes' intent was not to eliminate final causes completely but to remove them to the more elevated sphere of God's inscrutable purposes.[8] This tolerance did not last. Laplace's 1814 *Essai* (1902, p. 3) is quite forthright in denouncing the supposition of final causes for events. They result from "ignorance of the ties which unite such events to the entire system of the universe." Laplace reports triumphantly, "these imaginary causes [including chance] have gradually receded with the widening bounds of knowledge and disappear entirely before sound philosophy."

The central hero of the new era in physics, Isaac Newton, was not immune to failed causal pronouncements. His mechanics provided no mechanism mediating the gravitational attraction of the sun, millions of miles away, upon us here. Nonetheless, he denounced the idea of unmediated action at a distance (Newton, 1761, pp. 25-26):

> [T]hat one body may act upon another at a distance through a vacuum, without the mediation of anything else, by and through which their action and force may be conveyed from one to another, is to me so great an absurdity, that I believe no man,

[8] As Osler (1996) has shown, this relocation of final causes from the simple mechanics of bodies to the larger theological sphere was then a common view. Robert Boyle (1688), who coined the term "mechanical philosophy," held a version of it.

> who has in philosophical matters a competent faculty of thinking, can ever fall into it.

Yet for all Newton's self-assurance, his theory of gravity soon came to be established as the premier example of an action-at-distance theory. As late at 1901, in his authoritative Teubner Encyclopedia article on gravitation, Zenneck (1901) reported the failure of empirical efforts to see any effect of an intervening medium on the propagation of gravitation (§17, pp. 42-43) and to discern a finite speed of propagation of gravitation (§20, pp. 44-46). Against all efforts to show otherwise empirically, the best account of gravitation remained that it is an instantaneous, unmediated action at a distance.

These are just two of many causal dicta widely accepted in these earlier times, but which do not survive scrutiny. I can continue here only with a brief mention of more of them. In his *Treatise*, Hume (1739, p. 78, his emphasis), for example, reported: " 'Tis a general maxim in philosophy that *whatever begins to exist, must have a cause of existence.*" Hume proceeded to attack the maxim from the perspective of his rather severe inductive skepticism. Russell (1912-13, pp. 9-12) in his then inflammatory "On the Notion of Cause," listed five maxims that, he urged, "have played a great part in the history of philosophy." He proceeded to scrutinize and impugn them (Russell's emphasis and quote marks):

> (1) "Cause and effect must more or less resemble each other."
>
> (2) "Cause is analogous to volition, since there must be an intelligible *nexus* between cause and effect."
>
> (3) "The cause *compels* the effect in some sense in which the effect does not compel the cause."
>
> (4) "A cause cannot operate when it has ceased to exist, because what has ceased to exist is nothing."
>
> (5) "A cause cannot operate except where it is."

3.2 The Nineteenth Century Dogma of Determinism

In the course of the nineteenth century, causal theorists freed themselves from most of the complications of these maxims, save one. The notion of causation was reduced to one idea, determinism: the present state fixes a unique future state. The causal character of the world inhered in the fact of its determinism. Laplace's 1814 *Essai* gave an early definitive statement of it (1902, p. 4):

> We ought then to regard the present state of the universe as the effect of its anterior state and as the cause of the one which is to follow.

Laplace made the determination vivid through his fictitious calculator who could inerrantly calculate the unique future from a full specification of the present. Laplace concluded "for it, nothing would be uncertain and the future, as the past, would be present to its eyes." Probabilities, the subject of Laplace's *Essai*, merely reflected our ignorance of the immense complications of the full determination of things.

John Stuart Mill, perhaps the leading author on methodology in the nineteenth century, based his analysis of causation on this same identification of causation with determinism. He wrote in his *System of Logic* (Book III, Chapter V, §2. p. 236.)

> The Law of Causation … is but the familiar truth, that invariability of succession is found by observation to obtain between every fact in nature and some other fact which has preceded it; independently of all considerations respecting the ultimate mode of production of phenomena, and of every other question regarding the nature of "Things in themselves."

This identification of causation with determinism had the immense advantage of giving a clear and simple account of the nature of causation. It soon became clear that this appeal of simplicity and clarity set the view up for disaster. In the early decades of the twentieth century, quantum theory emerged. The theory contradicted determinism. The best its laws could provide, in the general case, was merely probabilistic assurance of future states, given present states.

The quantum theory was rightly recognized as a catastrophic blow to the nineteenth century conception of causation. Friedrich Waismann was a core figure in the logical positivist movement of the early twentieth century. He reflected in dark tones on not just the demise of determinism but the notion, championed by Niels Bohr, that even classical descriptions in space and time fail. In his 1958 lecture, "the Decline and Fall of Causality," he lamented that (1959, Ch. V, p. 208) "…causality has definitely come to an end: atomic science has penetrated to a depth where an entirely new orientation is called for."

3.3 The Twentieth Century Dogma of Probabilism

This obituary for causality proved premature. Even as Waismann was lecturing, a resuscitated causality was on the rise. Causes no longer determine their effects uniquely, the resuscitation asserted. They only affect their probabilities. This view remains today a dominant view of causality. In his synoptic encyclopedia article, Hitchcock (2021, Section 1.1) summarizes the new tradition:

> The central idea behind probabilistic theories of causation is that causes change the probability of their effects; an effect may still occur in the absence of a cause or fail to occur in its presence.

There are too many treatments of this notion of causation to survey tractably. Instead, we can sample the literature in Mellor's 1995 *Facts of Causation*. The driving theme throughout is to reassure readers that the loss of determinism does not mean the loss of causation. He bases his case in Ch 6 and 7 on identifying certain "connotations" of causation: causes are evidence for their effects, causes can explain their effects and causes are means of bringing about their effects. Mellor's core claim is that these connotations do not require that causation is determinism. Rather they *necessitate* a probabilistic conception of causation (p. 67):

> I shall argue in this chapter and the next that causation's connotations require every cause to raise the chances of its effects.

For "chances" here, we are to understand probabilities (p. 21, Mellor's emphasis):

> The first point I must make about chances is that they are *probabilities*. My other claims about chance may be contentious, but not this one: everyone measures chances by numbers which satisfy the standard calculus of probabilities.

Note that the claim is very strong. Causal connotations "…*require every cause…"* No instances of causation escape this theory.

Mellor's conception of the relevant type of probability is correspondingly strong. He insists (p. 75) that "…chances must be more than evidential probabilities… They measure possibilities …that are properties of facts about the world, and not merely of facts about our knowledge of the world."

Here Mellor has chosen one particular way to implement the probabilistic notion of causation using a physical or objective notion of probabilities. Others working in the subjective tradition of de Finetti may prefer to employ only a subjective notion of probability. For my purposes it is immaterial which particular sense of probability is employed in the account of probabilistic causation. They are all versions of a dangerous dogma:

> *Probabilism*: All indefinitenesses, whether of physical determination of events or of epistemic uncertainty, are always properly represented by probabilities.

Probabilism is presumed throughout Mellor's volume. Chances are identified as probabilities. Evidential support is presumed always probabilistic. This presumption is reflected throughout this literature.

3.4 Problems of Probabilism

The problem of probabilism is not that probabilities never apply. There are numerous cases of probabilistic relations that we might comfortably call causal. The problem is the insistence that causal connections are necessarily probabilistic; and that this probabilistic conception of causation is universal and exhaustive. That insistence makes probabilism a dogma that fares no better than the dogmas that have come before it. Just as nineteenth century thinkers had an unquestioned but misplaced confidence in determinism, so twentieth century probabilists have an unquestioned, misplaced confidence in probabilism.

The case against this dogmatic version of probabilism has been made at some length in Chapters 10-16 of Norton (2021). The central idea is that probabilities cannot be presumed by default as the way to represent all indefinitenesses. Rather, the applicability of probabilities to some particular instance of indefiniteness must be established by displaying what in the factual background authorizes the probabilities.

One strand of the case against probabilism reviews the range of arguments for the necessity of probabilities. They are, I find, all circular. For the applicability of probabilities to some system is an empirical matter. Thus, any deductive proof of their universal applicability must covertly assume contingent facts logically as strong as or stronger than the assumption of probabilism itself. Once one knows to look for it, it becomes straightforward to identify the presumption of probabilism hidden in each argument's premises.

Another strand of the case identifies systems in which, demonstrably, probabilities cannot be used to represent indefiniteness. These are cases of indeterminism in some physical theory in which the physical theory itself entails indeterminism only, but provides no measure to weight the different possible futures. To represent the indeterminism probabilistically is to add illicitly physical content to the physics applicable. These sorts of systems have been in the philosophy of science literature since Earman's (1986) path-breaking *Primer on Determinism*. Simple examples are collected in Norton (2021, Ch. 13, 15). Most of these cases arise in artificial contexts.

The simplification afforded by the artificial contexts enables quick proofs of the inapplicability of probabilities. Some, however, come from present science. The enduring "measure problem" in present inflationary cosmology is that the theory can provide no unique probability measure over the different non-inflating universes that can be spawned by a larger inflating cosmos. (See Norton, 2021a.)

Familiar background facts that do warrant probabilities can be supplied by a physical theory, such as versions of quantum theory. In the social sciences, probabilities may be introduced by the assumption that individuals of interest in a large population are selected randomly, where the import of "random" is a selection with equal probability.

If there are no authorizing facts but probabilities are still applied, there is a significant risk that results are introduced spuriously as artifacts of an inappropriate representation. For examples, see Norton (2021, Chapter 10, Section 4) and also Norton (2010). In them, strong conclusions are wrestled implausibly from assumptions that lack factual content sufficiently strong to support them.

In one such case, Van Inwagen has argued that, antecedent to specific knowledge of the world, it is very probable that there is something rather than nothing. There are, he notes, very many ways that things might be, but only one way that things might fail to exist. We do not know which of the ways is the case. We represent that ignorance by distributing probabilities fairly uniformly over all the possibilities. The result is that the overwhelming bulk of probability is assigned to the many ways that things might be. We are to conclude that, with this overwhelmingly high probability, there must be something. The result is clearly spurious. The fallacy resides in assigning probabilities in circumstances of complete ignorance, for those circumstances provide no background facts to authorize the application of probabilities.

This fallacy arises in the literature in many forms. The "doomsday argument" considers a process that we observe has endured for some time t and we are otherwise in total ignorance as to how long T the process persists. The key step in the ensuing argument is to assume that, since we are in complete ignorance as to when the observed time t arose in the overall time T, we say it can arise with equal probability anywhere in the overall process time T. An application of Bayes' theorem then assures us that any smaller value of T is more probable than any larger value of T. There is no background fact warranting the probabilistic representation of this ignorance. The result is merely an artifact of applying probabilistic reasoning illicitly.

If a probabilistic notion of causation is applied to systems indiscriminately without attention to the background facts that authorize the probabilities, there is a significant danger of spurious results. In new work, Wysocki (manuscript) argues that at least some of the cases now routinely dealt with by probabilistic notions of causation can be analyzed without probabilities; and may be better treated that way.

3.5 A Failure of A Priori Science

This concludes a brief review of efforts over the centuries to base a metaphysics of causation in some factual restriction on processes.[9] What we have seen is a sustained record of failure. Time and again, we are offered some pronouncement on the nature of causation as a deep metaphysical truth. New discoveries in science or just closer scrutiny then reveal that the pronouncements are no truths at all, but merely a reflection of the prejudices and dogmas of the moment. That this is the outcome should, on reflection, be no surprise. The idea that metaphysicians can anticipate the form that all future science must take is doomed to fail. It amounts to an attempt to do science a priori. This is precisely the project that Nature seems to enjoy confounding.

We should not confuse this record of failure with the fallibilism of science. The history of science is replete with cases of theories confidently pronounced and subsequently retracted. The difference is that science has within it the means to corrects itself. The major corrections come as empirical investigations advance and reveal more of the world. A priori metaphysics has no corresponding mechanism within it. All it can do it pronounce and then, when the science advances, retract and move on to the next pronouncement.

For example, nineteenth century Newtonians were assured by their physics that the fall of bodies and the motion of planets are deterministic. They presumed this determinism prevailed on all scales, including the very small. That presumption was overturned in the twentieth century

[9] This survey does not exhaust the present candidates for a metaphysical necessity of causation. Curie's principle, for example, requires that any symmetry of a cause must be reflected in the effect. Closer examination in Norton (2016), however, shows that the principle exploits a malleability in how we identify causes and effects in some particular system. By careful identifications, we can render the principle true or false in many cases at our whim.

when new investigations were able to probe empirically the behavior of particles at the minute atomic and subatomic levels. The revised science now allowed that determinism was preserved for the domain in which the empirical evidence for Newton's mechanics was found: ordinary falling bodies, planetary motions and the like. The atomic and subatomic levels would now be governed by a new, indeterministic quantum physics. This is a story of science evolving as it should by responding to the demands of an expanding body of evidence.

As long as the causal metaphysics merely echoes what the latest empirical science asserts, then the transition is comparable to that of the science. However that benign circumstance depends on the emptiness of a priori, causal metaphysics: it has no independent grasp on causation prior to the empirical investigations of science. If we suppose that the nineteenth century metaphysical identification of causation with determinism derives from some independent, prior knowledge of causation in the metaphysics, then the diagnosis must be that this prior grasp was mistaken, yet again.

4. The Second Horn: Causal Metaphysics is Not Factual

The causal metaphysician's dilemma is that we must choose between two horns: EITHER conforming a science to cause and effect places a restriction on the factual content of a science; OR it does not. The discussion of the last section shows the first horn is untenable. As a matter of deductive logic, then, we must accept the second horn. Causal metaphysics adds nothing factual to what we already know by empirical means in our sciences. That is, causal metaphysics has failed to provide novel facts by any means other than through its repetition of what is discovered empirically in scientific investigations.

This result does not deprive causal analysis of its value. Rather it requires us to reconceive its nature. In so far as it is useful, causal metaphysics does not tell us about the world. It tells us about ourselves. There are many ways that this might be so.[10] In my view, the most

[10] Here many find Kantian approaches to causation helpful. In them, causal notions arise of necessity through the way that we interact with the world. These are attempts to secure a priori necessities in another way. However, their prospects are no better than the other a priori dogmas of causation sketched in the last section. We simply are not good at intuiting a priori how things might unfold in the world. For all his prodigious mental powers, Kant's thinking was still

significant is that causal conceptions can have pragmatic value. We have a naïve conception of simple causal processes derived from our common experiences. When we push (the cause), a body moves (the effect). We treat simple cases like this as a template to be imposed on processes more remote from direct experience to enable us to comprehend them better. The template is imposed by redescribing processes in suggestive causal language, while not adding anything factually to the narrative.

Consider, for example, Ohm's law in simple electrical circuit theory. The voltage drop *V* in some conductor equals *IR*, the product of the current *I* and the resistance *R*. That is the full content of the law. We are able to use the law better if we attach causal labels to its parts. The voltage difference *V* is the cause, we might say, that pushes the electric charges through the conductor, producing the effect, *I*, the current. We have added nothing factual to Ohm's law with this narrative. But it does give us a better grasp on the process: the voltage different pushes the charges similarly to the way we might push a body. The addition of such new conceptions to the original science by this causal language is described in Norton (2003) as a "folk science."

More abstruse examples are possible. Electrons in atoms in a radiation field, jump up and down among different energy levels, while absorbing and emitting electromagnetic radiation. Einstein's celebrated "A and B coefficients" paper of 1917 summarized the three component processes. (a) In the presence of a radiation field of the right frequency, an electron can, with some probability, absorb energy from the radiation field and jump to a higher level. (b) When at the higher level, it will with some probability emit radiation and drop back to a lower level. (c) This probability is greater in the presence of a radiation field of the right frequency.

We commonly add causal language to this bare but factually adequate account. In (a), the electron is said to be induced—caused—to jump to a higher energy level by the radiation field. It does so in much the same way as we may induce a guest to a second helping of pie. In (b), the electron is drawn back to its natural, lowest energy state, which we conceive as the applicable final cause. Finally, in (c), the electron is stimulated by the field to jump down to the lower

confined by the prejudices of his era. Contrary to the certainties of his time, the geometry of space does not have to be Euclidean. We have no reason to expect present day Kantian approaches to fare any better, in so far as they arrive at facts by a priori means.

energy state,[11] in much the same way that the cat might be stimulated by a tickle to jump down from the sofa. Once again, adding the causal language "stimulated" has not added factually to the description of the process, but it has enhanced our grasp of it.

Examples like this show how readily we can add causal language to some process without adding factually to it. As far as I can see, there is no limit to this process. We can, if we wish, always find a way to add causal descriptions to any science. Any process can be conformed to some conception of cause and effect, without placing a factual restriction on the process.

5. Causation is Not Factual Discovery, but Convenient Definition

The analysis of the causal fundamentalist's dilemma leads to this result: the identification of causal processes in the world is not one of factual discovery, but of the application of convenient definitions. When we consider some process, there is no fact of the matter as to whether this particular process is causal or not. Whether it is or not is just a matter of how we decide to use the term "causal." This is what we have learned from the fragility of the many causal dogmas. What were yesterday's truths of causation are today's errors. The indeterminism of quantum processes was received in the 1920s with alarm. "…causality has definitely come to an end." This was Waismann's appraisal, as reported above. A century later it is standard to think otherwise. That the connections of quantum theory are merely probabilistic places them comfortably within routine causal analysis. Nothing factual has changed. The only difference is our choice of which definitions are congenial.

My fear is that readers, familiar with Humean skepticism, will misread these last claims as just another version of Hume's causal skepticism. He was averse to such notions as "efficacy, agency, power, force, energy, necessity, connexion and productive quality." (Hume, 1739, p. 157) Where others saw them, all Hume saw was the habit of mind of expecting past regular connections to continue.

My view is not a Humean skepticism. I fully accept that things connect in a myriad of ways; that science has revealed all manner of hidden powers and productive qualities; and that here science enjoys uncommon success. What I deny is that causal metaphysics has anything

[11] This stimulated emission is the process of a LASER = Light Amplification through Stimulated Emission of Radiation.

factual to add to these successes. Tides on the earth's oceans correlate with the positions of the moon and sun. Hume regarded this merely as a constant conjunction. I do not. The tides are raised by the gravitational action of the sun and moon. Call that action "causal" if you want. In so doing, nothing factual is added. All I have learned is how you conceive of the process and where you find causal labels apt.

Under this deflationary view of causation, philosophical theories of causation amount to declarations about how the theorist will use a term. For example, regularity theorists announce that the terms "cause" and "effect" can be used whenever certain sorts of regularities manifest between events of specified types. Those advocating a counterfactual theory of causation augment these occurrent regularities with further conditions concerning what regularities might, counterfactually, have occurred but did not. In the Salmon-Dowe process theory of causation (Dowe, 2000), a process is causal if it transmits a conserved quantity. They are the attaching of labels to the behaviors of things in the world as already discovered empirically by the sciences.

While these labels may be pragmatically useful, none of them and the other theories like them report a factual discovery about things in the world, made antecedently to empirical investigation. There is no compulsion that the world should present us with behaviors of the requisite type, beyond that provided by the evidence unearthed by the sciences. We have powerful evidence that processes in the world convey the conserved quantity energy and, to the extent that this is true, the Salmon-Dowe theory will be able to identify processes to which its causal label can be attached. There is no a priori necessity that all processes must be such. Recent science has at least toyed with the idea of processes we might like to call causal but do not pass energy. In traditional accounts of quantum theory, the mere observation of a quantum system can trigger its collapse without the process of observation exchanging energy with the system. Other accounts of causation face more serious challenges. The enduring difficulty of counterfactual theories of causation is that there is no agreement on precisely which formulation is the right one, if the relationship is to be identified as causal.[12]

---

[12] For a recent survey, see Menzies and Beebee (2020). Frisch (2009, a, b) has argued for a causality principle in physics in which the effect can never precede the cause. One of the criticisms of the proposal (Norton, 2009) is that it fails to identify what counts as a cause and an effect and to distinguish which is which.

6. The Manipulability or Interventionist Conception of Causation

Designations of causality are mere matters of definition. That fact, however, does not make the designations worthless. Namings, even arbitrary ones, can have great utility. Without names for streets, towns, cities and nations, routine travel would be greatly inconvenienced. So far, above, I have suggested that causal naming can enhance understanding. That is a foundationally fragile benefit, for understanding is connected to the idiosyncrasies of the individual. Probabilistic accounts of causation offer some a sense of understanding since now causal relations are expressed in a precise, quantitative theory, the probability calculus. For me, however, the situation is reversed. I find handing over determination of these relations to a computational tool obscures my understanding. When I do the calculation, I can see that conditionalizing on some factor increases the probability of some other factor. However, I need to look behind the probabilities to understand why this increase obtains.

The manipulability or interventionist conception of causation is free of these complications. It identifies a process as causal when our intervention or manipulation of what we will designate the cause results in a corresponding change in what we will designate as the effect. Knowing which processes are causal in this sense has great pragmatic value. The obvious examples are in medicine. We learn that high blood pressure causes strokes and high blood cholesterol levels cause heart disease in this sense. We then know something pragmatically useful: that intervening on high blood pressure and on high cholesterol levels with suitable medications or changes in diet and lifestyle can ameliorate the unwanted effects.

There are many other prosaic examples outside medicine. Water and air cause iron to rust. By intervening to preclude the contact of one or both of water and air with the iron, we can prevent the effect of the rusting of the iron. There are more exotic examples. We noted above that an excited atom can be stimulated to emit radiation if it is immersed in a radiation bath of the right frequency. The manipulation of the surrounding field causes an effect, the emission of radiation. This is how we build LASERs.

The foremost proponent of this approach to causation is Woodward, whose (2003) work is definitive. A version of the account, adapted to variables governed by deterministic relations, is given in a later work (forthcoming, §2) as:

(M) Where *X* and *Y* are variables, *X* causes *Y* iff there are some possible interventions that would change the value of *X* and if were such intervention to occur, a regular change in the value of *Y* would occur.

The most common setting, however, for this interventionist account is sets of variables that are nodes in an acyclic, directed graph of relationships, where a probability measure is defined over the variables. The identification of the causal relations then requires some elaborate lawyering. For what is termed "actual causation," the best developed condition in Woodward (2003, p. 84) is the condition AC*. While most of its details will not be relevant to the discussion here and will not be elaborated, the definition is reproduced to give a sense of the care and richness of the specification. The notion is defined over the variables mentioned in the formulation in upper-case letters and their values in the corresponding lower-case letters:

(AC*1) The actual value of *X*= *x* and the actual value of *Y* = y.

(AC*2) For each directed path *P* from *X* to *Y*, fix by interventions all direct causes Z_i of *Y* that do not lie along *P* at some combination of values within their redundancy range. Then determine whether, for each path from *X* to *Y* and for each possible combination of values for the direct causes Z_i of *Y* that are not on this route and that are in the redundancy range of Z_i, whether there is an intervention on *X* that will change the value of *Y*. (AC* 2) is satisfied if the answer to this question is "yes" for at least one route and possible combination of values within the redundancy range of the Z_i.

X=*x* will be an actual cause of *Y* = *y* if and only if (AC*1) and (AC*2) are satisfied

What will matter for the further discussion here is that this definition and others like it is incomplete without a definition of "intervention." That is provided as IV (p. 98):

(IV)

I1. *I* causes *X*.

I2. *I* acts as a switch for all the other variables that cause *X*. That is, certain values of *I* are such that when *I* attains those values, *X* ceases to depend on the values of other variables that cause *X* and instead depends only on the value taken by *I*.

I3. Any directed path from *I* to *Y* goes through *X*. That is, *I* does not directly cause *Y* and is not a cause of any causes of *Y* that are distinct from *X* except, of course, for those causes of *Y*, if any, that are built into the *I*-*X*-*Y* connection itself; that is, except for (a) any causes of *Y* that are effects of *X* (i.e., variables that are causally between *X* and *Y*) and (b) any causes of *Y* that are between *I* and *X* and have no effect on *Y* independently of *X*.

I4. I is (statistically) independent of any variable *Z* that causes *Y* and that is on a directed path that does not go through *X*.

The essential notion here is that of a "switch." Other similar accounts employ a related notion. Pearl (2009, p. 70) employs a "do" operator that models an intervention as the setting of the value of some variable to a stipulated value. It necessarily breaks some dependencies, while leaving as many others unaltered as possible.

In my view, the interventionist account is probabilistic only by an accident of the context in which it is commonly implemented. In my version of interventionism, what is essential is the notion of an intervention or manipulation. Thus, Woodward's later deterministic version "(M)" above gives a clearer presentation of the core idea. Pearl's "do" operator shows most directly that the probability distribution prevailing over the variables is inadequate to capture the notion of intervention. The "do" is implemented by breaking some relations through means outside the probability distribution. Woodward's condition "(IV)" seeks to implement the notion within the existing probabilistic setting. We shall see below that this leads to one of the enduring problems of the interventionist account.

This interventionist account should be contrasted with other accounts that are essentially probabilistic. Unlike the interventionist account, such accounts cannot be formulated without the probability distribution. In them, at their simplest, a causal relation is identified merely when conditionalizing on what will be designated a cause within the distribution raises the probability of what will be designated the effect. Or, in more complicated cases, such as Reichenbach's (1971, p. 157) common cause principle, conditionalizing on a common cause eliminates probabilistic dependence among the effects. Unlike the interventionist account, such theories cannot be formulated without probability measures.

7. Solving Problems for the Interventionist Account

Two problems have routinely troubled the manipulability or interventionist account: a concern over its limited scope and a threat of circularity. These are long-standing problems and were already addressed at some length in an early development of the interventionist account by Menzies and Price (1993).

7.1 The Problem of Limited Scope

The interventionist account can only identify a causal process if we can intervene on one of its variables. There are numerous processes that we might like to call causal in which this is not possible. The earth's magnetic field derives from electric currents associated with the convection of molten metals inside the earth. Or, to use an example from Menzies and Price (1993, §5), earthquakes arise from a shift in the tectonic plates of the earth's crust. We might like to say that the convection currents cause the magnetic field and the plate shifts cause the earthquakes. However, we have no practical means of intervening or manipulating these causes.

It is tempting to extend the interventionist account by analogy. The process producing the earth's magnetic field is analogous to the production of a magnetic field by an electric current in a laboratory solenoid. That electric current can be manipulated by a simple switch in the circuit and is the cause, in the interventionist sense, of the magnetic field. We might reason similarly in the case of earthquakes by considering how we might manipulate a laboratory scale model of the tectonic plates.

Woodward (2016, §2) has already raised concerns with this extension. How are we assured that the target system and the model are sufficiently analogous for the causal claim to carry over? That would seem to require some alternative conception of cause that has already been found to apply to the target system. Worse, we can proceed to more extreme examples in which the analogy is implausible. Laplace, we saw above, declared the universe at one time to be the effect of its earlier state and the cause of its future state. Thus, we might now conceive the unimaginably hot early universe of modern cosmology as the cause of our present universal state. However, barring fictional fantasies, manipulations of this early state are implausible even in analogy.

These are practical obstacles to extending the interventionist conception more broadly. My view is that there simply is no reason to extend it and good reason not to. We might like to think that the early universe is the precursor of its present state; or that internal currents produce the earth's magnetic field. However, nothing factually is added to the existing cosmological theory and to the existing electrodynamics by the appending the appellation "cause." To feel a pressure to extend the interventionist conception to cases where no intervention is possible is to fall prey to defective causal metaphysics. We would fall into the illusion that there is some deeper fact of causation in the world that goes beyond the explicit content of the sciences. We would be accepting a mistaken view: that these are just two examples of an enormous repository of canonical cases of this true sense of causation that extend beyond cases in which interventions are possible; and that any account of causation must somehow accommodate all these cases.

On the contrary, the strength of the interventionist account is that it gives us a pragmatic reason to designate certain processes as causal. They are processes that we can deploy to our advantage. To seek to apply the theory to cases in which such pragmatically useful manipulation is not possible is to undermine the very strength of the theory. There is no point is trying to find an interventionist apology for these cases. For they have no special claim on us in the first place. Laplace found it apt to label the past as the cause of the present without that designation adding anything to the physical theory that relates the past to the present. Seeking an interventionist interpretation of the relation has no pragmatic value. It dilutes and confuses the interventionist theory by depriving it of exactly what makes it work: that causes are identified as things that can be manipulated.

This last analysis also responds to another objection discussed in Menzier and Price's (1997, §6). Is the interventionist conception unacceptably anthropocentric. The answer here is that the account, or at least my version of it, is essentially and irremediably anthropocentric. It is all about what we can manipulate in the world. That is what makes the theory useful and it needs to offer no apologies for it.

7.2 The Threat of Circularity

An enduring issue with the interventionist theory is the accusation that the theory is circular and thereby undermined. The concern appeared early in Menzies and Price (1993, §4) and still required discussion in Woodward's Stanford Encyclopedia article (Woodward, 2016, §6). The difficulty is that an intervention is itself routinely understood to be a causal process. Thus, any attempt to define causal processes in terms of interventions presumes the very thing that is to be defined. The difficulty is quite clear in Woodward's (2003, p. 98) definition IV above of what it is for *I* to intervene on *X*. Its first clause is just "I1. *I* causes *X*."

While the literature has many responses, it seems to me that the simplest follows directly from the essential anthropocentrism of the conception. My austere version of interventionism can be reduced to two slogans:

A manipulation or intervention need not be a cause; and at least some are not.

A causal process is the propagation of an intervention or manipulation through a network of dependencies.

That is, we take the notion of an intervention, at least in some cases, as a primitive, deriving from our routine understanding of human actions. Once we have introduced an intervention in this primitive sense into the analysis, we can, if applicable, use the causal processes that ensue as interventions within further causal process. This expands the scope of the interventionist theory. We can set up chains of interventions. We intervene on this variable to manipulate another variable; and that manipulation is an intervention that leads to manipulation of a third variable; and so on.

An intervention, in this primitive sense, might arise when we intervene on the electrical conductivity of some component in an electrical circuit by throwing the switch. The propagation of that intervention leads to the light illuminating. The propagation, but not the intervention, is causal. The temptation is to try to apply the interventionist theory to the human action of throwing the switch. Might we say that my willing my hand to move intervened in my nervous system in such a way as to lead my hand to throw the switch? Since I have no special expertise in psychology and neuroscience, I have no idea if something like this can be offered responsibly as an account of my action. However, I am confident that any such account will merely push back the point at which interventions as a primitive are needed in the theory.

Attempts to avoid a primitive notion of intervention have not fared well. We might temporarily manage if there is a cascade of causal processes such that each cause in the cascade serves as an intervention for the next. However, we eventually run into the problem of characterizing the first intervention in the cascade. I read Woodward's definition IV as successfully characterizing interventions that arise as intermediaries in a causal cascade. However, the intervention at the outset of the cascade eludes the definition. There is no earlier intervention to allow it to be identified as causal. Pearl's "do" operator is unapologetically a primitive that lies outside the existing probabilistic relations prevailing over the variables in his directed graphs.

To accept the notion that some interventions are an ineliminably human, primitive notion, as my version does, is no compromise. It is to accept that the theory derives all its value from the pragmatic benefits it affords humans. That the theory depends on a human notion, primitive and irreducible in the context of the theory, is simply a core element of the theory. To try to explicate further this human notion is to undertake a project that belongs elsewhere in science, perhaps in psychology and neuroscience. To regard explication of the notion as outside the scope of the theory is not to offer it as a mystery. It is merely to say that, if it is treated as a primitive, it is a notion well enough understood for the purposes of the interventionist theory; and that further investigation of it belongs elsewhere.

8. Downward Causation

The failure of determinism in the early twentieth century was a disaster for causal metaphysics. My impression had been that causal metaphysicians were sufficiently chagrined by it that they subsequently avoided strong factual claims about causation in favor of accounts that could be read as definitions. That meant that a factual principle of causality or other causal maxims of comparable factual import started to fade from the literature. Such assertions had not fared well in the past. Why risk more of them?

While such prudent hesitation may be true for the mainstream of metaphysicians constructing positive accounts of causation, it turns out not to be the case in adjacent literatures. A notable example concerns the recent debate over "downward causation." An enduring approach by critics of downward causation is to argue that downward causation violates one or more supposedly unimpeachable maxims of causation. However, a closer examination of the maxims shows that none of them are sustainable as factual restrictions. They are not expressions of some factual principle of causality that has for so long eluded causal metaphysicians. In keeping with the general approach of this paper, my view is that whether downward causation occurs is not a factual matter. Rather its occurrence depends entirely on the definition of causation one finds apt. In this section, I will argue that interventionist accounts of causation do support downward causation. In the following section, I will argue for the failure of attempts to impugn this interventionist analysis by means the causal maxims.

8.1 Downward Causation

Downward causation arises when causation acts from a higher level of description of some system to a lower level. Specifying it requires an account of levels. A common example is that thermodynamics provides a higher-level account of thermal systems; and the statistical physics of molecules and radiation modes provides a lower-level account. Another is a motivating example for this literature. It concerns mental causation. The higher-level relations are among mental states and actions, such as willings of some behavior; and the lower level concerns neuronal activity that implements the action. There is much more to say than can be said here. In their introduction to an eighteen-chapter collection of papers on downward causation, Paoletti and Orilia (2017, p. 9) draw a list of roughly thirty proposed examples of downward causation from the volume. They derive from physics, chemistry, biology, neurosciences, psychology and sociology. Ellis (2009, p. 64) identifies a hierarchy of eight levels that rise from particle physics to sociology/economics/politics. Each level harbors its own phenomenological theories and each provides the possibility of top-down causation. "Top-down causation is ubiquitous in physics, chemistry, and biology, …" he concludes (p. 66). As the size of the Paoletti and Orilia volume and other related works suggest, the literature on downward causation is large. My account here can only explore a small portion of it that pertains specifically to issues of this paper.

While the debate over downward causation is energetic and long-lived, from the perspective of the present analysis, the debate depends on its participants employing different definitions of causation; or even just vague definitions of causation. That definitions are so central explains a striking feature of the debate: the factual relations prevailing in disputed cases are often fully known, so that a determination of whether downward causation is present or not adds nothing to our factual knowledge. That is how debates dependent of different definitions manifest. In some cases, however, the prevailing facts are not fully known. The prime example is of mental causation and its foremost exponent is Jaegwon Kim (whose analysis is recounted in the following section). Do our wishes cause our neurons to trigger? What is striking in Kim's analysis is just how irrelevant are detailed factual matters concerning the relationship of mental states and neural states. It is generally supposed that the mental states supervene on the neural states and, merely by reflecting on this supervenience, we can determine whether a causal relations obtains. Notably, a decision in either direction seems to add nothing to our knowledge of the factual processes that connect wishes and neural states.

8.2 The Interventionist Account: The Evaporation of Water

Whether downward causation occurs is, according to the perspective of this paper, simply a matter of how one defines causation. Depending on the definition one finds apt, it may or may not occur. If, for example, one favors a Salmon-Dowe conserved quantity definition of causation, then likely downward causation is precluded. Exchanges of energy and momentum are most plausibly restricted to one level. This is one objection Gillett (2017, pp. 250-51) raises against downward causation. An interventionist account, however, seems quite amenable to downward causation. That downward causation can be understood in the interventionist conception has been defended by Woodward, most recently in his Woodward (forthcoming), and also by Kistler.[13] Ensuring that plausible candidates for downward causation fit with the precise details of the interventionist account has required some elaborate argumentation and delicate adjustments that can be found in their papers. My concern here is not with the details of these adjustments since their success seems assured by the existence of simple examples the conform well with the interventionist conception.

It is easy to see through these examples that the general interventionist idea is quite hospitable to plausible candidates for downward causation. An example I find striking illustrates how this works. The example arises in thermal physics. Warm water evaporates faster than cold water. We can speed up the evaporation of water by intervening on its temperature. We may bring a dish of water into a warmer room; or we may gently warm it on a hot plate; or we may warm it by exposing it to sunlight. These interventions are on the variable of temperature that resides within the higher thermodynamic level of description. On the lower molecular level, evaporation occurs when individual water molecules are energized enough thermally to break free from the forces that hold them in the bulk liquid state. Increasing the temperature of the water increases the speeds of the water molecules and, with it, the probability that individual water molecules exceed this escape threshold. This probability is the lower-level variable affected by manipulation of the higher-level variable, the water temperature.

Two aspects of this instance of downward causation are notable. First, examination of the relation is not just qualitatively useful but also gives useful quantitative results. What follows is a

[13] The proposal has engendered a small debate in the literature. The course of the debate is summarized in Woodward (forthcoming) and Kistler (2017). Ellis' (2009) account also employs interventionist notions, although he does not label them as such. He writes (p. 66) "How do you demonstrate top-down causation? You show that a change in high-level variables results in a demonstrable change in lower-level variables in a reliable way, after you have altered the high-level variable."

very rough and ready illustration. The energies of water molecules are Boltzmann distributed. That is, the probability that a molecule has energy E is proportional to the Boltzmann factor, $\exp(-E/kT)$, where k is Boltzmann's constant. If we extract from this distribution just the kinetic energy associated with the vertical component of velocity v of each molecule (towards the water surface), we can approximate well enough the probability density over the vertical component of the velocity v of molecules of mass m as

$$f(v) = \left(\frac{m}{2\pi kT}\right)^{1/2} exp\left(-\frac{mv^2}{2kT}\right)$$

A first crude model is that a water molecule escapes the surface when its vertical speed v exceeds some value V. The probability that a water molecule exceeds this vertical speed in water of temperature T is

$$\Pr(T) = \int_{v=V}^{\infty} \left(\frac{m}{2\pi kT}\right)^{1/2} exp\left(-\frac{mv^2}{2kT}\right) dv = \int_{x=X}^{\infty} \left(\frac{1}{2\pi}\right)^{1/2} exp(-x^2/2) dx$$

where $x = (m/kT)^{1/2}v$ and $X = (m/kT)^{1/2}V$. This integral extends over an exponentially decaying tail, so that most of the probability will be massed close to the lower bound, X. Hence, as a rough approximation, we can assess the effect of varying temperature through the proportionality:

$$Pr(T) \sim \exp(-X^2/2) = \exp(-mV^2/2kT)$$

Taking logarithms, the temperature dependency of the probability can be expressed as

$$\log Pr(T) \sim \text{const.} - \text{const.}/\text{T}$$

Since the rate of evaporation is determined by this probability, we find a "-1/T" increase of the rate with temperature.

It is striking that this relation agrees with a corresponding result located entirely within the thermodynamic level. At this level, the rate of evaporation is determined by the vapor pressure P of the liquid. The Clausius-Clapeyron equation relates this vapor pressure P to the system temperature T. Under suitable idealizing assumptions, it reduces to[14]

---

[14] The Clausius-Clapeyron equation is $dP/dT = L/(T\Delta v)$, where L is the molar specific heat of vaporization and Δv is the molar specific volume change on evaporation. We assume that L is constant over the temperature range of interest and that the vapor phase is an ideal gas, so that $\Delta v = RT/P$, where R is the ideal gas constant and the small liquid volume is neglected. The Clausius-

$$\log P \sim \text{const.} - \text{const.}/T$$

The agreement of the rough molecular calculation with this result is actually too good considering the approximations employed. However, all that matters for our purposes is that this example of downward causation supplies useful quantitative results that could be refined with more realistic assumptions.

The second aspect pertains to a common rejoinder to examples of downward causation like this. There is, we are assured, necessarily a causal relation that resides fully within the lower level; and this renders the downward causation redundant. If the water is warmed on a hot plate, the relation would involve the thermally excited energy of molecules of the hot plate, how that energy is communicated to the water vessel and then to the water. If the water is warmed radiatively, however, the relation would pertain to the energetic coupling of radiation modes in incident sunlight with degrees of freedom of the individual molecules.

No doubt, some definition of cause would apply to these lower-level processes. For example, they involve the transfer of energy and momentum in conformity with the Salmon-Dowe account. However, identifying these relations as causal would have little utility. They do not aid in understanding but merely bury us in masses of superfluous molecular-scale details. They do not enhance our ability to affect the rate of evaporation. That capacity is already provided by the relation of downward cause: to increase the rate of evaporation, increase the temperature by whatever means you like. Indeed, if anyone were to attempt the fully molecular account, they would likely end up noting that that conductive or radiative heating merely serves to increase the parameter T in the Boltzmann distribution. In effect, in the essential part, they would end up agreeing with the downward causal account. It would be in a less informative way since the T in the account would merely be a parameter divorced from a higher-level, thermodynamic interpretation.

That is, in terms both of understanding and pragmatic application, the downward causal relation is superior to the one instantiated solely at the lower level.

Clapeyron equation becomes the differential equation $dP/dT = LP/RT^2$ and its integral solution is the “-1/T” dependency indicated.

9. Causal Maxims are Reawakened

Here I will review attempts to impugn downward causation by the suggestion that it contradicts supposedly secure maxims about causation. These attempts fail, I will argue, because of the dubious status of the causal maxims. To summarize what we will see below, there are two difficulties with the maxims.

The first difficulty is that the viability of a causal maxims depends sensitively on the sense of causation employed. A maxim that works for one sense may not for another. Since the Salmon-Dowe approach requires the passage of a conserved quantity from a cause to an effect, it is inherently asymmetric. However, the causal relations in the interventionist account are between variables and thus can go in either direction. Again, we may find that causal closure or completeness fail in quantum mechanics if our conception of causation is deterministic. It may succeed, however, if we posit a probabilistic conception. The result is that no single maxim can assure a blanket preclusion of downward causation. A causal maxim adapted to one sense of causation may not conform with another. This multiplicity of notions of causation is central to Ellis' (2009) identification of five different types of top-down causation that arise according to the context. In reflecting on them, he writes (p. 64, his emphasis):

> There could be others, but I claim that these can all be regarded as well established. In brief: *there are other forms of causation than those encompassed by physics and physical chemistry*. A full scientific view of the world must recognize this fact, or else it will ignore important aspects of causation in the real world, and so will give a causally incomplete view of things…

The second difficulty is more serious. These causal maxims are deployed by critics of downward causation as deep truths, not mere matters of comfortable definition. Thus, they need to be secured by appropriate empirical investigation. Yet these causal maxims are most commonly just declared as self-evident truths or perhaps just what everyone thinks. This is a fallacy of "argumentum ad populum" and not a secure basis for maxims that are to control a major, foundational debate. There are cases, we will see, in which a factual basis is identified. Commonly, the factual basis resides in carefully chosen examples. Yet, as we shall see, we can find other examples that contradict the maxim. More curiously, such identifications are often accompanied by the concession that the maxim has been disputed and has known counterexamples.

The result is that efforts to use causal maxims as a means of impeaching downward causation fail; and, in particular, they fail as objections to the interventionist account of downward causation.

9.1 Causal Exclusion and Causal Closure

Perhaps the best-known arguments against downward causation are offered by Jaegwon Kim in several papers and his two works (1998, 2005). They arise in the context of the philosophy of mind. His "supervenience argument" against downward causation is based on two maxims. The first is the "causal exclusion principle" (2005, p. 42):

> *Exclusion*. No single event can have more than one sufficient cause occurring at any given time—unless it is a genuine case of causal overdetermination.

The second is the causal closure principle (2005, p. 43):

> *Closure*. If a physical event has a cause that occurs at *t*, it has a physical cause that occurs at *t*.[footnote]

These principles are applied to the two levels in the philosophy of mind: the mental and the physical. If some event at the physical level has a cause, then, by *Closure* it has a physical cause. We may try to offer a mental cause as well. *Exclusion* precludes this mental cause and with it the possibility of downward causation. Kim (2005, p. 44, his emphasis) concludes "…*the assumptions of causal exclusion and lower-level causal closure disallow downward causation*."

Both these principles are inadequate from the perspectives of this paper. Both proceed without drawing anything of substance from the empirical investigation of psychology and neuroscience. Hence it is hard to see how these principles can be securely based. Both proceed without giving a clear definition of what is meant by the term "cause." Without it, it is hard to know precisely what they say. We might accept that the ontology of the lower level is exhaustive. That is, in this context, we might accept that all mental activity simply is neuronal. However, that view of the ontology is not sufficient to establish the causal closure principle. The notion of cause at issue must also be specified and then it must be shown that the closure principle applies to it. Might it be that the causes of an effective analysis are only properly expressible at some higher level so that there is no description of the cause in the lower level? Ellis (2009, p. 66) argues that just this does happen since, in his view, top-down causation depends on higher level context variables:

Altering the high-level context alters lower-level actions; this is what identifies the effect as top-down causation. In such cases the high-level context variables are not describable in lower-level terms, and this is what identifies them as context variables.

These ambiguities make the principles fragile. Kim's footnote to the above statement of *Closure* illustrates it. The footnote calls the reader's attention to two further sources for discussion. One is Lowe (2003) who points out that quantum theory prevails at lower levels. Under a deterministic conception of causation, probabilistic quantum events such as the death of Schroedinger's cat, have a cause but, he argues, not a sufficient cause. He continues to argue that replacing the deterministic notion of causation with a probabilistic one does not save the analysis. We are assured that: "we have every reason to suppose that (CCP) [causal closure principle] is empirically false, its falsity being testified to by the empirical data which support the principles of modern quantum physics." (p. 142) and finally "I am firmly convinced that (CCP) is false." (p. 145)

It is perhaps germane to add that this supervenience argument does not succeed in impugning most of the cases of downward causation proposed in the present literature. If it works at all, it would only apply to cases of downward causation such as would appear in a dualist theory of mind or other theories like it with a dualist ontology. Then the downward causation is from a mental world to the physical world that is distinct from it.

Aside from these dualist cases and others functionally similar to them, Kim's framework actually affirms downward causation for the common examples in the literature. This affirmation arises because of the way Kim implements the causal exclusion principle. The principle is developed in an earlier paper, Kim (1989). We are asked to consider two explanations for some event (p. 89)

Explanation A cites C as a cause of E

Explanation B cites C* as a cause of E

We are then taken through five cases. The first four are: C and C* are the same; C is distinct from C* but reducible to it; neither C nor C* are individually sufficient as a cause; and C and C* are different links in the causal chain that leads to E. The causal exclusion principle applies to none of these. For none provide independent causal explanations. The causal exclusion principle applies only to the fifth case in which C and C* are distinct and individually sufficient as causes of E. This fifth case includes dualist ontologies, such as a dualist theory of mind.

These first four non-dualist cases are the cases implemented in common claims of downward causation. There we typically have a higher-level description of a cause and lower-level description of a cause. They are commonly the same thing just described at different levels; or something close to it along the lines of Kim's first four cases above. At the higher level of description, transferring heat to water has the effect of increasing its rate of evaporation. At the lower level, energizing the individual water molecules has the effect of increasing the probability that they escape the liquid phase and evaporate. These two causes—transferring heat and energizing molecules—are just the same thing described at different levels. In such cases, if the thing identified at the lower level is a cause of the effect, then *necessarily* the thing identified at the higher level is also a cause of the effect. It has to be that way since they are the same things, just described at different levels.

We have the curious result that, if we accept Kim's maxims and his understanding of them, his work provides a vindication of downward causation in most of the cases considered in the Paoletti and Orilia (2017) volume.

9.2 Whole to Part Causation, Synchronic Causation

A different causal maxim that is used to criticize downward causation pertains to wholes and parts. Woodward (forthcoming, p. 12) reports:

> A very common criticism of the idea of downward causation is that this requires that "wholes" act downward on their "parts" and that the relation between a whole and its parts cannot be a causal relation of any kind.

He identifies Craver and Bechtel (2007) as an instance of it. In their work, the preclusion is justified by dual considerations: that wholes and parts are not sufficiently distinct: and that whole to part causation is synchronous. That is, it violates a second causal maxim that prohibits instantaneous causal actions. Their initial formulation is interesting for present purposes (2007, p. 551):

> Many common assumptions about the nature of causation preclude the possibility of causal relations between parts (components) and wholes (mechanisms). To start with an especially clear example, consider the view that all causation involves

> transmitting something such as a mark (Salmon 1984) or a conserved quantity (Dowe 2000) from one event, object, or process to another.

The reference to "common assumptions" reflects the concern above of an argumentum ad populum. The attachment to particular accounts also reflects the fragility of the maxim. For what holds in one account may not hold in another. Attached to this maxim is the second maxim that precludes synchronic causation, that is, causation in which the cause and its effect happen at the same time. Their concern is that we might have two system each instantly causing each other to acquire some causally efficacious property. They continue (p. 553):

> To avoid this problem, one might assume that causal transactions across levels take time: the effects of changes to a component alter the behavior of the mechanism as a whole at some later time, and vice versa.

We see a similar recourse to popular ("generic") views in Gillett's suppositions that could be used to preclude whole to part causation (2017, p. 250):

> I simply use generic features of the causal relation to assess the claim that FDR [Fundamental Downward Relation] is downward causation in such cases. I therefore assume causation to be a relation holding between wholly distinct individuals that is temporally extended and where its *relata* are usually at different locations.

Woodward's (forthcoming, p.8) defense of the interventionist account against these objections seems to me to be decisive:

> Scientifically plausible examples of downward causation do not involve wholes acting on parts but rather involve variables (as all causal relations do) and these need not stand in part/whole relationships, even when entities of which they are predicated do.

My concern here, however, is the cogency of the maxims that prohibit whole to part causation and synchronic causation. They do perhaps fit with the examples Craver and Bechtel favor. However, they do not fit with others. Laplace, we saw above, regarded the present state of the universe as a cause of its future state. Presumably the future state includes its parts. Thus, the totality of the position and velocities of the sun and other solar system bodies today determines their whole configuration tomorrow. Within that whole will be a part, such as a particular, localized event. It might be an eclipse tomorrow. Thus, today's state of the whole is the cause, in Laplace's sense, of that part of the tomorrow's whole.

This last whole to part relation is not synchronous. Today's whole causes something in tomorrow's part. Yet that a causal action is synchronous has historically not precluded it from serious consideration. For hundreds of years after Newton, a paradigm of causation was the gravitational action exerted among bodies. In it, the Newtonian forces acting on a body are the cause of the effect of the body's acceleration. Since Newtonian gravitational forces propagate instantly, the cause and its effect are synchronous. In the early nineteenth century, Laplace tried to place a bound on the speed of propagation of gravity and found such a high speed that he concluded it could be assumed infinite.[15]

If gravitational actions are causal, we may take this example as a case of whole to part, synchronous causation. The cause of a celestial body's acceleration is the entirety of the gravitational action of all the bodies in the universe. If we think in terms of fields, the force on the body is due to the prevailing gravitational field, which is in turn determined through instantaneous action by the masses and positions of all the bodies in the universe.[16]

9.3 The Asymmetry of Causation

Two related maxims concerning causation have also been applied in the literature critical of downward causation. They are:

---

[15] Laplace (1805, p. 526): "… one finds [from the secular motion of the moon] that the speed of the gravitational fluid is about seven million times greater than that of light … one may suppose for the gravitational fluid a speed at least a hundred million times greater than that of light. … Therefore, surveyors can, as they have done so far, assume that this speed is infinite."

[16] A loophole here is that this entirety may exclude the body in question. Since Newtonian theory is linear, altering the mass of the body does not affect the force acting on it. However altering its mass will alter the effect described, its acceleration, all else held constant. If this is still a concern, we can consider a non-linear Newtonian-style gravitation theory in which the body's own "self-field" contributes to the acceleration. Such self-field contributions arise in the classical theory of accelerating electric charges.

> The causal relation is asymmetric: causes act on their effects, but effects do not act on their causes.
>
> The causal relation is time ordered: causes precede their effects.

The difficulty with both maxims is that they have been challenged extensively in the present literature. The challenges are so pervasive that the favorable reporting of the maxims commonly include the concession that they are challenged.

Here is Craver and Bechtel's (2007, p. 552) introduction of the maxim that causes temporally precede their effects:

> Many theories of causation assume that causes precede their effects. This feature of causation is often disputed (see Faye 2005), and some accounts of causation (e.g., Reichenbach 1958) are designed as the foundation of an account of the temporal order, and so do not assume the temporal asymmetry of causation.

The Faye source cited (here, Faye, 2005) is to the then latest, 2005 Stanford Encyclopedia article "Backward Causation." (The article has since been revised in several later editions.) It constitutes a rather substantial dispute of the maxim. The article responds to the suggestion that backward causation is somehow paradoxical and reports several physical theories that propose backward causation.[17]

In another context, Frisch (2009a, §2) has argued for a similar principle in the context of dispersion processes in classical electrodynamics: "…an effect cannot temporally precede its causes…" A difficulty with the proposal is noted in Norton (2009). It is that the physics of dispersion in classical electrodynamics is time reversible. That means that, if we have a process in which state or event A precedes B in time, then we can also have the time-reversed process in which B precedes A. If we judge that A causes B in the first process, then we must also judge that, A causes B in the second time-reversed process. But that relation in the time-reversed case would be causation backward in time. An unsuccessful escape from this conclusion is the

[17] The Reichenbach citation is harder to characterize as a dispute. Reichenbach (1956, §21) introduces the "causal theory of time," in which one event is *defined* as later in time than a second event, if that second event can causally influence the first event. A unidirectional sense of time arises in this theory only if the causation relation is asymmetric.

empirically unfounded assumption that there are time-directed causal facts in classical electrodynamics that elude its laws.

After affirming the maxim that causes must precede their effects, Craver and Bechtel (2007, p. 553) turn to the second maxim of the asymmetry of causation:

> [The maxim that causes must precede their effects] raises a related worry about the asymmetry of causation.[18] It is a widely accepted condition on accounts of causation that they account for the asymmetry of causal dependency. The sun's elevation causes the length of the shadow, but the length of the shadow does not determine the elevation of the sun. The virus produces the spots on the skin, but the spots on the skin do not cause the infection with the virus. Causes produce their effects, and (at least in many cases) not vice versa. Examples such as these have the staying power that they do because the asymmetry of causation is so fundamental to our very idea of causation.

The work of Price (1996) referred to in the footnote to the text does more than merely question the asymmetry of cause and effect. It proposes that it is an artifact of our human conceptions and not a factual matter in the world (p. 161): "On the perspectival view, causal asymmetry reflects an asymmetry in us, not an asymmetry in the external world." Price (1996, pp. 183-87) also uses the time reversibility of elastic collisions of billiard balls in same way as I did above with the time reversibility of classical electrodynamic processes. Since the collisions are time reversible, there is nothing in the physics itself that distinguishes which state is the cause and which the effect.

More generally, the asymmetry of cause and effect appears sustainable only as long as we choose favorable examples. It is hard to retain this asymmetry when one considers equilibrium chemical reactions. For example, hydrogen and iodine combine to form hydrogen iodide:

$$H_2 + I_2 \rightarrow 2HI$$

Conversely, hydrogen iodide can decompose into hydrogen and iodine:

$$2HI \rightarrow H_2 + I_2$$

---

18 Craver and Bechtel's footnote: "Again, this principle has been questioned. See Price (1996) for a lengthy review."

We may fill our reaction vessel with hydrogen and iodine and the first reaction will start to produce hydrogen iodide. If we hold the temperature at 445^{o}C, the production will cease at 79% hydrogen iodide and 21% elemental hydrogen and iodine. If instead we had initially filled our reaction vessel with hydrogen iodide, the second reaction will start to form hydrogen and iodine. In neither case will the reactions go to completion. Rather holding the temperature at 445^{o}C, they will approach a dynamic equilibrium at the percentages above in which both reactions proceed but the rates of formation of reaction products of one reaction will match the rates of consumption in the other.[19]

One can describe this dynamical equilibrium as the occurrence of two asymmetric causal process at the same time. In one, the cause, hydrogen and iodine, leads to the effect, hydrogen iodide, through the first reaction. In the other, the cause, hydrogen iodide, leads to the effect, hydrogen and iodine in the second reaction. That is an artificial division. In this dynamical equilibrium, the reagents hydrogen, iodine and hydrogen iodide are both causes and effects at the same time. The dynamical equilibrium is the continuing execution of the two reactions above. However, if any causal account is appropriate, it would be one in which we say that these reagents are engaging in a continuing causal interaction, without the asymmetry of the distinction of cause and effect playing any role. All the reagents are at once both cause and effect.

Craver and Bechtel are concerned that allowing causal processes to proceed in both directions would lead to troublesome causal circles (2007, p. 552): “the possibility of bottom-up and top-down influence ‘propagated’ simultaneously across levels results in problematic causal circles.” Perhaps this may happen in specific cases. However, as a general matter, there can be no maxim that prohibits causal circles. For the chemical equilibrium described above is just such a causal circle. Equilibrium systems with analogous causal circles are pervasive in chemistry and in other sciences. Woodward (forthcoming, pp. 15-16) finds nothing incoherent in such causal cycles and reports their existence in biological, social and economic systems. They are essential, he notes, in control systems that employ feedback loops.

---

[19] These reactions and conditions are as reported in Senter (1913, p. 164). Senter notes also that: “The investigation of this reaction has been of considerable importance in the development of chemistry…”

10. Conclusion

This paper has presented a sustained critique not of causation, but of the metaphysical literature on causation. Its basis is empirical. Things connect with other things in the world in a myriad of interesting ways. It is the job of empirical science to find out just how they connect. Empirical science has enjoyed immense success in these efforts; and we can expect its successes to continue without any foreseeable limit.

As these successes continue, one feature of them persists: we cannot know just what the next discoveries will be. Non-empirical attempts to predict them have not and will not end well. Yet just such attempts have been the enduring substance of the metaphysics of causation. Its history is one of the repeated failure of efforts to impose enduring factual restrictions. By demanding that future science must respect the present conceptions of causation, they are attempts to circumscribe what empirical science can and cannot discover next. They have been, in effect, ill-fated attempts at a priori science.

The repeated history of errors and corrections in the metaphysics of causation is quite unlike the history of corrections in science. The power of discovery and correction derives from the empirical probing of the world that is lies at the center of science. It has within itself the means to correct its own errors. Metaphysical investigations into causation have no corresponding power. They can only pronounce and then accept their errors when subsequent discoveries in science reveal them.

If the analysis of causation cannot add anything factual to what science tells us, what can it provide? I have argued here that the identification of some things as causes and effects and some processes as causal is a matter of conventional definition. While this conception has lowered ambitions, it can be valuable if the designations are useful to us. Time and again, conceiving a process in causal terms calls up analogies that aid our limited human brains in comprehending novel processes. The interventionist approach provides special benefits. For if we know that a process is causal in the interventionist sense, we know that we can intervene in it and use it to manipulate things in the world. Since the account is purely one of definition, there is no presumption that Nature must present us with processes that admit interventions. But when they are found, the valuable discovery is promulgated efficiently by the simple expedient of attaching the label "causal" to them.

References

Audi, Robert (ed.) (1999) *Cambridge Dictionary of Philosophy*. 2nd ed. Cambridge: Cambridge University Press.

Boyle, Robert (1688) *A Disquisition about the Final Causes of Natural Things*. London: John Taylor.

Craver, Carl F. and Bechtel, William (2007) “Top-down Causation without Top-down Causes< *Biology and Philosophy*. **22**, pp. 547-63

Descartes, René (1982) *Principles of Philosophy*. Trans V. R. Miller and R. P. Miller (Dordrecht: Kluwer.

Dowe, Phil (2000) *Physical Causation*. Cambridge: Cambridge University Press.

Earman, John (1986) *Primer on Determinism*. Dordrecht: Reidel.

Einstein, Albert (1917) “On the Quantum Theory of Radiation,” pp. 63-77 in B. L. van der Waerden, *Sources of Quantum Mechanics*. Amsterdam: North-Holland, 1967.

Ellis, George F. R. (2009) “Top-Down Causation and the Human Brain,” Ch. 4 in N. Murphy, G. F.R. Ellis, and T. O’Connor, eds., *Downward Causation and the Neurobiology of Free Will*. Berlin: Springer.

Faye, Jan (2005) “Backward Causation,” *The Stanford Encyclopedia of Philosophy* (Fall 2005 Edition), Edward N. Zalta (ed.). https://plato.stanford.edu/archives/fall2005/entries/causation-backwards/

Fisher, Ronald A. (1934) “Indeterminism and Natural Selection,” *Philosophy of Science*, **1**, pp. 99-117.

Fleming, William (1860) *The Vocabulary of Philosophy*. 2nd ed. Philadelphia: Smith, English and Co.

Frisch, M. (2009a) “ ‘The’ Most Sacred Tenet”? Causal Reasoning in Physics”, *British Journal for the Philosophy of Science*, **60**, pp. 459–74.

Frisch, M. (2009b) “Causality and Dispersion: A Reply to John Norton,” *British Journal for the Philosophy of Science*, **60**, pp. 487–95.

Gillett, Carl (2017) “Scientific Emergentism and Its Move beyond (Direct) Downward Causation,” Ch. 14 in Paoletti and Orilia (2017).

Hitchcock, Christopher, "Probabilistic Causation," *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/spr2021/entries/causation-probabilistic/

Hume, David (1739) *A Treatise of Human Nature*. L. A. Selby-Bigge, ed. Oxford: Clarendon, 1896.

Kim, Jaegwon (1989) "Mechanism, Purpose, and Explanatory Exclusion," *Philosophical Perspectives*, **3** , pp. 77-108.

Kim, Jaegwon (1998) *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, MA: Bradford, MIT Press.

Kim, Jaegwon (2005) *Physicalism, or Something Near Enough*. Princeton: Princeton University Press.

Kistler, Max (2017) "Higher-Level, Downward and Specific Causation," Chapter 4 in Paoletti and Orilia (2017).

Laplace, Pierre-Simon (1805) *Traité de Mécanique Céleste*. Vol. 4. Paris.

Laplace, Pierre-Simon (1902) *A Philosophical Essay on Probabilities*. London: John Wiley & sons.

Lowe, E. J. (2003) "Physical Causal Closure and the Invisibility of Mental Causation," Ch. 6, pp. 137-54, in, Sven Walter and Heinz-Dieter Heckmann, eds., *Physicalism and Mental Causation*, Exeter, UK: Imprint Academic.

Mackie, John L. (1980) *The Cement of the Universe: A Study of Causation*. Oxford: Clarendon Press.

Mellor, David Hugh (1995) *The Facts of Causation*. London: Routledge.

Menzies, Peter and Price, Huw (1993) "Causation as a Secondary Quality," *The British Journal for the Philosophy of Science*, **44**, pp. 187-203.

Menzies, Peter and Helen Beebee (2020) "Counterfactual Theories of Causation," *The Stanford Encyclopedia of Philosophy* (Winter 2020 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/win2020/entries/causation-counterfactual/.

Mill, John Stuart (1882) *A System of Logic.* 8th ed. New York: Harper & Bros.

Newton, Isaac (1761) *Four Letters From Isaac Newton to Doctor Bentley*. Pall Mall: R. and J. Dodsley, 1761, pp. 25-26.

Norton, John D. (2003) "Causation as Folk Science," *Philosophers' Imprint* Vol. 3, No. 4.

Norton, John D. (2009) “Is There an Independent Principle of Causality in Physics?” *British Journal for the Philosophy of Science*, **60**, pp. 475-86.

Norton, John D. (2010) “Cosmic Confusions: Not Supporting versus Supporting Not-“. *Philosophy of Science*. **77**, pp. 501-23.

Norton, John D. (2014) “A Material Dissolution of the Problem of Induction.” *Synthese*. **191**, pp. 671-690.

Norton, John D. (2016) “Curie's Truism.” *Philosophy of Science*. **83**, pp. 1014-1026.

Norton, John D. (2021a) “Eternal Inflation: When Probabilities Fail.” *Synthese* **198** (Suppl 16), S3853-3875.

Norton, John D. (2021) *The Material Theory of Induction*. BSPSopen/University of Calgary Press.

Norton, John D. (manuscript) The Large-Scale Structure of Inductive Inference. http://sites.pitt.edu/~jdnorton

Osler, Margaret J. (1996) “From Immanent Natures to Nature as Artifice: The Reinterpretation of Final Causes in Seventeenth-Century Natural Philosophy,” *The Monist* , **79**, pp. 388-407.

Paoletti, Michele Paolini and Orilia, Francesco (2017) (eds.) *Philosophical and Scientific Perspectives on Downward Causation*. New York: Routledge.

Pearl, Judea (2009) *Causality: Models, Reasoning, and Inference.* Second Edition. Cambridge: Cambridge University Press.

Price, Huw (1996) *Time's Arrow & Archimedes' Point*. Oxford: Oxford University Press.

Reichenbach Hand (1958) *The Philosophy of Space and Time*. Dover, New York.

Reichenbach, Hans (1971) *The Direction of Time*. Ed. Maria Reichenbach. Berkeley: University of California Press.

Russell, Bertrand (1912-13) “On the Notion of Cause,” *Proceedings of the Aristotelian Society*, **13** (1912-1913), pp. 1-26.

Senter, George (1913) *A Text-Book of Inorganic Chemistry*. 2nd ed. London: Methuen & Co.

Waismann, Friedrich (1959) *Turning Points in Physics*. Amsterdam: North-Holland.

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

Woodward, James F. (2016) “Causation and Manipulability,” *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/win2016/entries/causation-mani/

Woodward, James F. (forthcoming) “Downward Causation Defended,” *Top-Down Causation and Emergence*, Eds. J. Voosholz and M. Gabriel. Springer

Wysocki, Tomasz (manuscript) “Underdeterminsitic Causation: A Proof of Concept.”

Zenneck, Jonathan (1901). “Gravitation.” *Encyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen*. Vol 5. Part 1 (1903-1921), pp. 25–67.