

Understanding: The One in the Many

Catherine Z. Elgin

Abstract

Building on my *True Enough* (2017), I argue that understanding is both one and many. It is one in that, regardless of field or subject matter, genuine understanding satisfies the same generic requirements. It is many in that choices are integral to the satisfaction of those requirements, and within limits multiple divergent choices may be acceptable. It may seem obvious that understanding black body radiation is different from understanding voting patterns in a democracy. My point is that there are multiple ways of understanding topics like black body radiation and voting patterns. To make my case, I discuss scientific, ethical, and aesthetic understanding.

1. Introduction

Question: is understanding one or many? Answer: It's complicated. *True Enough* argues for a conception of understanding that seems monistic. To understand is to reflectively endorse a comprehensive, systematically linked network of information that is a) grounded in fact, that is b) duly responsive to evidence or reasons, and that c) enables non-trivial inference, argument, and perhaps action regarding to the phenomena it pertains to (Elgin 2017). The complication arises because this is highly schematic. The criteria can be realized in multiple divergent ways – not only across disciplines, but within the same discipline. That suggests that the position is pluralistic. If the diversity is indicative only of the fact that different disciplines address different topics, it is uninteresting. Botany seeks a systematic understanding of plants. Orthopedics seeks a systematic understanding of bones. Musicology seeks a systematic understanding of music. This is so – even trivial – but it does little if anything to advance our understanding of understanding.

The problem is reminiscent of the issue that confronted Meno at the beginning of Plato's dialogue. Socrates seeks a definition of virtue. Meno

replies with a list (Plato 1976: §71e-72a). There's nothing inaccurate about the list, but it fails to get to the heart of the issue. What makes the several items on the list virtues? What unifies them? Similarly, one wants to know, what unifies the various types of disciplinary, interdisciplinary, and everyday understanding so that at a suitable level of abstraction they qualify as the same sort of epistemic accomplishment. There's no guarantee, of course, that anything does. Maybe 'understanding', like 'game', is a family-resemblance concept. In that case, although each type of understanding has something in common with each of the others, there is no unifying thread that runs through them all (see Wittgenstein 1953). That remains to be seen.

2. Access

My focus is on multiplicity as it bears on epistemic access. Avenues of access may vary with topic, but they also vary in approaches to understanding the same topic. We can begin to appreciate the roles of epistemic agency and choice in framing what is understood when we attend to this multiplicity.

A priori, it might seem that there is exactly one optimal way to understand a given range of phenomena – the way that gets things right. The problem, though, is that reality is complicated. Indeed, it is so complicated that what is plausible a priori is implausible a posteriori. William James described a newborn's experience as a 'blooming buzzing confusion' (1890: I, 488). A vast and motley collection of inputs impact the baby's central nervous system. She has to sort them out. To do so, she needs to ignore most of them and impose order on the rest. Part of her response may be innate; part may be directed by her caregivers; but with so many diverse inputs, it would be remarkable if there were just one way, or even just one optimal way, to make sense of things. Nevertheless, we do make sense of things. We develop and refine expectations that are often borne out. We formulate goals and sometimes achieve them. We eventually get some things right. And this, we think, is no coincidence. Matters work out when they do because we are in touch with the way things are. Indeed, it is not a coincidence. As we proceed, we refine our understandings so that our success rate improves.

This might encourage the idea that understanding consists in mirroring the facts. If so, an understanding of a topic reflects the way the phenomena are. On such a view, knowledge mirrors individual facts; understanding mirrors broader ranges of facts. Richard Rorty (1979) disparages the view that the mind is the mirror of nature. He maintains that this position is a shibboleth of analytic philosophy or a pipe dream of inquiry. I disagree. I think it captures an aspect of the way ordinary people (that is, people who are not analytic philosophers) think about knowledge of the world. When you are trying to find your keys, a

good way to proceed is to try to remember – that is, form a picture in your mind – of where and when you last saw them. When you are trying to understand the layout of the campus, a good way to proceed is to construct a mental map. Often enough, such a strategy works. By mirroring correctly, we achieve our epistemic ends. Nevertheless, there are problems with the mirroring metaphor.

Unlike a mirror, understanding omits. To understand a range of phenomena, we have to selectively disregard irrelevancies. James's 'blooming buzzing confusion' afflicts us all, not just babies. Despite our more sophisticated conceptual structures, we too are bombarded by useless information. But selective disregard requires identifying features that matter – focusing on them and sidelining the rest (see Elgin 1995b). Moreover, this requires adopting a suitable orientation toward those features. One inquiry might be interested in how customers for a given product are alike; another might seek to understand the relevant differences they display. The first, for example, might investigate whether there is a market for electric bicycles in a given city by looking at the full pool of potential customers. The second might seek to understand demographic differences among potential customers. Both are legitimate. But they would encourage focusing on different aspects of the target class.

Moreover, we need to identify appropriate modes of description. That involves choosing a taxonomy. In terms of what kinds will the phenomena be classified? At what grain or level of detail should they be investigated, represented, modeled? Should all the predators in an ecosystem be treated as alike? Or should we differentiate among them according to their species or the species they prey on or their population size or what? Depending on the questions we want to answer, any of these might be appropriate. And depending on the question, one might be far preferable to the others.

An orientation is a perspective on things. Every perspective occludes some things in order to afford epistemic access to others (see van Fraassen 2008). So the choice of an orientation involves a willingness to block some information from view – to sacrifice some information in order to gain access to other information.

Considerable pruning, shaping, and streamlining is thus required just to delineate the topic to be understood. The process is, roughly, Carnapian explication (Carnap 1950). We take a concept that has rough, ill-specified, or not entirely suitable boundaries, and refine it so that it better aligns with the function for which we seek to use it (see Brun 2020). The process is, as Carnap emphasizes, pragmatic. The cognitive end – the problem we seek to solve – justifies the means – the revisions we make in our pretheoretic concept.

More needs to be done. We need to identify what counts as evidence and what counts as sufficient evidence. We need to identify what counts as sound reasoning, and what counts as going too far afield. Opinions can diverge over these questions. Some might restrict evidence to quantitative measurements; others might prefer (or tolerate) qualitative evidence. Some might favor one threshold for statistical power and significance; others might set different limits. All, presumably, insist on satisfying the requirement on variety of evidence. But there may be differences of opinion as to what ranges of evidence answer to that requirement. Then there is the issue of weighing evidence. Even if all agree that a datum is a bit of evidence, different parties may responsibly assign it different weights, some considering it highly significant; others, less so.

Nor are we free from concerns about measurement. What measurements should we take? What measurements *can* we take? The latter concerns available measuring devices, and the extent to which we should trust them. This raises broader questions about how confident we should be about our assumptions, measurements, instruments, and calculations. The questions become more acute if the measurements and calculations are off-loaded onto self-learning AI. Even though the device continually revises and updates its own program, and does so in ways we are not privy to, we need to keep it on task. We don't want its learning curve to bend away from the function we expect it to perform.

To arrive at a genuine understanding of a topic, we need to establish and satisfy norms of acceptability. To avoid conscious or unconscious bias, they should be intersubjectively agreed upon. But by whom? It would be epistemically counterproductive to cast too wide a net. Presumably the community of particle physicists should set the standards of acceptability for particle physics, and the community of art historians should set the standards for art history. We might worry that the standards set by a community or sub-community are or could be skewed. The standards should not be arbitrary or idiosyncratic. They should stand up to scrutiny. This does not solve the problem; it merely identifies it. Different communities of inquiry have different methods at their disposal, different priorities, different resources, and different ends in view. All are subject to rules of coherence and consistency. All agree that like cases should be treated alike. Beyond that, there is little hope for one-size-fits-all standards. Nevertheless, the inquiries in one field should interface with those in adjacent fields, so the standards across related fields should either mesh or should fruitfully challenge one another.

It may seem that standards of acceptability should be indexed to truth. Then methods, criteria of evidence, and the like would be acceptable just in case

they, or their incorporation into an epistemic practice, were truth-conducive. This is overly optimistic. An inquiry takes place because we do not know what the truth is. We may design our investigations in hopes that they will arrive at the truth, but despite the fact that they are well designed, the truth may elude us. Moreover, disciplinary, interdisciplinary, and more informal inquiries make use of approximations, idealizations, distortions, and models that we know to be, sometimes drastically, inaccurate. We justify our use of them by saying (or preferably, showing) that in the respects that matter in the current context, the inaccuracies don't matter.

We might think that the factors I have mentioned are only features of inquiry – the route to understanding; they do not figure in understanding itself. But understanding is holistic. It cannot be divorced from the considerations that sustain it. We cannot claim to understand the genetic structure of the fruit fly without implicitly sanctioning the methods, standards, and evidence that underwrite that understanding. And if, as I have maintained, there are multiple choice points in arriving at an understanding of a range of phenomena, then the roads not taken would likely lead to somewhat different understandings of the same phenomena.

The features I have identified are characteristic of disciplinary understanding – the sort of understanding that explicitly depends on inference and argument – but they apply to everyday understanding as well. Avid sports fans develop a sophisticated, systematic understanding of their favored sport. They mount rigorous, evidentially informed arguments for their views. They dispute the weight of evidence as well. They argue, for example, that a strong defense is (or is not) more important than a strong offense, or argue that last year's standing is (or is not) relevant to assessing this year's prospects. So the weights different parties assign to a strong defense or a previous year's winning season diverge. With a bit of tweaking, the features also apply to discussions in the arts. I will have more to say about this below.

For now, there are several aspects of understanding that are worth noting. First, they are products of choice – choice of grain, orientation, methods, measures, and weights assigned to bits of evidence. Within limits, there are multiple acceptable choices. These choices will yield different verdicts about particular cases. The world does not mandate that it be described in certain terms, or at a certain grain. Nor does it mandate that it be assessed using certain standards of assessment. Moreover, when we set a threshold on evidence or statistical significance or intersubjective agreement, we commit ourselves to sacrificing information that does not reach that threshold. An uncorroborated report may be true; but, we think, it ought not be credited. This is entirely

reasonable, but it shows that the selection of the threshold is doing considerable work in framing the understanding that emerges.

Choices within acceptable limits, and in some cases, choices about what the acceptable limits are, are keyed to the questions we want to answer. One choice may be favored over another because different inquirers want to answer different questions, and thus need to focus on different aspects of the phenomena. Internal medicine, for example, may focus on the effects of a medication on an individual patient; public health may focus on its effects on the population at large. Still, the choices are not arbitrary. Neither is the wisdom of those choices. Both are grounded in the interests, priorities, and resources of epistemic agents and communities.

3. Scientific Understanding

Helen Longino's *The Study of Human Behavior* (2013) underscores the ways these points engender epistemological pluralism. Her discussion of studies of aggression focus on antisocial aggressive behavior by individual actors, not soldiers at war. Many focus on male aggression. It is not clear whether this emphasis is due to men being more aggressive than women, men's aggression being worse than women's, investigators simply being more interested in men's aggression than women's, there being more funding for studies of men's aggression than women's, or something else entirely. (The choice was not Longino's. It was made by those who study aggression.) Whatever the reason, a choice was made to make antisocial male aggression the object of study. It marked out the domain of inquiry. Somebody, or some group, made that choice. The delineation of the topic to be investigated was not dictated by the phenomena.

Longino discusses five distinct approaches to the study of aggression: 1) quantitative behavioral genetics; 2) molecular behavioral genetics; 3) social environment-oriented developmental psychology; 4) neurophysiology and anatomy; 5) ecological approaches, which focus on populations and environments rather than on individual agents. It is plausible that each of these approaches can discover something worth understanding about aggression. It is likely that genetics plays a role, that the configuration of the brain plays a role, that the social milieu in which a child develops and the one where an adult lives are significant, that characteristics of populations and both the physical and social environments matter. So it might seem that we could readily conclude that each approach tells part of the story. Then we could just conjoin the accounts and get a comprehensive understanding of male antisocial aggression.

Understanding: The One in the Many

The problem is that the approaches make different assumptions about the phenomenon being studied and the way it is to be studied. They rely on divergent idealizations. Each idealization brackets features that other approaches consider relevant. They make different assumptions about the causal factors involved. Nature/nurture issues loom large. Genetic approaches assume that inborn features explain aggression. Their mission is to determine which ones. Social and ecological approaches assume that environmental factors, either in upbringing or throughout life, play a dominant role in explaining aggressive behavior. They seek to ascertain what the environmental factors are and how significant they are. The approaches make different simplifying assumptions, yielding disparate models. They rely on different methods and measurements. Their time frames diverge. Some focus on individual acts of aggressive behavior; others look at aggression over a lifespan. Ecological approaches do not even investigate what makes an individual aggressive. They ask about what configurations of populations in what environments lead to different proportions of aggressive men.

The studies deploy different criteria for what counts as aggressive behavior. In some, any violent behavior counts as aggressive; in others aggressive behavior is restricted to violent impulsive behavior. In some, the criterion for being violent behavior is behavior that is more violent than the (supposed) base rate for the relevant population. In others, it is assumed that violent behavior is violent behavior regardless of what the base rate is. Thus a study that focuses on violence per se would consider a man who beats his wife violent. Another study, which focuses on impulsive violence, might exclude him because he is deliberative rather than impulsive in his domestic abuse. A third might exclude both the deliberative and the impulsive wife-beaters on the grounds that wife beating is widespread in their cultural milieu. Manifestly, these three studies are investigating different phenomena and quite intricate arguments will be required to bring their findings into alignment.

Nevertheless, there are a number of ways a monist might accommodate the apparent plurality of approaches to the understanding of male aggression that Longino has uncovered. Perhaps they are all reducible to a single base. If so, one might argue that to the extent that they are true they are not really in tension with one another. Alternatively, perhaps, one is right and the others are wrong. In that case, there really is only one way to understand male aggression. Even if the other theories have their shares of apparent successes, they are in the end unfounded. A third possibility is that they all get something right. In that case, we should identify what they get right, and conjoin whatever is right in them. To arrive at a plausible conjunction will require eliminating some material and devising a translation function to arrive at a common denominator.

Unfortunately, the matter is trickier than any of the suggestions supposes. The approaches, at their best, are and are acknowledged to be, incomplete. Each purports to provide a partial understanding – one that identifies an important factor in the explanation of aggression. Moreover, there is no hope of straightforward reduction, since each focuses on factors that the others sideline. One, for example, looks at neurotransmitters; another looks at social environments. The difference between looking for the causes of aggression inside the head or outside the head seems significant. But each simply marginalizes factors that the other considers crucial.

In the choices of the several approaches about what to focus on, measure, and model, they make different tradeoffs. If a researcher decides, for example to focus on neurotransmitters, he has implicitly or explicitly decided that a genetic explanation of why these transmitters behave as they do is relatively unimportant. He sets genetic factors aside. Similarly, if a researcher decides to focus on social factors, such as whether a child was raised in a single parent home, concerns about which neurons are firing are sidelined. It is reasonable to suppose that all the approaches, when done well, get something right. Their investigations disclose something important about aggression. But the issue that confronts us is not their several piecemeal successes. It is how, if at all, to bring them into accord.

We might think that the failure of the several approaches to align is due to the fact that they are trying to study a phenomenon that is characterized socially rather than scientifically. Aggression is the sort of behavior that is reported in the newspapers. It is not a natural (or scientific) kind. So perhaps it should be no surprise that different scientific approaches, trying to get some sort of handle on this important social kind, explicate it in different ways. This may be true, but it simply brings us back to my earlier point. Molecular genetics explicates aggression in such a way that it lends itself to being investigated by the methods of molecular genetics. Developmental psychology explicates aggression in such a way that it lends itself to being investigated by the methods of developmental psychology. The several sciences frame the objects of their study so that their methods, or perhaps refinements of their current methods, are suitable for studying those objects. Their framing commitments – interests, measures, orientations, and so forth – figure in the demarcation of the topic they seek to understand. But such framing commitments do not just draw a boundary between what is really there, and what is somehow infused with human interests. Human interests delineate what, among the things and arrangements that are really there, deserve systematic attention.

In any case, the issue is not just a difficulty that emerges from attempts to get a scientific understanding of a social phenomenon. It won't be resolved by

insisting that we focus on natural kinds. Teller (2018:403) discusses *water* – a philosophical paradigm of a natural kind. How should scientists model it? His answer: it depends. If they are interested in the flow of water and in wave propagation, scientists should represent water as a continuous, incompressible fluid. If they are interested in diffusion, they should model it as a collection of discrete particles in thermal motion. There are additional complications. If they are interested in water waves of a meter or more, they can ignore surface tension forces. If they are interested in ripples of a millimeter or so, surface tension is an issue that their representations should reflect. All of these interest-relative variations are consonant with water’s being a natural kind with determinate properties. If we want to understand something about water – maybe how its waves propagate, or how it diffuses, or how it contributes to sustaining life – we need to select certain among its properties, and devise an understanding that focuses on them. That involves devising models that set aside real but, for current purposes, irrelevant properties of water.

4. Ethical Understanding

I have been discussing scientific cases because science tends to make criteria and the reasons for them public. But the same sorts of considerations apply to understanding in other realms. Let’s turn to ethics.¹ Much moral reasoning is widespread. We do it everyday, and probably are inconsistent in the reasoning requirements we set for ourselves and for one another. We seem to vacillate. Sometimes we allow appeals to intuition, sometimes we do not. Sometimes we allow arguments from analogy, sometimes we resist them. Maybe there is an underlying principle that vindicates such apparent vacillation, but unless we can identify and justify it, we are, we feel, on shaky ground. We move from everyday moral decision-making to something more systematic when we suspect that our everyday practice, as it stands, is inadequate. It is vulnerable to well founded charges of inconsistency and incompleteness. Not infrequently, it leaves us at a loss about how to reason in ways we think we ought. It lacks the resources to solve our current moral problems. If we understood the topography of the moral realm, we would have greater confidence in our judgments. Issues of taxonomy arise in the demarcation of moral kinds such as virtue, vice, exception, and excuse. These are subject to something like Carnapian explication. Is lying by omission really lying? Is intentionally misleading a type of lying? Is it possible to promise something that is known to be impossible? Is it possible to have obligations to inanimate objects? Do we, for example, have obligations to the environment or to the

¹ Here I use ‘moral’ and ‘ethical’ interchangeably.

Grand Canyon? To answer such questions involves fine-tuning everyday concepts of lying, promising, and obligation; and lines could be plausibly drawn in different places, leading to divergent answers to these questions.

There are issues of relevance as well. A good deal of our moral understanding involves settling questions of the moral relevance or irrelevance of various factors. In deciding whether x 's behavior toward y is morally permissible, does it matter that x is y 's mother, or that x is y 's supervisor, or that x is the same race as y ? If the answer is 'It depends', then the question is what it depends on, and how we determine that.

We set criteria for what counts as a reason, what counts as a good reason, and what counts as a sufficient reason in moral assessment. These determine what is worthy of moral consideration and why. Opinions differ, for example, over whether groups have rights that are independent of, and can diverge from the rights of the individual group members. They differ over whether and how to assign proxies to protect or exercise the rights of those who cannot do so for themselves. Should we, for example, recognize proxies for children, for future generations, for tribes that have no conception of rights? If so, who should those proxies be? What standards should they be subject to?

We issue inference licenses. These determine what information we can appeal to and what inferences we can make on the basis of that information. Should we allow appeals to intuition or to tradition? Should we recognize the existence of moral expertise? If so, what qualifies someone as an expert? Then we need to determine what sorts of inferences are permissible. If *ought implies can*, we will face the question of the scope of the *can*. If weighting of alternatives is permissible or mandatory we need to identify criteria for assigning weights.

Idealizations figure in moral understanding. The core of a moral worldview may be that a moral agent should be like the Buddha, or be like Christ, or be like an Aristotelian phronimos. There is no expectation that a moral agent can actually achieve this state. But such an ethical idealization functions both as something to aspire to and in much the way that an idealization in science does. It exemplifies significant features and marginalizes features that (so the account supposes) do not matter. Even without such a core figure, we are often advised to act as some familiar and accessible role model would act. Ask yourself, 'What would Jon do?' and you may have a good guide to how you should behave in this sort of situation. On such a picture exemplars play a major role in orienting thought and guiding behavior. To be sure, other approaches are rule-based. You should maximize expected utility, or act in such a manner that you could will your maxim to be a universal law, or follow the Golden Rule, or whatever. But even in cases where the rules are thought to

be fundamental, they are apt to be illustrated by idealized examples that show how to apply them.

Our understanding of the moral realm is framed by the choices we make about these issues and the orientations we adopt to them. Minimally, some border-line cases will be treated differently, depending on what (reasonable) choices are made. Something that counts as, and is condemned as, a lie under one explication will not count as, and therefore ought not be condemned as a lie, under another, equally reasonable explication. But more importantly, different choices infuse our understanding of ourselves and one another as moral agents. They affect what we think we and others are doing. If, for example, we believe that giving to charity is morally mandatory, we will not consider ourselves particularly worthy of praise for making a charitable donation. Such a gift is only to be expected. If we consider it optional, we might think the donation makes us praiseworthy. A person's self-understanding of herself as a moral agent is thus keyed to her understanding of what morality requires.

As in the scientific cases, choices are correlated with our interests and goals. In some cases, of course, moral considerations are in tension with non-moral ones. Our aspirations, desires, and/or the councils of prudence may pull us away from where we think our duties lie. Nevertheless, we are told, we should do our duty. But that doesn't automatically get us far enough. Sometimes we confront tensions among our duties. When a person cannot do everything she thinks she is morally obliged to do, how should she determine which obligation takes precedence? What commitments should underwrite her decision?

It might seem that the moral facts, whatever they are, settle these questions. But the difficulty is that we need to figure out how to access those facts and how to assess our views of what those facts are. As is well known, there are a variety of approaches. As we saw with Longino's discussion of aggression, each seems to take a plausible stance, and each seems to have strengths and limitations. One major issue is that ethical theories are abstract. We need to be able to figure out how to put them into operation. A consequentialist maintains that the agent should do whatever alternative will produce the greatest good. There are familiar objections to this proposal, not the least of which is that it is imprecise. But even if we have settled on, for example, act consequentialism rather than rule consequentialism (or conversely), and even if we have restricted our range to readily foreseeable consequences, there remains the question of how to identify and weigh the relevant alternatives. It is not hard if one is to save a life and the other is to eat a candy bar; but if the alternatives are either to devote the medical budget to saving a few lives or devote it to

vaccinating thousands of children against diseases that are unlikely to kill them, things get tricky. Nor is this just a problem for consequentialists. Non-consequentialists recognize that it makes sense to recognize that foreseeable consequences should play a major role in policy-making decisions, and often in the choice between individual acts. They also recognize that the distribution of benefits, as well as the aggregate good, can be morally important. So both consequentialists and non-consequentialists face the problem.

Virtue theorists hold that the moral acceptability of an action depends on the virtue or vice embodied in performing that action. They face the question of how to assess virtuous actions that lead to undesirable results, or vicious actions that lead to desirable ones. Arguably non-virtue theorists do not have that problem, particularly if they maintain that the assessment of the act is more basic than the assessment of the actor. But even non-virtue theorists acknowledge that the cultivation of virtues should dominate our thinking about character development.

Deontologists maintain that the agent's motive for action (the maxim under which he acts) is the proper unit of moral assessment. This, of course, gets complicated if the agent's tunnel vision prevents him from a reasonable assessment of his action's fallout or from recognizing the relevant alternatives to the action he opts to perform. A reasonable assessment of the likely outcomes and the available alternatives seems mandatory. And arguably, even non-deontologists recognize that it may make sense to consider what the perpetrator thought he was doing (roughly, what his maxim was) in order to understand his crime, and to consider what the hero thought he was doing when he embarked on a risky rescue mission in order to assess his heroism.

In ethics, as elsewhere, our interests and resources frame the issues we address, the means we take, the goals we pursue. Since *ought implies can*, we are restricted by the means at our disposal. Ethics is not just concerned with understanding human behavior somehow. It is concerned with understanding human actions and opportunities in a way that makes certain concerns salient, and makes certain problems addressable and potentially solvable. It sets aside ways of looking at human behavior (e.g., as matter in motion) that do not put thinkers in a position to address its concerns.

This is consonant with moral realism, moral constructivism, moral subjectivism. The questions I focus on is how do we frame our inquiries to access and understand the moral realm. Issues of taxonomy, orientation, standards of relevance and of inference arise regardless of what the nature of moral facts is. And the understanding we gain from moral deliberation seems to have much the same form regardless of the underlying ontology.

5. Aesthetic Understanding

It might seem that the arts provide a greater challenge to systematic understanding than science, ethics, or even everyday opinion does. It is true that the arts do not display the sort of regimentation I discussed above. But this is not a critical difference. Like Goodman (1968), I think that works of art are symbols that are dense and replete. Insofar as they are dense, every difference in certain respects makes a difference. In that case, there is always another way of looking at the work that has the potential to enhance or challenge the prior understandings the viewer had of and through the work. Insofar as they are replete, works of art function along multiple dimensions. It follows from density and repleteness that there is more to the work than meets any particular eye.

How does this affect understanding? When dealing with symbols that are dense and/or replete, our understanding should be tentative and open to revision. A replete symbol conveys multiple messages along multiple axes. A work with a lilting melody may also have darker undercurrents in the bass line. A work might juxtapose tonality and atonality, or take seemingly forever to resolve a chord. Such configurations can disconcert. It can call into question and perhaps upset standard ways of thinking. Encounters with the arts, because they challenge complacent views, encourage fallibilism. They undermine stereotypes and the propensity to rely on stereotypes. They stretch our minds, by promoting and rewarding new ways of think about matters that we might previously have considered settled.

In discussing ethics, I spoke about the role of idealizations. Many of these are drawn from the arts. Works of art exemplify characteristics – often complex or subtle characteristics for which there is (as yet) no literal label. We identify the cluster of features as distinctive of, for example, King Lear or Emma Bovary or Don Quixote. This enables us to see in ourselves and our fellows features or patterns that we might otherwise have overlooked. The notion of an idealization here is not that of a moral ideal; rather it is, like the ideal gas, a clear case, unencumbered by confounding features. This is not to say that such ideals are uncomplicated. It is to say that the complications are not confounds. They matter. Since the ancient Greeks, philosophers have asked how (or whether) weakness of will is possible. We gain insight into the issue through a character like Pierre Bezukhov, who really intends to stop procrastinating but never gets around to doing so.

I said earlier that an understanding is a systematic network of commitments that is grounded in fact, is answerable to evidence, and enables non-trivial inference and perhaps action vis à vis the phenomena it concerns. This raises

two questions. One concerns understanding *of* the work itself; the other concerns understanding of things *through* the work.²

The facts an understanding of the work are grounded in are facts about the work. These may include facts about the genre, author, period, or style, as well as facts about the work's own content and stance. So the resources we can draw on to generate an understanding of the work may extend beyond the boundaries of the work itself. Still, the understanding is answerable to evidence in the work. It should therefore be possible to identify features of the work that underwrite the understanding. If a reader understands Gregor Mensa as a Christ figure, she should be able to point to features of *The Metamorphosis* that backs her claim. If a listener understands a Taylor Swift song as a sonata, he should be able to show how it exemplifies the sonata form. The understanding of the work should enable inferences that connect and make sense of salient elements of the work, and perhaps enable the interpreter to explain why other features, which might have been expected, were omitted.

Understanding through the work is a matter of projection. One can extrapolate from the understanding of the work to afford an understanding of features beyond the work. If Pierre Bezukhov's character shows how weakness of will is possible, we are in a position to see how the possibility is realized in the world outside the fiction – that is, how some sorts of procrastination that we engage in or see around us are manifestations of weakness of will, and perhaps how others are not. The extrapolation affords a basis for inference about what happens in the world – what we should expect, given what we have seen in and through the work.

It is worth noting that the same bifurcation of understanding occurs in other areas. There is understanding *of* the molecular genetics account of aggression and understanding of aggression *through* the molecular genetics account. There is understanding *of* consequentialism and understanding of domestic policy *through* consequentialism. Understanding of a phenomenon then implicates a (perhaps tacit) understanding of how we understand the phenomenon.

I will not go into detail about mundane understandings, such as the understanding displayed by sports fans or the understanding displayed by care givers of young children. These too involve explications that fine tune everyday concepts, set criteria of relevance and grain, determine standards of methods of assessment and so forth. They too are tested to see if they make sense of the phenomena – if, for example, their explanations mesh and their predictions are often enough borne out. What I have described as disciplinary

² Malfatti has drawn this distinction with respect to theories. See Malfatti 2019.

understanding is more methodologically self-conscious than other forms of understanding, but there is a continuum from the mundane to the scholarly.

6. Conclusion

I have construed an understanding as a network of substantive, methodological, and normative commitments. These commitments are not dictated by the phenomena, but reflect choices about how to conceptualize, investigate, infer, and extrapolate. An understanding frames the phenomena, reflecting our interest in them, the epistemic resources we have or can devise, the assumptions we can build on, and the tradeoffs we are willing to make. It is always incomplete. There is more to the phenomena than we have yet accommodated. It is apt to be inadequate as well. So a desideratum for a good understanding is that it admits of, and has the resources to promote, its own expansion and revision. Because choices frame and configure the understandings we develop, epistemic agents are active rather than passive (see Elgin 2025a). What we find depends on what we look for and how we look for it. It depends on what resources we can use and recognize that we can use in our investigations. The importance of epistemic agency does not threaten objectivity. What we find is a function of where and how we look, but it is also a function of what is there to be found. Understanding is grounded in fact, and our views about the facts may be wrong. It is answerable to evidence, and the evidence may tell against our expectations. It enables inferences which may fail to pan out and actions which may fail to achieve their aims. Success is not guaranteed. What we see depends on where we look and how we look. It also depends on what is there to be seen. Although agency is deeply implicated in understanding, the world pushes back.

References

- Brun, Georg (2020). 'Conceptual Engineering: From Explication to Reflective Equilibrium' *Synthese* 197:925–954.
- Carnap, Rudolf (1950). 'On Explication' *Logical Foundations of Probability*. London: Routledge and Kegan Paul, pp. 1–18.
- Elgin, Catherine (2025a). *Epistemic Ecology*. Cambridge MA: MIT Press.
- Elgin, Catherine (2025b). 'Selective Disregard' *Abstraction in Science and Art*, edited by Chiara Ambrosio and Julia Sánchez-Dorado. London: Routledge, pp. 17–29.
- Elgin, Catherine (2017). *True Enough*. Cambridge MA: MIT Press.
- Goodman, Nelson (1968). *Languages of Art*. Indianapolis: Hackett Publishing Co.
- James, William (1890). *Principles of Psychology*. New York: Henry Holt.

Catherine Z. Elgin

- Longino, Helen (2013). *Studying Human Behavior*. Chicago: University of Chicago Press.
- Malfatti, Federica (2019). 'Can Testimony Generate Understanding?' *Social Epistemology* 33: 477–490.
- Plato (1976). *Meno* translated by G. M. A. Grube. Indianapolis: Hackett Publishing Co.
- Rorty, Richard (1979). *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.
- Teller, Paul (2018). 'Measurement Accuracy Realism' *The Experimental Side of Modeling* ed. Isabel Peschard and Bas van Fraassen. Minneapolis: University of Minnesota Press, pp. 273–298.
- Van Fraassen, Bas (2008). *Scientific Representation*. Oxford: Clarendon.
- Wittgenstein, Ludwig (1953). *Philosophical Investigations*. Oxford: Basil Blackwell.

Catherine Z. Elgin
catherine_elgin@harvard.edu