

Dimensions of explanation

Eric Hochstein

University of Victoria

Abstract

Some argue that the term “explanation” in science is ambiguous, referring to at least three distinct concepts: a communicative concept, a representational concept, and an ontic concept. Each is defined in a different way with its own sets of norms and goals, and each of which can apply in contexts where the others do not. In this paper, I argue that such a view is false. Instead, I propose that a scientific explanation is a complex entity that can always be analyzed along a communicative dimension, a representational dimension, and an ontic dimension. But all three are always present within scientific explanations. I highlight what such an account looks like, and the potential problems it faces (namely that a single explanation can appear to have incompatible sets of norms and goals that govern it). I propose a solution to this problem and demonstrate how this account can help to dissolve current disputes in philosophy of science regarding debates between epistemic and ontic accounts of mechanistic explanations in the life sciences.

Keywords

evaluative dimension, communicative concept of explanation, representational concept of explanation, ontic conception of explanation; mechanistic explanation.

At its metaphysical core, what exactly is a scientific explanation? Is it a communicative act whereby we make some phenomenon intelligible or understandable to an audience (e.g. Nancy explains the

orbits of the planets to John)? Is it a type of theory, model, or representation (e.g. *The theory of evolution* explains why elephants have long trunks. Einstein's *General Theory of Relativity* explains why light bends when it passes a large gravitational body)? Or is it a collection of entities and relations out in the world (e.g. The car's velocity and the cut break line explains why the car crashed into the storefront)? Historically, philosophers have disagreed about which of these provides the best metaphysical account of what a scientific explanation is. For example, Peter Achinstein insists that "the illocutionary concept of explaining is fundamental and that other explaining concepts are explicable by reference to this" (1984, p.22). Conversely, Ruben insists that "explanations work only in virtue of the determinative relations that exist in the world" (Ruben, 1990, p.231).

In recent years, some have suggested that perhaps there are simply distinct concepts of "explanation" at work. Each is employed in different contexts within scientific practice, but we must be cautious not to equivocate between them. This idea has strong support. The goals and norms that seem to govern the various concepts of explanation appear to be different, and indeed can even be at odds with one another. Moreover, it seems we can often have one concept of explanation apply in a context where the others are not applicable at all. All this suggests that we may simply mean different things by the word "explanation" in different scientific contexts.

Despite the intuitive strength of such a position, in this paper I argue that this is not the case. Instead, we should think of a scientific explanation as an extremely complex and multifaceted entity. Every explanation can be analyzed along different dimensions. When each dimension of explanation is considered in isolation of the others, it provides a distorted account of what explanations are, and how they

function. But understanding how these different dimensions constitute one and the same explanation can help to clarify current disputes in the philosophy of science.

In order to make this argument, I start in section 1 with an account of why it is plausible to think that there are at least three distinct concepts or definitions of explanation: explanation defined as a *communicative* act (hereafter the “communicative” concept of explanation), explanation defined as an appropriate kind of model or theory (hereafter the “representational” concept of explanation), and explanation defined in terms of ontic structures, causes, and relations in the world that produce and sustain the phenomenon in question (hereafter the “ontic” concept of explanation). In section 2, I demonstrate why we should not consider these as distinct concepts of explanation after all, and why explanation always requires the interaction of all three elements. In section 3, I highlight why this shift from “concepts” of explanations to “dimensions” of explanation is not a trivial one, and explore how this new account can shed light on debates in the philosophy of science. Lastly, in section 4, I highlight the complications that this new account faces, and how they might be addressed.

1. Different Concepts of Scientific Explanation

An important insight that some have emphasized in recent years is that the term “explanation” may itself be ambiguous, denoting entirely distinct concepts. Waskan et al. (2014), for instance, argue that “‘Explanation’ appears to be ambiguous between a representational-artifact, an objective, and a doxastic sense” (Waskan et al., 2014, p.3090). Meanwhile, Craver (2014, p.29) argues that we should “disambiguate four ways of talking about explanation: as a communicative

act, as a representation or text, as a cognitive act, and as an objective structure”. Likewise, Gilpin et al. (2022, p.3) argue that “the term ‘explanation’ is a classic example of what Marvin Minsky referred to as suitcase words: words that contain multiple meanings which are interpreted in different ways for different people in different contexts”.

For the purposes of this paper, I want to focus on three different concepts of explanation, what I will call the “communicative” concept, the “ontic” concept, and the “representational” concept.¹ To begin, consider the communicative concept of explanation. This concept defines explanation as a social activity we engage in whereby we try to make some phenomenon understandable or intelligible to a given audience. As Achinstein (1983, pp.16–17) puts it, “S explains q by uttering u only if S utters u with the intention that his utterance of u render q understandable”. This definition of explanation can also be seen in Craver’s example of Jon, a professor trying to *explain* the action potential of the neuron to a classroom full of students (Craver, 2014, p.30). Similarly, Andrea Woody describes the Communicative concept of explanation in the following way:

There is always, at least implicitly, a request (question) and a response, and these are typically negotiated or shared among individuals or groups. Explanations frequently pass between individuals or groups with different levels of expertise. A person requesting an explanation often is not in a good position

¹ This is not intended to be an exhaustive list, and other concepts of explanation may well be worth including. For instance, both Waskan et al. and Craver acknowledge a *psychological* or *cognitive* concept of explanation. In this paper, I collapse the “communicative” and “cognitive/doxastic” sense of explanation into one (since the goal of the communicative concept of explanation is taken to be providing *understanding* to an audience, which is itself a psychological/cognitive phenomenon). That being said, if one wishes to differentiate these then the argument in this paper will still go through. For the sake of brevity, however, I will focus on these three concepts of explanation.

to judge the adequacy of the response, and must rely on the authority of the responder to accept the explanation as legitimate (Woody, 2015, p.81).

When thought of in this way, we can view explanation as having particular goals and norms that dictate its proper application. In Craver's example of Jon giving a lecture to a classroom of students, he tells us that if all goes right, then "the audience comes to understand how action potentials are produced." (Craver, 2014, p.30). Or consider Achinstein's definition above in which an explanation involves an utterance from a speaker *with the intention that their utterance renders some phenomenon understandable*. In this respect, we can see *understanding* or *intelligibility* as an essential goal of explanation, and the norms of good explanation can be thought of in terms of the norms of good pedagogy.

The communicative concept of explanation can be contrasted with another: the ontic concept of explanation. This concept of explanation is most commonly attributed to José Alberto Coffa (1974), but has been advocated for by many others since (Salmon, 1989; Machamer, Darden and Craver, 2000; Craver, 2006; Izadi, Anandakrishnan and Onufriev, 2014; Kaplan, 2011; Strevens, 2011; Waskan, 2011; Povich and Craver, 2018; Craver and Kaplan, 2020). Salmon describes the ontic concept of explanation as follows:

Proponents of this conception can speak in either of two ways about the relationship between explanations and the world. First, one can say that explanations exist in the world. The explanation of some fact is whatever produced it or brought it about. [...] [I]t seems entirely appropriate to say such things as the gravitational attraction of the moon explains the tides, or the drop in temperature explains the bursting of the pipes. The gravitational attraction and the drop in temperature are out

there in the physical world; they are neither linguistic entities, nor abstract entities. Second the advocate of the ontic conception can say that an explanation is something—consisting of sentences or propositions- that reports such facts (Salmon, 1989, p.86).

The idea is that a good explanation must identify the appropriate *causes, structures, or features out in the world* responsible for the occurrence of the explanandum phenomenon in order to be genuinely explanatory. To borrow an intuitive example from Craver & Kaplan (2020, pp.299–300), the explanation for why sea levels are rising is because of *global warming*. Note, the explanation of sea levels rising is not *my saying the words* “global warming” to an audience, nor is it a model or a theory of global warming by itself. What explains the sea levels rising out in the world is the actual warming of the planet. A scientific model of global warming is thought to be explanatory in the ontic sense in virtue of identifying the relevant dependencies in nature (i.e. the changing climate) which genuinely produces and thus explains the rising sea levels.

While many think that the ontic concept of explanation tends to presuppose a largely mechanistic view of explanation (i.e. that good explanations must identify *ontic mechanisms*), this need not be so and may include other kinds of entities/relations as well. Craver and Povich, for example, claim that “we adopt a more inclusive understanding of the ontic that embraces any natural regularity, e.g., statistical relevance, natural laws, or contingent compositional relations might also figure fundamentally in explanation” (Craver and Povich, 2017, p.32; see also: Povich and Craver, 2018). If one adopts this ontic definition of explanation, then there are different goals and norms that present themselves. The goals associated with this definition of

explanation are tied to things like identifying the relevant or appropriate natural regularities, laws, mechanisms, or dependencies out in the world that produce or sustain the phenomenon.

Traditional explanatory virtues such as *precision* and *depth* are often associated with this definition of explanation, as the more precise our descriptions of the ontic structures, causes, and regularities becomes, the deeper the explanation becomes. This idea is explicitly defended by Mackonis, who claims that “[explanation] H1 is *deeper* than H2 if H1 explicates a causal-nomological mechanism that produces the abductive trigger and H2 does not” (my emphasis), and that H1 is deeper than H2 if the mechanism posited by H1 is more “specific, precise, [and] fundamental” than one posited by H2 (Mackonis, 2013, p.985). Others have likewise associated the explanatory virtues of depth and precision with the ontic definition of explanation (see, Lipton, 2004; Strevens, 2011; Craver and Kaplan, 2020, p.313).

According to this definition of explanation, an essential norm of good explanation is that the more we can identify the ontic dependencies relevant to the production of the phenomenon in the world, the better the explanation becomes. As Craver (2014, p.41) argues, “the norms of scientific explanation fall out of a prior commitment on the part of scientific investigators to describe the relevant ontic structures in the world”.

Lastly, let us consider the representational concept of explanation. We frequently talk of explanations in terms of scientific *models* or *theories*. For example, we might talk about how G.A. Parker’s (1978) *model of dung fly copulation* explains why dung flies copulate for approximately 36 minutes (see, Rice, 2015); or how Heeger’s (1992) *Normalization Model* explains cross orientation suppression (see, Chirimuuta, 2014). As Waskan et al. (2014, p.3090) argue, “On one

common manner of speaking (e.g., ‘There is an explanation for the odd trait on page 25’) the noun ‘explanation’ does seem to refer to a set of representational-artifacts”.

When dealing with this concept of explanation, we again have different norms and goals that govern our explanatory practices. The goal of explanation under this definition involves our ability to represent phenomena in the appropriate sort of way to gain relevant insights into it. These insights are traditionally cashed out in terms of our ability to predict when the phenomenon occurs, describe its behaviour, identify constraints on the phenomenon, or the patterns/principles it obeys (e.g., Batterman, 2000; 2002; Batterman and Rice, 2014; Chirumuuta, 2014). This definition of explanation is often associated with the traditional explanatory virtues of *breadth* and *unification*. Breadth and unification are often characterized as the ability of an explanation to account for a wide range of disparate phenomena. The more phenomena that can be explained by the model or theory, the greater its explanatory breadth and ability to unify (see, Mackonis, 2013; Lombrozo, 2016; Mantzavinos, 2016). For instance, Mantzavinos characterizes unification in terms of explanations that seek “laws and principles of high generality with the aim of constructing a coherent world picture and fitting particular facts within this framework” (Mantzavinos, 2016, p.6, footnote). Thus, under this definition of explanation, a model that describes a set of patterns or principles that we can subsume a greater range of phenomena under, or which identifies specific constraints that apply to a greater range of phenomena, will be better explanations than models that do not. For instance, a continuum model is taken to be a good explanation of the flow behaviour of liquids given that it can successfully predict and describe the behavioural patterns of a wide range of different fluids using the same formalism (see, Batterman and Rice, 2014; Izadi, Anandakrishnan and Onufriev,

2014; Bokulich, 2018). Different kinds of predictions, constraints on phenomena, patterns that phenomena adhere to, and principles that govern phenomena, require different kinds of models or theories to be properly identified.

All this suggests that the distinct concepts of explanation (“communicative”, “ontic”, and “representational”) involve different goals and norms that govern good explanation. Moreover, it also seems to be the case that one concept of explanation can apply in situations where the others do not. For instance, let us return to our previous example of rising sea levels. Global warming explains the rising sea level even if no one understands that the globe is warming, or whether we have a scientific model or theory which can adequately represent it. In this case, it seems like the ontic definition of explanation applies even though the communicative definition, and the representational definition, do not. Conversely, we might claim that Darwin’s Theory of Evolution is a *good* explanation of why elephants have long trunks even if the audience I try to explain this to doesn’t understand it (if, for instance, I am explaining the theory to a group of toddlers), and even if we do not yet know what the underlying ontic dependencies are that influence selection and genetic transmission. Lastly, we might do pedagogical research about the best ways to explain scientific principles to non-experts, and discover that certain kinds of techniques (e.g. frequent quizzes, having class discussions, one-on-one tutoring sessions, etc) will allow us to better explain. In which case, the communicative concept of explanation would apply, but the representational and ontic concepts do not (as I can discuss good explanatory techniques without invoking any particular scientific model or theory, or identifying any particular ontic structures and causes in nature).

To sum up, there are good reasons to think that the word “explanation” is ambiguous, referring to at least three distinct concepts

that are invoked in science. Each brings with them distinct goals for explanation, and distinct norms that govern good vs bad explanation. Moreover, it seems that each definition of explanation can apply in situations where the others do not. Working scientists thus need to be careful not to equivocate between the different concepts. While this sounds intuitive and plausible, I propose that this is, in fact, the wrong story to tell. Explanation is far more complex than it may initially seem, and these different concepts of explanation are in fact far more interconnected and interdependent than they may appear. In the section that follows, I want to propose an alternative way of thinking about scientific explanation and motivate this new way of approaching it.

2. One Explanation with Multiple Dimensions

Instead of assuming that we are working with three distinct *concepts* of explanation (each defined in a different way), I propose that a single explanation always has a communicative *component*, a representational *component*, and an ontic *component*. In other words, these are not distinct concepts of explanation as much as they are different evaluative dimensions along which we can analyze a single explanation. This point is not a trivial one and has the potential to help reframe current disputes within the philosophy of science. In order to demonstrate this, let us begin with the hypothesis that these *are* intended to be entirely distinct concepts of explanation, and then explore how each concept of explanation in fact requires or incorporates the others.

Let's begin with the ontic concept. To say that we can have an ontic concept of explanation apply independent of a communicative and representational one would be misleading. To illustrate, consider

the phenomenon of phase transitions in statistical mechanics. When we boil a pot of water, it transitions from a liquid state to a gaseous one. Yet our ability to explain this phenomenon is notoriously difficult since we can only model this sort of transition if we mathematically treat the pot of water as being infinitely large, allowing the molecules an infinite degree of freedom (Batterman, 2002). Now, we know that ontically the transition from liquid to gas is somehow accomplished by the interaction of the finite molecules that make up the water, and the pot containing it. Yet, it would seem unhelpful to suggest that we've *always* had a good scientific explanation of phase transitions by merely *pointing to* a pot of boiling water. Here, we are denoting the ontic structures and causes out in nature that produce the phenomenon. Yet, we are no better at *scientifically explaining* the phenomenon than we were before. This is precisely why William Bechtel (2008, p.18) notes that “the problem with this ontic view is that mechanisms do not explain themselves”.

This has led several theorists to point out that those who claim to be working strictly within the ontic concept of explanation do not, and cannot, eliminate the communicative and representational dimensions from how they talk about explanations (Bechtel, 2008, p.18; Wright, 2015; Bokulich, 2018). For instance, Bechtel argues that:

Even the advocates of the ontic perspective are unable to avoid invoking epistemic notions, although they try to minimize them. Machamer et al. sometimes refer to “giving a description of the mechanism” (p. 3) and “revealing . . . productive relations” (p. 21), and Salmon uses such words as exhibiting. But these terms understate the cognitive labor involved (Bechtel, 2008, p.18).

Or consider Cory Wright (2015, p.27), who notes that “Salmon offered a putative example of ontic explanation that centers on an

epistemic agent who makes discoveries, and so is thereby positioned to elaborately communicate a causal story”. Many who endorse the ontic definition of explanation have granted that there is an essential representational and communicative component to explanation that is intertwined with the ontic one. Mark Povich, for example, points out that:

Craver’s ([2014]) most recent formulation of the ontic conception backs away from the metaphysical claim that explanations are ontic structures in the world and focuses on demarcatory and normative constraints on explanation. Craver ([2014]) writes that according to the ontic conception, ‘in order to satisfy these two objectives [of explanatory demarcation and explanatory normativity], one must look beyond representational structures to the ontic structures in the world’ (Povich and Craver, 2018, p.129).

Similarly, Waskan (2011, p.4) points out that “ontic theories might take many forms, so long as what they propose is that explanations (primarily) reveal something about objective states of affairs. [...] Explanations are, on this view, representations—objective facts are not in the business of revealing. Specifically, they are descriptions”. This point has likewise been emphasized by Phyllis Illari, who argues that advocates of ontic explanation acknowledge the essential role that representation and modeling plays in scientific explanation, and that the issue instead is one of normative constraints. In other words, that “ontic explanation is essential for marking several crucial normative dimensions by which scientific explanations are and ought to be evaluated” (Illari, 2013, p.243).² This provides evidence that there is

² This shift from ontic structures being explanations in and of themselves, to the idea that identifying the appropriate ontic structures provides essential normative constraints on explanation may lead some to conclude that we’ve weakened the ontic

not a distinct *ontic concept* of explanation, there is instead an *ontic dimension* to a scientific explanation that comes along with both representational and communicative dimensions.

To further highlight this idea, let us turn to the representational concept of explanation. It can often be the case that various theories or representations are treated as good explanations even when they are not understandable to a given audience (which runs counter to the communicative definition of a good explanation), or when no ontic dependencies are being identified by the representation in question (which runs counter to the ontic definition of a good explanation).

account of explanation to the point of triviality. If the claim merely reduces to some form of epistemological realism (i.e. that there are objective features of the world that we can learn about), then the claim is not particularly informative. Even those who explicitly reject the ontic definition of explanation would be willing to grant such a claim. Such a criticism would be uncharitable, however. The ontic concept of explanation does not merely insist that our models must describe some objective world, but that a good explanation must identify the *appropriate* ontic dependencies in the world necessary to account for the explanandum phenomenon. Specifically, those responsible for the production and sustaining of the phenomenon. This claim is directly at odds with those who insist that highly idealized models which do not identify any such ontic dependencies can still count as genuine explanations (e.g. Batterman, 2002; Batterman and Rice, 2014; Chirimuuta, 2014; Rice, 2015; Rice and Rohwer, 2020). Put another way, the ontic definition need not be committed to the idea that the ontic dependencies in the world are somehow self-identifying or self-explaining. It instead need only be committed to the idea that there are relations that exist between events/dependencies in the world such that some set produces and sustains the other, and that appropriate scientific explanations are those that correctly highlight such relations and dependencies. This idea has always been implicit in the ontic definition of explanation, with Salmon (1989, p.86) telling us that “the advocate of the ontic conception can say that an explanation is something—consisting of sentences or propositions- that reports such [ontic] facts”, and that “one can properly say either that the explanandum-fact is explained by the explanans-facts or that the explanans-statements explain the explanandum-statement”. Although it is worth noting that debates about whether this weaker interpretation of the ontic view is still problematic are ongoing. For details, see: (Illari, 2013; Wright and van Eck, 2018).

This seems to provide compelling evidence for the idea that there are distinct concepts of explanation at play. Let us examine the intuitive strength of these claims before responding to them.

Suppose a theoretical neuroscientist uses a computational model to provide a scientific explanation of the behaviour of a particular neural circuit. If the audience of the explanation happens to be a group of kindergarten children, then the computational model in question will fail to provide them with any sort of understanding of neural circuits. Thus, it would fail to be a good explanation under the communicative definition of explanation. However, such a failure on the part of the children to understand the model does not undermine the idea that the model is still considered a good explanation by the scientific community, suggesting that a different concept of “explanation” is at work.

We must be cautious with such a conclusion, however. This is because when dealing with *scientific* explanations, the audience who must understand the model or theory is not *any* audience. Put another way, the computational model is a good scientific explanation because it provides *working scientists* with an understanding of the neural circuit. Indeed, studies have shown that “high-science participants were less likely to regard a passage of text as an explanation when told that the representation, because of qualitative barriers, lacks intelligibility” (Braverman et al., 2012, p.1372). In other words, if a model or theory cannot provide an explanation *in accordance with the communicative definition* to working scientists, then scientists themselves do not consider the model or theory to be genuinely explanatory. The importance that understanding-to-a-scientific-audience has for the representational concept of explanation can be seen in the way that scientists explain to non-scientists. As Woody (2015, p.81) notes, “explanatory discourse often involves the communication of exemplary

explanations among members of a given scientific community, and one aim of explanatory practice seems to be training novice practitioners to recognize typical explanatory patterns.” And so to suggest that a scientific model or theory is explanatory independent of whether it provides understanding to an audience is deeply misleading, and one cannot tease apart a representational concept of explanation from a communicative one in this regard.

But what about the relationship between the representational concept and the ontic concept? There seem to be many examples in which a model is considered explanatory despite not identifying (or indeed deliberately misrepresenting) the ontic variables in nature responsible for the production of the explanandum phenomenon. Optimality models in evolutionary biology, for example, deliberately ignore or distort causal features in the evolutionary process in order to mathematically determine what sorts of traits would be locally optimal for an organism to have, thus explaining why such traits were selected for (Potochnik, 2010; Woods and Rosales, 2010; Rice, 2015). Likewise, continuum models in physics are used to explain the flow behaviour of liquids despite deliberately saying nothing as to the molecular components and causal processes that produce such behaviours (Batterman and Rice, 2014; Izadi, Anandkrishnan and Onufriev, 2014; Bokulich, 2018, p.803). In these instances, it appears like there is no ontic component to the explanation, suggesting a distinct concept of explanation.

Yet these cases are not as straightforward as they initially seem either. For instance, take optimality models. While it may be true that an optimality model does not identify the ontic structures and causes in nature responsible for the production of phenotypic traits, it is not true that a commitment to what those structures and causes are is not an essential part of the explanation. This is because the application

of optimality models *only* works when scientists already know essential ontic structures and processes at work in the production of the phenomenon. Angela Potochnik, for example, notes the following regarding the use of optimality models in evolutionary biology:

Because optimality models use highly simplified assumptions as placeholders for complex dynamics, their successful use depends upon evolutionary dynamics that the models themselves do not explicitly represent. In other words, optimality models are epistemically dependent on unrepresented dynamics. Information about these unrepresented dynamics helps establish whether an optimality model's simplifying assumptions are problematic, and thus how successful the model is (Potochnik, 2010, p.226).

And so the ontic dimension of the explanation is essential here. The model in question cannot function appropriately without the relevant ontic structures and processes being identified and understood by scientists, even if they are not explicitly stated in the model or theory itself. While a focus on the *representational dimension* of explanation does not highlight the ontic features of the explanation, they are a necessary part of the explanation itself, since they are required for the representations to function appropriately.

For another example, consider continuum models. As Bokulich (2018, p.803) rightly points out, "at large scales, continuum representations of water and the Navier-Stokes equations are typically more relevant" than other kinds of scientific models or representations for characterizing flow behaviour. The Navier-Stokes equations predict how fluids will behave even though they do not identify any of the underlying molecular components of the fluid. Instead, they treat the

fluid as a single continuous entity. Such models would appear to explain why fluids behave as they do despite saying nothing as to the ontic variables responsible for it.

Just like the case with models in evolutionary biology however, such models in physics are likewise *still dependent on scientists being committed to ontic variables that are not included in the model itself*. In this case, the ontic properties of the molecules that compose the fluid must be correctly understood by working scientists for proper application of the Navier-Stokes equations. Mark Povich, for example, points out that the Navier-Stokes equations cannot be used to account for the behaviour of fluids like liquid crystals. This is because “their often rod-shaped particles result in directional preference and lack of symmetry. Liquid crystals thus cannot be accurately modelled using the unmodified Navier-Stokes equations. The addition of a stress tensor or coupling with a Q-tensor system is required to take into account the anisotropy of liquid crystals” (Povich, 2018, p.124). Note that while the Navier-stokes equations themselves do not identify the molecular properties of the fluid, scientists must be committed to various ontic properties of the molecules that make up the fluid in order to tell if and when the equations will work, or if and when they will need to be modified.

Some, like Potochnik (2010), acknowledge the importance that these ontic commitments play in constructing and applying these idealized representations, but insists that it is the idealized model that is explanatory independent of these commitments. This is because while these ontic commitments are necessary to construct and apply something like an optimality model, the model itself can only identify the relevant evolutionary patterns or principles when these ontic structures and causes are excluded from the representation itself. By trying to add such details to the representation, it becomes worse at

identifying the patterns and principles we seek to represent. Thus, the model is explanatory independent of those ontic details. Similar arguments are also made in the context of highly idealized models in physics (see Batterman, 2002; Batterman and Rice, 2014).

The problem with this style of response is that it equates an explanation with a scientific model or theory. In this case, since the model or theory does not identify the relevant ontic variables, then neither does the *explanation*. However, a scientific explanation should not be conflated with a model or theory. A scientific representation is an integral *part* of an explanation, but the explanation itself includes elements that go beyond what is included within any single representation.

We have already seen evidence of this regarding the role that understanding plays in explanation. A model or theory may, for instance, correctly identify the pattern we care about, but if working scientists cannot understand the model or theory (i.e. if it is abstracted away from the communicative dimension of explanation), then scientists will not consider it explanatory. This already suggests that what constitutes an explanation goes *beyond* what is explicitly stated by any particular model or theory. Just as a communicative dimension is essential for any particular representation to count as a *part* of the explanation, so too is a commitment to essential ontic features of the world (even if they are not stated by the model or theory itself). A commitment to what the actual ontic variables in nature are plays an essential role in the creation of both optimality and continuous flow models, their boundary conditions, their applications, and what inferences we are licensed to draw from them (see Hochstein, 2019). In this regard, these ontic commitments are intertwined with the representational content of the models to instantiate the explanation. To pull them a part is to strip the representation of its ability to carry

explanatory content. In other words, we cannot cleanly demarcate a *representational* concept of explanation from an *ontic* concept of explanation.

Lastly, let us turn to the communicative concept of explanation. While we can talk about an act of explaining abstracted away from any particular representation or ontic commitments, the representational and ontic components of explanation are always present and ineliminable in any particular instance of scientific explanation as a communicative act. As Illari notes, when talking about scientific explanation (even as a communicative act), one cannot ignore the ontic dimension of explanation given that “explanations cannot ignore worldly things” (Illari, 2013, p.251).

To illustrate, recall that the primary goal of the communicative definition of explanation is to provide *understanding* in an audience. However, our understanding of a phenomena cannot be divorced from our ontic commitments about essential structural and causal features in the world since part of what it is to understand is to correctly account for such features. Waskan, for example, points out that “understanding” is frequently treated as a “success verb much like (a sense of) ‘see’—that is to say, in order to do it you must be successful at it. We might say, for instance, that whereas alchemists *felt* that they understood combustion, chemists do *genuinely* understand combustion” (Waskan, 2011, emphasis in the text). Bechtel (2008, p.14), Strevens (2011, p.3), and Illari (2013, p.245) make similar claims. Likewise, both Elgin (2004) and Potochnik (2015) argue that truth is an essential *threshold* concept when it comes to understanding. In other words, understanding requires that one correctly accounts for at least *some* of the ontic dependencies out in the world that are directly responsible for the production of the phenomenon. As Potochnik (2015, p.73) puts it, “a claim must be ‘true enough’ in

order to be epistemically acceptable; that is, any divergence from the truth must be negligible, or safely neglected". While both Elgin and Potochnik argue that the amount of truth required for an account to provide genuine understanding will vary based on our pragmatic needs, it highlights how the ontic dimension of explanation cannot be conceptually *divorced* from any account of understanding, making it an eliminable part of the communicative concept of explanation.

A very different kind of argument for the essential connection between the communicative, ontic, and representational definitions of explanation can also be found in the way that the various goals associated with scientific explanation are unavoidably intertwined and dependent on one another for their explanatory power (Hochstein, 2017). Consider that the primary goal associated with the ontic concept of explanation is the identification of relevant ontic dependencies existing in the world. Meanwhile, one of the primary goals associated with the representational concept of explanation is the ability to represent certain behavioural patterns that a given phenomenon adheres to. But what justifies the explanatory status of such goals? Put another way, *why* is the identification of ontic dependencies explanatory? What *makes it* explanatory? The same with the identification of behavioural patterns, or with prediction. Why consider such goals relevant to explanation? Typically, theorists will justify the explanatory status of one explanatory goal *by appealing to the other goals*.

For instance, we often justify the explanatory status of identifying ontic dependencies by appealing to the fact that knowing such dependencies allows us to better control and manipulate the phenomenon, which in turn allows us to better *predict* and *understand* it. Thus, we justify the explanatory goal of the ontic concept of explanation by appealing to the fact that it allows us to better satisfy the goals of the communicative or representational concepts of explanation. Similarly,

we might ask why predicting the occurrence of the phenomenon in a range of cases counts as explanatory. In such cases, it is common to justify the explanatory status of prediction on the grounds that “predictions help us check whether our accounts of the world have any veracity” (Douglas, 2009, p.453). Here we justify the explanatory status of a goal identified by the representational concept of explanation by appealing to how it helps us to better attain a goal identified by the ontic concept. This shows that the various goals of explanation are in fact inter-defined and inter-dependent. Building on this idea, I propose that instead of having distinct concepts of explanation, we have different evaluative dimensions of an explanation that are similarly inter-defined and inter-dependent.

All of this provides substantial evidence for the idea that there is always an unavoidable ontic, communicative, and representational dimension to one and the same explanation. While we can talk about these different evaluative dimensions abstracted away from one another, we must be cautious not to assume that these dimensions denote entirely distinct concepts of explanation that are autonomous and that apply in different contexts.

3. A Distinction Without a Difference?

Before proceeding further, it is important to highlight why this alternative view of scientific explanation is worth articulating. After all, are “dimensions” of explanation really that different from “concepts” of explanation? Am I merely drawing a distinction without a difference? Why is it important to demarcate these two positions? There are important implications to the view being defended here.

To assume that there are simply distinct concepts of explanation leads to the deeply problematic assumption that the goals and norms of a particular concept of explanation, in the appropriate context, is always sufficient by itself to account for a good scientific explanation. However, this would be to distort how explanations work in science, and why they are successful. To illustrate, consider Mantzavinos's example of neoclassical economics. While Mantzavinos does not distinguish explanations in terms of communicative, representational, or ontic concepts, he does imply such a distinction when insisting that models from neoclassical economics count as genuine scientific explanation in virtue of meeting the criteria associated with the representational concept, but not those associated with the ontic one. Specifically, he notes that neoclassical models are considered good scientific explanations within the field of economics because they unify a range of economic phenomena under a particular set of principles (Mantzavinos, 2016, p.12). However, he notes that such models *do not need* to identify relevant ontic structures or causes in nature to count as explanatory, since in the context of economics, this explanatory goal is simply not what is adopted by working scientists.

For example, he tells us that “microeconomics textbooks never pay any attention to ‘causal processes, causal interactions, and causal laws’, in the analysis of any type of market, be it the competitive market, a monopoly or an oligopoly”, and that such models do not identify the sort of ontic dependencies in the world that we can intervene on or manipulate because “it operates at a level of abstraction that makes it extremely difficult to test the theory empirically” (Mantzavinos, 2016, p.13). He ultimately concludes that the “decision calculi of neoclassical economic theory are clearly argument patterns that can be only accommodated by the unification model of explanation” (Mantzavinos, 2016, p.13). For our purposes, we can reasonably interpret this

as saying that for economists, it is the representational concept of explanation that matters (since it is the unification of various market phenomena under a single model or theory that makes it explanatory), not the ontic concept (since identifying ontic structures, causes, and variables that can be manipulated, are not taken to be explanatorily relevant in this context).

But we have good reasons to challenge Mantzavinos on this point. Whether such neoclassical models identify genuine principles and regularities in market forces, or merely notes contingent correlations that do not genuinely unify market phenomena at all, cannot be determined *unless we appeal to ontic dependencies and causal variables in the world*. Neoclassical economics has built into it certain implicit ontic commitments about the structures and dependencies in the world that must be true in order for the model to genuinely provide a unifying account. If these sorts of ontic dependencies are false, then the model's ability to unify successfully, and thus meet its own standards of good explanation, are undermined. Similarly, if the model does not identify ontic variables we can intervene upon, we have no way of testing or determining whether the neoclassical model is actually identifying unifying market forces since there is no way to manipulate relevant features of the market to see whether its underlying principles apply or not. This is precisely why *economists themselves* have questioned whether such models should be taken to provide genuine scientific explanations.

For a rather stark example, consider Alfred Eichner's now famous argument for why neoclassical economics fails to count as a science, and to provide scientific explanations. He notes that neoclassical economics is built on numerous implicit ontic assumptions about the underlying structures and dependencies in the world which it uses as the foundation for its unifying models, and that many such ontic

assumptions have proven false (for discussion of such incorrect ontic comments, see: Eichner, 1983, p.511). In virtue of getting these ontic dependencies wrong, Eichner argues that the explanatory value of such models is undercut given that they fail to account for the market forces that genuinely exist.³ In other words, the model fails to unify the actual range of phenomena in the world it is attempting to explain. As he puts it:

Economists as a group have adopted the view that formal, or mathematical, proofs are sufficient to establish the validity of a theory rather than just being necessary. [...] Thus it is common for the “theorists” to set up their models in such a way that the postulated behavior runs counter to all that is known about actual economic systems without this fact impugning either the argument or the economists’ reputation. So sharp a distinction between theory and empirical research is unknown in the natural and biological sciences, and for good reason. It leads to an outpouring of useless theories that waste the time and energy of empirical researchers (1983, p.517).

Similarly, the lack of intervention that such models provide is likewise taken to undermine the scientific and explanatory status of such models. This is because “the theory must be shown to make a difference to society, when translated into one or more public policies, that will lead to certain clearly distinguishable results. The policies must then be adopted and the predicted effect confirmed. This is the praxis test of a social science theory” (Eichner, 1983, p.510). Based on these very reasons, Eichner concludes that neoclassical economics fails to be scientific, and its models fail to provide genuine scientific explanations. Note that the issue here is not whether economists think

³ For others who likewise emphasize this point, see: (Hall et al., 2001; Hausman, 2008; Keen, 2022).

that unifying market forces within a model or theory is explanatory. The issue is that the ability of the model to successfully meet this criterion *requires* that the norms and goals of the ontic dimension of explanation be adopted. Mantzavinos suggests that economists themselves don't consider the goals of the ontic dimension relevant to explanation, but this is untrue. Even Milton Friedman himself grants that keeping track of these implicit ontic commitments are "extremely valuable in suggesting leads to follow in accounting for divergences between predicted and observed results; that is, in constructing new hypotheses or revising old ones" (Friedman, 1953, p.31n). Put simply, *economists themselves* have emphasized the importance of the ontic dimension of explanation when evaluating whether neoclassical models are successful at unifying phenomena under an appropriate scientific representation.

To assume that there is a representational concept of explanation that is distinct from the ontic concept and the communicative concept, and that the application of this one concept is sufficient for explanation in the right contexts, ignores that each "concept" in isolation does not, and cannot, do the explanatory work that it is expected to do unless we smuggle in the other concepts of explanation with it. It is important to distinguish "concepts" of explanation from "dimensions" of explanation since one implies that we can satisfy one concept of explanation without having to say anything about the other concepts, and that *good* scientific explanation is determined entirely by the particular concept of explanation we employ. If we understand these as different dimensions of the same explanation, on the other hand, then we can better understand the interplay between the different dimensions, and why we do not have sufficient resources for determining a good explanation when focusing only on one dimension in isolation of the others.

By changing our perspective from concepts of explanations to different evaluative dimensions of an explanation, we can also better clarify some confusions that have been at the heart of certain debates in the philosophy of science. Take the ongoing debate amongst mechanists in the context of biology and neuroscience. These theorists argue that explanation in the life sciences, especially biology and neuroscience, primarily proceeds by identifying the physical mechanisms responsible for the target phenomenon (often characterized in terms of organized parts and operations that produce regular change). However, some argue that mechanisms are out in the world, and that explanations merely describe or reveal such mechanisms (Machamer, Darden and Craver, 2000; Craver, 2006; Kaplan, 2011; Strevens, 2011; Craver and Darden, 2013). Meanwhile others argue that mechanisms are interpretations of the world (by way of mechanistic *models*), and that we should think of mechanistic explanation as cognitive and representational products we create to understand the world (Chirimuuta, 2014; Bechtel, 2015; Wright, 2015; Austin, 2017). The first group argues that mechanistic explanations are *ontic*. Meanwhile, the second group argues that mechanistic explanations are *epistemic*. So how ought we to understand mechanistic explanation?

Here the dispute rests on which dimension of explanation that theorists focus on at the expense of the others. For instance, Bechtel (2015) notes that causal variables in the world responsible for a cognitive phenomenon go well beyond the structures and processes occurring within the brain, and include events in the distant past, as well as all kinds of environmental structures and processes that are not typically treated as part of the biological system. In this respect, he argues that the boundary of a cognitive mechanism is often not determined by the world itself, but by the interests of working scientists who focus on only a subset of the causal variables responsible

for a given cognitive process. Since every mechanistic explanation idealizes numerous causal variables in the world and are dependent on the interests of working scientists, this suggests that we are not merely describing the ontic joints of mechanisms in the world but interpreting the world in mechanistic terms. In his words,

It is the scientists who impose boundaries around entities and activities in nature and impose a time scale on which their functioning is characterized. For different explanatory purposes researchers may draw these boundaries in different locations or at different time points. These choices, though, while not simply responsive to pre-existing boundaries, are not entirely arbitrary. [...] The networks of entities found in nature commonly exhibit small-world organization as well as being scale-free. This entails that while real-world networks are highly inter-connected, there are clusters within them that are semi-independent of the rest and productively posited to be the mechanisms responsible for specific phenomena (Bechtel, 2015, p.85).

Conversely, those who defend the ontic account note that even in such cases, there *are* “real-world networks” that cluster in ways that are semi-independent and productively posited to be mechanisms. Thus even if such models are highly idealized, they still carry explanatory content in virtue of revealing genuine ontic structures, organizations, and processes in nature that play an essential part in the production of the phenomenon.

How then should we settle such disputes? I propose that the framework presented in this paper can help untangle the confusion here. Those who advocate for *ontic* accounts of mechanistic explanation focus on the ontic dimension of explanation at the cost of the representational or communicative dimension. Meanwhile, those who argue for epistemic explanation focus on the representational and

communicative dimensions at the cost of the ontic dimension. Yet if we consider that explanation always has a communicative, representational, and ontic dimension to it, then the two accounts don't so much represent conflicting accounts of mechanistic explanation, as much as highlight why we should care about, and pay attention to, different dimensions of the same explanation.⁴

Illari (2013) similarly argues that the contrast between ontic and epistemic views of mechanistic explanation ignores the inseparable and intertwined nature of the various goals and norms that make up the two positions. For instance, she argues that many of the traditional

⁴ This account of explanation may also help to disentangle some confusions regarding the dispute between explanatory monists, and explanatory pluralists, in the philosophy of science. Explanatory monists argue that there is a single set of criteria for determining what counts as a good scientific explanation that applies across fields and contexts. Meanwhile, explanatory pluralists argue that there is a plurality of types of explanations in science, and that there is no single set of criteria that unities them all (e.g. the sorts of criteria that determine a good explanation in the context of physics may be distinct and unrelated from the sorts of criteria that determine a good explanation in the context of biology). We may be able to help clarify some of these issues by noting that if we focus on a single dimension of explanation, like the representational dimension, then we seem to find a plurality of types of models that can and do count as explanatory (e.g. statistical, mechanistic, topological, dynamical, etc). However, if we focus more generally on the essential set of norms that all three dimensions impose on each other, we may be able to find a general monistic set of criteria that all explanations require. For instance, a good explanation must be understandable to the scientific community. It must represent the phenomenon in a way that either unifies phenomena under a set of principles, identifies constraints on its behaviour, or makes predictions. Since a model that fails to allow us to predict what the phenomenon will do, provides no constraints on what the phenomenon cannot do, and provides no information as to what the various instances of the phenomenon have in common, may strip the explanation of its ability to say anything of use to scientists. Lastly, an explanation must identify essential ontic dependencies and regularities in nature, since otherwise we cannot tell if our representation is genuinely identifying principles the phenomenon obeys, or constraints on its behaviour, as opposed to mere contingent correlations. And so explanations may turn out to be pluralistic in one regard (say, in terms of model types), but monistic in others (say, the set of general constraints that guide communicative, representational, and ontic practices).

virtues of explanation, such as unification, simplicity, elegance, and intelligibility, in fact depend on the *interaction* of the various goals and norms that make up both ontic and epistemic accounts of explanation. As she puts it:

This means that what looks to us intelligible, simple and unified is not a static feature of human psychology, but is affected by our empirical engagement. Newtonian action at a distance used to seem quite impossible to us; so did quantum mechanical indeterminism, and non-locality. Physicists sincerely describe quantum mechanical equations as elegant, a claim that can generate hilarity in those unused to working with such theories. If empirical engagement continually forms what we find intelligible, simple and unified, then epistemic constraints on explanation, even the more ‘psychologistic’ constraints, are deeply entangled with ontic ones (2013, p.254).

In this regard, the ontic and epistemic camps are not providing distinct or incompatible views of mechanistic explanation, so much as emphasizing different essential features that a given explanation must have.

4. Implications and Concerns

On the surface, such a view may appear either obvious (“of course explanation involves all three elements!”) or naïve (“It’s too simplistic to assume that we can easily merge these accounts”). And so it is worth exploring the complexities and difficulties that such an account now presents us with. As noted previously, different dimensions of explanation emphasize different norms for evaluating good and bad explanations. Problems begin to arise, however, when we discover

that these norms can often conflict with one another. For instance, if we focus on the ontic dimension of explanation in isolation, it seems like the more we can identify and describe the ontic dependencies in the world, the better our explanation becomes. It is this intuition that underlies Kaplan's claim that it is a "highly plausible assumption that the more accurate and detailed the model is for a target system or phenomenon the better it explains that phenomenon, all other things being equal" (Kaplan, 2011, p.347). Yet this can run counter to some of the other norms of explanation associated with the other dimensions. For instance, adding too many accurate details to a model or theory can make it more difficult to understand, and thus makes it a worse explanation according to the norms associated with the communicative dimension of explanation.

Indeed, conflicts between the various explanatory goals of the different dimensions can be seen explicitly in scientific practice itself. Consider explanations of the action potential of the neuron in neuroscience. The Hodgkin & Huxley model of the action potential, developed in the early 1950s, mathematically characterized electrochemical features of the action potential (specifically: the ion flow of sodium and potassium channels). However, it did not provide an account of the ontic variables responsible for the production of these features. In this regard, it satisfied the explanatory goals of the *representational* dimension (by identifying constraints or patterns that the phenomenon adheres to), but not of the *ontic* dimension (by failing to denote the relevant ontic variables that actually produce these constraints and patterns). This led Hodgkin & Huxley to make inconsistent claims regarding whether their model is explanatory or not. In their original paper, they claim that their model provides "a sufficient explanation of the wide range of phenomena that have been fitted by solutions of the equations" (Hodgkin and Huxley, 1952, p.541).

Yet they also claim in the same paper that their model is in fact *not* a good explanation of the action potential, since it merely provides an “empirical description” of the phenomenon and not an account of what produces it (Hodgkin and Huxley, 1952, p.541; for the discussion, see also: Bogen, 2005; Craver, 2006).

So how then do we determine whether a model (like the Hodgkin & Huxley model) counts as a good scientific explanation or not given that the norms that govern explanation are inconsistent or contradictory across evaluative dimensions? This is a serious concern. One potential solution, recently proposed by Rice & Rohwer (2020), is to claim that a scientific model or theory need only satisfy *a sufficient number* of the norms that govern all the various dimensions of explanation to count as genuinely explanatory in a given context, as opposed to satisfying *all* of them. They claim that “many single models are sufficient to explain because the information they provide will satisfy a sufficient subset of the conditions included in the cluster concept” (Rice and Rohwer, 2020, p.1043).

The problem with this solution is that it makes the common mistake of conflating *explanation* with *representation* (i.e. the idea that explanation is always constituted by a particular model). However, as I have argued, features beyond the representational texts are also constitutive parts of an explanation. For instance, scientists have certain implicit commitments to ontic dependencies that are essential to how they construct, interpret, and draw inferences from their models despite such dependencies not being explicitly stated in the model itself. Without these commitments, the model cannot function. In this respect, the background beliefs and commitments of scientists are a component part of an explanation in addition to the particular mod-

els and theories being used (see: Bokulich, 2018; Hochstein, 2019). Similarly, how we convey models and representations to others is also an essential component to explanation:

The history of science, meanwhile, provides rich empirical support for the claim that production of explanations serves to constitute, rather than merely communicate, “intelligibility” for a scientific discipline. Precisely because explanatory discourse inculcates particular patterns of reasoning, it functions to sculpt and subsequently perpetuate communal norms of intelligibility. In effect, explanations encode the aims and values of particular scientific communities, telling practitioners what they should want to know about the world and how they should reason to get there (Woody, 2015, p.81).

This means that a single explanation may be constituted by collections of *representations*, *social practices*, *psychological processes*, and *accurate commitments to the relevant ontic features of the world*. Likewise, determining whether an explanation is good or bad will involve an evaluation of *all* these components.

At first glance, this would appear to make the evaluation of scientific explanations a herculean affair! This idea is not without precedent however. Mantzavinos (2016) proposes an account of explanation that makes a similar claim. Mantzavinos proposes that instead of talking about explanations as entities, we should understand them as collections of inferential practices governed by the norms of the scientific community, which he calls “explanatory games”. Most notably, he tells us that determining “which rules of representation guide the explanatory activities is fundamentally important for the quality of explanations generated during the game” (Mantzavinos, 2016, p.42). Here, the *representational* dimension of explanation is emphasized. However, he also notes that:

No explanatory game can take place in a metaphysical vacuum. The metaphysical assumptions act as constraints on the generation of the other rules, and belong therefore to the constitutive rules. The structure of the game is predicated on prior assumptions concerning the way the world is and by what means it is explainable in principle. These rules can be implicit or explicit and they can vary from stone-age metaphysics to highly refined metaphysical assumptions (2016, p.41).

Here, the *ontic* dimension is emphasized. Lastly, he tells us that explanatory practices are “a process of social interaction unfolding within the given institutional rules.” (Mantzavinos, 2016, p.68). In this way, the *communicative* dimension of explanation is highlighted. For Mantzavinos, we evaluate whether an explanation is good or bad by evaluating whether all of these inferences are appropriate and in accord with the inferential rules of the explanatory game.

The problem, of course, is that we must now evaluate a huge number of variables (communicative, representational, and ontic) in order to determine whether any explanation is good or bad, which is no small feat. Yet I propose that a great deal of how we evaluate explanations in science already implicitly does this. For instance, knowing which parts of a model count as idealizations and which don't require that working scientists already have ontic commitments regarding the variables in the world that the model is deviating from, and the wrong sorts of ontic commitments are often taken to be signs of a problematic explanation. As Anya Plutynski (2013, p.472) points out:

Good modelers are careful to be clear about when an assumption used to construct a model is deliberate simplification or simply false, and when it is a hypothesis supported by evi-

dence. Unfortunately, what starts as deliberate simplification may often be confused with actual hypothesis and latter reified into theory.

Thus we often evaluate good explanations in terms of whether one's ontic commitments are appropriate and being used correctly in the construction of our models. Similarly, the way that knowledge is communicated throughout the scientific community is also evaluated when determining good explanation. As Woody (2015, p.85) points out,

Explanations must be minimally intelligible to be accepted by individuals. But explanations, especially in educational contexts, condition reasoning, and thus sculpt social norms of intelligibility. Explanations intelligible to members of one community are frequently opaque to outsiders, even other scientists.

Therefore the appropriate means of communication can be essential to whether something is adopted by the scientific community as a good explanation. Likewise, various constraints on the communicative practices we have (such as passing peer review in order to be published in scientific journals, or accepted at conferences and workshops) are also considered essential to whether something should be accepted or rejected as a good *scientific* explanation.

Like evaluating any complex system, we may have to trust that some parts are in better working order than others or look for faults only when something has gone astray. Instead of attempting to evaluate all parts of the explanation simultaneously, we focus our attention on certain features of it depending on our interests and needs, or evaluate the particular features of the explanation that we think may be problematic.

By analogy, imagine evaluating whether a car is in good working order. We might have particular concerns about a given type of car (“I hear it gets bad gas millage!”, “The engine is known to frequently stall!”), and use this to gauge if the particular car we want to purchase has these problems. If not, we may conclude it is in good working order. Yet if someone comes to us and points out that the break line is cut, or that the spark plugs need replacing, or that the trunk is welded shut, then those concerns will point our attention to parts of the car we assumed were in working order, and now causes us to doubt whether the car is worth purchasing. But of course, of the thousands of things that could potentially be wrong with a car, we do not and cannot evaluate all of them simultaneously. Instead, we focus our attention on some features we consider to be most important or worrisome, and assume the rest of the car is in good working order *until given reason to doubt it*. This is because there are certain standards we assume the car must have met to be sold on the lot.

The same will apply to scientific explanations. We might question some aspects of an explanation, while assuming that others are in working order in virtue of scientific communal standards. For instance, if a paper has been published in a respected peer reviewed journal, then we might assume that it is intelligible to the relevant scientific audience. Of course, if we discover that a paper was accepted due to a case of fraud, then it would immediately cast into doubt that the paper has in fact met the appropriate standards of the communicative dimension. Similarly, if scientists working in various labs report being able to control and manipulate episodes of depression by changing the concentrations of various gut bacteria, then we may be justified in assuming that the gut microbiome is part of what explains those episodes of depression. However, if we learn that their ability to control and manipulate depressive episodes in this way was greatly

overstated, or far more limited than they let on, then it will call our attention to the ontic features of the explanation that we assumed were in good working order.

Scientists often have little reason or incentive to identify all component features of an explanation during their everyday activities. To ask whether a particular model is a good explanation for example, is not to suggest that the communicative or ontic dimensions to that explanation *do not apply*, or that *we mean something different by explanation*. It is merely the case that the other dimensions are either obvious to those involved, not points of contention, or not their primary focus. Sometimes it is the ontic component of the explanation that we may disagree about, at other times the best way to convey models to an audience. We must not confuse the fact that we can analyze explanations along different dimensions with the idea that these different dimensions correspond to distinct concepts of explanation, or denote entirely metaphysically distinct explanatory entities.

Conclusion

Traditionally, philosophers have adopted either a representational definition, communicative definition, or ontic definition of scientific explanation. This has given way more recently to a pluralistic view which argues that each concept of explanation may be scientifically appropriate in particular contexts and for particular purposes. In this paper, I have argued for an alternative view which better highlights the complex ways in which representational, communicative, and ontic features of a scientific explanation interact and depend on one another. Instead of distinct explanatory concepts, we have different evaluative dimensions along which we can analyze any instance of a scientific

explanation. The upshot of this view is that it not only better accounts for why scientific explanations work when they do, but can also help to clarify disputes within philosophy of biology and neuroscience between ontic and epistemic views of mechanistic explanation.

This new account is not without its own set of challenges, however. Evaluating whether an explanation is good or bad is substantially more difficult under this new framework. Evaluations must be done in a piecemeal way, and we may have to trust that some features of an explanation are better justified than others until given reason to think otherwise. While this makes the evaluation of explanations more difficult, it is in-line with how scientific practice in fact works.

Bibliography

- Achinstein, P., 1984. The pragmatic character of explanation. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1984(2), pp.274–292. <https://doi.org/10.1086/psaprocbienmeetp.1984.2.192509>.
- Achinstein, P., 1983. *The Nature of Explanation*. New York ; Oxford: Oxford University Press. Available at: <<http://archive.org/details/natureofexplanat0000achi>> [visited on 29 September 2023].
- Austin, C.J., 2017. The philosophy of biology. *Analysis*, 77(2), pp.412–432. <https://doi.org/10.1093/analys/anx032>.
- Batterman, R.W., 2000. Multiple realizability and universality. *The British Journal for the Philosophy of Science*, 51(1), pp.115–145. <https://doi.org/10.1093/bjps/51.1.115>.
- Batterman, R.W., 2002. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, *Oxford studies in philosophy of science*. Oxford: Oxford University Press.
- Batterman, R.W. and Rice, C.C., 2014. Minimal model explanations. *Philosophy of Science*, 81(3), pp.349–376. <https://doi.org/10.1086/676677>.

- Bechtel, W., 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Routledge.
- Bechtel, W., 2015. Can mechanistic explanation be reconciled with scale-free constitution and dynamics? *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 53, pp.84–93. <https://doi.org/10.1016/j.shpsc.2015.03.006>.
- Bogen, J., 2005. Regularities and causality; generalizations and causal explanations. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2), pp.397–420. <https://doi.org/10.1016/j.shpsc.2005.03.009>.
- Bokulich, A., 2018. Representing and explaining: The eikonic conception of scientific explanation. *Philosophy of Science*, 85(5), pp.793–805. <https://doi.org/10.1086/699693>.
- Braverman, M. et al., 2012. Intelligibility is necessary for scientific explanation, but accuracy may not be. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 34(34), pp.1368–1373. Available at: <<https://escholarship.org/uc/item/32m3q2sp>> [visited on 29 September 2023].
- Chirumuuta, M., 2014. Minimal models and canonical neural computations: the distinctness of computational explanation in neuroscience. *Synthese*, 191(2), pp.127–153. <https://doi.org/10.1007/s11229-013-0369-y>.
- Coffa, J.A., 1974. Hempel’s ambiguity. *Synthese*, 28(2), pp.141–163. <https://doi.org/10.1007/BF00485232>.
- Craver, C.F., 2006. When mechanistic models explain. *Synthese*, 153(3), pp.355–376. <https://doi.org/10.1007/s11229-006-9097-x>.
- Craver, C.F., 2014. The Ontic Account of Scientific Explanation. In: M.I. Kaiser, O.R. Scholz, D. Plenge and A. Hüttemann, eds. *Explanation in the Special Sciences: The Case of Biology and History*. Vol. 367, *Synthese Library*. Dordrecht: Springer Netherlands, pp.27–52. https://doi.org/10.1007/978-94-007-7563-3_2.
- Craver, C.F. and Darden, L., 2013. *In Search of Mechanisms: Discoveries across the Life Sciences*. Chicago, IL: University of Chicago Press. Available at: <<https://press.uchicago.edu/ucp/books/book/chicago/I/bo16123713.html>> [visited on 29 September 2023].

- Craver, C.F. and Kaplan, D.M., 2020. Are more details better? On the norms of completeness for mechanistic explanations. *The British Journal for the Philosophy of Science*, 71(1), pp.287–319. <https://doi.org/10.1093/bjps/axy015>.
- Craver, C.F. and Povich, M., 2017. The directionality of distinctively mathematical explanations. *Studies in History and Philosophy of Science Part A*, 63, pp.31–38. <https://doi.org/10.1016/j.shpsa.2017.04.005>.
- Douglas, H.E., 2009. Reintroducing prediction to explanation. *Philosophy of Science*, 76(4), pp.444–463. <https://doi.org/10.1086/648111>.
- Eichner, A.S., ed., 1983. *Why Economics is Not Yet a Science*. Armonk, N.Y: M.E. Sharpe.
- Elgin, C.Z., 2004. True enough. *Philosophical Issues*, 14(1), pp.113–131. <https://doi.org/10.1111/j.1533-6077.2004.00023.x>.
- Friedman, M., 1953. *Essays in Positive Economics*. Chicago, IL: University of Chicago Press.
- Gilpin, L.H. et al., 2022. "Explanation" is Not a Technical Term: The Problem of Ambiguity in XAI. arXiv. <https://doi.org/10.48550/arXiv.2207.00007>.
- Hall, C. et al., 2001. The need to reintegrate the natural sciences with economics: Neoclassical economics, the dominant form of economics today, has at least three fundamental flaws from the perspective of the natural sciences, but it is possible to develop a different, biophysical basis for economics that can serve as a supplement to, or a replacement for, neoclassical economics. *BioScience*, 51(8), pp.663–673. [https://doi.org/10.1641/0006-3568\(2001\)051\[0663:TNTRTN\]2.0.CO;2](https://doi.org/10.1641/0006-3568(2001)051[0663:TNTRTN]2.0.CO;2).
- Hausman, D.M., 2008. Why look under the hood? *The Philosophy of Economics: An Anthology*. Cambridge: Cambridge University Press, pp.183–187. <https://doi.org/10.1017/CBO9780511819025.012>.
- Heeger, D.J., 1992. Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9(2), pp.181–197. <https://doi.org/10.1017/S0952523800009640>.
- Hochstein, E., 2017. Why one model is never enough: a defense of explanatory holism. *Biology & Philosophy*, 32(6), pp.1105–1125. <https://doi.org/10.1007/s10539-017-9595-x>.

- Hochstein, E., 2019. How metaphysical commitments shape the study of psychological mechanisms. *Theory & Psychology*, 29(5), pp.579–600. <https://doi.org/10.1177/0959354319860591>.
- Hodgkin, A.L. and Huxley, A.F., 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4), pp.500–544. <https://doi.org/10.1113/jphysiol.1952.sp004764>.
- Illari, P., 2013. Mechanistic explanation: Integrating the ontic and epistemic. *Erkenntnis*, 78(S2), pp.237–255. <https://doi.org/10.1007/s10670-013-9511-y>.
- Izadi, S., Anandakrishnan, R. and Onufriev, A.V., 2014. Building water models: A different approach. *The Journal of Physical Chemistry Letters*, 5(21), pp.3863–3871. <https://doi.org/10.1021/jz501780a>.
- Kaplan, D.M., 2011. Explanation and description in computational neuroscience. *Synthese*, 183(3), pp.339–373. <https://doi.org/10.1007/s11229-011-9970-0>.
- Keen, S., 2022. The appallingly bad neoclassical economics of climate change. In: B.K. Gills and J. Morgan, eds. *Economics and Climate Emergency, Rethinking globalizations*. London; New York: Routledge, pp.79–107.
- Lipton, P., 2004. *Inference to the Best Explanation*. 2nd ed., *International library of philosophy*. London; New York: Routledge/Taylor and Francis Group.
- Lombrozo, T., 2016. Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, 20(10), pp.748–759. <https://doi.org/10.1016/j.tics.2016.08.001>.
- Machamer, P., Darden, L. and Craver, C.F., 2000. Thinking about mechanisms. *Philosophy of Science*, 67(1), pp.1–25. <https://doi.org/10.1086/392759>.
- Mackonis, A., 2013. Inference to the best explanation, coherence and other explanatory virtues. *Synthese*, 190(6), pp.975–995. <https://doi.org/10.1007/s11229-011-0054-y>.
- Mantzavinos, C., 2016. *Explanatory Pluralism*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781316440599>.

- Parker, G., 1978. Searching for Mates. In: J.R. Krebs and N.B. Davies, eds. *Behavioural Ecology: An Evolutionary Approach*. Oxford: Blackwell Scientific, pp.214–244.
- Putynski, A., 2013. Cancer and the goals of integration. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 44(4), pp.466–476. <https://doi.org/10.1016/j.shpsc.2013.03.019>.
- Potochnik, A., 2010. Explanatory independence and epistemic interdependence: A case study of the optimality approach. *The British Journal for the Philosophy of Science*, 61(1), pp.213–233. <https://doi.org/10.1093/bjps/axp022>.
- Potochnik, A., 2015. The diverse aims of science. *Studies in History and Philosophy of Science Part A. Special Section: Formal Epistemology and the Legacy of Logical Empiricism*, 53, pp.71–80. <https://doi.org/10.1016/j.shpsa.2015.05.008>.
- Povich, M., 2018. Minimal models and the generalized ontic conception of scientific explanation. *The British Journal for the Philosophy of Science*, 69(1), pp.117–137. <https://doi.org/10.1093/bjps/axw019>.
- Povich, M. and Craver, C.F., 2018. Because without cause: Non-causal explanations in science and mathematics (review). *Philosophical Review*, 127(3), pp.422–426. <https://doi.org/10.1215/00318108-6718870>.
- Rice, C., 2015. Moving beyond causes: Optimality models and scientific explanation. *Noûs*, 49(3), pp.589–615. <https://doi.org/10.1111/nous.12042>.
- Rice, C. and Rohwer, Y., 2020. How to reconcile a unified account of explanation with explanatory diversity. *Foundations of Science*, 26(4), pp.1025–1047. <https://doi.org/10.1007/s10699-019-09647-y>.
- Ruben, D.-H., 1990. *Explaining Explanation*. 1st ed., *The problems of philosophy: their past and present*. New York: Routledge.
- Salmon, W.C., 1989. Four decades of scientific explanation. In: P. Kitcher and W.C. Salmon, eds. *Scientific Explanation, Minnesota studies in the philosophy of science*, 13. Minneapolis: University of Minnesota Press, pp.3–219. Available at: <<http://conservancy.umn.edu/handle/11299/185700>> [visited on 4 October 2023].

- Strevens, M., 2011. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- Waskan, J., 2011. Intelligibility and the CAPE: Combatting Anti-Psychologism about Explanation. Available at: <<http://philsci-archival.pitt.edu/8530/>> [visited on 4 October 2023].
- Waskan, J., Harmon, I., Higgins, A. and Spino, J., 2014. Three senses of 'Explanation'. In: P. Bello, M. Guarini, M. McShane and B. Scassellati, eds. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society, pp.3090–3095.
- Woods, J. and Rosales, A., 2010. Virtuous Distortion. In: L. Magnani, W. Carnielli and C. Pizzi, eds. *Model-Based Reasoning in Science and Technology: Abduction, Logic, and Computational Discovery, Studies in Computational Intelligence*. Berlin, Heidelberg: Springer, pp.3–30. https://doi.org/10.1007/978-3-642-15223-8_1.
- Woody, A.I., 2015. Re-orienting discussions of scientific explanation: A functional perspective. *Studies in History and Philosophy of Science Part A*, 52, pp.79–87. <https://doi.org/10.1016/j.shpsa.2015.03.005>.
- Wright, C., 2015. The ontic conception of scientific explanation. *Studies in History and Philosophy of Science Part A*, 54, pp.20–30. <https://doi.org/10.1016/j.shpsa.2015.06.001>.
- Wright, C. and van Eck, D., 2018. Reconciling ontic and epistemic constraints on mechanistic explanation, epistemically. *Ergo, an Open Access Journal of Philosophy*, 5(38). <https://doi.org/10.3998/ergo.12405314.0005.038>.