Beyond Explanation: Understanding as Dependency Modeling

Finnur Dellsén

Forthcoming in *The British Journal for the Philosophy of Science*
(please cite published version when available)

Abstract

This paper presents and argues for an account of objectual understanding that aims to do justice to the full range of cases of scientific understanding, including cases in which one does not have an explanation of the understood phenomenon. According to the proposed account, one understands a phenomenon just in case one grasps a sufficiently accurate and comprehensive model of the ways in which it or its features are situated within a network of dependence relations; one's degree of understanding is proportional to the comprehensiveness and accuracy of such a model. I compare this account with accounts of scientific understanding that explicate understanding in terms of having an explanation of the understood phenomenon. I discuss three distinct types of cases in which scientific understanding does not amount to possessing an explanation of any kind, and argue that the proposed model-based account can accommodate these cases while still retaining a strong link between understanding and explanation.

Keywords: objectual understanding; scientific explanation; non-explanatory understanding; mental modeling; dependence relations.

1 Introduction

Objectual understanding is a type of cognitive achievement canonically expressed by sentences of the form 'S understands P', where P is a phenomenon. *Understanding-why*, by contrast, is the type of achievement canonically expressed by sentences of the form 'S understands why E', where E is an explanandum (something to be explained). Understanding why E is clearly closely associated with having an explanation of E, but it is less clear that objectual understanding requires having or knowing an explanation. While such *explanatory accounts* of objectual understanding enjoy considerable popularity (e.g., Strevens, 2013; de Regt, 2017; Khalifa, 2017), others have argued that objectual understanding of P does not require that one has an explanation of any aspect of P (e.g., Lipton, 2009; Kvanvig, 2009; Gijsbers, 2013). Two non-explanatory views have gained traction, viz. a *manipulationist account* that construes objectual understanding in terms of an ability to manipulate a representation of P for various purposes (Wilkenfeld, 2013, 2015), and an *epistemic account* that construes objectual understanding as an unusually comprehensive type of propositional knowledge (Kelp, 2015, 2017).

This paper develops and argues for a new account of objectual understanding, based on the idea that understanding a phenomenon amounts to having a model of its dependence relations. The account differs from extant non-explanatory accounts in retaining a crucial insight of explanatory accounts, viz. that understanding consists in locking onto the underlying dependence relations that typically undergird correct explanations. According to the proposed account, one understands a phenomenon P just in case one grasps a sufficiently accurate and comprehensive model of the network of dependence relations in which P, or its contextually relevant parts, is situated; and one's degree of understanding of P is proportional to the comprehensiveness and accuracy of such a model. As we will see, it's possible to increase both the accuracy and the comprehensiveness of such a dependency model of P without learning an explanation of any aspect of P. Accordingly, this account of understanding accommodates the possibility of achieving understanding through means other than explanations.

I proceed as follows. I start by clarifying the explicandum, viz. what's called 'objectual understanding' (§2), and then present and elaborate on my preferred model-based account (§3). The rest of the paper compares this account with explanatory accounts (defined in §4). In particular, I discuss three types of cases in which understanding comes apart from explanation, thus illustrating the limitations of explanatory accounts (§§5-7). For each type of case, I show how the model-based account I favor allows us to attribute understanding to the agent, thus delivering the more plausible verdict about such cases. I conclude with a brief discussion of how the key insight of explanatory accounts – that understanding is closely related to explanation – can be accommodated within a plausible non-explanatory account (§8).

2 Objectual Understanding

The topic of this paper is the type of understanding that is canonically attributed to someone with a sentence of the form 'S understands P', where P is a phenomenon. This type of understanding is usually referred to as *objectual understanding*. The term is slightly misleading since the target phenomenon P need not be a single object; rather, it may be a complex system that is itself most naturally described as being composed of several interacting objects. Thus, for example, one can be said to understand complex systems like machines, economies, and human minds – all of which are naturally thought of as systems with interacting parts. In using the term 'phenomenon' for the target of objectual understanding, I do not mean to imply that these targets must be observational or empirical. Atoms, quarks, and electric fields can be understood, even though they cannot be (directly) observed. Indeed, I take it that everything that exists is, or is part of, some phenomenon or other – and can thus conceivably be the target of someone's understanding in the relevant sense.

I am following Kelp (2015, 2017) in characterizing objectual understanding as understanding of phenomena. Some have instead characterized objectual understanding as understanding of a 'topic' (Brun and Baumberger, 2017) or 'subject matter' (Carter and Gordon, 2016; Khalifa, 2017). However, that encourages the slip from understanding something to understanding a discipline that studies it, as in 'she understands quantum physics'. Since a discipline could afford a very poor understanding of the phenomena it studies (as in the case of alchemy, for example), this distinction should not be overlooked. For similar reasons, it is important to distinguish objectual understanding from *understanding of theories* – what Newman (2017) helpfully calls 'theoretical understanding.' After all, it is one thing to understand a phenomenon itself and quite another to understand the theories that may or may not be true of those phenomena. For example, I can have perfect understanding of the phlogiston theory of heat, but that gives me a limited understanding of heat itself (since the theory is radically false).[1]

Objectual understanding is often contrasted with the type of understanding that is canonically attributed to someone with a sentence of the form 'S understands why

[1]With that said, understanding of theories may or may not be a necessary requirement, or perhaps a species of, objectual understanding. For reasons of space, I won't discuss the relation between objectual understanding and understanding of theories further in this paper.

E', where E describes an explanandum.[2] This is known as *understanding-why* – or sometimes, since having understanding of this kind clearly involves having or knowing an explanation, *explanatory understanding*. Although much of the recent philosophical enthusiasm about understanding has explicitly concerned this kind of understanding (e.g., Hills, 2009, 2016; Pritchard, 2009, 2010; Grimm, 2010, 2012, 2014), it's also clear that understanding-why has already received a great deal of attention in the literature on scientific explanation and explanation-seeking why-questions. Thus, if understanding-why is the only kind of understanding that is of philosophical interest, or if other kinds of understanding can be reduced to understanding-why without loss (as is argued by Khalifa, 2013a), then the topic of understanding may not be as novel or interesting as many epistemologists and philosophers of science have suggested (see esp. Khalifa, 2012, 2017). This paper in effect pushes back against attempts to reduce objectual understanding to understanding-why by arguing that, in contrast to understanding-why, objectual understanding does not necessarily involve having or knowing an explanation of the understood phenomenon or any of its features.

Let me highlight two further features of (objectual) understanding[3] that matter to my arguments below. First, understanding is a matter of degree in a way that propositional knowledge, for example, is not. It's not just that one can understand more or fewer phenomena; rather, one can have more and less (or, if you prefer, 'better' and 'worse') understanding of a single phenomenon P. Thus it makes sense to say, for example, that an experienced physicist has a greater understanding of Brownian motion than her graduate student; who in turn has a greater understanding of this phenomenon than the freshmen taking her introductory courses; and so forth. A corollary of this gradability of understanding is that our understanding of a phenomenon can (and typically does) *improve incrementally* over time. Second, I will assume that paradigmatic cases of understanding are to be found in the natural sciences – thus, what follows can be viewed as a discussion of *scientific* understanding. Importantly, this means that a central adequacy condition for an account of understanding is that it makes sense of how understanding is in fact achieved in the

[2]Objectual understanding should also be distinguished from *linguistic understanding*, e.g. understanding a sentence in a particular language (see Longworth, 2009). Furthermore, it is unclear whether objectual understanding is the type of understanding one can have of other persons (see Grimm, 2016). Since my concern is with understanding in natural science, I don't take a stand on that issue here.

[3]From now on, I will often omit 'objectual' from 'objectual understanding' to save space. To prevent misunderstandings, I will not do this for 'understanding-why'.

empirical sciences.[4]

3 Understanding as Dependency Modeling

The account of understanding that I propose holds that to understand a phenomenon is to grasp a specific kind of model of that phenomenon's dependence relations. I will develop this *dependency modeling account* of understanding in stages, starting with the term 'model'.

For my purposes, a *model* is simply an information structure of some kind that is interpreted so as to represent its target. These information structures can be *concrete*, as in Watson and Crick's original model of DNA, or *abstract*, as in mathematical and computational models like the Lotka-Volterra model of predation in ecological systems. In both cases, the structures are associated with an intended interpretation that specifies how the different parts of the structure correspond to different elements and relations in the phenomenon – a 'key' (Frigg and Nguyen, 2016). In the Lotka-Volterra model, for example, its differential equations contain variables that are meant to correspond to determinables of the real system, such as the sizes of the predator and prey populations. Note that these equations are not themselves a model; rather, they become part of a model once the variables therein have been interpreted.[5] The same is true for other information structures – they become a model only once different parts of the structures have been associated with specific parts of the phenomenon in question.

Now, since understanding is something that happens in our minds, a model-based account of understanding requires that the relevant kind of models must somehow be present in thought. There is indeed an influential, though somewhat non-standard, psychological account of mental representation which holds that agents construct so-called 'mental models' to use in reasoning and deliberation (Craik, 1943; Johnson-Laird, 1983, 2013; Nersessian, 2007; Thagard, 2010). Some theorists even propose to replace the traditional picture of the mind as occupied by propositional attitudes, with a picture on which thought is based on mental models (Waskan, 2006). However,

---

[4]Indeed, understanding is arguably an important cognitive aim of science (Potochnik, 2015, 2017; de Regt, 2017), and scientific progress at least sometimes consists in increasing understanding of natural phenomena (Bangu, 2015; Dellsén, 2016; Stuart, 2016).

[5]Here I am roughly following Weisberg (2013) in thinking of models as interpreted structures. I should note that the kind of modeling I have in mind is what Weisberg refers to as 'target-directed modeling' (Weisberg, 2013, esp. ch. 5).

it is ultimately an empirical question whether the mind should be thought of as containing such models (in addition to, or instead of, propositional attitudes), and I would prefer to stay neutral on this issue. Fortunately, as far as the current account of understanding is concerned, all that's required is that models can be *represented in* the mind, e.g. by sets of propositions that describe these models. As a shorthand for the relation between the mind and the models – whatever it turns out to be – I will use the term 'grasp'.[6] Thus we can say that, on the model-based account I have in mind, understanding consists of grasping a certain kind of model of the understood phenomenon.

What kind of model must an understanding agent grasp? By their nature, all models are incomplete or inaccurate representations of their targets. As Borges's "On Exactitude in Science" famously illustrated, only an exact copy of something can represent all its aspects in an accurate way. So which aspects of a phenomenon must be represented by an understanding agent's model? My answer to this question is inspired by previous and related work on understanding, which connects understanding-why with *dependence relations* (Kim, 1974; Greco, 2014; Grimm, 2014).[7] Building on this idea, I say that the aspects of a phenomenon that matter for understanding are the dependence relations that the phenomenon, or its features, stands in towards other things. The most well-known kind of dependence relation is *causality* – an effect depends on its cause – but there are arguably other kinds of dependence relations as well. One of these arguably non-causal dependence relations is *grounding* – i.e., the in-virtue-of relation (Fine, 2001; Schaffer, 2009; Audi, 2012). There might very well be other non-causal dependency relations that do not reduce to either causation or ground, e.g. the relation between a foundational mathematical axiom and a theorem it explains (see, e.g., Steiner, 1978), the kind of relation there is between Socrates's death and the widowing of Xanthippa (Kim, 1974), or the kind of relation there is between the curvature of space time and the shape of an undisturbed electron cloud traveling through that space (Nerlich, 1976). At any

[6]I am not alone in using the term 'grasp' as a placeholder for a relation constitutive of understanding that remains to be further specified (see, e.g., Strevens, 2013, 511-512). See Bourget (2015) for a systematic discussion of the relevant notion of 'grasping'.

[7]With that said, the position developed in this paper differs from previous accounts of understanding in terms of dependence relations in showing why and how objectual understanding can come apart from having an explanation (see §§5-7). Indeed, this possibility would seem to be excluded by both Greco and Grimm, who construe understanding as *knowledge of causes* (where 'cause' is defined broadly so as to include grounds and other factors on which something might depend). After all, someone who possesses knowledge of the kind required by these accounts would seem to have an explanation of the understood phenomenon or its relevant features.

rate, the kind of model that I think is involved in understanding is one that aims to capture the network of dependence relations that a phenomenon stands in, whatever these relations turn out to be. I will refer to this as a *dependency model*.

To get a sense of what I mean by a dependency model, consider causal graphs (also known as 'causal maps'). Causal graphs are pictorial representations of models that are meant to capture the causal relations in some given system. Specifically, they purport to tell us which elements of the system are causes and effects of which other elements, and also which elements are not either causes or effects of specific other elements. In other words, they are meant to contain both 'positive' information about the causal relations that *are* present within a system, and also 'negative' information about which parts of the system *are not* causally related. Now, since causal graphs are only concerned with one kind of dependence relation, they can only afford a partial understanding of the target phenomenon (and, of course, they do so only when the agent has indeed grasped the model depicted by the graph). However, it's easy to see how a more complete model can be constructed by representing other dependence relations, such as grounding. Thus a pictorial representation of what I am calling a dependency model would be a graph which purports to depict how each element depends, or does not depend, on each other element – causally or otherwise.

Of course, to have understanding of a phenomenon P, it is not enough to grasp any old dependency model of P. Rather, the model must in some sense be a *good representation* of the relevant dependence relations. So what makes such a model better or worse qua representation? My answer is simple: A dependency model better represents P to the extent that the network of dependence relations that P stands in is correctly depicted by the model. Since a dependency model can thus fail either by incorrectly representing (i.e. misrepresenting) some aspect of this network, or by not representing it at all, we can identify two separate criteria here, viz. *accuracy* and *comprehensiveness*. These two criteria can come into conflict: Increasing comprehensiveness may require us to to sacrifice accuracy, in which case we engage in a kind of *idealization*; conversely, to increase accuracy we may have to sacrifice our model's comprehensiveness, in which case we engage in *abstraction*. Since understanding is a function of how well a model does on both accuracy and comprehensiveness at once, it is sometimes possible to increase one's understanding by sacrificing either accuracy or comprehensiveness respectively, provided that doing so brings sufficient benefits with regard to the other criterion. Incidentally, this helps to explain why idealizations and abstractions can both increase understanding.

I have noted that understanding is a gradable notion – that one can have various degrees of understanding of the same phenomenon. In a model-based account of the sort I am proposing, this is explained by the fact that the two aforementioned criteria (accuracy and comprehensiveness) are both gradable.[8] Thus one's degree of understanding of some phenomenon P is proportional to the comprehensiveness and accuracy of one's dependency model of P. Of course, we commonly use the term 'understanding' in what seems to be a binary way. Such a binary concept of understanding can be explicated in terms of a degree of understanding that exceeds some threshold t, where the t is presumably contextually-determined (e.g. by the speaker's context of utterance).[9] Thus when someone says 'S understands P', we can take that to mean that S has a degree of understanding that is higher than such a threshold t, i.e. that S grasps a model of P that is sufficiently accurate and comprehensive to exceed the thresholds determined by the speaker's context of utterance.

Context also enters into our usage of the term 'understanding' in a couple of other ways. First, context plausibly determines which parts of a complex phenomenon need to be understood to a significant degree in order for it to be felicitous to say that the phenomenon itself is understood. For example, what it means to understand human mating will differ substantially depending on whether one is speaking to a physician, a psychologist or an ecologist. Relatedly, context plausibly determines which other phenomena (or parts thereof) are sufficiently salient for their dependence relations to the target phenomenon P, or lack thereof, to be relevant to whether someone understands P. For example, most events will have enormous causal histories tracing back to the beginning of the universe, but only a small subset of the causes will be sufficiently salient in a given context for one's dependency model to contribute more understanding of such events if one correctly represents the dependence relation between the events and the relevant cause. Similarly, only a small subset of the phenomena that a given target phenomenon P does *not* stand in a dependence relation towards will be sufficiently salient in a given context for the lack of such a

[8] Additionally, the gradability of understanding can be explained by the gradability of 'grasping', i.e. by the fact that one can have a stronger or weaker grasp of the dependency model in question. I won't discuss this feature further in this paper, since I am not focusing here on the psychological aspects of understanding.

[9] The context-sensitivity of understanding is also discussed by Wilkenfeld (2013, 1007-9) and Kelp (2015, 3813-14)

relation to contribute anything to one's understanding of P.[10]

With these points in mind, the dependency modeling account of understanding I am proposing – *DMA*, for short – can be summarized as follows:

> **DMA:** S understands a phenomenon P if and only if S grasps a sufficiently accurate and comprehensive dependency model of P (or its contextually relevant parts); S's degree of understanding of P is proportional to the accuracy and comprehensiveness of that dependency model of P (or its contextually relevant parts).

Since dependence relations are, at least normally, what undergirds explanations, there is still a strong connection between understanding and explanation on this view. So there is a kernel of truth in explanatory accounts of objectual understanding. However, whereas explanatory accounts take explanation to be a necessary component of such understanding, the current account allows for understanding to come apart from explanation. We will examine a number of cases that illustrate this in various ways below – cases that support DMA and undermine any account that makes explanation a requirement for understanding.[11]

4 Explanatory Accounts

Having sketched my dependency modelling account (DMA) of objectual understanding, I now turn a class of accounts that I will take as my main target in this paper, viz. explanatory accounts. Roughly speaking, an *explanatory account* of understanding takes explanation to be a necessary condition on understanding. A particularly clear example of an explanatory account is Strevens's 'simple view', where having a

[10] For example, there is almost no imaginable circumstances in which the number of grains of sand on one of the Maldives Islands is salient in the context of understanding the 2016 US presidential election. That is why grasping the independence between the two phenomena would almost never constitute any degree of understanding of the 2016 US presidential election. I say 'almost never' because it is *possible* to force the number of grains of sand on one of the Maldives Islands to become salient, e.g. if the number of grains of sand on the Maldives Islands were (for whatever strange reason) known to have been a significant causal factor in the previous (i.e. 2012) US presidential election. In that case, grasping the independence between the numbers of grains of sand and the 2016 elections would indeed constitute *some* understanding of the 2016 elections. (Thanks to a reviewer for pressing me to clarify my account in response to this example.)

[11] I lack the space here to discuss other non-explanatory accounts of objectual understanding, such as Wilkenfeld's (2013; 2015) manipulationist account and Kelp's (2015; 2017) epistemic account. However, it should be clear that these accounts differ from the current account since DMA does not define understanding in terms of manipulability or knowledge.

correct explanation is taken to be both necessary and sufficient for understanding:

> An individual has scientific understanding of a phenomenon just in case they grasp a correct scientific explanation of that phenomenon" (Strevens, 2013, 510).[12]

However, I will also classify as explanatory a wide range of accounts that posit a less tight connection between explanation and understanding. These include accounts on which there are non-explanatory requirements for understanding, so that having or grasping an explanation is not sufficient for understanding. They also include accounts on which the explanations need not be correct, but instead merely 'adequate' in some broader sense (e.g. empirically adequate, or well-confirmed by scientific evidence).

To make this demarcation between explanatory and non-explanatory accounts more precise, I will count as explanatory any account of understanding that satisfies the following condition:

U→E: S understands P only if S grasps enough of an adequate explanation of P (or its relevant features); *other things being equal*, S has more understanding of P to the extent that S grasps more of an adequate explanation of P (or its relevant features).

This requirement is accepted by many – probably most – philosophers who have written about scientific understanding and explanation (e.g., Salmon, 1993; Nounou and Psillos, 2012; Khalifa, 2012, 2013a, 2017; Strevens, 2013; de Regt, 2017; de Regt and Gijsbers, 2016). Indeed, U→E is even accepted by some who do not particularly highlight the role of explanation in understanding, but who nevertheless make explanation a necessary condition on understanding. Consider, for example, de Regt's most recent account of scientific understanding:

> A phenomenon P is understood scientifically if and only if there is an explanation of P that is based on an intelligible theory T and conforms to the basic epistemic values of empirical adequacy and internal consistency (de Regt, 2017, 92).

[12]Given that this is meant to specify when someone "has understanding of a phenomenon" (rather than, e.g. when someone *understands why* something is the case), I take this to be an account of objectual understanding rather than (or perhaps in addition to) an account of understanding-why. To be fair, however, Strevens does not discuss the distinction between objectual understanding and understanding-why.

While the focus of de Regt's work has largely been to develop and defend his account of understanding as involving intelligibility and conformism to empirical adequacy and consistency (as opposed to truth), his account is clearly explanatory in the sense of making explanation a necessary requirement for understanding (as per U$\rightarrow$E).

One might even wonder whether my own account of understanding – DMA – satisfies U$\rightarrow$E and thus counts as an explanatory account by these lights. There is some truth in this thought, since a dependency model typically contains the kind of information that can be used in adequate explanations, such as information about causal relations. But in order for DMA to count as explanatory by the lights of U$\rightarrow$E, it would have to be the case that any information about a phenomenon's dependency relations would also be part of an explanation for the phenomenon (or its relevant features). Below I discuss various cases in which that is not so. To foreshadow, this is roughly because the information about dependence-relations that serves to increase our understanding in these cases consists of, or implies, the fact that the understood phenomenon cannot be explained – either by some specific other fact, or at all. The upshot, then, will be that it's possible to gain understanding of a phenomenon without gaining any information that would serve to explain that phenomenon or its relevant features.

In what follows, I discuss three distinct ways in which U$\rightarrow$E fails. In each case, scientists gain understanding of a phenomenon without obtaining any information that would serve to explain that phenomenon (or its relevant features). I will argue that DMA can easily account for cases of these types without invoking any additional theoretical machinery. By contrast, explanatory accounts can at best attempt to explain away these types of cases by adding theoretical baggage that brings its own set of problems for these accounts. So while the types of cases I discuss below do not conclusively refute explanatory accounts of understanding, I do maintain that they shift the overall argumentative burden so as to favor non-explanatory accounts such as DMA.

5 Explanatory Bruteness

The first type of cases in which U$\rightarrow$E fails involves what I will call *explanatory bruteness*. These are phenomena that have no explanation at all – phenomena that are not merely unexplained, but unexplainable.

An interesting subset of the kind of cases I am interested in here has previously been discussed by Kvanvig (2009) in the context of arguing against a reduction of objectual understanding to understanding-why. Kvanvig asks us to imagine a system in which it is indeterminate whether an electron will go to the left or to the right, with an equal chance of each option materializing. Kvanvig says that the path of the electron cannot be explained in such a system, since its path will by construction be "the result of chance rather than of causation" (Kvanvig, 2009, 101). I agree with Kvanvig on this point, but a fuller discussion of the example would need to include a story of how exactly the system *is* understood.[13] I believe DMA provides exactly that story: The dependence relations in an indeterministic system can be *modeled*, no less than those in a deterministic system. Thus, if we grasp the appropriate kind of model, i.e. a sufficiently accurate and comprehensive dependency model, then we understand it. Importantly, note that a model is more comprehensive if it includes rather than excludes the information that the electron's going right rather than left is a 50-50 chance event, the outcome of which does *not depend* on anything else in the system. Thus gaining this crucial information increases one's understanding of the system according to DMA.

So in my view, Kvanvig's example provides a nice illustration of the benefits of DMA over explanatory accounts. However, the dialectical potency of Kvanvig's example has been contested at length by Khalifa (2013a), who argues that the electron's going left rather than right can in fact be explained after all. Khalifa's line of argument starts by pointing out that the non-contrastive fact that the electron went to the left can be explained probabilistically according to some influential theories of probabilistic explanation, such as Railton's (1978; 1981). Khalifa further argues that even the contrastive fact that the electron went *left rather than right* can be explained in the same way. This is implausible on its face, since explaining a contrastive fact – that A rather than B – seems to require an explanans that makes A more likely to occur than B, and by stipulation there is no such explanation in the present case. But Khalifa responds that this

> [...] confuses the source of explanatory relevance or difference making. An explanation needn't make the probabilities between contrasted outcomes different from each other; rather, these probabilities must be different than they would be had the explanans been different (Khalifa, 2013a,

[13]Kvanvig is surprisingly silent on this point, though he does say of understanding generally that it involves having a "grasp of the structural relationships (e.g. logical, probabilistic, and explanatory relationships) between the central items of information regarding which the question of understanding arises" (Kvanvig, 2009, 97).

1163).

But I see no confusion here, for the claim that the probabilities for different outcomes in contrastive indeterministic explanations must be different is perfectly compatible with the claim that an explanans must make a difference to (i.e. raise or lower) the probabilities of the outcomes. Indeed, some of the most influential discussions of contrastive explanations, such as Sober (1984) and van Fraassen (1980), explicitly make the stronger requirement that such explanations must make the explanandum more probable than the alternatives in the relevant contrast class.

Moreover, I don't think Khalifa's defense generalizes to understanding of other indeterministic systems. Consider a modification of Kvanvig's example in which there is a much greater (e.g., a trillion times greater) chance that the electron goes right than left; nevertheless, the electron went left. By the same token as before, the system can be understood. But can the fact that the electron went left (rather than right) be explained? Such an explanation would have to reveal how extremely unlikely the electron is to have gone left, thus making the event even more surprising than it would have seemed before. But such a small probability, if it can be said to explain at all, surely provides a worse explanation than a larger counterpart. This is not just 'intuitively' so, but also strongly suggested by scientific practice – in particular, by statistical mechanics – in which conveying a larger probability on events is taken to provide a better explanation for those events (Strevens, 2000). So in this modified Kvanvig-style case we would at best have an extremely poor explanation for the electron's trajectory. Contrary to Khalifa (2013b, 169), understanding thus clearly cannot be identified with having a "reasonably good" explanation.[14] And of course it is even less plausible that the contrastive fact that the electron went left rather than right can be adequately explained by citing a probability that so strongly favors the contrasting event over the explanandum.

Proponents of explanatory accounts of understanding could of course simply insist that the relevant (contrastive) outcomes can be adequately explained after all. But this raises the issue of what exactly is required for indeterministic (contrastive) explanation by their lights. As we have seen, Khalifa is explicit on this point, saying that the "probabilities [of contrasted outcomes] must be different than they would be had the explanans been different" (Khalifa, 2013a, 1163). Specifically, Khalifa

[14]In his most recent work, Khalifa (2017) does not explicitly talk about "reasonably good" explanations, though he does say that the kind of understanding he views as fundamental – explanatory understanding – is "characteristic of good explanations" (Khalifa, 2017, 2).

appeals to Hitchcock's (1999, 587) account of contrastive indeterministic explanation, which holds that A is 'explanatorily relevant'[15] to the contrastive outcome that E rather than F, relative to background conditions B, just in case:

$$P\big(E|(A\&B)\&(E \vee F)\big) \neq P\big(E|B\&(E \vee F)\big).$$

This is an extremely liberal conception of contrastive explanatory relevance; it counts A as explanatorily relevant to E rather than F if there is *any* change in the probability of E when we conditionalize on A – including a *lowering* of this probability. Nevertheless, let us grant it for the sake of the argument. The question is whether this conception makes it impossible to construct Kvanvig-style cases in which understanding comes apart from explanation (or explanatory relevance).

It does not, for Kvanvig's case can be modified so that the probabilities of the contrasted outcomes are the same before and after conditionalizing on A. For instance, consider a variant on Kvanvig's original case in which (it is part of the background conditions B that) the electron *either* goes through the original indeterministic process[16] (whatever that process was) *or*, alternatively, goes through another indeterministic process for which the probabilities of the electron going left and right are exactly the same as in the first process.[17] So, for example, if there was an equal chance of the electron going left rather than right in the first process, then there is also an equal chance of these two events in the second process. Now, in such a case, the information that the electron went through the first process, for instance, would not be explanatorily relevant to the electron going left rather than right according to Hitchcock's criterion, since the probability of the event would not change at all when we conditionalize on this information.[18] Nevertheless, by the same token as before, it would surely increase our understanding of the event to find out which of

[15] I will set aside the fact that Hitchcock's condition is meant to explicate what it is for something to be 'explanatorily relevant', not what it is for something to explain something else.

[16] In using the term 'process', I do not mean to assume that later events in the process depend on earlier events (as an effect depends on its cause). Indeed, for my purposes here, 'process' can be replaced with 'sequence of events (which may or may not depend on each other)'.

[17] It can be assumed that it is itself indeterministic whether the electron goes through the first or the second process, although I don't think that is necessary for the case to work as it should.

[18] Let A be the claim that the electron went through the first process; and let E and F be the claims that the electron went left and right respectively. Given our background conditions B, $\neg A$ is equivalent to the claim that the electron went through the second process. By construction, our case is such that:

$$P\big(E|(A\&B)\&(E \vee F)\big) = P\big(E|(\neg A\&B)\&(E \vee F)\big).$$

the two indeterministic processes the electron actually went through.

So I have my doubts about Khalifa's response to Kvanvig's example of understanding of indeterministic systems. However, I will now set these doubts aside, because I believe that an argument similar to Kvanvig's can be made by appealing to cases that do not in involve any indeterminacy. What I have in mind are *explanatorily brute* facts – facts that aren't merely currently unexplained, but for which there is no (correct) explanation, deterministic or otherwise.[19] There are numerous plausible candidates for such explanatorily brute facts, ranging from various mundane coincidences to fundamental physical truths (for some examples, see Owens, 1992; Barnes, 1994; Fahrbach, 2005; Lando, 2017). For any such fact, there is a phenomenon that is better understood if one's representation of the phenomenon includes a depiction of the very fact that makes it unexplainable, viz. that its occurrence or existence does not depend on anything else. Any such case undermines U→E. Indeed, note that in so far as we are seeking an account of understanding that is not merely extensionally adequate in the actual world, but is also adequate to counterfactual scenarios (such as when we ask whether a a given false theory, e.g. the phlogiston theory of combustion, would have provided scientific understanding if it were true), we have here an argument against explanatory accounts if it is merely *possible* that there are explanatorily brute facts.[20]

Of course, it is very much an unsettled empirical question whether a given fact is explanatorily brute. However, just for illustrative purposes, consider what physicists often consider to be a plausible candidate for explanatorily brute facts, viz. the values of the dimensionless fundamental physical constants. These are constants that play an indispensible role in fundamental physical theory and whose values cannot themselves be explained by any currently accepted scientific theory. Accordingly, the only known way to determine the values of these constants is via measurement. Fortunately, since these constants are dimensionless, their values are the same in

---

Using Bayes's Theorem, it is easy to prove that this entails:

$$P\big(E|(A\&B)\&(E \vee F)\big) = P\big(E|B\&(E \vee F)\big),$$

which negates Hitchcock's criterion.

[19] Barnes (1994, 62) refers to facts of this kind as "ontologically brute".

[20] I don't see any reason to reject the possibility of explanatorily brute facts, except perhaps on the grounds of some sort of *a priori* commitment to the principle of sufficient reason (PSR). While I suppose proponents of explanatory accounts could avoid the argument of this section by claiming that the PSR is (necessarily) true, that would be a very high price to pay for avoiding just one of three arguments against U→E.

any system of measurement. For example, one such constant is the *electromagnetic coupling constant* (also known as the 'fine structure constant'), $\alpha \approx 1/137$, which describes the strength of electromagnetic forces in our universe. Now, suppose there is no (correct) explanation for this value – that the value of α is explanatorily brute.[21] If so, it seems clear that we can increase our understanding of a number of phenomena – most obviously, of electromagnetic interactions – by discovering that this is the case. To see this most clearly, contrast a scientist (or a scientific community) who has made this discovery with one who is mistaken or in the dark about whether the value of the electromagnetic coupling constant can be explained. Clearly, the latter's understanding of electromagnetic interactions is wanting or defective in a way that the former's is not.

So brute facts are a problem for explanatory accounts of understanding. Admittedly, they do not constitute a knock-down argument against explanatory accounts, since their proponents can always bite the bullet and deny us understanding in these cases. At the very least, however, explanatorily brute facts count in favor of the non-explanatory account proposed here, DMA, as against its explanatory counterparts. After all, note that in cases of explanatorily brute facts, understanding comes from realizing that the phenomenon in question is not caused, grounded, or otherwise dependent on anything else. Such cases are easily accommodated by DMA, since a dependency model that depicts such a phenomenon or its features as not dependent on anything else would be more accurate than an otherwise identical model that represents them as dependent on something else, and more comprehensive than an otherwise identical model that abstracts away from the issue. Indeterministic systems of the kind discussed by Kvanvig (2009) and Khalifa (2013a) are a special case of this type: In such systems, the explanatorily brute feature is an outcome of an indeterministic chance process, such as the electron's going left rather than right. This outcome is not dependent on anything else in the system, even if the outcome's probability of occurring clearly is.

6 Explanatory Targetedness

I turn now to a second class of cases in which understanding comes apart from explanation. In these cases, we come to understand through grasping an explanation,

[21] Given the current state of physical theory, this may well be true; although again I emphasize that it is an open empirical question and one that may never be settled with certainty.

but the explanation helps us understand the explanans rather than the explanandum. Thus, in these cases, the target of one's understanding differs from the target of one's explanation in a way that separates understanding of P from grasping an explanation of P.

As a case in point, consider the transition from Rutherford's planetary model of the atom to Bohr's quasi-quantum model. In both models, the atom is depicted as consisting of negatively charged electrons that orbit a positively charged nucleus. But whereas Rutherford's model did not specify what possible locations and energy levels the electrons could occupy, Bohr proposed that the orbits of the electrons had certain fixed radii associated with specific energy levels. Of course, Bohr's model is not entirely accurate by our lights, since the contemporary fully quantum mechanical model of the atom depicts the electrons as spread out in a probabilistic electron cloud, not as located on orbits with fixed radii. Nevertheless, since the fully quantum mechanical model does imply that the electron cloud is densest around Bohr's orbitals – in the sense that electrons are most likely to be found near these orbitals – Bohr's model is reasonably accurate by the lights of contemporary physical theory. Thus it certainly seems right to say that Bohr's model increased our understanding of the atom as compared with Rutherford's previous model.

Now, Bohr's model certainly contained a great deal of explanatory information. Consider what is perhaps the most striking explanatory achievement of Bohr's model, viz. the explanation of the Rydberg formula for spectral lines of several chemical elements, such as hydrogen. A crucial part of this explanation is the assumption, built into Bohr's model, that electrons would have to be located on fixed orbits around the nucleus. Given this, an electron can only gain or lose energy in fixed discrete quantities (when it 'jumps' between orbitals), which explains why the radiation from each atom has certain fixed wavelengths, viz. those described by Rydberg's formula. So Bohr's model provides a reasonably correct explanation for the wavelengths described by Rydberg's formula. Moreover, this explanation could not have been made using Rutherford's model, since the latter does not specify that the electron orbits should have fixed radii. Thus, one might think, explanatory accounts of understanding can account for how the transition from Rutherford's model to Bohr's increased our understanding of the atom.

But this tempting line of thought is mistaken. To see why, note that the spectral patterns described by Rydberg's formula are not a feature of any atom, but a feature of the radiation that's omitted from such atoms. So the phenomenon that is being explained in the above explanation – the explanandum – is not a feature of the atoms as described by Bohr's model at all. It's true that a feature attributed to the atom by Bohr's model, viz. the fact that the electron orbits have fixed radii, occurs *in* this explanation. But this feature serves not as the explanandum but as part of the explanans. So while Bohr improved our understanding of the atom with his explanation of the Rydberg formula, he did not thereby explain the atom or its relevant features. Nor did Bohr provide some other explanation for the atom or it's relevant features, such as for the stipulation that the electron orbits had fixed radii; rather, these features are simply posited by the model. What Bohr's understanding of the atom allowed him to do, explanation-wise, was to increase our capacity to explain other related phenomena, viz. the atoms' spectral patterns. But that point is neither here nor there with regard to U→E, which would require that Bohr had some explanation of the the atom itself.

Now consider this episode from the point of view of my dependency modeling account of understanding, DMA. Here the question is not whether Bohr managed to explain any specific feature of the atom, but whether Bohr's model provided a more accurate and/or comprehensive model of the network of dependence relations that the atom, or its features, stand in towards other things. Bohr's model, in contrast to Rutherford's, reveals that the spectral patterns described by Rydberg's formula are caused by radiation emission in fixed quantities corresponding to the fixed radii of the electrons' orbits. Thus the transition from Rutherford's model to Bohr's provides a more comprehensive model of the dependence relations in which the atom stands towards spectral lines. In this way, DMA validates the judgment that Bohr's model really did increase our understanding of the atom, despite the fact that the model did not provide an explanation of any of the atom's features.

By contrast, explanatory accounts would have to demote Bohr's achievement to something less than an increased understanding of the atom, such as to a mere *description* of the atom, or to an increased understanding *of spectral patterns* (and not also of the atom itself). This has to be considered a significant cost of accepting explanatory accounts, for Bohr's model is frequently described in terms of an understanding of the atom in scientific practice. For example, a widely used textbook describes Bohr's achievements by noting that "even though his model later proved to be incorrect, Bohr remained a central figure in the drive to understand the atom" (Zumdahl and DeCoste, 2010, 331). The most natural interpretation is surely that Bohr is taken to have advanced our understanding of atoms, albeit to a non-maximal degree. Of course, this is not a knock-down argument for DMA as against explanatory accounts by any stretch, since proponents of the latter may reinterpret such statements in line with their accounts, or simply reject them as confused. However, DMA has the advantage over explanatory accounts in that it does not require rejecting or reinterpreting statements that seemingly attribute to Bohr some degree of understanding of the atom.

I conclude that, in some cases, understanding of a phenomenon plausibly does not come from enabling us to explain that phenomenon (or its features), but in modeling the phenomenon in a way that enables us to explain other phenomena. There is a broader lesson in the vicinity here. Explanatory accounts of understanding can seem plausible, perhaps even irresistible, because understanding does tend to bring increased capacities to explain. In that sense, explanation and understanding are indeed closely linked. However, at least in some cases, the targets of understanding and the associated explanation are not the same, for the understood phenomenon (or its features) sometimes functions as the explanans, not as the explanandum. This crucial difference is sometimes forgotten when understanding is simply associated with 'explanatory information' – as if there was some special sort of information that's eligible for explanatory work. In truth, any piece of information can be used in an explanation, and so any information is 'explanatory' in the right context. What matters for the purposes of evaluating explanatory accounts is whether the understanding-increasing piece of information is part of an explanation of the phenomenon that is being understood. I have argued that this is not always so.

7 Explanatory Disconnectedness

A third and final type of case in which understanding comes apart from explanation is when our understanding is increased by virtue of discovering that some object or feature *does not* explain another object or feature. This is illustrated by a version of an example that Lipton (2009) introduced into the debate, viz. Galileo's reductio

[22] It is worth noting that Lipton himself and many of his interlocutors who defend explanatory accounts (e.g., Khalifa, 2012, 2017) are primarily concerned with what this example tells us about *understanding-why*. In particular, Lipton argues that the example shows that understanding-why can come apart from explanation, while Khalifa objects that Lipton has failed to establish this. In contrast to Lipton, my use of this example is only meant to illustrate that *objectual understanding* can come apart from explanation. Indeed, as I note in the next footnote, I disagree with Lipton that the example illustrates that understanding-why can come apart from explanation as well. Thus, for me, Galileo's reductio exemplifies an important difference between objectual understanding and understanding-why, viz. that the former can come apart from explanation in cases of this sort while the latter cannot.

of the idea that heavier objects have greater gravitational acceleration than lighter objects.[22] As Lipton presents it, the reductio is a thought experiment in which we suppose that a lighter object is fastened to a heavier object. If lighter objects accelerate slower, then the lighter object should slow down the heavier object, so the two objects should accelerate slower than the heavier object would by itself. However, the two objects can also be considered together as one larger object, which is thus heavier than either of the objects that it is composed of, so this composite object should accelerate faster than the heavier object. But since the two objects cannot both accelerate faster and slower than the heavier object would by itself, the idea that heavier objects accelerate faster than lighter objects cannot be correct.

Lipton claims that Galileo's reductio provides understanding even though it is not an explanation. I agree with Lipton *that* the reductio provides us with understanding, but I don't accept Lipton's analysis of *why* it provides understanding.[23] According to Lipton, the reductio reveals a specific kind of necessity:

> Galileo's thought experiment is a wonderful demonstration of the necessity of the mass independence of gravitational acceleration and, since seeing necessity is one form of understanding, it is a source of understanding [...] (Lipton, 2009, 48).

However, I fail to see how the *necessity* of the fact that gravitational acceleration is independent of mass is responsible for our understanding in this case. In my view, Galileo's reductio instead shows that understanding can be increased by grasping that two factors are independent, whether by necessity or as a contingent matter. In other words, Galileo's reduction provides understanding not by showing necessity, but by showing a certain kind of *independence*.

To bring that out, consider an otherwise similar case in which we come to realize that two seemingly related factors are in fact independent, where this independence is a contingent matter. When we look up at the sky, the sun and the moon appear to be almost exactly the same size. Accordingly, it is not unreasonable to suspect – as generations of astronomers did suspect – that somehow the apparent sizes of the moon and sun are dependent upon one another, or both dependent on a third object or feature (as when two events have a common cause). As a matter of fact, however, these phenomena are entirely independent. The sun and the moon only appear to

[23]I also disagree with Lipton that Galileo's reductio enables one to "understand why acceleration must be independent of mass" (Lipton, 2009, 47). In my view, Galileo's reductio provides objectual understanding only, not also understanding-why.

be equally sized because while the sun's diameter is approximately 400 times larger than that of the moon, the sun is also approximately 400 times farther away from the earth. This information gives us some understanding of the relative apparent sizes of the sun and the moon, in that we now realize that the apparent sizes of the moon and the sun are independent after all. In this case, it would surely have been possible for the apparent sizes of the sun and the moon to have depended on one another, so the independence in question is contingent (not necessary). Now, since this example is similar in other respects to Galileo's reductio (both involve seemingly dependent variables that are shown to be independent), it strongly suggests that it is not the fact that the reductio 'shows necessity' that is responsible for our increase in understanding. Rather, our understanding increases because we come to realize that seemingly-related factors – viz. mass and gravitational acceleration – are independent.

These types of cases are easily accommodated by DMA, since one way in which a dependency model of a phenomenon can be made more comprehensive is by specifying that two of its features – e.g. an object's gravitational acceleration and its mass – do not depend on each other. By showing that an object's gravitational acceleration is independent of its mass, we are able to model its dependency relations more comprehensively. Of course, this increase in understanding is by itself rather modest according to DMA, since it does not tell us what factors gravitational acceleration does depend on, only that a particular contextually salient factor, viz. mass, is not one of these. Again, this appears to be a correct prediction, since the understanding of gravitational acceleration provided by Galileo's reductio is indeed rather incomplete (a point emphasized by, e.g., Khalifa, 2013b; Strevens, 2013).

Proponents of explanatory accounts have provided a number of rebuttals of Lipton's argument, with Khalifa (2017, ch. 5) providing an especially detailed critique.[24] It is worth noting at the outset that Khalifa's critique, like Lipton's original argument, is concerned with understanding-why rather than objectual understanding. However, let us see if Khalifa's responses hold up against the version of the argument that I have given here (which is exclusively concerned with the latter type of understanding – see footnotes 22 and 23). Khalifa's critique consists of two arguments: First, Khalifa (2017, 138-9) argues that if the understanding brought about by Galileo's reductio is not due to us having or grasping an explanation of some

[24] I will focus on Khalifa's most recent discussion (Khalifa, 2017), as opposed to a previous discussion on which it is partly based (Khalifa, 2013b).

sort, then Lipton lacks any resources for classifying it as genuine understanding, since Lipton himself relies on a conception of understanding as the cognitive upshot of explanation. This argument is clearly not applicable to my use of Galileo's reductio, since I do not share Lipton's explanation-based conception of what understanding consists in. On my view, (objectual) understanding is not the cognitive upshot of explanation, but instead roughly the possession of a model of the understood phenomenon's dependence relations (see §3 above).

Khalifa's second argument (2017, 135-8) is that even if Galileo's reductio does not provide us with genuine understanding, it does give us what he refers to as 'proto-understanding'. Khalifa's definition of 'proto-understanding' is complex (see Khalifa, 2017, 86-7 & 129), but for our purposes it suffices to say that one way in which an agent can come to have 'proto-understanding' is through obtaining information that 'plays an explanatory role' with regard to a given proposition. Importantly, however, information can 'play an explanatory role', in Khalifa's sense, without amounting to, or even being part of, a correct explanation. Rather, it suffices that the information in question figures in comparisons between competing potential explanations of the relevant proposition.[25] Indeed, it is crucial for Khalifa's response to Lipton that one way for a piece of information to 'play an explanatory role' with regard to a proposition p is when it consists of "justified true beliefs that potential explanations [of p] are incorrect" (Khalifa, 2017, 136). Khalifa refers to information that plays this particular kind of role in a comparison of potential explanations as 'critical information'. Khalifa's point is then that Galileo's reductio does provide 'critical information' in this sense (and thus also 'proto-understanding'), since it does rule out some potential explanations of an object's gravitational acceleration, viz. all those explanations in which gravitational acceleration does depend on mass.

There is an important commonality between Khalifa's treatment of Galileo's reductio and that suggested by DMA. In both treatments, the central thought is that Galileo's reductio conveys a kind of negative information, viz. that an object's acceleration is not dependent on its mass. However, the two accounts characterize this information in slightly different ways. According to Khalifa, this negative information contributes 'proto-understanding' because it rules out a specific type of explanation of gravitational acceleration. According to DMA, it contributes some

[25] On Khalifa's view, these comparisons are made with reference to the so-called 'explanatory virtues', e.g. simplicity and explanatory scope. Thus the comparison in question resembles that which occurs in Inference to the Best Explanation (e.g., Harman, 1965; Lipton, 2004).

degree of understanding because it increases the accuracy or comprehensiveness of one's dependency model of gravitational acceleration. While this difference in how the two accounts treat this case might seem negligible, I now want to suggest that DMA provides a more plausible way to characterize Galileo's reductio. This is so for two related reasons.

First, while DMA treats Galileo's reductio as providing some degree of understanding, Khalifa's account implies that the reductio does not provide any degree of understanding. This might seem surprising since we have just seen how Khalifa's account classifies the reductio as providing 'proto-understanding'. However, 'proto-understanding' turns out to be compatible in Khalifa's system with lacking even a minimal degree of understanding: Khalifa says that "S has minimal understanding of why p if and only if, for some q, S believes that q explains why p, and q explains why p is approximately true" (Khalifa, 2017, 14). By this definition, Galileo's reductio doesn't provide minimal understanding, since the reductio doesn't require or enable us to even conceive of an approximately true q at all; much less do we come to believe that such a q explains why p.[26] Now, since we don't have minimal understanding according to Khalifa, we don't have any greater degree of understanding either. Hence Khalifa's account implies that Galileo's reductio does not actually provide any degree of understanding of gravitational acceleration. While I have no independent argument for Lipton's claim that Galileo's reductio provides some understanding of gravitational acceleration, this judgment has seemed highly intuitive and is widely accepted (a point acknowledged by Khalifa, 2017, 138). Thus denying that we gain even a minimal degree of understanding in this case seems to me to at least be a cost of accepting Khalifa's treatment. By contrast, DMA implies that Galileo's reductio does provide some degree of understanding, since representing gravitational acceleration as independent of mass improves the accuracy or comprehensiveness of the relevant dependency model.

A second reason I think my DMA-based treatment of Galileo's reductio is preferable to Khalifa's concerns whether Galileo can in fact be said to have 'proto-understanding' in all relevant versions of the example. In elaborating on how his account would handle Galileo's reductio, Khalifa claims that "Galileo's demonstration provides critical information [...] at the stage of comparison" (Khalifa, 2017, 136-7), i.e. when one is comparing one potential explanation to alternatives. Khalifa (2017, 135-6) argues

[26] Of course, we can imagine cases in which we have conceived of this q, and in which we do believe that q explains why p, but those conditions are not part of Lipton's example.

convincingly that this was true of the actual Galileo, who was comparing Aristotelian and non-Aristotelian potential explanations of gravitational acceleration. However, *in Lipton's own* presentation of Galileo's reductio, there is no reference to any *comparison* between alternative explanations; rather, Galileo's thought experiment simply "demonstrates that gravitational acceleration is independent of mass" (Lipton, 2009, 47). Contrary to what would be required for someone to have 'proto-understanding' in Khalifa's sense, there are no hypotheses or potential explanations that are being compared in Lipton's presentation of Galileo's reductio. This points to an important limitation to Khalifa's treatment, because it means that it is applicable only to Galileo's actual argument in the *Dialogue Concerning Two New Sciences* (Galilei, 1638) and not also to Lipton's (admittedly somewhat anachronistic) version of the same argument. In so far as we agree with Lipton that we gain some form of understanding from *his* presentation of the reductio as well as from Galileo's own, Khalifa's treatment thus fails to deliver the desired verdict. By contrast, DMA straightforwardly counts the reductio, in either Galileo's own version or in Lipton's anachronistic version, as providing us with some degree of genuine understanding.

8 Conclusion

I have discussed three types of cases in which (objectual) understanding of a phenomenon comes apart from explaining that phenomenon. These cases conflict with my criterion for classifying accounts of understanding as explanatory, U→E. Furthermore, in all these cases, the information in question does serve to increase the accuracy and/or comprehensiveness of our models of the understood phenomenon's dependence relations. So these cases serve double duty as positive arguments for the dependency modeling account of understanding presented above. Although this account is not an explanatory account of understanding, it does preserve the kernel of truth in explanatory accounts in so far as a sufficiently accurate and comprehensive dependency model contains the sort of information about a phenomenon that is required to explain it and related phenomena, provided that they can be explained at all. This is so for the simple reason that the dependence relations that these models must correctly represent in order to provide understanding (e.g., causal and grounding relations) are precisely the sort of relations that form the basis for correct explanations.

One benefit of separating understanding from explanation in this way is that it enables us to identify instances of understanding without already having settled on an account of scientific explanation. Despite being the focus of decades of intense intellectual efforts, the nature of explanation is still highly contested and poorly understood. Accordingly, accounts of understanding in terms of explanation run the risk of replacing a hard question, 'What is understanding?', with an even harder question, 'What is explanation?'. Of course, any account of understanding will have to appeal to notions that could themselves be brought into question, but history suggests that 'explanation' is particularly unsuitable in this regard. By contrast, the account of understanding I have proposed here promises to specify when, and to what degree, someone possesses understanding in a way that doesn't rely on any contentious account of explanation. This may, in turn, make it easier to answer a number of important philosophical questions about understanding, such as that regarding the value of understanding (e.g. Elgin, 2006; Kvanvig, 2009; Riggs, 2009; Pritchard, 2010; Grimm, 2012; Ahlstrom-Vij and Grimm, 2013) and its distinctive role in science (e.g. Bangu, 2015; Potochnik, 2015; Dellsén, 2016; Elgin, 2017).

Acknowledgements

For helpful comments on earlier drafts of this paper, I am very grateful to Mike Bertrand, Insa Lawler, Mike Stuart, Elanor Taylor, two anonymous reviewers for this journal, and the audience at the Explanation and Understanding workshop at Aarhus University in May 2016.

References

Ahlstrom-Vij, K. and Grimm, S. (2013). Getting it right. *Philosophical Studies*, 166:329–347.

Audi, P. (2012). Grounding: Toward a Theory of the In-Virtue-Of Relation. *Journal of Philosophy*, 109:685–711.

Bangu, S. (2015). Progress, Understanding, and Unification. In Toader, I. D., Sandu, G., and Parvu, I., editors, *Romanian Studies in Philosophy of Science*, pages 239–253. Springer.

Barnes, E. (1994). Explaining Brute Facts. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pages 61–68.

Bourget, D. (2015). The role of consciousness in grasping and understanding. *Philosophy and Phenomenological Research*, DOI: 10.1111/phpr.12208.

Brun, G. and Baumberger, C. (2017). Dimensions of objectual understanding. In Grimm, S., Baumberger, C., and Ammon, S., editors, *Explaining Understanding: Perspectives from Epistemology and Philosophy of Science*, pages 165–189. Routledge, London.

Carter, J. A. and Gordon, E. C. (2016). Objectual understanding, factivity, and belief. In Grajner, M. and Schmechtig, P., editors, *Epistemic Reasons, Norms and Goals*, pages 423–442. De Gruyter, Berlin.

Craik, K. (1943). *The Nature of Explanation*. Cambridge University Press, Cambridge.

de Regt, H. W. (2017). *Understanding Scientific Understanding*. Oxford University Press, Oxford.

de Regt, H. W. and Gijsbers, V. (2016). How false theories can yield genuine understanding. In *Explaining Understanding: Perspectives from Epistemology and Philosophy of Science*, pages 50–75. Routledge, New York.

Dellsén, F. (2016). Scientific Progress: Knowledge versus Understanding. *Studies in History and Philosophy of Science*, 56:72–83.

Elgin, C. Z. (2006). From Knowledge to Understanding. In Hetherington, S., editor, *Epistemology Futures*, pages 199–215. Oxford University Press, Oxford.

Elgin, C. Z. (2017). *True Enough*. MIT Press, Cambridge, MA.

Fahrbach, L. (2005). Understanding Brute Facts. *Synthese*, 145:449–466.

Fine, K. (2001). The question of realism. *Philosopher's Imprint*, 1.

Frigg, R. and Nguyen, J. (2016). The Fiction View of Models Reloaded. *The Monist*, 99:225–242.

Galilei, G. (1914/1638). *Dialogues Concerning Two New Sciences*. Macmillan, New York. Translated by Henry Crew and Antonio de Salvio.

Gijsbers, V. (2013). Understanding, explanation, and unification. *Studies in History and Philosophy of Science*, 44:516–522.

Greco, J. (2014). Episteme: Knowledge and Understanding. In Timpe, K. and Boyd, C. A., editors, *Virtues and their Vices*, pages 285–302. Oxford University Press, Oxford.

Grimm, S. (2010). The goal of explanation. *Studies in History and Philosophy of Science*, 41:337–344.

Grimm, S. (2012). The Value of Understanding. *Philosophy Compass*, 7:103–117.

Grimm, S. (2014). Understanding as Knowledge of Causes. In Fairweather, A., editor, *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of Science*, pages 347–360. Springer, Dordrecht.

Grimm, S. (2016). How understanding people differs from understanding the nautral world. *Philosophical Issues*, 26:209–225.

Harman, G. (1965). The Inference to the Best Explanation. *The Philosophical Review*, 74:88–95.

Hills, A. (2009). Moral testimony and moral epistemology. *Ethics*, 120:94–127.

Hills, A. (2016). Understanding Why. *Nous*, 50:661–688.

Hitchcock, C. (1999). Contrastive Explanation and the Demons of Determinism. *British Journal for the Philosophy of Science*, 50:585–612.

Johnson-Laird, P. N. (1983). *Mental Models*. MIT Press, Cambridge, MA.

Johnson-Laird, P. N. (2013). Mental models and cognitive change. *Journal of Cognitive Psychology*, 25:131–138.

Kelp, C. (2015). Understanding phenomena. *Synthese*, 192:3799–3816.

Kelp, C. (2017). Towards a Knowledge-Based Account of Understanding. In Grimm, S., Baumberger, C., and Ammon, S., editors, *Explaining Understanding: Perspectives from Epistemology and Philosophy of Science*, pages 251–271. Routledge, New York, NY.

Khalifa, K. (2012). Inaugerating understanding or repackaging explanation? *Philosophy of Science*, 79:15–37.

Khalifa, K. (2013a). Is understanding explanatory or objectual? *Synthese*, 190:1153–1171.

Khalifa, K. (2013b). The Role of Explanation in Understanding. *British Journal for the Philosophy of Science*, 64:161–187.

Khalifa, K. (2017). *Understanding, Explanation, and Scientific Knowledge*. Cambridge University Press, Cambridge.

Kim, J. (1974). Noncausal Connections. *Nous*, 8:41–52.

Kvanvig, J. (2009). The Value of Understanding. In A. Haddock, A. M. and Pritchard, D., editors, *Epistemic Value*, pages 95–111. Oxford University Press, New York.

Lando, T. (2017). Coincidence and Common Cause. *Nous*, 51:132–151.

Lipton, P. (2004). *Inference to the Best Explanation*. Routledge, London, 2nd edition.

Lipton, P. (2009). Understanding without Explanation. In de Regt, H. W., Leonelli, S., and Egner, K., editors, *Scientific Understanding: Philosophical Perspectives*, pages 43–63. University of Pittsburgh Press, Pittsburgh.

Longworth, G. (2009). A plea for understanding. In Sawyer, S., editor, *New Waves in Philosophy of Language*. Palgrave Macmillan, London.

Nerlich, G. (1976). *The Shape of Space*. Cambridge University Press, Cambridge.

Nersessian, N. (2007). Mental Modeling in Conceptual Change. In Vosniadou, S., editor, *Handbook of Conceptual Change*, pages 391–416. Lawrence Erlbaum, Mahwah.

Newman, M. P. (2017). Theoretical understanding in science. *British Journal for the Philosophy of Science*, 68:571–595.

Nounou, A. and Psillos, S. (2012). Book Review of Scientific Understanding: Philosophical Perspectives. *Studies in History and Philosophy of Modern Physics*, 43:72–74.

Owens, D. (1992). *Causes and Coincidences*. Cambridge University Press, Cambridge.

Potochnik, A. (2015). The diverse aims of science. *Studies in History and Philosophy of Science*, 53:71–80.

Potochnik, A. (2017). *Idealization and the Aims of Science*. University of Chicago Press, Chicago, IL.

Pritchard, D. (2009). Knowledge, Understanding, and Epistemic Value. In O'Hear, A., editor, *Epistemology (Royal Institute of Philosophy Lectures)*, pages 19–43. Cambridge University Press, Cambridge.

Pritchard, D. (2010). Knowledge and Understanding. In *The Nature and Value of Knowledge: Three Investigations*, pages 3–90. Oxford University Press, Oxford.

Railton, P. (1978). A Deductive-Nomological Model of Probabilistic Explanation. *Philosophy of Science*, 45:206–226.

Railton, P. (1981). Probability, explanation, and information. *Synthese*, 48:233–256.

Riggs, W. D. (2009). Understanding, Knowledge, and the Meno Requirement. In haddock, A., Millar, A., and Pritchard, D. H., editors, *Epistemic Value*. Oxford University Press, Oxford.

Salmon, W. (1993). The Value of Scientific Understanding. *Philosophica*, 51:9–19.

Schaffer, J. (2009). On what grounds what. In Manley, D., Chalmers, D. J., and Wasserman, R., editors, *Metametaphysics*, pages 347–383. Oxford University Press, Oxford.

Sober, E. (1984). *The Nature of Selection*. MIT Press, Cambridge, MA.

Steiner, M. (1978). Mathematical explanation. *Philosophical Studies*, 34:135–151.

Strevens, M. (2000). Do Large Probabilities Explain Better? *Philosophy of Science*, 67:366–390.

Strevens, M. (2013). No Understanding Without Explanation. *Studies in History and Philosophy of Science*, 44:510–515.

Stuart, M. T. (2016). Taming theory with thought experiments: Understanding and scientific progress. *Studies in History and Philosophy of Science*, 58:24–33.

Thagard, P. (2010). How brains make mental models. In Magnani, L., Carnielli, W., and Pizzi, C., editors, *Model-Based Reasoning in Science and Technology*, pages 447–461. Springer, Berlin Heidelberg.

van Fraassen, B. C. (1980). *The Scientific Image.* Clarendon, Oxford.

Waskan, J. A. (2006). *Models and Cognition.* MIT Press, Cambridge, MA.

Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World.* Oxford University Press, Oxford.

Wilkenfeld, D. A. (2013). Understanding as Representation Manipulability. *Synthese*, 190:997–1016.

Wilkenfeld, D. A. (2015). MUDdy understanding. *Synthese*, DOI: 10.1007/s11229-015-0992-x.

Zumdahl, S. S. and DeCoste, D. J. (2010). *Introducing Chemistry.* Cengage Learning, Boston, MA, 7th edition.