

The Ontology of Causation: A Carnapian-Pragmatist Approach

Zili Dong

Department of Philosophy, Western University

Abstract: Metaphysicians of causation have long debated the existence of primitive causal modalities (e.g., powers), with reductionists and realists taking opposing stances. However, little attention has been given to the legitimacy of the metaphysical question itself, despite our longstanding awareness of Rudolf Carnap’s critique of metaphysics. This article develops a (broadly) Carnapian-pragmatist approach to causation as an alternative to existing metaphysical approaches. Within this pragmatist approach, metaphysical questions about causation are reinterpreted as practical questions about the choice of causal frameworks. To motivate and justify this new approach, I argue that, in emphasizing the priority of ontology over methodology, metaphysical approaches to causation fail to adequately capture the interplay between causal ontology and causal methodology in scientific practice. In contrast, the Carnapian approach provides a more appealing alternative that emphasizes the mutual dependence and “balance” between the two in an ongoing process of scientific inquiry. I use the recent controversy over “What counts as a cause” in statistical causal inference as a case study to demonstrate how the Carnapian approach can help us better understand the role of ontological issues in methodological practices.

Keywords: Causation, Causal inference, Metaontology, Rudolf Carnap, Pragmatism

Acknowledgements: I am grateful to Wayne Myrvold and two anonymous reviewers for their invaluable feedback on earlier versions of this paper. I also thank Xiuyuan An, Weixin Cai, Eric Desjardins, Yousuf Hasan, Yichen Luo, Zihan Qu, Shuguo Tang, and Shimin Zhao for their helpful comments.

1. Introduction

In “Empiricism, semantics, and ontology”, Rudolf Carnap (1956/1988) proposes a pragmatist meta-ontology in an attempt to settle various traditional metaphysical disputes. Despite remaining a minority view, Carnap’s meta-ontology has received increasing attention in contemporary debates (e.g., Blatti & Lapointe, 2016; Eklund, 2009, 2013; Price, 2009; Thomasson, 2014). Thus far, these debates have focused primarily on general ontological issues. However, I believe a Carnapian meta-ontology can also shed light on a broad range of concrete and domain-specific issues, such as causation and modality, chance and probability, natural kinds, scientific models, and social ontology.¹ This article is an attempt to apply a Carnapian meta-ontology to causation.²

Philosophers, especially analytic metaphysicians, have paid a great deal of attention to metaphysical issues about causation—for example, whether there are primitive causal modalities in the world. However, little has been said about the legitimacy of this kind of metaphysical inquiry. Given the recent revival of interest in Carnap’s meta-ontology, it is time philosophers of causation take Carnap’s critique of metaphysics seriously. In this article, I shall develop a Carnapian-pragmatist approach to causation as an alternative to existing metaphysical approaches. In this pragmatist approach, metaphysical questions about the “external reality” of causation are recast as pragmatic questions about the choice of causal frameworks. This project has various implications for our understanding of causation; for instance, as I will endeavour to show, it depicts a more sensible picture of the role of ontological questions (e.g., “What can be a cause”) in causal inference practice.

Since it is not my goal here to contribute to Carnap scholarship, I will only provide a brief introduction to Carnap’s meta-ontology. A key element of Carnap’s meta-ontology is the notion of a

¹ Danks (2015) and Ludwig (2016) can be interpreted as taking (broadly construed) Carnapian approaches to scientific ontology. Antoniou (2021) defends a Carnapian pragmatist approach to the ontology of scientific models. Lauer (2022) proposes a pragmatic approach to social ontology that seems to be Carnapian in spirit.

² Fischer (2023) develops a Carnapian approach to the ontology of actual causation which emphasizes the role of goals and context in improving or explicating our concept of actual causation. In addition, broadly pragmatist approaches to causation can be found in Eagle (2007), Hitchcock (2012), Price (2001, 2007), and Woodward (2014, 2015, 2017, 2021).

(linguistic) *framework*: a system of vocabularies and rules that can be used to formulate assertions about a system of ontological postulates. For example, to introduce natural numbers to our ontology, we need to construct a framework that provides the necessary vocabulary and rules for talking about natural numbers. Based on this idea, Carnap makes the famous distinction between internal and external ontological questions. *Internal* questions are questions of existence raised within a framework that can be answered by either logical or empirical methods. For this reason, they are theoretically meaningful: answers to these questions can be judged as true or false within the given framework. For example, within the framework of ordinary objects, I can ask, “Is there a desk in my room?” and the answer is “Yes”, given the empirical evidence I have. In contrast, *external* questions are concerned with the existence of a whole system of ontological postulates, such as “Does the system of ordinary objects exist?”.

Carnap distinguishes between two ways of understanding external questions: they can be understood as *theoretical* questions about the nature of reality or as *practical* questions about the choice of frameworks. Traditional metaphysicians, including many contemporary analytic metaphysicians, treat external questions as theoretical. For them, truth values of assertions about ordinary objects or natural numbers are not contingent on the frameworks we choose to formulate the assertions. They often use “really” (e.g., “Are there *really* X?”) to emphasize that these metaphysical questions should be understood in a framework-independent way. Carnap contends that traditional metaphysicians have misunderstood the nature of external questions. External questions are not theoretical but practical: they are questions about whether we should accept or reject a framework, and such decisions should be made on *pragmatic* grounds:

[The acceptance of the new entities] does not mean for us anything more than acceptance of the new framework ... The acceptance cannot be judged as being either true or false because it is not an assertion. It can only be judged as being *more or less expedient, fruitful, conducive to the aim for which the language is intended*. Judgments of this kind supply the motivation for the decision of accepting or rejecting the kind of entities. (Carnap, 1956/1988, p. 214; emphasis added)

In summary, Carnap's meta-ontology includes not only a denial of the theoretical significance of external questions but also an affirmation of the practical significance of these questions as questions about *framework choice*. In other words, in addition to deflating the metaphysical significance of external questions (this is also why Carnap's meta-ontology is often said to be "deflationary"), Carnap's meta-ontology also inflates their practical significance. Indeed, it is exactly because external questions about framework choice have important practical consequences, especially in the scientific context, that Carnap warns us against conflating practical external questions with theoretical ones.³

The rest of the article is structured as follows. Section 2 outlines and compares three approaches to the ontology of causation: reductionism, realism, and Carnapian pragmatism. Section 3 explains in detail how the Carnapian-pragmatist approach works and what implications it has for causal pluralism and the objectivity of causation. Section 4 offers some motivations and justifications for this Carnapian-pragmatist approach to causation. I argue that compared to metaphysical approaches to causation, Carnapian-pragmatism provides a better understanding of the relationship between causal ontology and causal methodology in scientific practice. I use the recent controversy over "what counts as a cause" in statistical causal inference as a case study to illustrate how the Carnapian approach helps us make sense of the role of ontological controversies in causal inference practice. Section 5 is a brief conclusion.

For clarity, I will use "ontology_M" and "ontology_C" to refer to the metaphysician's and the Carnapian's ontology, respectively; when "ontology" is used without subscripts, it refers to any kind of ontological inquiry, broadly construed.⁴

³ At the end of the paper, Carnap (1956/1988, p. 221) says, "[t]o decree dogmatic prohibitions of certain linguistic forms instead of testing them by their success or failure in practical use ... is positively harmful because it may obstruct scientific progress."

⁴ I thank two anonymous reviewers for suggesting this important clarification. My distinction here draws inspiration from Woodward's (2015) distinction between ontology₁ and ontology₂. Our distinctions are essentially the same except for one important difference: my ontology_C subsumes Woodward's narrower notion of ontology₁. An "ontology₁" (e.g., a gene ontology) refers to a system of "basic" entities, properties, and structures, together with ways of classifying them, in a particular scientific domain. Ontology_C, in contrast, encompasses any kind of

2. Approaches to the ontology of causation

Contemporary discussions on the ontology of causation have been dominated by the long-lasting debate between reductionists and anti-reductionists (aka realists).⁵ Despite being rivals, reductionists and realists share an important presupposition: whether causation is reducible, and more broadly, whether the world is Humean, are legitimate theoretical questions.

Causal reductionists deny the ultimate need to postulate irreducible or fundamental causal connections in nature. Causation can and should be reduced to something *non-causal*. Consequently, reductionists prohibit themselves from using any causal-laden concepts (e.g., “manipulation”) in their definitions of causation. There have been various attempted reductive definitions of causation. One attempt is to define causation in terms of probabilistic relevance (Suppes, 1970). The core idea is that an event C is a *prima facie* cause of another event E , if and only if, C raises the probability of E (i.e., $P(E|C) > P(E)$). This simple definition is far from adequate; refinements have been made by Suppes and later philosophers, but it is fair to say we still do not have a tenable probabilistic reduction of causation. Another attempted reduction is the Lewisian-counterfactual approach. Lewis (1973) defines causal dependence in terms of counterfactual dependence: for two (actual) events c and e , c causally depends on e , if and only if, had c not occurred, e would not have occurred. Lewis further analyzes counterfactual dependence in terms of comparative similarities between possible worlds. Unfortunately, this counterfactual analysis of causation is still afflicted by counterexamples despite numerous attempted remedies (Paul & Hall, 2013).

Unsatisfied with these reductionist attempts, some philosophers advocate a realist or primitivist approach to causation. Realists complain that reductionists mistakenly presuppose a “Humean” world in which things and events are merely loosely correlated without genuine connections (Mumford, 2009). Given this Humean ontology_M, realists argue, it is simply impossible to make sense of causation,

ontological commitments, classifications, frameworks, and inquiries in science or philosophy that are Carnapian in nature.

⁵ In this paper, “causal realism” refers specifically to a metaphysically inflationary thesis about causation. Therefore, my rejection of causal realism doesn’t imply that causation cannot be “real” in some deflationary sense.

especially its distinctive modal character. Realists contend that causal claims are committed to something more, namely, the existence of irreducible *causal modalities* of some sort—powers, dispositions, or necessities. Some quotes from causal realists will be useful (emphasis added):

Realism about causation requires two things. First, ... causation is *objective*, meaning that it is something that occurs in an “*external reality*” ... Second, ... causation involves some sort of *necessity* ... (Chakravartty, 2005, p. 8)

According to causal realism, causation is *a fundamental feature of the world*, consisting in the fact that the properties that there are in the world ... are *dispositions* or *powers* to produce certain effects. (Esfeld, 2012, p. 157)

Aristotelian *powers*, we maintain, are part of the *basic ontology of nature* ... Powers are the best way to make sense of familiar methods for inferring and testing causal claims in contemporary science, from physics to economics. (Cartwright & Pemberton, 2013, p. 93)⁶

Similar to Mumford (2009), the above-cited authors also motivate causal realism by rejecting causal reductionism or its Humean ontology_M. For example, according to Chakravartty (2005, p. 7), “causal realism is the view that accounts of causation in terms of mere, regular or probabilistic conjunction are unsatisfactory”; Esfeld (2012) defends his realism about causal-dispositional properties in fundamental physics by rejecting the Humean view of categorical properties. Moreover, many causal realists share optimism about the prospect of a unified and basic causal ontology_M. Although Esfeld chooses fundamental physics as his domain of focus, he does not restrict his causal-realist thesis to

⁶ I agree with Cartwright (1989, esp., 2007) on many things she has said about causation; nevertheless, I think Cartwright and Pemberton (2013) have gone too far in making unnecessary metaphysical commitments to Aristotelian powers. I will say more about this in section 4.

physics. Instead, Esfeld (p. 167) claims that causal realism provides “a unified ontology_[M]” for both fundamental physics and special sciences such as biology. Similarly, Cartwright and Pemberton (2013) also promise us a “basic ontology_[M] of nature” that can be used to make sense of causation across various scientific domains. This is not a coincidence. According to causal realism, causation is objective and framework-independent: a realist ontology_M of causation should not be contingent on the domain or context in which the ontology_M is applied, nor should it be relativized to a particular framework. Although the promise of a unified and basic causal ontology_M may sound attractive, it can cut both ways, as we shall see later.

From a Carnapian perspective, both reductionists and realists have gone the wrong way. They both take questions on the existence of causation as theoretical questions about the alleged “objective” and “external” reality of causation. However, these are simply pseudoquestions, says the Carnapian. The Carnapian offers a pragmatist alternative to causation, according to which an ontological commitment to causation means nothing more than accepting a framework of causation that proves fruitful, expedient, and conducive to our goals in daily life or science.⁷

Carnap’s (1956/1988) meta-ontology_C has sometimes been confused with either metaphysical nominalism or realism. To avoid similar confusion about my view, I shall explain in more detail how the above three approaches to causation differ from each other. The divergence between causal reductionism and causal pragmatism lies in the fact that the former asserts that there are no genuine causal connections in nature, whereas the latter refuses to make such a strong metaphysical claim. Importantly, the Carnapian does not deny the possibility that in the distant future, we might obtain a reductive framework of

⁷ Readers may be interested in what Carnap himself thinks about the ontology of causation. Unfortunately, he was somewhat equivocal on this. On the one hand, Carnap (1966, p. 201) was sympathetic to reductionism: “A statement about a causal relation ... describes an observed regularity of nature, nothing more.” Therefore, it seems that here, even Carnap failed to resist the temptation to make a metaphysical (i.e., a Humean) claim about causation. On the other hand, Carnap also acknowledged that “I do not deny the possibility of introducing a necessity concept, provided it is not a metaphysical concept but is a concept within the logic of [causal] modalities” (p. 208). The idea expressed in this latter quote is clearly a Carnapian one: if a framework needs to postulate causal necessities to fulfill our goals, we are justified to admit them to our ontology_C.

causation that proves fruitful and expedient; such a possibility, however tiny it is, should not be ruled out a priori.

The disagreement between realism and Carnapian-pragmatism is more subtle. Both the realist and the Carnapian have an open-minded attitude towards ontological commitments. The Carnapian is happy to acknowledge that causal and modal concepts play an important role in daily life and science, and at least for now, it is neither feasible nor necessary to reduce or eliminate them. For this reason, it is pragmatically justified to posit irreducible causal relations in our ontologies_C. On the other hand, unlike realists, the Carnapian refuses to make commitments to any sort of framework-independent “external reality” of causation or modality. Therefore, although a realist and a Carnapian may appear to accept the same system of ontological commitments in the same context, they disagree on the nature of these commitments. For the Carnapian, the existence of a system of ontological postulates is *relative to* a framework that is accepted and used in a particular context; moreover, the acceptance such a framework is goal-dependent. If it is found later that we have good reasons to adopt another framework with a different ontology_C, the Carnapian will not hesitate to revise their ontological commitments accordingly.

Realists may contend that if we agree that causal notions are indispensable in formulating many true claims or approximately true theories about the world, this should count as evidence for the external reality of causation. This is an instance of the indispensability argument. However, I do not think the argument succeeds since it begs the question by assuming that assertions about the external reality of causation have truth values.⁸ If we pause and think from a Carnapian perspective instead, we will immediately realize that such assertions about external reality have no truth values to begin with, and hence, there is no question of “evidence” for them.

Indeed, a key insight of the Carnapian meta-ontology_C is the recognition that the relationship between our conceptual frameworks and the world is rather intricate. It is unlikely that there is a simple

⁸ Here “truth” is understood in terms of some sort of correspondence relation between assertions and (some portion of) external reality. I assume that causal realists, and more generally speaking, metaphysical realists, typically adopt a *correspondence* theory of truth which says that truth consists in some kind of correspondence between statements/propositions/frameworks and external reality (see David, 2022).

and straightforward correspondence between (fragments of) conceptual frameworks and (portions of) external reality that necessitates the “truth” of some fundamental ontology_M of the world. For this reason, it is too hasty to posit powers or necessities as “fundamental features” of the world, claiming that they constitute the “true” ontology_M of causation. Although we humans find frameworks with causal and dispositional concepts useful, this does not exclude the possibility that intelligent beings with different cognitive capacities or interests may prefer conceptual frameworks with different causal vocabularies or rules (or even without causal concepts at all). It is highly plausible that causal ontologies_C accepted by us, especially in ordinary languages and high-level sciences, are contingent upon features of ourselves.

3. The Carnapian-pragmatist approach to causation

To further explain how this Carnapian-pragmatist approach works, we need to take a closer look at some basic elements of a Carnapian-pragmatist ontology_C:

- *Frameworks*: Linguistic systems (e.g., a causal inference framework) that provide necessary terms and rules for formulating ontological commitments and principles.
- *Users*: Agents who choose and use frameworks (e.g., us).
- *Context*: A context or domain in which the chosen framework is intended to be used (e.g., epidemiology).
- *Goals*: Goals and criteria based on which a framework is accepted or rejected (e.g., fruitful causal control).

3.1 Users, context, and goals

Carnap (1956/1988) does not explicitly mention users and context, but their relevance is evident. The decision to accept or reject a certain framework must be made by agents like us, who are looking for a framework to use. Moreover, the decision is always made relative to a certain context or domain since the goals or criteria based on which the decision is made may vary from context to context. I take these points to be uncontroversial. Nevertheless, there is a separate controversy over whether (human) agency is *constitutive* of the notion of causation, which I will discuss briefly below.

Agency theorists, such as Price (2017), believe that causation is conceptually tied up with human agency, whereas his opponents (notably, Woodward) deny that human agency plays any special role in our notion of causation. I believe the Carnapian can remain neutral on this issue. To see why, note that the role of users or agents in framework choice is a generic one that is not unique to causal frameworks. Users also play a role, for example, when mathematicians choose a framework of numbers or when biologists choose a framework of species. This means the kind of pragmatism I advocate here does not lend direct support to Price’s agency theory of causation; Carnapian pragmatism, as a meta-ontological doctrine, does not directly imply that reference to agency within a causal framework is necessary for the framework to be fruitful and expedient. Of course, agency theorists may want to further contend that reference to “agency” is necessary for a causal framework to be useful at all; but this should be distinguished from my point about the generic role of agency in framework choice.

Turning to the goals in causal framework choice, there are a few of them that are often considered important by the Carnapian:

- *Prediction, explanation, and control* (i.e., *methodological goals*): A causal framework can be used to obtain predictions, explanations, or control of the world.
- *Explication*: A causal framework can provide an explication of causal notions used in other causal frameworks, hoping to make the latter more precise.⁹

Not all these goals need to be considered when evaluating a framework. Methodological goals are important in assessing causal frameworks in the sciences. Explication as a goal is particularly important for evaluating causal frameworks in philosophy, although sometimes it can also be of direct relevance in scientific practice. It is worth repeating that for the Carnapian, the “truth” of a whole framework is not a legitimate goal for framework choice.

Given the goals we would like to achieve, we can evaluate a causal framework based on how fruitfully and expediently it can help us achieve these goals. There is often a trade-off between

⁹ According to Carnap (1962, p. 3), “explication consists in transforming a given more or less inexact concept into an exact one.” See also Fischer (2023) for how Carnapian explications work in the context of actual causation.

fruitfulness and expediency in the evaluation (Carnap, 1956/1988): since fruitful frameworks also tend to be more complicated and cumbersome, sometimes we may need to sacrifice fruitfulness for expediency or vice versa. Take the ordinary-language framework of causation as an example. Its main goal is to let us efficiently accomplish daily tasks of causal prediction and control in a rapidly changing environment, and for this goal, expediency matters more. This is why we have not and probably never will abandon the ordinary-language causal framework even if it is far less fruitful in prediction and control than its scientific alternatives.

What about the objectivity of causation? Readers may have the following concern about my Carnapian approach to causation: if the Carnapian thinks the acceptance of a causal framework and its ontology_C is a matter of choice based on pragmatic considerations, it seems what causes what is *up to us!* That would compromise the objectivity of causation; some might say. What causes what should not be a subjective matter; as Chakravartty (2005) points out, causation should be *objective* in the sense that it is mind-independent, which presumably implies that the reality of causation should not depend on framework choice.

In response, the Carnapian suggests a more fine-grained conception of objectivity in place of the simple dichotomy of subjectivity and objectivity. First, note that in the Carnapian approach, one does not get to decide the truth of a statement *within* the chosen framework. This means, for the Carnapian, there is an important sense in which causation *is* objective: within a chosen causal framework, whether there is a causal relation between two given variables becomes an internal question that needs to be answered by empirical means. That is, once we have accepted a particular causal framework with causal variables *X* and *Y* (e.g., smoking and lung cancer), whether *X* causes *Y* is then a matter of objective fact.¹⁰

¹⁰ A similar point has also been made by Woodward (2007, p. 90): “Relative to a specification of system and a level of description or graining for it ... one fixes the variables one is talking about, it is [an] ‘objective’ matter whether and how [the variables] are causally related”, and by Eagle (2007, p. 167): “even if the variable and their ranges are chosen for pragmatic and context-sensitive reasons, the truth of the resulting counterfactuals will be a perfectly objective feature of those variables.”

Second, although our choice of causal frameworks is goal-dependent and context-sensitive, this is not to say the choice is arbitrary or subjective. The world still plays a crucial role in regulating our choices, even if it does not necessitate the “truth” of any single framework. For example, in medicine, we choose causal frameworks that can enable us to achieve better control of public health; the fact that these frameworks deliver better control of the world suggests that they must be somehow “latched onto” the world. Therefore, the fact that causal frameworks and the ontologies_c of these frameworks have pragmatic dimensions does not prevent causation from being “objective” in important ways. It is just that the Carnapian has a weaker and less absolute notion of “objectivity” than the metaphysician.

For the metaphysician, saying that “causation is objective” implies that causation is some sort of mind-independent “external reality” (see, e.g., Chakravartty, 2005). The Carnapian wants to resist this overly strong conception of objectivity: from the Carnapian perspective, metaphysicians have overemphasized the “worldly” side of causation and neglected its “human” side. I do not deny that causation must be in some sense grounded in the external world. As Weinberger, Williams, and Woodward (2023) point out, our success in causal inference has support in the “worldly infrastructure” of causation.¹¹ However, this worldly infrastructure should go hand in hand with the conceptual framework we choose to frame the infrastructure. There is no such thing as a “causal relationship/structure/process” unless we have chosen a conceptual framework to frame certain aspects of the world *as* a causal relationship/structure/process.

3.2 The plurality of causal frameworks

I have been using the notion of a “(causal) framework” in a quite flexible manner; a causal framework need not be a general, axiomatic, or mature theory of causation. From a Carnapian perspective, causal frameworks will necessarily be *domain-specific*, since different domains may require different

¹¹ What is this infrastructure? Weinberger et al. (p. 4) write: “there are certain generic features of our world that license and support the application of causal thinking and inferences to causal conclusions ... (i) some variables are statistically independent of others ... (ii) interventions, in the sense of unconfounded manipulations, are often possible ... (iii) the macroscopic, coarse-grained behavior of many systems is largely independent of variations in their microscopic realizing details ...”

causal frameworks. In addition, I use the notion in a relatively fine-grained way: two causal frameworks count as different as long as they differ in their ontological commitments or principles such that a choice between them is deemed necessary in a certain context. In fact, two different causal frameworks may even postulate seemingly incompatible ontologies_C (e.g., omissions are causes within some philosophical frameworks of causation but not in others). This need not necessarily be a problem: these “incompatible” frameworks may very well be accepted in different contexts or be used for different purposes. Of course, problems may arise if one attempts to unify the two frameworks or to retain only one of them; however, from a Carnapian perspective, one is not obliged to pursue a single, unified framework of causation unless we have reason to believe that doing so will serve our goals better.¹²

This explains why we find a plurality of causal frameworks—and a plurality of causal ontologies_C—in the sciences and philosophy. The Carnapian approach to causation, therefore, offers a philosophical ground for the increasingly popular thesis that causation is in some sense *pluralistic*. Note, however, that the idea of “causal pluralism” itself is a pluralistic one, and the Carnapian does not embrace all versions of the idea.¹³ If causal pluralism is formulated as a metaphysical thesis (as in Hall, 2004), then Carnapian pragmatists will have to reject it. The version of causal pluralism supported by the Carnapian is close to what Hitchcock (2007) calls “methodological pluralism”, but it also bears similarity to Cartwright’s (2007) pluralism about (thick) causal concepts.¹⁴

At the same time, it is no less important to note that the Carnapian also acknowledges that there should be a (perhaps very minimal) *single* core meaning or linguistic function that is shared by plural causal concepts used in different frameworks. After all, there should be a reason why they are all “causal” concepts. It seems to me that this core meaning has to do with the modal or dispositional character of causation, which appears to be shared by all our currently used causal notions (even including domains such as fundamental physics where causal notions differ significantly from those in other domains).

¹² Fischer (2023) reaches a similar conclusion about frameworks of actual causation.

¹³ Hitchcock (2007) helpfully identifies “a plurality of causal pluralisms”; see references therein.

¹⁴ Although I am not opposed to Cartwright’s thesis of causal pluralism, I argue in section 4.1 that her thesis about causal diversity does not square well with Cartwright and Pemberton’s (2013) desire for a unified causal ontology_M.

Therefore, causal realists are certainly onto something when they insist that causation involves some sort of irreducible modalities. (Again, it is crucial to note that the disagreement between a Carnapian and a realist lies in whether we should take a step further and regard these causal modalities as “objective”, framework-independent entities).

Scientific frameworks of causation. Scientists from various domains have developed a variety of causal frameworks. For example, in Newtonian mechanics, we have a framework of causation that postulates deterministic causation: under a well-controlled (e.g., isolated) condition, the microscopic state of a system at a time uniquely determines the system’s microstate at any later time. In high-level sciences (e.g., medicine), however, what we often choose is a *probabilistic* (or pseudo-indeterministic) framework of causation in which causes work probabilistically and statistically (e.g., smoking increases the risk of developing lung cancer). The reason for this choice is obvious: deterministic causes will have to include a lot of minuscule details, making deterministic relationships practically impossible to be formulated or discovered in terms of macroscopic predicates used in special sciences; even if we can, they are too idiosyncratic to be useful for future predictions or control.

Let us look more closely at frameworks of statistical causal inference used in the medical and social sciences, using the *graphical* or *structural causal modelling* (SCM) framework as an example (Pearl, 2009; Spirtes et al., 2000).¹⁵ Typically, the framework uses directed acyclic graphs (DAGs) to represent causal structures of target systems. In a causal DAG, a direct causal relation between two variables is represented as an arrow between two nodes. Direct causation can be defined in terms of Pearl’s (2009) notion of atomic intervention or Woodward’s (2003, p. 98) notion of *ideal intervention*. X is a direct cause of Y (relative to a DAG G) means it is possible to change the value of X through some ideal interventions such that the probability distribution of Y will change accordingly (when all other variables in the graph are held fixed by ideal interventions). A DAG only represents a qualitative causal structure; more detailed information about the causal structure is encoded in the joint probability

¹⁵ Another popular statistical causal inference framework is the potential-outcomes framework (also known as “the Rubin Causal Model”; see Rubin, 1974; Holland, 1986), which is discussed in section 4.3.

distribution over the DAG. A DAG and its corresponding joint probability distribution are assumed to satisfy the causal Markov condition (a modern successor to Reichenbach's principle of common cause).

Why do many scientists accept the SCM framework? The motivation is arguably *not* metaphysical. It is true that in accepting the framework, scientists are committed to the existence of causal structures or data-generating mechanisms underlying observed associations. This is what it means for one to adopt causal methods in place of traditional regression methods. However, this does not mean that scientists have to hold that causation "is a fundamental feature of the world" and that causation must be grounded on an ontology_M of modalities. In contemporary causal inference practice, the metaphysical debate between causal reductionists and realists seems irrelevant; at least to my knowledge, metaphysical arguments for causal realism were not a major reason why practitioners started to embrace causal inference methods. In fact, statisticians have long known that causation is not correlation. Many statisticians had been shying away from causal talk, especially in nonrandomized studies, not because they were reductionists about causation, but because they lacked a rigorous framework for causal talk and a reliable and fruitful tool for statistical causal inference.

The framework of SCMs and other causal methods are gradually being accepted exactly because they could fill in this gap: these methods offer statisticians ways of making sufficiently reliable and unbiased causal inference, even in purely observational studies. Pearl (2009), for example, lays theoretical grounds for using DAGs to address biases in observational causal inference. Relying on graphical rules (e.g., the back-door criterion) derived from the causal Markov condition, Pearl demonstrates that DAGs can be reliable and expedient in adjusting for confounding and other forms of biases, a problem that has haunted statistical scientists for decades.¹⁶ In the context of empirical research, such as medical research, it has been shown that DAGs can do a better job in identifying and adjusting for confounding biases in causal inference, compared to traditional regression-based methods (see, e.g., Shrier & Platt, 2008; cf. Greenland et al., 1999; Suzuki et al., 2020; more recently, Cinelli et al., 2022). Confounding adjustment is

¹⁶ I don't have the space to elaborate on how the backdoor criterion works, but see Pearl (2009, pp. 79-80) and Pearl et al. (2016, p. 61-64) for detailed explanations.

essential for fulfilling the goals of statistical causal inference, such as having effective infection control during a pandemic.

Philosophical frameworks of causation. There is a variety of frameworks of causation in philosophy as well. The complexity here is that not all philosophers agree with the Carnapian on what the goals of a philosophical framework of causation should be. The Carnapian thinks the primary and distinctive goal of a philosophical framework of causation is to offer *explications* of causal notions used in other domains. Of course, philosophical frameworks of causation can contribute to methodological goals as well. Such a contribution can be achieved by directly engaging with methodological studies on causation, but more often philosophers contribute to causal methods through explications of (imprecise) causal notions used by scientists. For example, Woodward’s (2003) interventionist framework of causation can be seen as an explication of causal notions used in biomedical and social sciences (especially the notion of intervention or manipulation). Through explication, his work makes an indirect contribution to causal methodology (see also Woodward, 2015, for examples of how such explication works). Here, we see a substantial continuity between philosophical and scientific frameworks of causation, which I consider to be a merit.

Explication can also be a goal of non-Carnapian metaphysical frameworks of causation such as Lewis’s counterfactual theory. To be fair, Lewis’s theory can do a good job in explicating how causal notions are used in a range of ordinary-life cases; however, it is also true that the theory has not been very fruitful for explicating causal notions used in the sciences, mainly because Lewis’s possible-world semantics of causal counterfactuals lacks empirical testability and hence has limited scientific relevance (especially if we compare it with Woodward’s interventionist semantics of counterfactuals). It is unclear how talk of similarity between “possible worlds”—whatever the ontological status of these “worlds” are—can help us interpret and understand the results of observational causal studies or clinical trials.

Metaphysicians such as Lewis would likely reject my claim that explication is the primary goal of a philosophical framework of causation. What they have in mind is a more ambitious *metaphysical* goal; namely, a philosophical framework of causation should solve “deep” metaphysical questions about

causation, such as whether the world is truly Humean, whether there are really powers... Explication is a useful tool for answering these questions, but it is not the ultimate goal we should be after. Lewis, for instance, would certainly not be satisfied with his counterfactual analysis of causation being seen merely as an explication of the ordinary notion of “cause”. What he tries to get at through the explication is the following deeper metaphysical truth: namely, the notion of causation is not really primitive even if it appears to be so. The Carnapian, of course, is strongly opposed to this metaphysical approach. In the next section, I will explain why I think the metaphysician’s approach to causation is flawed and then present some arguments in favour of the Carnapian approach.

Philosophers are all very familiar with the Quinean meta-ontological maxim that “the world contains whatever our *best scientific theory* (or theories) says it contains”. This maxim is in an important sense on par with Carnapian pragmatism: both conceive ontological commitments in a language-dependent way. Nevertheless, the Quinean maxim is far stricter on what kind of language frameworks are eligible for determining our ontological commitments; unlike the Quinean, the Carnapian does not require frameworks to be our best theories in order for them to have ontological implications. An immediate difficulty with the Quinean approach is that it is not clear at all whether we have “theories” in many special sciences (e.g., epidemiology and econometrics), not to mention the “best” one(s). What we have in these areas are often piecemeal models used to solve specific causal questions. Due to causal heterogeneity, it is usually impossible and counterproductive to integrate these models into a unified theory with a consistent set of ontological commitments.¹⁷ Now, one may contend that “the best theories” should be confined to physical theories; high-level sciences are simply not among “our best scientific theories”. This would lead to the consequence that ontological commitments (to high-level causal relations or mechanisms) in these high-level sciences are not to be taken seriously; I do not think this is a good way of interpreting the practices of special sciences. Fortunately, this problem does not bother the

¹⁷ An analogous problem is also noted in Danks (2015). Danks shows that depending on our goals, we may need different models with incompatible ontological commitments even about the same target system.

Carnapian who sets a much lower bar on when ontological commitments of a framework can be justified and accepted.¹⁸

4. Ontology and methodology: Motivating the Carnapian approach

While I believe there are good reasons for favouring the Carnapian approach to causation over metaphysical approaches, I must admit that I do not have knockdown or conclusive arguments against the latter. This may be disappointing but is hardly surprising. As we have seen in the last two sections, pragmatists and metaphysicians disagree on so many things (especially on the goals of ontological inquiry) that probably any argument favouring one will rely on premises or intuitions that the other finds unjustified.¹⁹ For this reason, the “arguments” I shall provide in this section are not aimed at settling the debate or disagreement; instead, the goal is to stimulate further exchanges between pragmatists and metaphysicians.

4.1 What is wrong with the metaphysical approach?

The Carnapian and the metaphysician have very different understandings of the relationship between ontology and methodology. For metaphysicians, the study of ontology_M (i.e., metaphysics) is a distinct project from that of methodology or epistemology: ontological issues are concerned with what

¹⁸ In a broadly Quinean spirit, Schurz and Gebharter (2016) argue that causation is a “theoretical concept” explicated by “axioms” of the “theory” of causal Bayes nets (aka SCMs). My project and theirs share important common grounds, but I disagree that the SCM framework alone is sufficient to explicate the concept (or rather *concepts*) of causation. In making this claim, they seem to assume that SCMs can capture everything important about causation and that it is the single best causal framework we have. There are good reasons to doubt these assumptions. First, the framework of SCM is merely concerned with the “thin” concept of causation. But there are plenty of “thick” causal concepts used in the sciences that the formalism simply cannot capture (cf. Cartwright, 2007). For example, in classical mechanics, the detailed dynamics of a system’s evolution needs to be described using Hamiltonian equations which we do not find in SCMs. Additionally, axioms of SCMs presuppose idealized assumptions (such as modularity) that may be violated in biological sciences. Moreover, even in domains where SCMs plays a significant role, causal notions used there are often richer and messier than what the axioms of SCMs can tell us. For example, actual interventions conducted in clinical trials are much more complicated than the kind of atomic or ideal interventions assumed in the SCM framework. Additionally, the formalism of SCM does not capture other peripheral causal notions like invariance, proportionality, and stability, which play an indispensable role in causal inference practice.

¹⁹ I am not the only one who is skeptical about the possibility of having conclusive arguments against the metaphysician. Woodward (2017, p. 193), for example, found that “putting everything into an ordinary ‘linear’ argument [against the metaphysician] was impossible”; so, he organized his article as an imaginary dialogue with “Professor Metafisico”.

exists in the world, whereas methodological issues are concerned with how to obtain *knowledge* about things that exist and how to *use* the knowledge. Ontological and methodological inquiries are of different nature: they have distinct goals, distinct methods, and most evidently, distinct jargon; for this reason, the former cannot be assimilated to or replaced by the latter. This is not to say they are unrelated, but it does imply that ontology_M is an autonomous enterprise, and it is, in important ways, free of interference from methodological considerations.

According to the metaphysician, ontology_M enjoys some kind of *priority* over methodology. For example, Anjum and Mumford (2018, pp. 25-26) claim that “once a distinction is established, between ontology_[M] and epistemology, there can also be a question of priority”, and they think “ontology_[M] has priority”. In a similar spirit, Paul (2012, pp. 5-6) claims that “many concepts of metaphysics are conceptually prior to the concepts of science” and that “metaphysics involves questions about features of the world that are prior to those described by science.” Woodward (2017, p. 195) also notices that “what we too often get [from metaphysicians] are claims that work on [ontology_M] is logically required if one works on [methodology], that the former is ‘deeper,’ and so on.” To be clear, the priority of ontology_M should be *logical* or conceptual: logically speaking, the existence of X is not contingent on what method we use to discover X, whereas asymmetrically, the validity of methods for discovering X depends on the ontology_M of X. This kind of logical order between ontology_M and methodology does not imply an epistemic order; metaphysicians need not deny that methodologists can reliably discover X in ignorance of the “true ontology_M” of X. What the logical order does imply is that methodological successes should not be allowed to directly enter metaphysical theorization or justification. This is why metaphysicians will deny that we can directly infer the true ontology_M of the world from the methods we use to obtain knowledge of the world.

Applying the above picture to causation, it follows that causal ontology_M and causal methodology are distinct research projects, and the former has a (logical) priority over the latter. In this picture, methodological issues about causation do not directly concern metaphysicians, nor can the fruitfulness of a causal method be cited as direct evidence for an ontology_M of causation since it is possible that fruitful

causal methods may have no implication about the “true nature” of causation. For example, even if metaphysicians are willing to acknowledge that Woodward’s (2003) interventionist framework of causation is useful in explicating how causal notions are used in methodological practices, they will refuse to admit that this reveals any genuine insight about the ontology_M of causation.

There is something deeply problematic with this metaphysical picture. By assuming that causal ontology_M has priority over causal methodology, the picture encourages metaphysicians to focus their attention on ontological questions that are posed in a highly *domain-general*, *context-insensitive*, and *goal-independent* way (whereas, in contrast, methodological issues are typically domain-specific, context-sensitive, and goal-dependent). Indeed, metaphysical debates on causation tend to strip away the (scientific or methodological) context of a legitimate ontological question and transform it into an *abstract* metaphysical question.²⁰ It is commonly assumed by metaphysicians that such abstract metaphysical questions can be solved with methodological issues being completely set aside. Moreover, many metaphysicians of causation, especially reductionists, explicitly reject that causal methods have any implications for seeking the true causal ontology_M. For those realists who do care about causal methods, many of them fail to give a satisfactory depiction of how ontological and methodological issues (should) interact with each other in good scientific practice due to their faulty presumption about the priority of some “basic” causal ontology_M. Let me expand on the latter two points.

Reductionists typically take the reduction of causation as a separate and deeper issue than methodological issues concerning causal inference; consequently, they are not motivated to engage with ideas or discussions from the methodological literature. This explains why the reductionist project, especially the Lewisian approach, is relatively isolated from the methodological literature on causal

²⁰ E.g., consider downward causation in science (e.g., biology). The existence of downward causation is often accepted without questioning by biologists due to the evident usefulness of this ontological postulate; what biologists are interested in are typically internal questions about downward causation. Metaphysicians, however, are interested in the following metaphysical question: “Is there *really* downward causation?” Surprisingly, the metaphysician Kim (1993) has “compellingly” argued—in the sense that the premises of the argument are widely accepted among metaphysicians—that downward causation is impossible! This sharp discontinuity between ontological postulations in scientific practice and our metaphysics seems baffling and counterintuitive, to say the least.

inference, even if Lewis' approach to causation and the potential-outcomes approach to causal inference (Rubin, 1974; Holland, 1986) share similar conceptual grounds (e.g., both use counterfactuals to define token causation). More importantly, since thus far none of our best causal methods (including the SCM and potential-outcomes frameworks) has attempted to reduce causation to something non-causal, the methodological literature on causation bears little relevance to the reductionist project (see Paul & Hall, 2013). For these reasons, it is hardly surprising that the reductionist project contributes little to causal methodology or the explication of causal notions used in the sciences, which I take to be an undesirable consequence. Given this situation, also considering the fact that we are far away from having a reductive definition of causation, one may wonder whether it is worth the effort to continue pursuing the reductionist ideal.

Is causal realism better off? In an important sense, it is. Taking causal notions as primitive, causal realism does not run afoul of our currently accepted causal methods in which causation is assumed to be primitive. Some causal realists have explicitly argued that their realism thesis is motivated by the fact that our best methods presuppose primitive causal notions. In particular, Cartwright (2007, p. 132) claims that “if our metaphysical account does not tie in with our best methods for finding out, we should be suspicious of our metaphysics”. Until now, I am in total agreement with Cartwright. The main divergence between Carnapian-pragmatists and realists lies in the fact that causal realists (e.g., Cartwright & Pemberton, 2013; Anjum & Mumford, 2018) go one step *further* and claim that their powers ontology_M reveals the “true” nature of the world. For these causal realists, there is such a thing as a “basic ontology_M” of causation that is logically and metaphysically prior to methodological frameworks of causal inference. Note that the word “basic” here carries strong metaphysical commitments: powers are not just basic within a particular domain or framework but are basic *simpliciter*. For this reason, such a “basic ontology_M” of causation is believed to be able to serve as a foundation for a variety of causal methods.

For example, Cartwright and Pemberton (2013) start by pointing out that the concept of “powers” is indispensable in “making sense of contemporary scientific practice”; a pragmatist could totally agree with this claim as long as the notion of “powers” is not yet reified as a metaphysical concept. The

conclusion they draw from the indispensability of powers, however, is a metaphysical one. As Cartwright and Pemberton (p. 93) say, “Aristotelian powers ... are part of the basic ontology of nature.” By emphasizing that they are talking about “*Aristotelian* powers”, I believe they intend to make a metaphysical claim about powers.²¹ If this is the case, this must be the place in which pragmatists part ways with powers realists. For pragmatists, acknowledging the indispensability of causal-modal notions in scientific methods does not necessitate a metaphysical commitment to powers.

Moreover, Cartwright and Pemberton (p. 93; emphasis added) claim that “Powers are the best way to make sense of familiar methods for inferring and testing causal claims in contemporary science, *from physics to economics*”. The idea of having a domain-general and basic ontology_M of causation that serves as a foundation for a variety of causal methods is certainly attractive; however, I remain skeptical about this move due to an inherent tension within the idea. On the one hand, this basic ontology_M should not be too *specific* or restrictive on what causation is. If it includes any informative principles or even domain-specific constraints about causation (e.g., causation must be deterministic or temporarily asymmetric), it will fail to accommodate a variety of causal frameworks used in the sciences “from physics to economics” (e.g., general relativity, neurobiology, and econometrics). In fact, Cartwright (2007) emphasizes the diversity of causation in different domains and proposes a version of causal pluralism based on the diversity of thick causal concepts. According to her, the abstract and domain-general notion of causation is a “thin” concept that carries very little detail about the meaning of “thick” causal concepts in various specific domains.²² Note that the Carnapian has no objection to Cartwright on this; as I said earlier in section 3.2, the Carnapian supports a version of causal pluralism that bears similarity to Cartwright’s thesis about causal diversity.

²¹ I am not saying it is incoherent to be a pragmatist Aristotelian, but I doubt that anyone who are sympathetic to a pragmatist meta-ontology_C would label themselves an Aristotelian without adding any caveat.

²² Cartwright (2007, p. 19; emphasis added) writes: “We think of causation as a single monolithic concept. But that is a mistake ... there is no single thing of much detail that [causal laws] all have in common, something they share that makes them all causal laws.”

On the other hand, if the powers ontology_M merely captures a “thin”, detail-lacking notion of causation, then I fail to see how it can be of help in making sense of various causal methods that often invoke much thicker causal concepts. In other words, if the realist’s basic causal ontology_M imposes no informative or substantive constraints on causation, how can it provide a good *explication* of causal methods in a variety of domains? I do not see how it can. Suppose we agree with Carnap (1950) that an important desideratum for explication is the similarity of the explicatum to the explicandum.²³ However, if we accept Cartwright’s thesis of causal diversity, then it is not true that causal notions in physics and economics can share the same explicatum. As a result, the realist must admit that the generic and thin concept of powers cannot provide a good explication of causal notions in both physics and economics; otherwise, they must give up the thesis of causal diversity (which I do not think they should). Now, if the former is true, that is, if powers cannot provide a good explication of plural causal notions, then how on earth powers are the *best* way to make sense of causation in the sciences from physics to economics? The burden of proof now lies on realists such as Cartwright and Pemberton (2013). Note that the mere fact that a detail-lacking powers ontology_M is *compatible* with a variety of causal methods does not make it the best way to make sense of these methods.

4.2 The pragmatist alternative

As I have shown, the main problem with the metaphysician’s approach to causation lies in granting a special kind of priority to the “true” causal ontology_M. With ontology_M being prioritized over methodology, it is only natural to expect that the schism between them will cut deep. I believe opting for a pragmatist meta-ontology_C may help restore the “balance” between ontology and methodology. The Carnapian denies that we need a causal ontology_M that is in some sense “basic” and serves as a “foundation” for causal methods. From the Carnapian perspective, legitimate ontological questions always arise within a certain methodological context and should not be completely abstracted away from

²³ “The explicatum is to be similar to the explicandum in such a way that, in most cases in which the explicandum has been so far used, the explicatum can be used” (p. 7).

the context. This is also why, as we have seen in section 3, in scientific practice, what we need is a variety of causal ontologies_C presupposed by a variety of causal methods.

The relationship between causal ontology_C (or ontologies_C) and causal methodology in this Carnapian picture is not a matter of priority but a matter of balance: in (good) scientific practice, there is often mutual dependence or a (moving) *state of balance* between causal ontology_C and causal methodology. On the one hand, a causal method must presuppose a causal ontology_C. A methodological framework of causation is always committed to the existence of a system of causal dispositions, causal relations, or causal processes. On the other hand, granted that our causal ontology_C regulates and warrants our causal methodology, this does not mean we cannot adjust our causal ontology_C or choose between alternative ontologies_C in response to methodological considerations. In this picture, causal ontology_C does not enjoy any special kind of logical priority or fundamentality over methodology.

Such a state of balance between ontology_C and methodology provides a good characterization of the interplay between the two in the sciences. Scientists themselves often have in their minds a similar conception of ontology_C; for a quick example, just consider how our cognitive ontologies_C had undergone revisions (from rejecting mental causes to embracing them again) in response to methodological progress in psychological sciences since the 1910s. Indeed, it is commonly held by scientists that in scientific practice, if introducing a new (set of) ontological postulate(s) to our causal method will make it more fruitful, expedient, and conducive to our goals, we are obliged to do so. If a postulate is no longer needed in our new methods, we can simply remove it from our ontology_C. Ideally speaking, in the end, a balance can be achieved between ontology_C and methodology. A stabilized causal ontology_C in a certain domain will be practically useful since it frees later scientists from worrying about ontological issues. Of course, as science progresses, scientists may need to adjust their ontologies_C based on newly emerged methodological considerations.

To sum up: by emphasizing the “balance” between ontology_C and methodology in an ongoing process of scientific inquiry, the Carnapian approach provides a much-improved picture of how the two interact with each other. Of course, things are not always so ideal and simple as I have said here. In some

cases, for instance, when we try to merge two causal methods with conflicting causal ontologies_C, it may be difficult to find a balance between ontology_C and methodology. The nuances and complexities in the process toward such a balance should therefore be appreciated. Let us now consider a case study.

4.3 “What can be a cause?” A case study

Although less known among philosophers, the potential-outcomes framework for causal inference (Rubin, 1974; Holland, 1986) is as widely used in health and social sciences as the SCM framework. The two frameworks, however, originated from very different traditions; the potential-outcomes framework was first proposed by Neyman (1923/1990) in the context of randomized field experiments in agriculture (and was later extended to nonrandomized studies by Rubin), whereas the SCM framework was first proposed around the 1920s by Wright (1920) in the context of genetics. The difference in historical contexts has also led to more substantial differences between the two frameworks. The most salient difference is that the potential-outcomes framework builds its conceptual basis on the idea of a (hypothetical) randomized controlled trial (RCT), whereas the SCM framework is more concerned with the causal structure or mechanism underlying a causal query. Indeed, the primary contribution of Rubin (1974) is to demonstrate that we can estimate causal effects from nonrandomized data if the data can be carefully controlled such that “as if the study had been randomized” (p. 698).

In an influential article, Holland (1986) restates the conceptual basis of the potential-outcomes framework on (hypothetical) experimentation. “What can be a cause?” is a question asked by Holland against this theoretical context.²⁴ His answer is that “causes are only those things that could, in principle, be treatments in experiments” (p. 954). Therefore, for Holland, the scope of causes is rather narrow: biological or social characteristics of objects or individuals (what he calls “*attributes*”) are not causes since we cannot (hypothetically) assign these characteristics as treatments in randomized experiments. For example, attributes such as scholastic performance, obesity, and race are not causes because it is

²⁴ The question, especially of whether attributes or characteristics should be seen as causes in causal inference, has stimulated much debate in both philosophy and science (Holland, 2003; Woodward, 2003, 2016; Greiner & Rubin, 2011; Sen & Wasow, 2016). My discussion below will unavoidably be brief and oversimplified.

impossible, even in principle, to randomly assign these attributes as treatments in (hypothetical) RCTs.²⁵ This might sound surprising to philosophers, but Holland’s position, or some more nuanced version of his position, has been widely accepted among the potential-outcomes community (see, e.g., Hernán & Taubman, 2008; Greiner & Rubin, 2011).

The SCM community, including philosophers who favour the SCM framework, tends to reject Holland’s restriction of causes to treatments. For example, Pearl (2018) argues that physical, biological, and social attributes can be causes (see also Glymour & Glymour, 2014). To support his claim, Pearl distinguishes between *knowledge*-based and *policy*-based causal statements. In Holland’s logic, diet can be a cause of death since it can be assigned as a treatment. This type of causal statements has an immediate implication for policymaking since they can offer clear guidance on how to control the number of deaths by changing diets. This is what Pearl calls policy-based causation. However, Pearl emphasizes that scientists also have an epistemic interest in causal mechanisms and causal explanations, such as “obesity causes deaths”, even if these causal statements may be much more ambiguous than policy-based causal statements (as argued by Hernán & Taubman, 2008). For Pearl, attributes like obesity can legitimately be seen as causes in a knowledge-based causal framework.

My goal here is not to settle this complicated controversy but to argue that we should make sense of the controversy from a Carnapian-pragmatist perspective. Although “what can be a cause?” or whether attributes can be causes is arguably an *ontological* question, neither Holland nor Pearl takes it as an abstract metaphysical question and turns to metaphysics for help. In fact, it is doubtful that metaphysical discussions on causation can illuminate this controversy. For example, Lewis would probably suggest that only (perfectly) natural properties (cf. Lewis, 1983) can be causes. However, for reasons that have been pointed out in Woodward (2016), this suggestion is not helpful for answering “what can be a cause” in a way that can be of service to scientific practice. Realists cannot offer much help on this question either

²⁵ Note that you can of course change these attributes indirectly by assigning treatments to an individual. For example, you can improve a student’s scholastic performance by helping them with their homework. But this is not the same thing as assigning high scholastic performance as a treatment to that student. What is assigned as treatment here is homework help, not high scholastic performance.

since their “basic” powers ontology_M is not sufficiently informative on this rather domain-specific issue (see section 4.1). Are obesity and race causes, according to the realist’s powers ontology_M? Do they have “real” causal powers? It seems that either the realist has to admit that the powers ontology_M cannot directly answer this type of domain-specific questions, or they need to deny that this is a fair question.

In my view, Carnapian pragmatism offers a more sensible way to make sense of this controversy. Holland’s answer to “what can be a cause” is no longer so surprising if we notice that it is a direct consequence of the causal framework he has accepted. He restricts the scope of causes to treatments because the potential-outcomes framework he has adopted takes randomized experiments as the paradigm of causal inference. In RCTs, the causes being considered are, by default, treatments. Now, Holland does have good reasons to accept the potential-outcomes framework. First, randomized experiments have been considered one of the most reliable methods of causal inference. Based on the idea of an RCT, the potential-outcomes framework provides a conceptually clear and practically convenient definition of causal effects. In particular, as Hernán and Taubman (2008) argue, if a variable cannot be conceived as a treatment variable in a randomized experiment, its effects on other variables may be ill-defined. Besides, the primary goal of the potential-outcomes framework is to inform public policy in medical and social sciences. The framework is a suitable tool for achieving this goal. In view of all this, it starts to sound reasonable that causes should be restricted to treatments.

Pearl has a quite different answer to “what can be a cause” given his acceptance of another causal framework, namely, SCMs. Causes are not, and cannot, be restricted to treatments in this framework since to represent the underlying causal structure or mechanism of a target system with a graphical model, we often need to include various attributes as causes. The SCM framework is hence in sharp contrast to the potential-outcomes framework in the sense that the latter takes the target system as a black box and only cares about the effects of treatments on the system. But why does Pearl choose the SCM framework? An important motivation is that Pearl’s causal inference framework wants to represent the causal structure of a target system rather than just informing policymaking. This can be seen from Pearl’s (2018) distinction

between policy-based causation and knowledge-based causation I mentioned above. This divergence in goals partly explains why Pearl thinks Holland’s restriction of causes to treatments is too limited.

What makes the problem even more complex and interesting is that, recently, there have been attempts to show that the two causal frameworks are actually theoretically equivalent or at least “productively compatible” concerning their key concepts and assumptions (Weinberger, 2021; Ibeling & Icard, 2023). A natural question that follows from such conciliatory approaches is: what would this imply for the apparent ontological disagreements between the two frameworks on “What can be a cause”? Do such ontological disagreements pose a direct challenge to these conciliatory attempts? Or do such conciliatory attempts (if successful) prove that the ontological conflict between the two frameworks is not as serious as it seems? I tend to think the latter is the more reasonable option: a conciliation in the two frameworks should lead to mitigation of the ontological conflicts between them. After all, in my Carnapian approach to causation, our ontology and our methods should be kept in a state of balance. But it remains to be seen how such a balance can be achieved.

5. Conclusion

According to the Carnapian approach advocated in this paper, the metaphysical debate between causal realists and causal reductionists leads nowhere and the ontological status of causation should be understood from an entirely different perspective. Although the Carnapian acknowledges that causation has its basis in the external world, the worldly infrastructure of causation alone does not constitute the thing we call “causation”. For there to be causation, the infrastructure must be framed by an appropriate linguistic system—a causal framework. Moreover, although the worldly basis of causation is not up to our choices, we do have a say in the frameworks we use to conceptualize the infrastructure. The world does not uniquely determine one single correct way of framing the infrastructure. Importantly, as I have argued in this paper, what counts as an appropriate framing is a pragmatic issue. This implies that causation has a “human face” (Woodward, 2021), so to speak; or put differently, causation is not objective in the absolute sense. Consider, for example, a simple causal DAG, “ $X \rightarrow Y \rightarrow Z$ ” (smoking causes lung cancer which, in

turn, causes death). Does it really make sense to say this DAG refers to some framework-independent causal structure in the world? In fact, the very idea of a “causal structure” already presupposes a conceptual framework that we choose to frame the relevant worldly infrastructure.

It is worth mentioning that the separation of causation into worldly infrastructure and conceptual framing assumed in this paper is just a useful simplification; in fact, the two aspects of causation are so deeply intertwined that in some cases it may be indeterminate whether something (e.g., the causal Markov condition) belongs to the worldly or the human side of causation. This is an intricate issue that could be explored in the future.

References

- Anjum, R. L., & Mumford, S. (2018). *Causation in science and the methods of scientific discovery*. Oxford University Press. <https://doi.org/10.1093/oso/9780198733669.001.0001>
- Antoniou, A. (2021). A pragmatic approach to the ontology of models. *Synthese*, 1-20. <https://doi.org/10.1007/s11229-021-03085-9>
- Blatti, S., & Lapointe, S. (Eds.). (2016). *Ontology after Carnap*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199661985.001.0001>
- Carnap, R. (1956). Empiricism, semantics, and ontology. Reprinted in *Meaning and necessity: A study in semantics and modal logic*, enlarged edition (first published in 1950). University of Chicago Press.
- Carnap, R. (1962). *Logical foundations of probability* (2nd edition). University of Chicago Press.
- Carnap, R. (1966). *Philosophical foundations of physics: An introduction to the philosophy of science*. Basic Books.
- Cartwright, N. (1989). *Nature's capacities and their measurement*. Oxford University Press. <https://doi.org/10.1093/0198235070.001.0001>
- Cartwright, N. (2007). *Hunting causes and using them: Approaches in philosophy and economics*. Cambridge University Press.
- Cartwright, N., & Pemberton, J. (2013). Aristotelian powers. In Groff, R., & Greco, J. (Eds.). *Powers and capacities in philosophy: The new Aristotelianism* (93-112). Routledge.
- Cinelli, C., Forney, A., & Pearl, J. (2022). A crash course in good and bad controls. *Sociological Methods & Research*. <https://doi.org/10.1177/00491241221099552>
- Danks, D. (2015). Goal-dependence in (scientific) ontology. *Synthese*, 192(11), 3601-3616. <https://doi.org/10.1007/s11229-014-0649-1>
- Eagle, A. (2007). Pragmatic causation. In Price, H., & Corry, R. (Eds.). *Causation, physics, and the constitution of reality: Russell's republic revisited*. Oxford University Press.
- Eklund, M. (2009). Carnap and Ontological Pluralism. In D. Chalmers et al. (eds.). *Metametaphysics: New essays on the foundations of ontology*. Oxford University Press.
- Eklund, M. (2013). Carnap's metaontology. *Noûs*, 47(2), 229-249. <https://doi.org/10.1111/j.1468-0068.2011.00830.x>
- Esfeld, M. (2012). Causal realism. In *Probabilities, laws, and structures* (157-168). Springer, Dordrecht. https://doi.org/10.1007/978-94-007-3030-4_11
- Fischer, E. (2023). Actual causation and the challenge of purpose. *Erkenntnis*, 1-21.
- Glymour, C., & Glymour, M. R. (2014). Commentary: Race and sex are causes. *Epidemiology*, 25(4),

- 488–490. <https://doi.org/10.1097/EDE.0000000000000122>
- Greiner, D. J., & Rubin, D. B. (2011). Causal effects of perceived immutable characteristics. *Review of Economics and Statistics*, 93(3), 775-785.
- Greenland, S., Pearl, J., & Robins, J. M. (1999). Causal diagrams for epidemiologic research. *Epidemiology*, 37-48.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and Counterfactuals* (pp. 225-276), MIT Press.
- Hernán, M. A., & Taubman, S. L. (2008). Does obesity shorten life? The importance of well-defined interventions to answer causal questions. *International Journal of Obesity*, 32(3), S8-S14. <https://doi.org/10.1038/ijo.2008.82>
- Hitchcock, C. (2007). How to be a causal pluralist. In P. K. Machamer & G. Wolters (Eds.), *Thinking about causes: From Greek philosophy to Modern physics* (pp. 200-221), University of Pittsburgh Press.
- Hitchcock, C. (2012). Events and times: A case study in means-ends metaphysics. *Philosophical Studies*, 160(1), 79-96. <https://www.jstor.org/stable/23262474>
- Holland, P. W. (1986). Statistics and causal inference. *Journal of American Statistical Association*, 81(396): 945–60. <https://doi.org/10.2307/2289064>
- Holland, P. W. (2003). Causation and race. *ETS Research Report Series*, 2003(1), i-21.
- Ibeling, D., & Icard, T. (2023). Comparing causal frameworks: Potential outcomes, structural models, graphs, and abstractions. *arXiv preprint arXiv:2306.14351*.
- Kim, J. (1993). The non-reductivist's troubles with mental causation. In Heil and Mele (eds.), *Mental causation* (189–210). Oxford University Press.
- Lauer, R. (2022). Motivating a pragmatic approach to naturalized social ontology. *Journal for General Philosophy of Science*, 1-17. <https://doi.org/10.1007/s10838-021-09581-3>
- Lewis, D. (1973). Causation. *Journal of Philosophy*, 73, 556–567.
- Ludwig, D. (2016). Ontological choices and the value-free ideal. *Erkenntnis*, 81(6), 1253-1272. <https://doi.org/10.1007/s10670-015-9793-3>
- Mumford, S. (2009). Causal powers and capacities. In Beebe, H., Hitchcock, C., & Menzies, P. (Eds.), *The Oxford handbook of causation*. Oxford University Press.
- Neyman, J. (1990). On the application of probability theory to agricultural experiments. Essay on principles. Section 9 (D. M. Dabrowska & T. P. Speed, Trans.). *Statistical Science*, 465-472. (Original work published 1923)
- Paul, L. A. (2012). Metaphysics as modeling: the handmaiden's tale. *Philosophical Studies*, 160, 1-29.
- Paul, L. A. & Hall, N. (2013). *Causation: A user's guide*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199673445.001.0001>
- Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd edition). Cambridge University Press. <https://doi.org/10.1017/CBO9780511803161>
- Pearl, J. (2018). Does obesity shorten life? Or is it the soda? On non-manipulable causes. *Journal of Causal Inference*, 6(2). <https://doi.org/10.1515/jci-2018-2001>
- Pearl, J., Glymour, M., & Jewell, N. P. (2016). *Causal inference in statistics: A primer*. John Wiley & Sons.
- Price, H. (2001). Causation in the special sciences: the case for pragmatism. In D. Costantini et al. (eds.), *Stochastic causality*. CSLI Lecture Notes.
- Price, H. (2007). Causal perspectivalism. In Price, H., & Corry, R. (Eds.), *Causation, physics, and the constitution of reality: Russell's republic revisited*. Oxford University Press.
- Price, H. (2009). Metaphysics after Carnap: The ghost who walks. In D. Chalmers et al. (eds.), *Metametaphysics: New essays on the foundations of ontology*. Oxford University Press.
- Price, H. (2017). Causation, intervention and agency—Woodward on Menzies and Price. In Beebe, H., Hitchcock, C., & Price, H. (Eds.), *Making a difference: essays on the philosophy of causation*. Oxford University Press. <https://doi.org/10.1093/oso/9780198746911.003.0005>
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized

- studies. *Journal of Educational Psychology*, 66(5), 688-701. <https://doi.org/10.1002/j.2333-8504.1972.tb00631.x>
- Schurz, G., & Gebharter, A. (2016). Causality as a theoretical concept: Explanatory warrant and empirical content of the theory of causal nets. *Synthese*, 193(4), 1073-1103. <https://doi.org/10.1007/s11229-014-0630-z>
- Sen, M., & Wasow, O. (2016). Race as a bundle of sticks: Designs that estimate effects of seemingly immutable characteristics. *Annual Review of Political Science*, 19, 499-522.
- Shrier, I., & Platt, R. W. (2008). Reducing bias through directed acyclic graphs. *BMC Medical Research Methodology*, 8(1), 1-15.
- Suzuki, E., Shinozaki, T., & Yamamoto, E. (2020). Causal diagrams: pitfalls and tips. *Journal of Epidemiology*, 30(4), 153-162.
- Spirtes, P., C. Glymour, and R. Scheines. (2000). *Causation, prediction and search*. MIT Press.
- Suppes, P. (1970). *A probabilistic theory of causality*. North-Holland Publishing Company.
- Thomasson, A. L. (2014). *Ontology made easy*. Oxford University Press.
- Weinberger, N. (2021). Comparing Rubin and Pearl's causal modelling frameworks: a commentary on Markus (2021). *Economics & Philosophy*, 1-9.
- Weinberger, N., Williams, P., & Woodward, J. (2023). *The Worldly Infrastructure of Causation*. PhilSci-Archive. <http://philsci-archive.pitt.edu/id/eprint/22125>
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.
- Woodward, J. (2007). Causation with a human face. In Price, H., & Corry, R. (Eds.). *Causation, physics, and the constitution of reality: Russell's republic revisited* (66–105). Oxford University Press.
- Woodward, J. (2014). A functional account of causation; or, a defense of the legitimacy of causal thinking by reference to the only standard that matters—Usefulness (as opposed to metaphysics or agreement with intuitive judgment). *Philosophy of Science*, 81(5), 691-713. <https://doi.org/10.1086/678313>
- Woodward, J. (2015). Methodology, ontology, and interventionism. *Synthese*, 192(11), 3577-3599. <https://doi.org/10.1007/s11229-014-0479-1>
- Woodward, J. (2016). The problem of variable choice. *Synthese*, 193(4), 1047-1072.
- Woodward, J. (2017). Interventionism and the missing metaphysics. In Slater, M. H., & Yudell, Z. (Eds.). *Metaphysics and the philosophy of science: New essays*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199363209.003.0010>
- Woodward, J. (2021). *Causation with a human face: Normative theory and descriptive psychology*. Oxford University Press. <https://doi.org/10.1093/oso/9780197585412.001.0001>
- Wright, S. (1920). The relative importance of heredity and environment in determining the piebald pattern of guinea-pigs. *Proceedings of the National Academy of Sciences*, 6, 320-332.