Grasp and Scientific Understanding: A Recognition Account

Michael Strevens

To appear in *Philosophical Studies*

Abstract

To understand why a phenomenon occurs, it is not enough to possess a correct explanation of the phenomenon: you must grasp the explanation. In this formulation, "grasp" is a placeholder, standing for the psychological or epistemic relation that connects a mind to the explanatory facts in such a way as to produce understanding. This paper proposes and defends an account of the "grasping" relation according to which grasp of a property (to take one example of the sort of entity that turns up in explanations) is a matter of recognitional ability: roughly, a property is grasped to the extent to which the would-be understander is capable of recognizing instances of the property.

Why don't magnets run down like batteries? Why do birds-of-paradise have such flamboyant plumage? Why does printing money cause inflation? What's sought by such questions is understanding of the way the world works—a kind of understanding that is often called explanatory understanding or simply "understanding-why" (Kvanvig 2003).

The vehicle of such understanding seems to be, as the terminology suggests, an explanatory model of some sort. A straightforward analysis of understanding-why and its connection to explanation is supplied by the "simple view" of understanding-why (Strevens 2013):

> To understand why a phenomenon occurs is to grasp a correct explanation of that phenomenon.

The simple view builds understanding-why out of two elements: explanation and grasp. Explanation has been scrutinized deeply by philosophers of science and, I would say, is now quite well understood, even if many of the details remain contentious (Woodward and Ross 2021; Ross and Woodward 2023). Grasp is, by contrast, a metaphor, a mystery. As conceived by the simple view, it is a relation between a would-be understander and an explanation—a relation that somehow acts as a conduit through which explanatory insight flows into the mind. But how to make sense of these figures of speech?

This paper sets out to give a philosophical account of the nature of grasp built around recognition, according to which grasp of an explanation consists in the ability to recognize the components of the explanation along with the relational structure that knits them into an explanatory whole.

"Grasp" is a word with many meanings, related to understanding in various ways. Throughout what follows, I am singularly focused on the meaning implicit in the simple view. The question I pose, then, is this: what relation must a thinker have to a correct scientific explanation, if that explanation is to confer understanding, in the "why" sense, of what is explained?

1. The Need for Grasp

Must we really introduce a new, apparently sui generis notion of "grasp" to capture the connection that brings explanatory information into the mind in a way that confers understanding? Why is it not sufficient, in order for an inquirer to understand some phenomenon, for that inquirer to believe, or accept, or know, or have good evidence for the correct explanation of the phenomenon?

The answer to that question has been laid out in part by writers such as Pritchard (2010) and Strevens (2013), but it will be useful to introduce some case studies of scientific understanding in which belief, knowledge, and other such relations to a correct explanation are present while understanding is lacking. We can then ask: what is missing? The answer to this question will point us to the true nature of grasp.

1.1 *Effective Population*

Let me begin with a case centered around the notion of "effective population" in population ecology and evolutionary biology.

Population models typically contain many idealizations, assuming for example that the organisms in the modeled group mate at random, with any member of the group being equally likely to be chosen. Sometimes this idealization makes no difference to the phenomenon that the model is used to represent, predict, or explain. In that case, the idealization is harmless, and typically makes the model easier to use and to understand. But in many cases, the idealization results in a significant deviation between what the model says and reality. The random mating assumption underestimates the rate of genetic drift and inbreeding, for example, in a group in which individuals typically mate only with their neighbors.

In such cases, the accuracy of the model might be improved by abandoning the idealization. In practice, however, it is common for modelers to retain the idealization but to alter other parameters of the model, in particular, the population size. In a system where organisms tend to mate with their neighbors, for example, the idealized model would be run not with the actual population size but with a smaller population size, the "effective population". It will then give answers to questions about drift, inbreeding, and things that depend on them that are far more accurate. (Think of effective population, then, as an abstract but real property of biological groups, as center of mass is an abstract but real property of groups of physical objects.)

There are a number of different formulas for calculating effective population; the right one to use is determined by the relevant organisms' lifeways and the nature of the phenomenon to be predicted or explained. Determining the correct formula is a matter of applying the general notion of an effective population to the scenario at hand.

It is possible to imagine—and perhaps among professional educators very little imagination is required—a biology student who, impatient to solve the exercises and to pass the exam, pays little attention to the concept of effective population in general but rather learns the formulas and their various conditions of application by rote. Such a student will have the ability, let's say, for any straightforward case, to build an appropriate model and to apply the relevant effective population formula to answer some given question about, say, the rate at which alleles will drift to extinction. From this model, they will derive a certain degree of understanding of the allele extinction rate in the group concerned.

Nevertheless, I suggest, their understanding is limited. They apply the notion of effective population accurately but blindly, without really knowing why they are doing what they are doing—why they are using the particular formula they use, and indeed, why any such formula is required at all, that is, why they use some number in the model that differs from the actual population.[1]

Such a student has at their fingertips a correct explanation of the allele extinction rate, or so I assume. They represent the explanation, they believe the explanation (so we may suppose), and indeed they know the explanation, in virtue of its impeccable epistemic pedigree—passed on as it is by their teacher, a skilled and well-informed practitioner of the modeling arts.

What lacuna accounts, then, for their failing to achieve complete understanding? I suggest it is their failing to adequately grasp a key component of the explanation, namely, the notion or property of a group's effective population. They put that notion to work, but they do not fully apprehend what they are doing. Their explanatory practice is in part opaque to them: effective population sits at arm's length, competently manipulated but imperfectly

1. Opinion might well be divided as to whether the student has any grasp of effective population whatsoever. The recognition approach developed in this paper takes a liberal view; they do have a little, and in particular they have more than the many people who have never heard of effective population.

intellectually assimilated.

1.2 *Third-Law Forces*

Let me next consider another kind of student, wrestling for the first time with a simple, high-school version of Newtonian physics, comprising the three laws of motion, the law of gravity, rigid objects, and perhaps some special forces such as friction. Using these notions or properties, they attempt to build Newtonian models of simple systems such as a satellite or a pendulum.

Perhaps the greatest obstacles to the novice Newtonian are certain forces that fall under the purview of the third law of motion. Some of the third-law forces are unexpected, but present no deep intellectual difficulties: it is surprising to be told that a falling stone exerts an equally great gravitational force on the planet earth as the earth does on the stone, but I think that most of us quickly accustom ourselves to the dynamical symmetry proclaimed by the third law.

A far more formidable difficulty emerges when analyzing even a very simple system such as a block's sliding down a frictionless inclined plane. What's to be explained, let's stipulate, is the direction of the block's movement under the influence of the force of gravity. Of course, we all know how the block will move—but why?

From the second law, the student knows that the direction is that of the net force on the block. From the law of gravity, they know that there is a force exerted downward on the block (and from the third law, an equal and opposite force on the planet, though with negligible effect). But then what? We need a net force that is not downward, but parallel to the plane. Yet to the Newtonian novice, it seems that we have run out of forces.

The answer is that the block exerts a force on the plane equal to the component of the gravitational force that is perpendicular to the plane, and that the plane therefore—by the third law—exerts an equal and opposite force on the block. Sum the two forces operating on the block, namely, the downward gravitational force and the plane's "third-law" force, and you obtain a net force parallel to the plane (figure 1). That is why the block slides down the plane.

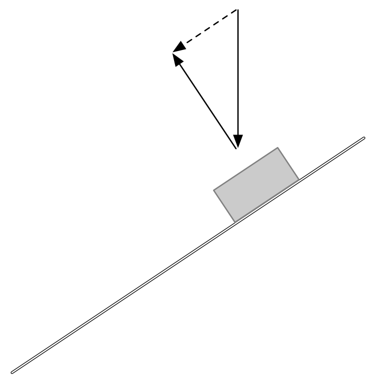

Figure 1: Two forces act on a block on a frictionless inclined plane: gravitational force (downward arrow) and a "passive third-law force" (upward arrow perpendicular to the plane). They sum to a net force (dashed arrow) that moves the block down the plane.

The student's teacher will tell them about this force, and other such third-law forces, but it takes some time to grasp where and when the forces are at work. This is surely because they are exerted by objects that appear to be completely passive: the plane, for example, does not seem to be in any obvious sense pushing up on the block. (I remind you that we are dealing with a simplified physics in which rigidity is a primitive property of the plane.) In any case, students at this stage in their education will tend to omit these passive third-law forces when asked to draw arrows representing the forces acting on objects, or to put them in the wrong places, even adding arrows willy-nilly in an unprincipled effort to produce something that will sum to an empirically viable net force.

Such a student may nevertheless have come, through practice and rote learning, to correctly identify the passive third-law forces for simple cases such as the inclined plane. Yet even in these scenarios, it seems, their understanding of the behavior of the system in question is compromised. This is because, though they draw the passive third-law arrow as instructed and therefore correctly, they don't fully understand what they are doing or why they are doing it—or in other words, because there is something about the passive third-law force that they do not yet fully grasp.[2]

This is not a failure to represent some aspect of the explanatory model, or a failure to believe that it is the correct model. It is not a shortfall of knowledge: the student is fully justified in their belief, and it is safe and secure. (They reliably produce the right answer to any exam question about inclined planes.) What's missing, rather, is—as in the case of effective population—a grasp of a key element of the explanation.

1.3 *Grasping Properties and Grasping Explanations*

Understanding-why is a matter, according to the simple view, of grasping an explanation. In the above examples, by contrast, what the would-be understander lacks is firm grasp of a property, namely, the property of effective population or passive third-law force.

The connection between grasping an explanation and grasping a property is, however, straightforward. To grasp an explanation is to grasp the components of the explanation and the way they are put together. The components of a scientific explanation are principally properties and the facts underlying property-involving generalizations or "laws": causal or other nomic relations between properties. What binds them into an explanation is a certain deductive or causal or other necessitation structure, which can itself, like any relation, be understood as a polyadic property. To grasp an explanation, then, is in large part to grasp an inventory of explanatory properties, some of which are dyadic or polyadic relations. (Besides properties, individuals also play

2. This case should be distinguished from one in which the student has mastered simple Newtonian physics, placing the passive third-law forces even in novel scenarios, but has no further explanatory story about the underpinnings of the passive forces, for example, no explanation of rigidity in solid objects. That, of course, was the situation of Newton himself. I am not claiming a deficit of grasp here.

a prominent role in scientific explanations, namely, those of or by singular events, such as the Big Bang, the K-T extinction, and the 2007–2008 financial crisis.)

This paper will focus wholly on grasp of properties, and for the most part simple, monadic properties. I take this sort of grasp to be central to the grasp of explanations more generally; it is something that an account of grasp must get right if it is going to be at all useful in the context of the simple view of understanding-why. Further, it is the "home base" for the recognition account of grasp to be proposed and defended in this paper, and so the natural arena within which to develop the recognition idea. As I will suggest in the paper's final section, however, the recognition idea promises also to be quite serviceable when it comes to the complex relational properties characteristic of nomic relations and causal structures, and to individuals too. Plausibly, then, it can give us an account of grasp for every element involved in a scientific explanation, including the relational aspect that joins the other elements into a unified explanatory structure. Insofar as the present paper is principally concerned with grasp of relatively simple properties, however, it is only a first step—though an extremely important one—toward a complete recognition theory of explanatory grasp.

There is one other respect in which this paper delivers something less than a full account of explanatory grasp. The shortcomings of the biology student described above might be described either as a failure to firmly grasp the property of effective population, or as a failure to firmly grasp the concept of effective population. I don't think it makes much difference in this particular case, but the conceptual characterization would seem to be more general, if we want to capture grasp of explanations that involve fictional idealized properties or false theories containing empty property terms, such as "phlogiston".[3]

3. Arguably, passive third-law forces are such a case: in a full-dress Newtonian physics, there are no passive forces; rather, all forces are consequences of force laws such as the law of gravity that implicitly conform to the third law of motion. The third law is thus to be understood not as a separate physical principle but as a constraint on force laws.

Ideally, then, an account of grasp would tell us what it is to grasp a property concept, and would construe grasping a property in terms of grasp of the corresponding concept. For simplicity's sake, however, I develop the recognition account in the guise of a theory of grasp of properties themselves. In the concluding section, I point to a way in which the account can be extended to apply to the grasp of property concepts.

2. Approaches to Grasp

The case of effective population may suggest that grasp is a matter of acquiring a definition: the biology student's partial deficit of understanding is due to their not knowing what effective population is, which is to say, not knowing the definition of "effective population"—or to put it more metaphysically, not knowing its nature or essence.

Can we say, then, that grasp of a notion or property consists in knowledge of the notion's definition or of the corresponding property's essence? This might be glossed as the Platonic approach to grasp.

It is surely true that grasp is sometimes attained by learning a definition. But that will not provide us with a fully general account of grasp, for several important reasons.

First, our biology student might turn out, in their avid pursuit of a passing grade, to have memorized a definition of effective population. They can repeat it when asked, but they don't know how to put it to use in their problem sets. They have, in the usual sense, propositional knowledge of the definition. Yet they have failed to incorporate it into their understanding. Something over and above knowledge of the definition seems to be needed—something that sounds suspiciously like grasp. That does not mean that we have made no progress: we have analyzed grasp of a notion in terms of grasp of the definition of that notion. But clearly, if definitions bottom out in primitive notions—concepts that cannot themselves be defined in terms of further concepts—we will need some route to grasp that does not go by way of definition.

Second—and this is perhaps no more than an alternative perspective on the first point—some notions that play an essential role in scientific understanding may very well have no definition. Examples that have been offered in the philosophical literature run from “species” to “temperature” to “water”—and to these might well be added the “passive third-law force” that features in my Newtonian case study.

Third and finally, even in cases where a definition exists, it seems possible to have a firm grasp on the notion in question—which is to say, sufficient grasp to extract all or almost all the available understanding from a correct explanation—without having any knowledge of that definition, or for that matter, while having beliefs about the definition that are quite wrong.

Suppose, for example, that philosophers of biology and theoretical systematists converge on the correct metaphysics for biological taxa such as the insects. Knowledge of the definition or essence of insecthood, as disclosed by this metaphysical theory, may well be, as the Platonic approach declares, sufficient for grasp of the notion of an insect. But surely the many entomologists who are ignorant of and uninterested in the philosophy of biology, along with those who lived before the philosophical breakthrough, have or had a rather good grasp of explanations involving insects—their evolution, their morphology, their ecology—all the same? Knowledge of the definition itself would seem to add little or nothing to their scientific understanding.

Or consider, in a similar but more mundane vein, ordinary thinkers’ grasp of what it is to be a dog. They may conceivably have a smattering of knowledge about the kinds of facts that figure in the definition or essence of doghood (if there is such a thing)—something to do with genes, perhaps (and not even that for the dog-lovers of past centuries). But insofar as they grasp explanations of the physiology and behavior of their favorite animals—as many surely do—it depends not on these vague metaphysical intimations, but on their real-world acquaintance with dogs and their habits. A regular dog-owner who comes across a correct explanation for some aspect of canine physiology or behavior—say, an explanation of dogs' acute sense of smell or their susceptibility to hip dysplasia—is not cut off from understanding such facts for the lack of an opinion about the ultimate foundation of doghood. It's good enough that they know a dog when they see one.

For these and related reasons, contemporary writers on scientific and other forms of explanatory understanding have tended to look not to definitions but to other cognitive achievements or abilities as the foundation of grasp.

One such route emphasizes knowledge, but not necessarily knowledge of definitions. The entomologist and the dog-lover grasp their respective subject matters, according to this sort of theory, in virtue of an extensive and secure body of knowledge about the creatures in question—more practical in the latter case, and more theoretical in the former.

The knowledge route is taken by Greco (2010), Kelp (2015), Riaz (2015), Sliwa (2015), and Khalifa (2017) (though Khalifa's notion of "scientific knowledge" is richer than the usual epistemological conception). It is identified as the standard approach in the philosophy of science by Pritchard (2010) (who rejects it). It might also be attributed to Kant, who, channeling eighteenth-century logic, writes "To understand what gold is I need nothing more than to know the properties of this metal, that it is, e.g., ductile, yellow, heavier than others, etc., that it does not rust" (Kant 1992, 106).

Another route emphasizes certain kinds of cognitive capacities, such as the ability to recognize instances of the property in question (e.g., individual dogs), to answer questions involving the property, to construct explanations involving the property, and so on. (Of course, such abilities in many cases depend on the thinker's propositional knowledge—Sullivan (2018) makes the case that they do so entirely. And even if they are in part more "knowledge how" than "knowledge that", the distinction between the two may be invidious (Stanley and Williamson 2001). For such reasons, it is perhaps unwise to insist on a clear dividing line between the "knowledge" and "ability" approaches to grasp, though I and many others find it to be a useful organizing principle.)

The notion that understanding is not merely manifested in, but constituted by, a suite of abilities is characteristic of the thought of the American pragmatists and the later Wittgenstein. In contemporary philosophy, it is endorsed by (among others) Elgin (1996, 123–4), Newman (2012), Wilkenfeld (2013), Hills (2016), de Regt (2017) (for whom a scientist finds a theory intelligible to the degree that they are able to put it to use), and Robert Brandom (e.g., "grasping or understanding a concept just is being able to place it in a network of inferential relations" (Brandom 2009, 118)). It is a natural home for philosophers who hold that understanding might not always be rooted in propositional attitudes (Zagzebski 2001; Lipton 2009). And it can also be discerned in Grimm's (2014) suggestion that causal understanding consists in a thinker's facility with the map of modal connections inherent in a causal theory, or Le Bihan's (2017) proposal that understanding is a matter of "[navigating] some of the possibility space associated with the phenomena" (p. 112).

The purpose of this paper is to advocate and develop a specific and intentionally narrow version of the ability route to grasp: as I have said, one on which grasp consists entirely in recognitional capacity.

Ideally, the recognition view would be motivated by a careful comparison and contrast with the numerous other views of grasp noted above, which are themselves only a sample of the literature on understanding. Such a survey of the literature would itself constitute a rather long philosophical essay, however, not only because of the wide range of opinion, but because much of that opinion is couched in terms other than the very particular notion of grasp that figures in the simple view of scientific understanding. With some careful exegesis, an implied view of the nature of grasp in my sense can be extracted from many authors' writing on a case by case basis, but it would be a elaborate exercise.

I think it is best, then, to use the space I have at my disposal to lay out my own proposal and to test it against a variety of challenging case studies.

Because the recognition theory attempts to do a lot with a little—analyzing grasp of a thing or property in terms of nothing over and above the ability to recognize the thing or instances of the property—I hope it will turn out to be, at the very least, a fruitful starting place or foil for philosophers of understanding.

* * *

Before I present the recognition theory of grasp, let me make some remarks about the desiderata for such a theory and a few simplifying assumptions.

First, grasp is itself a species of understanding. The simple view says that we understand why a phenomenon occurs just in case we grasp a correct explanation of that phenomenon, but it might equally as well have attributed understanding-why just in case we *understand* a correct explanation. (The partial understanders in my case studies can quite reasonably be said to lack a full understanding of effective population and of passive third-law forces, respectively.)

There are, in fact, many epistemic connections, states, or achievements that might be glossed as forms of "understanding" (which is one reason why the literature on understanding can be difficult to unravel). Understanding-why or explanatory understanding is one such state, and what I am calling grasp is another. Then there are "objectual understanding" (understanding of a subject matter), the "perspective-taking" characteristic of the *verstehen* tradition running through Dilthey and Collingwood, and so on. As Grimm (2021) remarks, "Understanding is a protean concept in philosophy".[4]

According to the simple view, grasp is a *precondition* for explanatory understanding, in the sense that explanatory understanding is acquired by grasping the right sort of thing (a correct explanation). As such, grasp is a simpler, prior form of understanding, a familiarity with the subject matter that constitutes a first step on the way to higher forms of understanding such

4. For these and other philosophically interesting varieties of understanding in everyday life and science, see Kvanvig (2003), Baumberger et al. (2017), and Grimm.

as explanatory understanding.

Second, let me emphasize that, because scientific understanding consists in grasp of an explanatory model, and such models in science frequently traffic in esoteric properties such as dark energy, mental representation, quantum superposition, and money, it must be possible to grasp such properties. In some cases, they may rightly be called unintuitive, even unintelligible: when Richard Feynman remarks "I think I can safely say that nobody understands quantum mechanics", we know what he is getting at. But if quantum mechanics affords us scientific understanding—if a skilled physicist has a far better understanding of phenomena such as quantum tunneling, the outcomes of double-slit experiments, and EPR phenomena than a beginning student—then in the technical sense central to this paper, the physicist must grasp characteristic quantum properties such as superposition. For that reason, we should not insist that intelligibility in the deepest sense is necessary for grasp—another sign that the sort of grasp proper to the simple view is more shallow or preliminary than other forms of understanding.

Third, grasp comes in varying degrees. A biological ignoramus has no grasp whatsoever of the property of effective population. The incurious biology student in my case study has relatively little grasp, but more than the ignoramus. An advanced student, or a biologist in a related area, may have considerably more grasp than that, though still rather less than an expert. A good account of grasp should accommodate and explain this graded quality.

Fourth: Is knowing that a state of affairs obtains a necessary condition for grasping that it obtains? Kvanvig (2003) argues that unjustified beliefs can provide a foundation for understanding; Pritchard (2010) disagrees. A number of other writers have joined the debate (Hannon 2021). I am inclined to agree with Kvanvig that justification, safety, and other forms of epistemic security associated with propositional knowledge are not a part of, and therefore are not necessary for, grasp. That is because I regard grasp of a thing not as a special kind of knowledge of that thing, but as a special kind of representation of that thing, one that is "direct" or that brings with it a familiarity, even an intimacy, with the subject matter. Putting the same point in a slightly different way: in order for true belief to attain the status of knowledge, it must have a certain degree and kind of security, whereas in order for true belief to attain the status of grasp it must have a certain degree and kind of *perspicacity*. Security and perspicacity, and therefore knowledge and grasp, are independent epistemic excellences of belief. Or so I contend—but considerations of length compel me to leave this opinion undefended.

Fifth and finally, grasp figures as the chief epistemic component of several varieties of understanding besides explanatory understanding. The notion of objectual understanding, for example—the understanding a thinker has of a subject matter, such as evolutionary biology or the Thirty Years War—can be glossed as a matter of grasping certain facts and relations central to the topic in question. Likewise, the notion of understanding another person can be glossed as a matter of grasping certain facts about that person: their values, their cognitive and emotional tics, perhaps what it's like to experience the world from their point of view. In these characterizations, as in the simple view, grasp operates as the epistemic connective tissue of understanding. Each variety of understanding is attained by a thinker's relating to a proprietary body of fact, but the relation is in every case the same: grasp.

Will the recognition account make sense of grasp across the board? I can appreciate a certain degree of skepticism on the part of the reader. Let me lay aside these more expansive claims, then, and focus on the task at hand: to develop an account of grasp in the narrow sense of the relation that a thinker must have to the properties specified in a correct explanation, in order to derive explanatory understanding of the thing explained.

3. Grasp and Recognition

3.1 A Recognition Theory

I wish to propose a theory of grasp that equates a thinker's grasp of a property with their ability to recognize its instances. In so doing, the theory conjoins two ideas: first, that grasping a property is a matter of knowing what it is to instantiate the property, and second, that something close to full grasp of a property is possible even for a thinker who lacks knowledge of its essence and other theoretical insight. What can it mean to know what it is to instantiate doghood, if it is not to know the essence of, or to harbor some deep theory of, doghood? It is to know dogs when you see them.[5]

The word "recognition" must be understood rather liberally, if this approach to grasp is to have any generality. Recognition need not go by way of the senses, and even when it does, it need not involve a face to face encounter with the recognized instance. You might recognize 41 as a prime number, you might recognize that trilobites flourished in the Cambrian period, or you might recognize that there is a black hole at the center of our galaxy. Recognition is, then, simply inference to a conclusion of the form, "This is one of those", where "this" might be something far away in space and time or entirely immaterial, and "those" might be a set of things that cannot be directly observed even up close. To put it a different way, a recognitional capacity is an ability to pick out things of a certain kind in a given domain, whether spatiotemporal, mathematical, or for that matter moral. Further, "things" need not be objects: I can recognize that a certain trend is regression to the mean, or that a certain system conforms to the ideal gas law (approximately or exactly); I thereby display my grasp of regression and the gas law.

5. There is—perhaps needless to say—a strong connection between a recognition approach to grasp and various recognitional approaches to the nature of representation, running from causal covariance theories (Fodor 1990) to what Evans (1982) called "Russell's principle", the precept that in order to think about a thing, you must be able to distinguish it from other things.

To contribute to grasp, recognitional capacities must not rely on help from other agents. If I recognize trilobites only because they are helpfully labeled as such, or if I can locate passive third-law forces only by asking for help from my physics teacher, that does not count toward my grasp. In this sense, my recognitional capacities must exist "in the wild", they must be (to borrow a term from computer science) "end to end".

Must they do without aids of any kind? Evidently not, if we are to allow grasp of theoretical properties and entities such as the PAX6 gene, a nation's GDP, or the weak nuclear force. Our capacity to measure and place such things depends on a wide range of detection and information-processing equipment; technological know-how, then, plays a critical role in our grasp of the hidden structure of the world—in a sense soon to be specified with more care.

Grasp is a matter of degree, and on the version of the recognition approach that I favor, its magnitude extends along two dimensions: accuracy and scope. Other things being equal, you have a better grasp of doghood the more reliably you are able to distinguish dogs from non-dogs; likewise, you have better grasp, the greater the variety of dogs for which you are able to exercise this capacity. That is all there is to degree of grasp: fluency in the exercise of recognitional capacities, for example, makes no contribution to grasp except insofar as it affects accuracy or scope. (That is why practicing long division, though it makes you faster, does not enhance your grasp of the division operation itself.)

Merely possible as well as actual instances count toward grasp; a thinker exhibits their grasp as surely when identifying fictional dog breeds or feigned black holes (as perhaps in an astronomy exam or a case concocted to test an instrument) as when they recognize the real thing. This matters when real instances are rare. There are only a handful of hunter-gatherer societies left on Earth; anthropologists' grasp of what it is to be a hunter-gatherer rests on a far more general ability, however, than the recognition of the hunter-gatherer lifestyle in those few remaining groups.

Perhaps unexpectedly, having more than one way to recognize a certain instance, or being able to recognize an instance under a wide range of conditions, does not contribute to scope for the purpose of assessing grasp.[6] Provided that I have some method that works under some circumstances for recognizing a given instance, my ability to recognize that instance contributes to scope as much as it possibly can.

To illustrate this point, consider my grasp of the property of being a forged artwork.[7] Wandering through the galleries, I may not be able to distinguish forgeries from the real thing. Does that diminish my grasp of what it is to be a forgery? No, because I can recognize any forgery in one particular circumstance, namely, by witnessing its creation, that is, by seeing the forger engaged in the act of creating their fake. My knowledge that this makes for forgery is the basis for my grasp. It is irrelevant that I will most likely never find myself in such circumstances. Grasp is a matter of in-principle, not practical, recognitional power.

Let me now apply the recognition account to the two case studies from my introductory pages. The biology student who fails to grasp fully the notion of effective population possesses a string of formulas for determining effective population in a range of different circumstances. That gives them a certain recognitional capacity, and so distinguishes them from the biological ingenue who knows nothing of effective population. But their ability is nowhere near fully general. There are many—indeed, infinitely many—pairs of ecosystems and questions about those ecosystems for which they have no idea how to calculate the effective population. Their ability to "recognize", that is, to determine, effective populations is therefore greatly limited by comparison with an expert who can apply the general idea of an effective population to any scenario. (Being able to recite the definition of effective population without putting it to use helps, of course, not at all.) Thus their grasp falls far short of

6. Having multiple methods for recognizing an instance may, however, enhance accuracy, thereby augmenting grasp along the other dimension.

7. Thanks to Johann Frick for suggesting this case.

what is possible, as does their understanding.

Much the same applies to the Newtonian sophomore. With a little instruction, they have come to recognize the passive third-law forces in a few situations, such as the inclined plane. Their grasp thus outstrips that of the complete beginner. But there are innumerable scenarios in which they are unable to determine the third-law forces; their shortfall in grasp, and hence understanding, is therefore profound.

With effort and instruction comes that moment when the notion of a passive third-law force "clicks". They get it at last, and soon enough they are drawing in the third-law arrows more or less accurately in an unlimited range of scenarios. With this great expansion in recognitional capacity comes a great augmentation of grasp. The recognition theory identifies the two: whatever psychological breakthrough gives the student their new capacity, also thereby gives them their newfound grasp, and with it, physical understanding.

Perhaps the most striking features of the recognition approach are, first, that it allows that a property may be grasped indirectly, through another property with which it is correlated, and second, that it makes grasp entirely a matter of detection and not at all of use. Let me comment in turn on these two provocations.

3.2 *Grasp May Be Indirect*

According to the recognition account, although it is possible to grasp a property by knowing its definition or essence, it is equally possible for grasp to come by way of a property that is merely reliably connected to the grasped property itself.

Most people grasp the property of doghood, to return to an earlier example, not by knowing what doghood really is, but by learning some reliable signs of its presence: those characteristic looks and behavior familiar to every dog-lover. Likewise, our grasp of theoretical properties, as I will emphasize in the following section, hinges not on direct acquaintance, nor in many cases on definition or indeed any inkling about the grasped property's intrinsic nature (think of "spin" in quantum mechanics, for example), but on knowing the characteristic causes or effects of those properties. Any sign of this sort, if sufficiently reliable in both actual and counterfactual circumstances, can give us firm grasp of the property in question in the relevance sense—the sort of "knowing what we're talking about" that is sufficient for explanatory understanding in science and elsewhere to get off the ground.

The case of doghood further shows, I think, that the features by which we recognize a property instance need neither be directly caused by, nor a cause of, that property. Doghood is arguably a historical concept, a matter of an animal's lineage rather than the physical properties that cause its characteristic appearances. Perhaps doghood causes nothing at all; it is, nevertheless, robustly connected to the observable properties by which most of us recognize dogs, and the reliability of this connection is enough to grant ordinary dog-lovers grasp.

What about a case of accidental, or at any rate rather fragile, correlation? The pink coloration of flamingos is due to their diet rather than their genetic makeup. Young flamingos, flamingos in captivity, and even adult flamingos in some unusual natural environments, are white. But I suggest that our knowledge of flamingos' characteristic pink color, by contributing to our recognitional ability, contributes to grasp of "flamingo-hood"—that is, to our having a firm grip on the subject matter when we engage in explanatory discourse about flamingos. The knowledge of pinkness offers a relatively minor boost to grasp, I should add, precisely because some actual and (so I suppose) many counterfactual flamingos lack that color. Yet like other characteristic flamingo appearances, it makes a contribution all the same.

This comports with the role that I have identified for grasp in the philosophy of understanding: grasp of a property or fact, in my sense, is only a first step toward the totality of understanding that is possible regarding that property or fact. We start out by grasping doghood or flamingo-hood; that is, by having some sense of what kinds of things dogs and flamingos are. Only later do we build the explanatory and theoretical structures in virtue of which we become experts in these and other biological matters.

3.3 *Grasp Is One-Sided*

According to the recognition approach, grasp is all about input: grasp of a property depends entirely on a thinker's ability to register its instantiation in the world, and not in the least on any further ability to exploit that information, to take advantage of an instance (or the lack of an instance) so as to advance other aims. By contrast, a more general ability-based approach to grasp, such as that suggested by Hills (2016), implies that a thinker's grasp is enhanced if they have the means to use information about instantiation to infer further facts or to perform further tasks.[8] Grasp, according to these views, is "two-sided": input and output, learning and exploitation, count equally in deciding the firmness of a thinker's grip on the world. The recognition theory of grasp rejects this precept.

Can that be correct? Surely a scientist's grasp of a theoretical property such as electrical charge, evolutionary fitness, or symbolic capital depends as much on what they do with their representations of such things—in particular, on the way that they put models incorporating charge, fitness, and symbolic capital to work so as to predict and explain—as it does on their disposition to attribute those properties to particular particles, organisms, or people?

It does indeed, and the recognition approach is quite capable of explaining why. Science's theoretical properties cannot be recognized directly. Their presence is, rather, determined either by prediction, as when I determine that a glass rod will carry a charge because I have rubbed it with a silk cloth and I know that such actions create a charge, or more often by way of something like inference to the best explanation, as when I recognize a charge by its

8. Likewise, according to the knowledge-based approach, grasp-constituting information is not limited to facts that have some recognitional use.

causal consequences: sparks, or a certain reading on a multimeter. Either way, recognition is indirect, and therefore requires some hypothesis that links the presence of the theoretical property to what can be directly observed. If recognition is to be reliable, those hypothetical links must themselves be reliable—not necessarily true, but a good guide to the presence of the property (as phlogiston theory, for example, gave its adherents reliable methods for detecting "dephlogisticated air").[9] Grasp hinges ultimately on input, but in these cases, input hinges on theoretical beliefs that may play just as large a role in exploitation as in detection. Or to put it another way, the recognition approach to grasp is in principle one-sided, but for the sorts of unobservable properties posited by scientific theories, successful recognition requires two-sided expertise.

The one-sidedness of grasp can be seen clearly in a different kind of case. Consider a child who has solid arithmetical skills and has recently learned the definition of a prime number. I take it that, in virtue of knowing the definition, the child has a firm grasp of the property of "primality" (as mathematicians sometimes call it). When they think about prime numbers, they know exactly what they are thinking about, namely, numbers that have no divisors other than 1 and themselves. (In this instance, grasp comes by way of a definition. The recognition approach happily accommodates such cases; in the right circumstances, definitions or essences are a powerful recognitional tool.) Typically, such a child can do almost nothing with their knowledge apart from checking for primality itself. They may not know the fundamental theory of arithmetic (that every natural number has a unique prime factorization), and even if they do, they are highly unlikely to know how to put it to use to derive any further results. They are unlikely, indeed, to have a clear conception of mathematical proof. In short, their grasp of primality is strong and secure, yet they have a negligible ability—in the absence of further tutoring—to deploy

9. Understanding that term to refer, in such cases, to oxygen, as suggested by Kitcher (1993).

primality in mathematical discovery or explanation.

The example shows that some properties, at least, can be grasped without much in the way of a capacity for exploitation. The correct account of grasp must therefore allow for one-sidedness; grasp, then, is not inherently two-sided.

It might be objected that because of the one-sidedness of their grasp, the child does not understand everything that there is to understand about prime numbers. That is quite true; it shows, however, not that grasp is two-sided, but that there are forms of understanding that go beyond grasp. Grasp of primality is necessary for, but very far from sufficient for, an understanding of the role of the primes in the system of natural numbers. This illustrates once again that grasp of the explanatory materials functions as nothing more than the first step toward full explanatory understanding.

3.4 The Role of Sense Organs and Other Technology

Our powers of recognition—our abilities to pick out things and properties in the world—are greatly enhanced by technology. Indeed, our grasp of some unobservables, as I noted above, turns wholly on our ability to use various detection and measurement technologies. Let me explain how.

Begin with a negative case. In the post-war years, the data from particle accelerators took the form of large numbers of photographs of particle interactions in cloud chambers and bubble chambers. The scientific teams involved could not possibly inspect them all, so they hired teams of "scanners": meticulous inspectors who typically lacked any background in physics but who learned to distinguish the characteristic patterns produced by protons, pions, and other particles involved in the photographed events.

A talented scanner with little interest in physics might reliably identify protons (say) without having much idea what sort of thing protons are. They presumably have the ability to think about protons courtesy of their physicist mentors (or the public-language term "proton"), but augmenting it with their considerable recognitional ability seems not to make much of a contribution to their grasp of protonhood—which remains scant even as they hone their scanning skills. Their recognitional capacity is extraordinary, their understanding nugatory. Has something gone wrong with the recognition account?

It is no use replying that the scanners' means of recognizing protons is a rather narrow technique, or that it is possible only under special circumstances, requiring access to a particle accelerator with the right sort of detector. I have declared that having a single, merely in-principle detection technique of great reliability contributes just as much to grasp as a panoply of practically available methods. (That, I proposed, is why we have a firm grip on the notion of a forged artwork, even if in most circumstances, we cannot recognize fakes.)

The key to solving the problem is rather to recall that recognitional capacities, if they are to contribute to grasp, must be "end to end": in determining grasp, their effectiveness is to be evaluated "in the wild", without help of any kind from other knowers. A scanner who has been trained simply to look for a certain pattern in the photographs does not have their recognitional power end to end, because—unlike the physicists for whom they work—they have no idea how to prepare the photographs in which they see protonhood manifested so clearly. Their scanning talent therefore no more contributes to their grasp of protonhood than does their ability to recognize trilobites, in the cushy scenario where every specimen is clearly marked either "trilobite" or "not a trilobite". The photographic pattern is for them, in effect, nothing more than a rather esoteric glyph supplied, through some unknown process, by their employers.

Now, it is important to distinguish two ways in which the scanners need help. First, someone must prepare the photographs for them. Second, someone must tell them which patterns indicate protonhood. It is the first need that undermines grasp. The second is quite compatible with full grasp: that you acquire your recognitional capacities from your teachers (as even experts normally do) does not compromise your grasp, provided that you can go on to exercise those capacities single-handedly, without any further epistemic assistance.

To see this point more clearly, compare the scanners to the dog-lovers who grasp doghood simply because of their highly reliable ability to know a dog when they see one. In a certain sense, the dog-lovers are like the scanners: they have been trained (in this case by parents, teachers, and so on) to "scan" for doggy looks. They don't have the metaphysical expertise to assess independently whether the things they recognize as dogs by looks alone really do possess the property of doghood. But by contrast with the scanners, once trained, they need no further help. The critical question is: when confronted by a putative instance of the property in the wild, can you put your recognitional ability into action without calling on your trainers or teachers? The dog-lover can: they simply look. The scanner cannot; they lack any sense of how to make a bubble-chamber image.

A positive view emerges from this discussion: grasp of an esoteric unobservable property requires the scientific expertise to reliably utilize a technique or device for detecting that property with no outside assistance. The source of this expertise is irrelevant—it may be, and typically is, passed on by teachers, books, employers—but once acquired, the learner must be able to put it to work in the wild, if it is to count toward grasp.

There may be, as a consequence, some aspects of the subject matter of contemporary science concerning which very few or perhaps no scientists have a firm grasp. Such grasp as exists is a collective achievement of various instrument-builders, theorists, and data analysts.

But it should not be thought that this is the norm. For certain esoteric properties, I do not need any detectors whatsoever; theory alone is sufficient for grasp. Provided that I have some rudimentary knowledge of the workings of general relativity, for example, I can detect the local curvature of spacetime simply by observing the trajectories of falling objects (not tremendously accurately, perhaps, but with a certain qualitative precision). That gives me grasp.

For other unobservable properties, I may need detectors, but the function and operating procedures of those detectors can be ascertained without sophisticated theoretical knowledge. I can determine that an instrument is a reliable microscope, for example, without any real expertise in optics, simply by experimentation: I point it at small things with a known structure, and see that structure magnified. That is enough to give me the capacity to recognize and use a microscope "in the wild" (sitting, perhaps, in a scientific laboratory, but without signage) and consequently the capacity to grasp various classes of microscopic entities—*E. coli*, diatoms, and so on—that I could detect using such an instrument.

I do not, indeed, need to be able to lay my hands on a microscope at all. Recognitional ability is a matter of potential: it is enough that, were I to come across a microscope—impossibly, perhaps, if I were an inhabitant of fourth-century Athens—I could determine its function and put it to use. That is why, when the ancient Greeks philosophized about the invisible structure of things, they grasped to a certain degree what they were (rightly or wrongly) thinking. Putting the point more generally, it is not so much actual technology, meaning physical devices and so on, as it is the capacity to recognize such technology—a purely intellectual matter—that allows our grasp to extend beyond the powers of detection that we possess in virtue of our innate sensory apparatus. Let me remark on two consequences of this precept.

First, once attained, grasp is not easily destroyed. I do not lose any of my grasp of the world when I take off my glasses, turn out the lights, or go to sleep. The recognitional ability is still there, although the circumstances for putting it to work are temporarily out of reach. Even various misfortunes that permanently deprive me of recognitional power do not deplete my grasp. For example, in cases where my recognitional ability is predominantly visual—of colors or faces, perhaps—I would not lose my grasp if I lost my sight. It is sufficient for grasp that I know how to use a functional pair of eyes to detect these things. I might, however, lose grasp through permanent mental impairment, which could rob me of precisely that kind of knowledge.

These observations suggest that the recognition approach to grasp might treat the sense organs and associated cognitive machinery as entirely on a par with artificial detectors. Innate biological equipment has no special epistemic status when it comes to grasp; its prominence in securing understanding owes entirely to our exiting infancy already knowing how to use it for the purposes of recognition.

Second, recognition in physically impossible circumstances can count toward grasp. Consider the following test case. Regular people have little grasp of dark energy. They may have heard the term before, and they may know that it plays an important role in cosmology, but that is all. The recognition approach to grasp offers a straightforward account of their lack of grasp: they have little or no ability to use advanced physics or astronomical observation to deduce the existence, in our universe, of dark energy.

Now contrast the case of dark energy with the case of exoplanets (that is, planets orbiting stars other than our own). The detection—the recognition—of exoplanets is a complicated business. As with the measurement of dark energy, it involves elaborate instrumentation, complex physical models, and a lot of computation. Yet ordinary people surely have a far firmer grasp of exoplanethood than of dark energy. They know what they're reading about when they read about exoplanets in a way that most of them don't when they come across some mention of dark energy.

Why the difference? I suggest the following explanation. Thanks to our experience with our own planet, our own solar system, and various quite accessible models and descriptions of exoplanets, an ordinary educated person is in a position to recognize an exoplanet close up—from, say, a spacecraft orbiting such a body. Now, of course, they will never in fact occupy such an orbit; indeed, it is physically impossible for them to do so: it would take more than a human lifetime to travel to almost any other star system, even in principle. This impossibility seems not, however, to matter in the least. If I were to travel the 50 odd light-years to the star 51 Pegasi, I would see its known planet for what it was, and that and the corresponding counterfactuals for numerous other exoplanets give me my rather firm grip on exoplanethood. Hence, when I read an explanation of a phenomenon involving exoplanets, I grasp what sort of thing it is talking about.

3.5 Rote Recognition

In the light of certain cases of "rote recognition", the recognition approach must be made just a little more sophisticated.

Suppose that the mathematical authorities define a term, 'thrime', referring to the property of being a prime between 15 and 30. We might then find a mathematical ingenue, and tell them that this word refers to a certain mathematical property (we don't say which one) having as its extension the set $\{17, 19, 23, 29\}$. Arguably, the ingenue can now, using public term 'thrime', think such thoughts as: because 17 is in the set, it is a thrime. That capacity, exercised competently, is sufficient for grasp of thrimality, according to the recognition theory. Yet this rote recognizer seems to have meager, if any, grasp of thrimality.

Or consider the philosophically famous chicken sexers. Unlike the rote recognizer of thrimality, they presumably have a solid grasp of the properties that they recognize so well—of being a pullet and of being a cockerel (young female and male chickens respectively). Yet their ability to distinguish pullets from cockerels, based (according to the story) on their detection of small physical differences of which they are consciously barely aware, plays little or no role in securing this grasp—or so it can plausibly be contended. (The case of the scanners, described in section 3.4, has a similar feel, though in that case the recognitional ability fails to contribute to grasp for the independent reason that it is not end to end.)

In these cases, impressive recognitional ability appears to confer rather little grasp. What is important instead, it might seem, is a more intellectual engagement with the properties in question. The rote recognizer of thrimality needs to learn its definition to attain grasp, and the chicken sexers have their grasp by way of their mature concepts of maleness, femaleness, and chickenhood, not their sexing expertise.

Might the recognition account have left the noetic side of grasp too far behind? Perhaps the knowledge-based approach is right to insist (as it does in its usual guises) that some sort of theoretical insight is of critical importance for grasp? Perhaps we ought even to reconsider the Platonic approach; how, after all, to firmly grasp the property of thrimality if not by learning what makes a number a thrime? The recognition approach need not be abandoned; it might rather be augmented, by requiring (say) that knowledge of essence or other theoretically deep aspects of a property must play a leading role in a recognitional capacity if it is to contribute in any substantial way to grasp.

Yet these proposals have great difficulty in dealing with certain cases that I have already discussed. An expert may have a very firm grip on protons or trilobites without knowing the corresponding definitions or essences, or while having deeply misguided beliefs about such matters. That would seem to put paid to any attempt to make knowledge of definition or essence mandatory.

And an ordinary dog-owner surely has a reasonable grasp of doghood—they are capable of understanding explanations of various properties of dogs perfectly well—though they may altogether lack a theory of doghood, whether metaphysical, physiological, or evolutionary. Their grasp seems to be based solidly on practical powers that give them a familiarity with the property in question: they know what dogs look like, how they behave, what they eat, perhaps a little about their training—and that is all. What distinguishes them, then, from the chicken sexer or the rote recognizer of thrimality? That is the question I now set out to answer.

The rote mathematical recognizer acquires the concept of thrimality by learning the term 'thrime'. They acquire their ability to recognize thrimality by mastering a list of four numbers. There is, I suggest, a disconcerting disconnect between, on the one hand, what enables them to think about the property of thrimality—the word—and on the other hand, what gives them their recognitional expertise—the list. The word in itself endows them with very little recognitional ability (outside a social context, which does not count for grasp), while the list greatly underdetermines what property they are thinking about, since there are numerous, indeed infinitely many, mathematical properties with the extension $\{17, 19, 23, 29\}$. I suggest that the rote recognizer's lack of grasp is due to this gap: their grasp of thrimality is minimal because the means by which they think about the property and the means by which they recognize the property are almost entirely disjoint. Significantly, the very same thing can be said of the chicken sexer (and indeed, the particle scanner).

Here, then, is my reformulation of the recognition theory. A thinker grasps a property to the extent that the aggregate of knowledge and perceptual and other abilities that empowers them to think about the property also empowers them to recognize the property (where, as before, recognition must notionally take place in the wild, that is, unaided). Whereas the original recognition view allowed any recognitional ability to count toward grasp, this new theory restricts the grasp-constituting abilities to those that play a substantive role in the thinker's capacity to mentally represent the property in question.

To frame the idea in slogan form—pithy if a little rough—representation is accompanied by grasp to the degree that it exists in virtue of discriminating power that is not only semantic, but also epistemic. Russell's Principle (note 5) is therefore wrong for representation in general but right for the kind of perspicuous representation that confers grasp.

This more subtle version of the recognition theory makes it very clear why definitions and deep knowledge are such powerful sources of grasp. A definition is an especially sure means both of thinking determinately about a property and of recognizing that property's instances—which is to say, it provides semantic and epistemic discriminating power in equal, and equally great, proportion. It is for this reason, and not because definitions have an exclusive claim to supply grasp, that knowledge of definitions or essences, when it exists, often seems so central to grasp.

The same is true (though in varying degrees), of other forms of "deep knowledge" about a property's place in the order of things. These are royal roads to grasp, but not the only roads—or even, in many cases, the roads most often taken.

4. Concluding Remarks

To fully integrate the recognition theory of grasp with the view that "understanding why" is a matter of grasping explanatory models, there is, as foreshadowed in section 1.3, plenty more work to do. Next on the agenda, I think, is to examine in detail what the recognition account has to say about grasp of explanatory relations such as causation. The outline of the story should be clear enough: grasp of a certain causal structure consists in the ability to recognize that structure in the wild, in one way or another (not infallibly, of course). Grasp of causation in general consists in the corresponding more general recognitional ability. What's needed is a careful investigation of the role that such abilities play in our attributions of causal understanding. (Subsequent work might develop a recognition approach to grasp of an individual such as the French Revolution or the Juan de Fuca plate.)

Another major item on the agenda is to extend the recognition account to the understanding of theories that are flawed in some way, such as phlogiston theory, or that for some other reason contain non-referring terms or make use of fictional entities. Again, the overall strategy seems clear enough: the recognition approach can be developed in a counterfactual mode, so that grasp hinges on the ability to recognize certain kinds of things in counterfactual circumstances in which the theory is correct or its postulated entities real.[10] As in the case of grasp of causal structure, I hope that the promise of the recognition theory is clear—but the development must be left to another time.

To conclude: I have argued that grasp is a measure of the extent to which a thinker's semantic connections to the world are also epistemic connections. It quantifies the thinker's ability to keep up with their own thinking, recognizing the things that they represent, discriminating what they are thinking about from what they are not. As such, it is a kind of precision in the exercise of thought, but it is not so much a matter of profound noetic insight as of epistemic dexterity. Of all the metaphors we have for this cardinal mind-world relation, the most apt is, indeed—grasp.

Acknowledgments

I have benefited greatly from the opportunity to present this material, in various incarnations, to audiences at the University of Michigan, New York University, the University of Oxford, Princeton University, and the online Understanding Working Group organized out of UCSD, and at the following events: "Explanation Across the Disciplines" at Middlebury College, "Understanding, in Science and Beyond" at the University of Iceland (online), the *SURe IV* conference at Fordham University, "Understanding Progress and Progress in Understanding" at the University of Iceland, and "Philosophy Meets Science" at NYU Shanghai. Thanks to Johann Frick for the case of the art forger, and to all of those who commented so perspicuously and so productively on my ideas.

10. More carefully: in circumstances in which the terms of the theory refer and the statements framed using those terms are true.

References

Baumberger, C., C. Beisbart, and G. Brun. (2017). What is understanding? An overview of recent debates in epistemology and philosophy of science. In Grimm et al. (2017).

Brandom, R. (2009). *Reason in Philosophy: Animating Ideas*. Harvard University Press, Cambridge, MA.

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

Elgin, C. (1996). *Considered Judgment*. Princeton University Press, Princeton, NJ.

Evans, G. (1982). *The Varieties of Reference*. Oxford University Press, Oxford.

Fodor, J. A. (1990). *A Theory of Content and Other Essays*. MIT Press, Cambridge, MA.

Gentner, D. and A. Stevens (eds.). (1983). *Mental Models*. Lawrence Erlbaum, Hillsdale, NJ.

Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press, Cambridge.

Grimm, S. R. (2014). Understanding as knowledge of causes. In A. Fairweather (ed.), *Virtue Epistemology Naturalized: Bridges between Virtue Epistemology and Philosophy of Science*. Springer-Verlag, Heidelberg.

———. (2021). Understanding. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Summer 2021 edition. URL = <https://plato.stanford.edu/archives/sum2021/entries/understanding/>.

Grimm, S. R., C. Baumberger, and S. Ammon (eds.). (2017). *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. Routledge, New York.

Hannon, M. (2021). Recent work in the epistemology of understanding. *American Philosophical Quarterly* 58:269–290.

Hills, A. (2016). Understanding why. *Noûs* 50:661–688.

Kant, I. (1992). *Lectures on Logic*. Translated by J. M. Young. Cambridge University Press, Cambridge.

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

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

Kitcher, P. (1993). *The Advancement of Science*. Oxford University Press, Oxford.

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

Le Bihan, S. (2017). Enlightening falsehoods: A modal view of scientific understanding. In Grimm et al. (2017), pp. 293–317.

Lipton, P. (2009). Understanding without explanation. In H. W. de Regt, S. Leonelli, and K. Eigner (eds.), *Scientific Understanding: Philosophical Perspectives*, pp. 43–63. University of Pittsburgh Press, Pittsburgh.

Newman, M. (2012). An inferential model of scientific understanding. *International Studies in the Philosophy of Science* 26:1–26.

Pritchard, D. (2010). Knowledge and understanding. In *The Nature and Value of Knowledge: Three Investigations*. Oxford University Press, Oxford.

Riaz, A. (2015). Moral understanding and knowledge. *Philosophical Studies* 172:113–128.

Ross, L. and J. Woodward. (2023). Causal approaches to scientific explanation. In E. N. Zalta and U. Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*. Spring 2023 edition. URL = <https://plato.stanford.edu/archives/spr2023/entries/causal-explanation-science/>.

Sliwa, P. (2015). Understanding and knowing. *Proceedings of the Aristotelian Society* 115, Part 1:57–74.

Stanley, J. and T. Williamson. (2001). Knowing how. *Journal of Philosophy* 98:411–444.

Strevens, M. (2013). No understanding without explanation. *Studies in History and Philosophy of Science* 44:510–515.

Sullivan, E. (2018). Understanding: Not know-how. *Philosophical Studies* 175:221–240.

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

Woodward, J. and L. Ross. (2021). Scientific explanation. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Summer 2021 edition. URL = <https://plato.stanford.edu/archives/sum2021/entries/scientific-explanation/>.

Zagzebski, L. (2001). Recovering understanding. In M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press, Oxford.