Varieties of Understanding

Michael Strevens

Talk given at the Pacific APA, 2010

Abstract

A number of writers have urged that the goal of inquiry is to acquire understanding rather than knowledge. The nature of understanding in science has long been subject to philosophical investigation, under the rubric of the philosophy of scientific explanation. But what about those fields of inquiry outside of science, such as mathematics, literary criticism, the interpretation of others' behavior, art appreciation, the interpretation of sacred texts, the rightness and wrongness of human actions, and inquiry into the meaning of life (insofar as these are not covert forms of scientific inquiry)? I will argue that insights from recent work on scientific explanation can be applied to unravel partly the nature of understanding in other realms.

Suppose that the principal goal of human inquiry is not knowledge but understanding (Kvanvig 2003). What then?

Innumerable philosopher-hours of thought have gone into the elucidation of knowledge, but not much has been said about understanding. High on the agenda, then, might be the following item: figure out what kind of epistemic state understanding is.

Perhaps the closest thing in the recent decades of philosophical literature to a sustained investigation of understanding is work on scientific explanation, which looks very much like an inquiry into the nature of a particular kind of understanding, namely, the kind of understanding that we extract

from science. Someone ought to ask the philosophers of science, then: What can you tell us about this particular kind of understanding? And how might what you tell us be generalized to understanding in other epistemic endeavors? What kinds of lessons can be drawn from the philosophy of scientific explanation for, say, moral understanding, or mathematical understanding, or metaphysical understanding, or literary understanding?

The aim of this paper is to provide the beginnings of an answer to this question, in three short steps. First, I will propose a simple principle to connect explanation to understanding. Second, I will give an overview of my own recent work on scientific explanation. Third and most important, I will explain how an important part of the framework that I use in my account of scientific explanation can be repurposed to give an account of explanation, and so an account of understanding, in other domains of human inquiry.

1. Explanation and Understanding

To understand something, I propose, is to grasp a correct explanation of that thing (Strevens 2008, 3). This is a view that is so simple and obvious that it might fittingly be called the naive view of the connection between explanation and understanding. Nothing called the "naive view" survives for the duration of a philosophy paper, and this one won't quite make it to the end of the present section, but the basic idea—that understanding consists in grasping correct explanations—will endure.

What is it to "grasp" an explanation? On most accounts of scientific explanation, including my own, an explanation is a set of propositions with a certain structure—in many cases, the structure of a logical argument (Hempel and Oppenheim 1948; Kitcher 1981; Strevens 2008)). To grasp an explanation, I suggest, is to grasp the truth of these propositions and also to grasp that they stand in the relations required by the correct theory of explanation.[1]

1. Alternatively, but somewhat less ecumenically, to grasp an explanation is to grasp that

OK, but what is "grasping"? It is a certain kind of understanding. There is no circularity here, however, because the kind of understanding that is constituted by grasping is not the kind of understanding that is supposed to be characterized by the "naive view". Let me borrow from Kvanvig's (2003) discussion of understanding to clarify what's going on here. Kvanvig notes that there are a number of related, but distinct, senses in which we use the term *understanding*. The two most important, he suggests—and I will follow him here—are the sense in which we understand a subject matter, such as classical electromagnetic theory or a friend's psychology, and the sense in which we understand that something is the case. It is the former, "objectual" understanding that Kvanvig takes to be the proper goal of human inquiry; it is the latter kind of understanding that is equivalent to the "grasping" that figures in the naive view. What the naive view is doing, then, is characterizing objectual understanding in terms of the other kind of understanding. There is no circle.

I have two further remarks on the naive view; the issues I mention will not play any role in the remainder of the paper, however. First, there are other views of scientific understanding in the philosophical literature, most notably De Regt and Dieks's (2005), which puts great emphasis on the notion of *intelligibility*. At least, I think that this account is an alternative to the naive view, though it might also be understood as an elaboration of the naive view that uses the notion of intelligibility to make the notion of "grasping" intelligible.

Second, the naive view concurs with Kvanvig (2003) that knowledge is not required for understanding. Indeed, on the naive view all that is required for understanding is true belief, whereas Kvanvig requires also "subjective justification"—though I suppose that it might be argued that sincere belief

---

the states of affairs asserted to hold by the propositions do indeed hold, and also to grasp the holding of some further states of affairs represented by the structure of the propositions. I am thinking here of my own account of explanation, on which the logical relations of the propositions represent certain causal relations between the corresponding states of affairs.

itself requires subjective justification.

⋆ ⋆ ⋆

Can the naive view of the connection between scientific explanation and scientific understanding be generalized to other forms of understanding? Or rather, can it be generalized to other forms of the kind of understanding that serves as the aim of inquiry, namely, *objectual* understanding? Can it be generalized to the case of understanding a friend, understanding a proof, understanding the moral complications involved in waterboarding, understanding the music of Bartók?

A positive answer to this question, or set of questions, would imply that for each of these feats of understanding, there is a set of explanations, the grasping of which constitutes the understanding in question. Are there such explanations? What set of explanations, when grasped, constitutes an understanding of Bartók's music?

You might think that difficult questions of this sort arise even in the case of scientific understanding. One of Kvanvig's examples of objectual understanding in science is "understanding the theory of relativity". On the naive theory, to understand something is to grasp its explanation, so to understand the theory of relativity would be (more or less) to have the ability to explain it. But what if, as seems quite possible, the theory of relativity is a part of fundamental physics? Then it has no explanation, or at least, no non-trivial scientific explanation.

A little reflection on this case suggests that, although relativity's fundamentality blocks a certain kind of understanding—an understanding that is apparently possible for non-fundamental laws, in virtue of their explicability in terms of fundamental laws—there is another kind of understanding with respect to which it clearly makes sense to say that one person can, while another person cannot, understand relativity.

In fact, I would distinguish two things that might be meant by such a claim. The first is simply that a person understands that the things asserted of

the universe by relativity theory hold. That is a case of "understanding-to-be-the-case", not objectual understanding, so I will put it aside. The second sense of understanding I will unpack in terms of scientific explanation: to understand relativity in this sense is to be able to explain some range of phenomena using the theory of relativity. This allows for a multi-dimensional gradedness in understanding, with the level of your understanding increasing as the range of phenomena that you can explain broadens. I won't say anything more about this, except to make the obvious claim that full understanding of the theory would be (more or less) constituted by the ability to use relativity to explain anything that it is capable of explaining.

The naive theory has become somewhat more complex, then, but it is still pretty straightforward. I have distinguished two kinds of scientific understanding, both connected in simple ways to explanation. In the first sense, to understand X is to grasp a correct explanation of X. In the second sense, to understand X is to grasp a range of correct X-involving explanations of other things.[2] The first sense applies to both phenomena and laws (or law-like stuff), the second only to laws. That is very rough, but it will do.

Enough about science. What about understanding the moral complications surrounding waterboarding? Here is my proposal. Such understanding consists in grasping correct explanations of facts of the following sort:

1. Waterboarding is always morally wrong.

2. Waterboarding is morally permissible in "ticking bomb" scenarios.

3. Although waterboarding is morally wrong even in "ticking bomb" scenarios, it is less wrong in such cases, or there are *pro tanto* reasons in its favor in such cases that do not exist in other cases.

2. You might weaken this formulation as follows: to understand X is to grasp a range of X-involving explanations of other things that *would be correct* if X itself were real/true/correct. This allows talk about understanding false theories—say, understanding Newtonian physics—but, though we do talk this way in syllabuses and so on, it is not clear that such understanding is an end rather than merely a (sometime) means of inquiry.

4. Waterboarding is morally permissible whenever lives are at stake.

Of course these cannot all be moral facts; you will have to take your pick. But you get the general idea. I should add that we can understand a fundamental moral principle—say, the principle of utility or the categorical imperative—in the same way that we understand a fundamental physical law, by grasping correct explanations in which it figures, in this case, correct moral explanations.

The case of understanding music is perhaps even more complex. One way that you might understand Bartók's oeuvre is by understanding the social and psychological considerations that explain the production of the music, but that is of course not the only or even the usual construal of what we mean when we talk about understanding music (or literature, or art, or philosophy). Unlike a scientific theory or a moral principle, a piece of music cannot be used to explain anything itself, so that way of connecting understanding and explanation is ruled out. What remains?

One option is to connect aesthetic understanding with aesthetic explanation, where aesthetic explanation is understood to be the kind of thing that critics and humanities professors engage in (I mean the real humanities, not philosophy). To understand a poem, for example, might be to grasp correct explanations of facts as diverse as:

1. That a typical reader has such and such a reaction to certain transitions in the poem,

2. That these otherwise quite different parts of the poem exemplify the same formal structure,

3. That the dilemma of this character is structurally similar to (or reminds the typical reader of) the dilemma we all face because of our mortality,

4. That this narrative trope fits a broader pattern in Western literature,

5. That this resolution serves to maintain a moral order conducive to the smooth operation of the dominant means of production,

6. That this line of the poem calls into question the fundamental premise on which the rest of the poem is predicated.

The question of what critics ought to be doing, and whether what they are doing can be construed as explanation, is of course tremendously controversial. Clearly, there is much more to say—things have only just begun to get interesting! But I need to move on. I hope I have said enough to establish the plausibility of the view that objectual understanding consists in grasping correct explanations, in which case I can turn, in the remainder of this paper, to the question of explanation itself.

2. Causal Explanation

I am going to present the outlines of my theory of scientific explanation without any of the usual philosophical niceties. I will not argue either for the account itself or for the causal view of the world presupposed by the account. I don't yet even care whether you accept the account. The point is not to convince you, but only to explain how the account works. The convincing—or at least, a few stirrings of philosophical sympathy—will come, I hope, in the next section, when I show how nicely the picture of scientific explanation generalizes to other kinds of explanation.

My theory of explanation is a theory of causal explanation; let me therefore begin with causality. Following more or less in the footsteps of Salmon (1984), Dowe (2000) and others, I suppose that fundamental physics gives us a picture of the world in which all matter is embedded in a dense and tangled causal web. The nature of the web is perhaps best appreciated by thinking in Newtonian terms; the move to a more modern physics will not change the qualitative picture very much (or so many of us hope).

In a Newtonian world there are a number of fundamental force laws, most notably laws of gravitational and of electromagnetic force, telling us when and to what extent one object causes another to accelerate. Force, then, is to be understood as a causal relation: in the kind of world I have in mind, one object is causally related to another if it exerts a force on the other. Call this relation *causal influence*. (I will be rather casual about the relata of the causal influence relation.)[3]

There is a lot of causal influence going around. Every object with mass exerts a gravitational force on, and so causally influences, every other object with mass, at any given time. The same goes for electrically charged objects, of which almost everything is composed. This is why I say that the web of causal influence is dense and tangled.

One further thing: causal influence is *the* fundamental causal stuff: insofar as there are high-level causal relations, they are "built out of", or "grounded in", or "nothing over and above" fundamental-level relations of causal influence. (Choose the locution that suits you; I have my own story about the reduction of high-level to fundamental-level causal relations, but the details do not matter here.)

Equipped with this story about the fundamental causal stuff, you could give a bare-bones theory of causal event explanation along the following lines: to explain an event is to specify a complete set of causal influences on the event. Or alternatively, it is (a) to specify some set of causal influences on the event, or (b) a complete set of causal influences operating at a time.

Call this the *minimal causal account* of explanation. The thought behind the minimal account is, I take it, something like this. Causal influence is a

3. Here is a somewhat more careful way of doing all of this (following Dowe). First, understand the causal influence relation as holding not between objects but between states of objects at times; in virtue of the force of gravity, then, the current state of the Earth is influencing the current state of the Moon. Second, as above that force is a causal influence. Third, stipulate that the state of an object at one time is a causal influence on its state at any later time. Fourth, stipulate that causal influence is transitive. Then you can have earlier states of things influencing later states of other things.

kind of metaphysical dependence relation: the way things are in the world, or in a certain part of the world, depends on the sum total of the causal influences operating on the region in question. A causal explanation specifies these dependencies; equivalently, understanding the way things are in some place at some time is a matter of grasping these dependencies.

Perhaps the most prominent advocate of the minimal account (or something close enough to it) was Salmon (1984). Salmon himself later concluded, however—quite rightly, in my view—that the minimal account was too minimal, suggesting that it needed a certain supplement (Salmon 1997). His argument was as follows.

Imagine that a group of philosophers are playing baseball. They are a boisterous lot, arguing loudly about the distinction between being and becoming as they play. They are not paying much attention to the game; a ball flies off in an unexpected direction and breaks a window. What explains the window's breaking? According to the minimal account, the answer is: all the causal influences on the breaking. But as Salmon points out, although some of these influences are certainly explanatorily relevant—the bat's striking the ball, the ball's striking the window—others seem to contribute nothing to our understanding of the breaking. In particular, the shouts of the philosophers as they play, though they causally influence the window throughout the shattering process, because they cause it (or its pieces) to vibrate in various ways, do not seem to belong in an explanation of the shattering. The ball's striking the window and the players' shouts are both causal influences on the window's breaking, then, but only the former is explanatorily relevant to the breaking. It seems that the minimal account is too liberal in its attribution of explanatory relevance; some further condition on relevance must be found.

Intuitively, what is going on is this. Although the state of the window as it breaks depends in part on the players' shouts, the shouts do not make a difference to *whether or not* the window breaks. When we explain an event, we are not interested in just any causal influence on that event; we are interested

only in difference-making influences. This is why, although the planet Mars exerts a slight gravitational influence on terrestrial events such as window-breakings, we do not consider it to be even slightly explanatorily relevant to (most) such events. What the minimal account needs, then, is to attend to these facts about difference-making. As Salmon puts it (and I will follow his lead), a state of affairs should be considered to be explanatorily relevant to an event only if it satisfies two criteria: (a) it causally influences the event, and (b) in virtue of its causal influence, it makes a difference to whether or not the event occurred.[4]

The importance of difference-making in explanation is widely appreciated, and there are a number of accounts of causal difference-making available. Perhaps the most familiar is the counterfactual account: event *c* makes a difference to event *e* if, had *c* not occurred, *e* would not have occurred. Salmon favors a statistical relevance account: *c* makes a difference to *e* if it changes the probability of *e*.[5] I favor a third alternative, which as you will eventually see is much better adapted to the needs of theories of explanation in other domains—mathematical explanation, moral explanation, and so on. It is convenient to have a name for this third account of difference-making; I call it the *kairetic* account. Let me explain how it works.

Go back to the ball-players and the broken window. How is it possible that the state of the window as it breaks depends on the players' shouts, but the players' shouts make no difference to the breaking? The answer, it seems, is as follows: some aspects of the state of the window as it breaks make it a window-breaking, while others do not. Fracture lines spreading through the

4. Or to presuppose a little less about the metaphysics of events: it makes a difference to whether or not the event fits the description specified in the explanatory request.

5. Although I am using these accounts to filter out irrelevant causal influences, they will be more familiar to many philosophers as the principal ingredients of self-standing theories of causation (Suppes 1970; Lewis 1973). In my view, the "is a cause" relation analyzed by Lewis and Suppes is identical to the relation of causal-explanatory relevance, so giving an account of the "is a cause" relation and giving an account of causal event explanation are one and the same endeavor (Strevens 2008, chap. 6).

glass, for example, belong in the former class, while movements of particular molecules in intact chunks of glass belong to the latter class. When we ask what causal influences make a difference to the window's breaking, we are asking about the things that influence the former class, but not the latter class. That is, we are asking what the aspects of the window's state that constitute its breaking depend on, but not what any other aspects of the window's state depend on. Only the former dependencies count as difference-making.

How, then, to extract just the difference-making causal influences? Here is my suggestion. Imagine a description of all the factors that causally influence the window's breaking, along with the physical laws in virtue of which they exert their influence. No, wait, that's too hard. Pick a time (say, 30 seconds before the breaking), and include in your description only those causal influences on the breaking present at that particular time—so, a time slice of the complete causal picture. I will suppose that the universe in question is deterministic. Then the factors in the time slice together with the laws will entail the precise state of the window as it breaks, down to the exact positions of every fundamental particle. Since this state constitutes a breaking, the factors and laws will also entail that the window breaks.[6] I call such a description a causal model for the window's breaking.

Now imagine taking a part of the model—a description of a particular states of affairs—and making it somewhat less precise. Where the model specified an exact position for a certain fundamental particle, substitute a specification that the particle's position fell into a certain range. (But don't substitute a falsehood for a truth: the range should include the particle's actual position.) The result of this operation is a new, more abstract causal model—a model that fits what actually happened to the window, but that also fits some alternative, non-actual physical processes in which things went slightly differently.

6. Maybe a "bridge principle" defining breaking in fundamental physical terms will have to be added to obtain this last entailment; if so, go ahead and add it.

The new, more abstract model will (almost certainly) not entail the exact actual state of the window as it breaks. But if you did not change too much, it will still entail that the window does break. To put it another way, the different physical processes that fit the model will all lead to the window's breaking. If so, I say that the details you removed from the model did not make a difference to the breaking. The way that the window broke depended on what was removed (or else it would not have been in the model to begin with), but the *fact* that the window broke did not. Or (one last formulation): the aspects of the window's breaking that depended on the removed details were not among the aspects that *made it* a breaking.

That is, very loosely, my test for difference-making: a causal detail makes no difference to the occurrence of an event if it can be removed from a causal model for the event without invalidating the model's entailment of the event.[7]

Let me emphasize that, when I say "very loosely" I mean *very* loosely. There are a number of amendments that must be made to this picture. First, it is incomplete: something needs to be said about difference-making in a fundamentally indeterministic universe, and about the explanation of laws as opposed to events. Second, even as an account of event explanation in a deterministic universe, it is has certain flaws.[8] Both kinds of deficit are taken care of in Strevens (2008)—but, for all its inadequacies, the picture presented above is good enough for my purposes here.

7. To pick up where note 5 left off: this account of difference-making may remind you of a third approach to "is a cause of" claims, Mackie's (1974) account.

8. I feel compelled to mention one at least: a causal influence that is removable from one model may not be removable from another model for the same event. In my full account, a causal influence is a non-difference-maker only if it is removable from all models in which it appears.

3. Varieties of Understanding

An account of causal explanation in science can make sense not only of scientific understanding (when conjoined with the view of the connection between explanation and understanding advocated in section 1), but also of everyday understanding that is based on the apprehension of causal relationships, such as understanding why refrigeration helps to preserve food or understanding why your car won't start.

Under this heading I would also place various kinds of psychological understanding: understanding a friend, the Republican Party, or your dog, I suggest, consists in grasping a range of correct causal explanations for the characteristic behaviors of said entities. Additionally, perhaps certain kinds of aesthetic understanding are partly based in causal explanation, if they involve the appreciation of the causal effects of artworks on their audience.

But there is clearly a limit on the degree to which a theory of causal explanation can elucidate understanding: understanding may consist in grasping correct explanations (that, at least, is my assumption in this paper), but it cannot consist only in grasping correct *causal* explanations. Mathematical understanding will consist in grasping correct mathematical explanations and moral understanding in grasping correct moral explanations, but it seems unlikely that either mathematical or moral explanations are causal. What is needed is an account of explanation in domains where causal relations are auxiliary or absent. Here, I think, my kairetic account of explanatory relevance is very suggestive.

A factor is causally relevant to the explanation of an event according to the kairetic account if

1. The factor causally influences the event, that is, the unfolding of the event causally depends on the factor in some way, and

2. The factor passes the difference-making test, that is, the dependence between the event and the factor makes a difference to whether or not

the event is an event of the type to be explained (e.g., whether or not it is a window-breaking).

The account of causal-explanatory relevance therefore has two parts. First, it identifies a metaphysical dependence relation, causal influence, and requires that the state of affairs to be explained depend on the explainer by way of that sort of relation. Second, it requires that the dependence make a difference to whether or not that state of affairs obtains: the contribution it makes to the state of affairs must be significant enough that it not only makes a difference to the way that the state of affairs is realized (e.g., the way the window breaks), but it also makes a difference to the fact that the state of affairs is realized at all (e.g., to the fact that what happens to the window constitutes a breaking).

I propose that explanation in every domain is structured in the same way; all that changes from area to area is the identity of the relevant metaphysical dependence relation. (On the connection between explanation and dependence across the board, see Rosen (2010).) Consider, for example, explanation in mathematics. If mathematical explanation is possible at all, there must be a dependence relation whereby some mathematical facts depend in some sense on other mathematical facts. It is in virtue of this dependence relation that some mathematical facts are fundamental, and others should be understood as holding in virtue of the fundamental facts. For example, if all facts about the natural numbers—say, the fact that there are infinitely many primes—depend on the properties of natural numbers articulated by the Peano axioms, then the Peano properties are the fundamental properties of natural numbers, and all other facts about natural numbers are to be explained by showing that they depend on the Peano properties. (I allow that the dependence relation may be metaphysically very lightweight; for a mathematical formalist, it may be no more than entailment by the axioms.)

A proof in mathematics may or may not be explanatory. In order to be explanatory, the entailment structure of the proof must mirror the relevant

mathematical dependencies, just as the entailment relations in an explanatory scientific model must mirror the relevant causal dependencies. (Failure to satisfy this criterion accounts for the fact that proof by reductio is typically not considered explanatory: in a reductio, the fact to be explained is deduced from, but does not mathematically depend on, the fact that its denial leads to contradiction.) But that is not a sufficient condition for explanatoriness. It is also important that an explanation contain no irrelevant information, and in particular, no fundamental facts that make no difference to the fact that is to be explained. For example, many of the fundamental properties of the real numbers, as articulated by the axioms for the reals, are irrelevant to explaining why there are infinitely many reals (by which I mean, why the reals are not finite, not why they have the particular cardinality that they do). Such properties ought not to appear in an explanation of the reals' infinitude. They cannot be disbarred on the grounds of non-fundamentality; rather, I suggest, it is the kairetic difference-making criterion that determines their irrelevance.

The application of the criterion exactly parallels its application in the causal case. Start with the complete set of axioms for the reals, and derive from this set the fact to be explained—the reals' infinitude—in such a way that the direction of entailment mirrors the direction of dependence.[9] Then try removing information from the set of axioms. You might do this in one of two ways: by removing an axiom altogether, or by substituting for an axiom something logically weaker (but with the same subject matter). Anything that can be removed without invalidating the proof is a non-difference-maker, and so is irrelevant to the explanation of the fact in question. What remains are the aspects of the axioms—hence, the aspects of the real numbers—that make a difference to their infinitude, or equivalently, on which their infinitude depends in a difference-making way. It is these properties that explain the reals' infinitude.

9. I pass over the fact that the usual proof is by reductio!

Observe that other accounts of difference-making that have a plausible flavor in the context of causal explanation are quite unappetizing when considered as determiners of mathematical difference-making. There is not much sense that can be made of probability-raising as a determinant of mathematical difference-making. A counterfactual test is slightly more promising, but requires an account of counterfactual conditionals on which the consequences of the falsehood of metaphysical necessities can be meaningfully investigated. Providing an "impossible world" semantics is not hopeless, but nor is it a prospect to relish. The kairetic criterion is by contrast relatively easy to carry over to the world of mathematics.

So much for mathematical explanation (though of course, many questions remain; I have barely scraped the surface). What about other forms of understanding? What about, say, moral understanding?

What is it to understand that, say, lying is usually wrong (or that lying is wrong except in such and such circumstances)? I will suppose that we can make a distinction between fundamental moral principles and derived moral principles, and that the truth of the derived moral principles depends on the truth of the fundamental principles. There is, then, a relation of "moral dependence". I won't speculate as it what it is (though it might be something as simple as logical entailment by the fundamental principles).

Suppose that the wrongness of lying is not fundamental but derived. Then to understand why lying is wrong is, I suggest, to grasp the way in which lying's wrongness depends on the fundamental moral principles (cf. Rosen 2010). An aspect of the fundamental moral principles is relevant to such an explanation if lying's moral status depends on that aspect in a way that makes a difference to whether or not lying is wrong. As in the case of causal and mathematical explanation, this criterion has two parts. First, lying's moral status must bear the right kind of dependence relation to the fundamental principle in question. Second, the dependence must make the right sort of difference. It is not enough, for example, that the dependence helps to

determine *how* wrong lying is (if moral depravity is a graded notion, as I very much hope it is); it must help to determine *that* lying is wrong.

The kairetic criterion for difference-making looks to be a satisfying judge of such matters. Take a derivation of lying's wrongness from the fundamental principles that exhibits the dependence of wrongness on the principles. Remove whatever you can from the premises of the derivation without invalidating the entailment of wrongness. What is removed are the elements of the fundamental principles that make no difference to lying's wrongness; what remain are the elements that do make a difference, along with a derivation that illuminates the aspects of the dependence relations in virtue of which they make their difference. To explain, and so to understand, lying's wrongness is just to grasp these facts.

So it goes for every variety of objectual understanding—that is my claim. Understanding a subject matter consists in grasping correct explanations of, using, or otherwise related to that subject matter. Grasping the correct explanation of a state of affairs means grasping the dependence relations that make a difference to whether or not that state of affairs holds. Different subject matters call for different dependence relations, but a single criterion for difference-making is explanatorily universal: whatever the subject matter, thus whatever the dependence relations, the kairetic criterion is able to distinguish between those aspects of the dependence structure that make a difference to some state of affairs' holding and those that do not.

To conclude, work in the philosophy of science can make three contributions to a general theory of understanding. First, the same simple theory of the connection between explanation and understanding that works in the case of science also works elsewhere. Second, the causal approach to scientific explanation leads by example in demonstrating that understanding is a matter of grasping the existence and structure of certain dependence relations. In science these are causal dependence relations; in other domains, they may be something quite different. Third, my own account of scientific

explanation provides a piece of the general theory of understanding, namely, an account of difference-making that tells you, for any kind of dependence relation, which aspects of dependence are difference-makers for a given state of affairs, and so which aspects explain that state of affairs.

References

De Regt, H. W. and D. Dieks. (2005). A contextual approach to scientific understanding. *Synthese* 144:137–170.

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

Hempel, C. G. and P. Oppenheim. (1948). Studies in the logic of explanation. *Philosophy of Science* 15:135–175.

Kitcher, P. (1981). Explanatory unification. *Philosophy of Science* 48:507–531.

Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press, Cambridge.

Lewis, D. (1973). Causation. *Journal of Philosophy* 70:556–567.

Mackie, J. (1974). *The Cement of the Universe*. Oxford University Press, Oxford.

Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale and A. Hoffman (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press, Oxford.

Salmon, W. C. (1984). *Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, NJ.

———. (1997). Causality and explanation: A reply to two critiques. *Philosophy of Science* 64:461–477.

Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.

Suppes, P. (1970). *A Probabilistic Theory of Causality*. North-Holland, Amsterdam.