CHAPTER 2

What Would the Community Think?

2.1 The Social Enterprise of Knowledge

In the ideal case, the output of scientific research is knowledge. Lots of research fails to rise to this level. Undoubtedly, there will be thousands of scientific research articles published this year that, for one reason or another, contribute nothing to human knowledge. There will be hundreds of thousands of research hours spent on investigations that yield no insight into nature, or into how to improve our investigation of it. But often enough, and often in astonishing ways, scientific research results in something that we can use to manage our relationship with nature; something on which future generations of researchers can build to refine and extend our general picture of the natural world; something that provides us with understanding. We call this product *scientific knowledge*.

Modern scientific knowledge is traditionally associated with distinctive methods of investigation. Scientific investigation characteristically involves the observation of natural phenomena. Frequently, this includes observing the results of careful and systematic experimentation of one kind or another. These experiments often incorporate the use of instruments specifically designed for producing, observing, and measuring effects. The production, observation, and measurement of certain effects are commonly carried out in relation to a hypothesis. At every step, it is customary for mathematics to play a central and indispensable role as scientists endeavor to clarify their ideas and evaluate them with precision. This specific medley of methods employed in the investigation of nature has, since the early seventeenth century, been associated with science's marked capacity for generating knowledge. We refer to the period marking the beginning of widespread uptake of these methods as "the Scientific Revolution." More than anything, the Scientific Revolution was a revolution in epistemology, a fundamental transformation of our basic conception of what it took to acquire knowledge of the world. Since that time, the use of these methods

https://doi.org/10.1017/9781009067508.003 Published online by Cambridge University Press

has been central to our understanding of what allows us to resolve, more or less definitively, the fascinating and multilayered set of puzzles presented by the natural world. An idea rises to the level of scientific knowledge when it solves one or more of those puzzles.

How do puzzles get solved? A natural response to this question is to invoke the use of the methods described above: when we make observations, formulate hypotheses, design experiments to test them, and then analyze the results (ideally, guided by mathematics), we eventually arrive at the solution to our puzzle. The reason that this answer has such a satisfying ring to it is because it is an abstraction of the Scientific Revolution's legacy – and because it is, in part, correct. Our ability to solve natural puzzles of all sorts – including really, really hard puzzles – has increased dramatically since the beginning of the seventeenth century; the methods of scientific investigation deserve much of the credit for that. But not all of it. While the production of scientific knowledge might characteristically *require* the use of certain varieties of investigative approach, it is not *reducible* to the use of those varieties.

Part of the reason we know this is because the history of each of these methods predates the advent of modern science. This historical fact has led some scholars to make the hasty inference that there was no such thing as the Scientific Revolution. If all there was to the emergence of modern science was the emergence of scientific methods, then that inference might hold up. However, during this fascinating time, we also see the diffusion of something distinctly *social*, a spirit or recognition among practitioners that their embrace of this suite of methods places them in relation to each other and to future practitioners in ways that seem to matter for developing a better understanding of nature. This recognition forms the basis of what will evolve into increasingly well-defined *intellectual communities*, of which our modern scientific communities provide instructive examples (to varying degrees).

Groups of this kind are themselves nothing new in the seventeenth century. Intellectual communities can be found in recognizably mature form since antiquity in cultures throughout the world. But, for reasons which I believe are (like the seventeenth century itself) still not very well understood, the intellectual communities that coalesced around scientific methods – and the intellectual cultures which came to define those communities – proved to be remarkably effective at producing knowledge of nature. When this suite of methods combined with a certain set of norms governing a culture of inquiry, something historically unprecedented was achieved. Precisely how or why this happens has proven to be a very difficult question to answer. Nevertheless, one of the most important developments in our understanding of the nature of science has been the scholarly consensus that has emerged around the idea that these communities are essential to the epistemic power of modern science.

Once you start looking at science from this socially oriented perspective, you quickly begin to see a number of ways in which the entire edifice of scientific knowledge – indeed, the scientific enterprise itself – rests on phenomena that are essentially social in nature. To take one example already mentioned, what does it mean to say that an idea becomes knowledge when it solves a puzzle? What I called the "natural response" to this question – the method-centered response – does not work, and the reason it does not work is because an investigation could still check all the "method" boxes and yet fail to be accepted as a solution by members of the scientific community. Another idea might be accepted as a solution despite failing to check several of the "method" boxes. Ultimately, what seems to matter is just whether the community accepts the idea as a solution to the puzzle. The conditions under which it is prone to do so vary across disciplines and across time within a discipline. What remains relatively invariant, though, are (1) the way in which widespread agreement within the research community results in the suspension of debate on an idea's acceptability, (2) the lifting of that suspension in the face of overwhelming pressure on the community to do so, and (3) the way in which widespread agreement *per se* fosters the community's ability to generate more knowledge. Each of these invariants plays a deep and significant role in the production of scientific knowledge. And for each invariant – the status of a puzzle as *solved*, the role of that status in the further development of inquiry, and the eventual erosion of that status – the prime mover is the set of social relations upon which a puzzle's solution ultimately rests.

When we look at the extraordinary achievements of the natural sciences, it is tempting to connect their ability to produce knowledge *per se* to the distinctive suite of methods for which the seventeenth century is credited as midwife. But it's crucial here to disentangle two dimensions of scientific knowledge that become deeply intertwined from around that time, right down to the present day: (1) the specific methods of investigation that practitioners employ as part of their efforts to refine our understanding of nature, and (2) the social mechanisms by which those efforts are weighed and filtered. While the former methods can plausibly be understood to be importantly related to what makes some knowledge *scientific*, they do not, I will show, provide us with an understanding of how some instances of scientific investigation become *knowledge*. The natural sciences produce knowledge. Not necessarily because they do experiments, or because they use precise measurement devices, or because they investigate reality, but because they have developed highly conservative epistemic cultures whose members are overwhelmingly concerned with what the community thinks. My purpose in this chapter is to support this claim as one component of a more general conception of disciplinary knowledge, a species of knowledge of which both the natural sciences and the humanities have historically been able stewards. If we use the natural sciences as a model for what real knowledge looks like, the question, "Do the humanities create knowledge?" turns not so much on the degree to which the humanities employ the Scientific Method, but on the degree to which they take part in the social processes by which disciplinary knowledge is achieved.

2.2 Alternatives to Knowledge

The eminent historian Stefan Collini rejects the notion that the humanities are primarily engaged with the production of knowledge: "…'knowledge' itself is surely less than ideal as a description of what we're after… 'Knowledge' is too easily thought of as accumulated stock, as something that doesn't need to be discovered again and is simply there for anyone who wants to use it." Rather, he suggests:

> [t]he contrast with 'understanding' indicates a lot of what it leaves out or misrepresents, and even a term like 'cultivation' has a claim here, or would do had it not come to be so closely associated with images of affected connoisseurship and simple snobbery….But 'understanding' underlines that it's a *human* activity, and so is inseparable from the people who do it. Notoriously, the possibilities of extending our understanding depend not just on what we already understand, but also on what sorts of people we have become. (Collini 1999, 237)

For Collini, this contrast underlies the distinction between "research" and whatever it is that humanists do:

> It has to be said – and has to be said now more emphatically than ever – that in many areas of the humanities 'research' can be a misleading term. It is difficult to state briefly how work in these areas should be characterized, but…we are at least pointed in the right direction by phrases like 'cultivating understanding', 'nurturing and extending a cultural heritage', and so on …. Publication in the humanities is, therefore, not always a matter of communicating 'new findings' or proposing a 'new theory'. It is often the expression of the deepened understanding which some

individual has acquired, through much reading, discussion, and reflection, on a topic which has been in some sense 'known' for many generations. (Collini 1999, 243)[1]

I imagine that Collini's putative alternative form of scholarly production will resonate with many humanists – as it does with me – because it reflects widely held views about the spirit and function of humanistic inquiry. It also identifies categories – *understanding*, *cultivation*, and *cultural heritage* – which most of us recognize as important and worthy of cultural investment. If our scholarly effort "cultivates understanding" or "preserves cultural heritage," we needn't trouble ourselves over the possibility that the humanities do not create knowledge. Perhaps they even create something more valuable than knowledge. Who would doubt that scholarly engagement with the dialogues of Plato, or the Homeric epics, or Alberti's writings on perspective, or Thucydides' account of the Peloponnesian War, or the Sistine Chapel, will outlast Einstein's General Theory of Relativity, or the Standard Model of particle physics? It is easy to imagine a world three hundred years from now in which, although physical science continues to grow, Einstein's theory is relegated to a footnote, while Plato and Confucius continue to enjoy tens of thousands of references annually. That is a world in which Einstein's theory has been superseded by a framework that better facilitates physical inquiry; we *hope* to see such a world. By contrast, it would seem that for the world to become Plato-free in three hundred years would require a cultural catastrophe of historic magnitude. To envision a world no longer in conversation with the wisdom of Socrates is to envision a profoundly impoverished iteration of human civilization.

For reasons which I articulate throughout the book, I do not think that distancing the humanities from the idea of knowledge is a very promising strategy. One very simple reason why it lacks promise is that the assertion is radically at odds with the history of humanistic inquiry, both in word and in deed. Humanists and their intellectual communities from antiquity to the present typically *do* relate to their scholarly endeavors in ways that are most naturally interpreted as epistemic; this is part of the reason why the social epistemology of the natural sciences appears to offer an intuitive framework with which to guide an examination of humanistic inquiry. Peer-reviewed research journals are a social epistemic phenomenon. Research grant review panels are a social epistemic phenomenon. Academic conferences are a social epistemic phenomenon. Each of these

[1] This distinction is endorsed by Small (2013, 2).

institutions is crucial to the development of knowledge in scientific communities. We have a fairly good understanding of how they function to promote the growth of scientific knowledge. As in the humanities, they do so through the social processes that make knowledge *disciplinary*.

Of equal significance is the fact that the kinds of scholarly categories which Collini takes to be distinctive of inquiry in the humanities are *also* serviceable vehicles for capturing the nature of scientific research. Mathematics provides a convincing illustration. It is common within mathematics to search for new ways of proving a theorem that has already been proven. While a new proof of an old theorem is a "new finding" in Collini's sense, that is not what motivates the mathematician's search, nor is it what is valued by the broader community of mathematicians. Instead, what would make the new proof valuable is if it offered us a better understanding of *why* the theorem is true.[2] This is literally how they talk. The focus on well-established theorems as subjects of mathematical research resembles in every way Collini's image of distinctively humanistic research as "expression of the deepened understanding which some individual has acquired, through much reading, discussion, and reflection, on a topic which has been in some sense 'known' for many generations."

This symmetry can also be found across the great works of scientific achievement that are venerated to this day, even outside the confines of scientific communities. Darwin's *Origin*, for example, reports but a few isolated and mostly insignificant "new findings," such as his observations about the ability of seeds to germinate after a long soak in the sea or after having been dried out for ages; there can't be more than a dozen references to such low-level findings in the nearly 500 pages that comprise the *Origin*. And it scarcely bears mention that these findings are in no way the anchor for the *Origin*'s intellectual contribution, nor have they ever been. The *Origin* is at its core a collection and highly persuasive arrangement of facts that had been uncovered by the last few generations of naturalists, along with an articulation of a handful of common-sense principles about selective breeding that had been in practice since who knows when. If ever there was an "expression of the deepened understanding which some individual has acquired, through much reading, discussion, and reflection, on a topic which has been in some sense 'known' for many generations," it is *On the Origin of Species.*[3]

2 See, for example, Tappenden 2005 and Lange 2015.
3 Ospovat (1981, chapter 4) argues that the bulk of Darwin's scientific research consisted of precisely these activities.

Although it is true that most modern scientific research publications do not take the form of great works of literature in the way that Darwin's do,[4] it is absolutely routine for them to take the form of "the expression of the deepened understanding which some individual has acquired, through much reading, discussion, and reflection, on a topic which has been in some sense 'known' for many generations." I will here offer two more well-known examples, with the promise that further illustrations will be introduced later on in the book. The fact that the sun was somehow the cause of the orbit of the planets around it had been known for a few generations before Newton (Cohen 1985, Chapter 6). Among the first handful of results that Newton proves in the *Principia* is the fact that a body moving in a straight line while simultaneously subject to a continuous impulsive force in the direction of a particular fixed point will trace an ellipse around that fixed point (Proposition 2, Theorem 2). Upon reading this proof, Newton's audience would have been treated to "the expression of the deepened understanding which some individual has acquired, through much reading, discussion, and reflection, on a topic which has been in some sense 'known' for many generations." Those who followed Newton's proof would, in the same sense to which Collini appeals, have found themselves to have gained a deep and profound understanding of the well-known fact of the sun's causal role in producing elliptical orbits.

Again from physics, Einstein's famous 1905 paper, "On the Electrodynamics of Moving Bodies," reports no "new findings," nor does it "propose a new theory." It merely (!) observes a tension between a few fundamental tenets of physical science. I would hesitate to describe it as an "expression of deepened understanding," except in the sense that it deepened our understanding of our own ignorance, which is arguably of greater value. However, Collini's "cultivating understanding" seems entirely appropriate: physicists of the nineteenth century had adopted a number of commitments in different physical contexts which, when brought together by Einstein, enabled them to see clearly that their views had unacceptable consequences regarding the difference between an electrical field and a magnetic field (Renn 2007).

Further illustrations from the history of science could be adduced, but they ultimately would only reinforce a point which I hope has by now been convincingly established – namely, that a qualitative distinction between the cognitive aims of scientific inquiry and those of the humanities is not easily made. The "nurturing, animating, revising, and extending

[4] *Descent of Man* (1871) and *On the Various Contrivances* (1862) are beautiful and engaging works.

[of] our understanding" is every bit as central a preoccupation of inquiry in mathematics and the natural sciences as it is in the humanities (Collini 1999, 238). The perception that it is *not* central – let alone so marginal as to constitute a passable metric for distinguishing between scholarly production in the sciences and the humanities – is based on a misleadingly limited conception of what actually goes on in science, one which envisions the essence of productive scientific inquiry to reside in the ahistorical accumulation or "cataloguing" of facts.[5]

Contrast that limited conception with this poignant characterization of inquiry by Collini, in which any practitioner of mathematics or the natural sciences would instantly recognize her own struggle to generate deeper scientific understanding:

> The truth is that there is often work by our predecessors which it may be right neither simply to repeat (even were that strictly possible) nor to repudiate and replace with something else. The proper response may be to acknowledge it, possess it, learn from it, and allow it to inform our understanding. One trouble with this way of putting it is that it may seem vulnerable to the charges of rigidity and passivity: any suggestion of merely handing on our cultural inheritance makes us seem like rather indolent museum curators … who are sure that everything worth preserving is already in the collection. But this is a misconception of what this kind of understanding involves. For each generation to repossess a cultural inheritance … is to modify and extend it. Apart from anything else, our understanding has to be different from that of previous generations just because it is ours: we fit it into the framework of other things we understand, we articulate it with our other concerns … and we restate it in our idiom and for our audience.

Each of the historical examples I provided above is well captured by this account of inquiry. Darwin, Newton, Einstein, along with every other scientist since the seventeenth century, have each taken the work of previous generations, "acknowledged it, possessed it, learned from it, and allowed it to inform their understanding." Far from employing it "rigidly and passively," they "fit it into the framework of things they understood, articulated it with their other concerns … and restated it in their idiom and for their audience." For instance, Galileo, Newton, and Einstein would each look anew at the phenomenon of an object in free fall. Galileo's understanding of this phenomenon as a physical process is strongly informed by his predecessor and chief antagonist, Aristotle.[6] Newton, with his sights

[5] Ibid., 238. "Cataloguing" is another term that Collini employs as an example of what the humanities are *not* about, in contrast to other scholarly endeavors.
[6] Westfall 1971, esp. Chapter 1.

https://doi.org/10.1017/9781009067508.003 Published online by Cambridge University Press

set on an inertial physics, has no use for Galileo's Aristotelian *physical* conceptualization of free fall. Nevertheless, he incorporates Galileo's *mathematical* description of free fall into his new physics, "articulating it with his other concerns," and "restating it in the idiom of" his Second Law of Motion.[7] More than two centuries later, Einstein would reinterpret the phenomenon of free fall – along with the entire edifice of physical science which had been erected upon it – within the framework of relativity, a framework which had grown out of his 1905 attempt to articulate the other concerns mentioned above. It is impossible to see Einstein as repudiating either Galileo's or Newton's work; surely he sees more of himself in them than they would. He "acknowledges it, possesses it, learns from it, and allows it to inform his understanding." He then "modifies and extends it." The observation below, made by Damerow et al. in reference to the period between Galileo and Newton, could with equal justice apply to the transition from Newton's "classical mechanics" to Einstein's mechanics:

> The conceptual development embodied in the transition to classical mechanics cannot be identified with any particular way station and is not to be found in any particular text. It is a process which begins with such figures as Descartes and Galileo and takes shape with the generation of their successors. These disciples or even adversaries read the old problems and arguments from the point of view of their new solutions, thus establishing classical mechanics, because their point of departure was now the concepts as they are implicitly defined within the derivations of the theorems, e.g., the law of free fall. Thus, while for the first discoverer, the law of free fall is achieved by applying and modifying an independently grounded, pre-existing conceptual system, for his disciples it is the law of fall that canonically defines key concepts in a new conceptual system. The very same reading of these theorems that establishes classical mechanics also obliterates the traces of its real historical genesis because the original problems and the concepts involved are now understood within a very different theoretical and semantic framework. But since the successors themselves derive the inherited theorems on the basis of the new concepts, they impute these concepts to the discoverers. (Damerow et al. 1991, 5)

The interesting thing here is that the poignant characterization of inquiry with which I began the previous paragraph is actually Collini's attempt to characterize inquiry in the *humanities*. It is unsurprising that his entirely apt characterization of humanities research so closely matches the description of the conceptual development of classical mechanics, along with

[7] Notes the Newton historian I.B. Cohen (1985, 155), "only a Newton could have seen [the Second Law] in Galileo's studies of falling bodies."

https://doi.org/10.1017/9781009067508.003 Published online by Cambridge University Press

countless other episodes in the history of science. The pursuit of scientific knowledge is an intensely historical endeavor, an underappreciated fact in itself and one which partly accounts for the immense power of modern science. Indeed, how could it be otherwise? Were it so ahistorical, each practitioner would be "forced to build his field anew from its foundations" (Kuhn 1962, 13). Although it is true that, as Collini remarks, "the humanities … are inherently 'conversational' subjects," this is equally the case for science. Scientists and mathematicians are in constant conversation with their predecessors, "and conversing … requires a constant, flexible, responsiveness" – precisely the kind of flexible responsiveness we see in the sciences as each researcher picks up the work of previous generations, reinterprets it from within her framework, connects it with her concerns, and restates it for her audience and in her idiom (Collini 1999, 238).

I have made a point of quoting Collini at length because it is important to see how a very capable account of scholarly effort in the humanities is able to function as an accurate and powerful encapsulation of scientific research. That is no accident. We in the humanities have adopted a certain vernacular for describing our creative output, I suspect with an eye toward contrasting our intellectual contribution with that of science for reasons that need no rehearsal (talking about university admin here). I do not believe that the intended contrast holds up to scrutiny, and would further wager that an alternate vernacular designed to facilitate similar ends would meet a similar fate. The problems we face in articulating what is distinctive and valuable about our work are not of the superficial sort that can be solved merely through artful redescription. They are substantive. And they can only be addressed by holding ourselves to a standard that is not designed to ensure our success.

The way we talk about the humanities – or the way in which we attempt to contrast them with the sciences – is a significant problem, one which is rooted partly in our defensiveness about our value in the current culture and partly in widespread ignorance about the nature of science. But it is not the deepest problem we face. That problem, too, can be brought to the fore by once again looking closely at Collini's able description of the humanities. We have seen that, in developing his contrast with knowledge, Collini emphasizes the alternative humanistic production of *understanding*: "Publication in the humanities … is often the expression of the deepened understanding which some individual has acquired"; "the possibilities of extending our understanding depend not just on what we already understand, but also on what sorts of people we have become"; and so forth.

https://doi.org/10.1017/9781009067508.003 Published online by Cambridge University Press

I now want to look at what sort of state of affairs we are asked to envision when we imagine a scenario involving "deepened understanding," or "extended understanding," as well as "what we already understand," etc.

Regarding a publication in the humanities, for example, how could we tell whether it "expresses deepened understanding which some individual has acquired"? For an individual to "acquire deepened understanding" sounds like an achievement of some sort. We're all familiar with the feeling of that kind of achievement. I can remember my wife explaining to me why Darcy concealed his past dealings with Wickham from Elizabeth. For me, understanding that was an achievement. I don't share Jane Austen's social sensibilities; the idea of family honor is not intuitive to me. Also, not so good when it comes to understanding people's emotions. But I got there eventually. As educators, we're often lucky enough to elicit that feeling in our students. I'll never forget the expression on my son's face when he realized why the area of a triangle *had* to be *bh*/2. Now, of course, neither of these events is appropriate for publication intended for an audience of scholarly peers. But not because they fail to qualify as "deepened understanding which some individual has acquired." In fact, it is precisely *because* they are such clear instances of deepened understanding that makes any announcement to the scholarly community unnecessary. We already have a good understanding of why the area of a triangle is *bh*/2. Anyone who is not burdened by my endearing brand of sociopathy is able to appreciate Darcy's reticence. In order to be appropriate for submission to a scholarly audience, there needs to be some outstanding question regarding whether or not the insight I think I've achieved is genuinely plausible. If there is no such question – if, say, Darcy's guarded behavior is a part of "what we already understand" – there's not going to be much demand for me announcing that I've finally caught up with the rest of you. However, if my thoughts do not reflect "what we already understand," then I can hardly claim to have achieved deepened understanding just because I say or think something. Perhaps what I'm expressing is not the acquisition of deepened understanding but rather a misguided, muddled, unenlightened, or superficial reflection. Whether my humanities publication expresses deepened understanding is something that has to be sorted out at the level of the *scholarly community*.

2.3 From Ideas to Knowledge

To begin framing the problem of humanistic knowledge, I want to set out the preliminaries for a distinction that will occupy us in one way or another for the remainder of our discussion. Traditionally – at least, within philosophy – we are prone to contrasting knowledge with "mere opinion" or "mere belief." That contrast has been one of the main foci of epistemology since it was explored with estimable clarity in Plato's *Meno*. Although there is instructive overlap between Plato's (and most epistemologists') conception of knowledge and the kind of knowledge we'll explore over the course of this book, they are nevertheless quite distinct. Whereas the appropriate contrast class with Plato's *knowledge* is "mere opinion," the distinction that concerns us is between knowledge and *ideas*.

In his study of the reception of scientific theories, *Making 20th Century Science*, historian of physics Stephen Brush tells the story of "how theories became knowledge."[8] For Brush, an idea becomes knowledge when that "idea is adopted [/accepted] by the relevant scientific community" (Brush 2015, 3). In Brush's account, an idea is adopted/accepted (henceforth "adopted") when it moves beyond the unruly realm of scientific disputation and into the comfortable dotage of warranted presupposition. His interest lay in understanding the kinds of reasons that tended to motivate members of a community of practitioners to eventually treat some idea as a relatively fixed point with which they ought to make their own research consistent going forward. Although the specific form taken by these motivating reasons is in some sense an artifact of their association with natural science, there are a few instructive lessons that we can derive, both from the goals of inquiry that these motivating reasons seem to service, as well as the conception of knowledge that Brush employs.

The first thing worth observing about Brush's study is that an idea does not need to be "true" in any sense in order to "become knowledge." Neither, interestingly (and relatedly), does the idea need to be *believed* by anyone. All that's required is that the idea come to be *used* by the community's members as part of the more or less uncontested background to research and education in the discipline. Knowledge in this sense and in this context thus differs in important ways from the kind of knowledge possessed by an individual agent who has a true belief that is also justified. To say that I know next week's lottery numbers are going to be 1, 2, 3, 4, and 5 is, in the traditional sense, to say that (a) the lottery numbers *are* going to be 1–5, (b) I believe that they are going to be 1–5, and (c) I have persuasive evidence that they are going to be 1–5 (let's just say I know a

[8] In an earlier, unpublished draft which Brush was kind enough to share with me, the book's title was, *How Ideas Became Knowledge*. I assume that my interest in the contrast between ideas and knowledge (or, at least, a certain conceptualization of that contrast) derives from this encounter with Brush's terminology.

guy…).[9] Satisfying only (a) and (b) won't get me there; that's just a lucky guess, a "mere opinion" which happens to be correct. Nor will satisfying (b) and (c) suffice; I'd just be *wrong*. If anything is incompatible with the traditional Western philosophical understanding of knowledge, it is the notion that someone could know something that is not true.

When we think of knowledge in scientific contexts, however, we need to treat the communal function of scientific knowledge as paramount. And this is a function which ideas can perform whether or not they are true, as well as whether or not they are believed to be true. It is a function which ideas can perform even when there is no persuasive evidence in favor of their truth. This is because the role that ideas play in science depends more on what the community of practitioners *agrees to use* to propel the study of nature than it does on what mind-independent nature is fundamentally like. But, surely practitioners would not adopt an idea unless they believed it to be true, right? Right?! I think it is very far from clear whether that is the case, and I think that lack of clarity says something of profound significance about the peculiar nature of *scientific* knowledge.

Looking across the history of science, we find countless instances of ideas which we would regard as literally false nevertheless serving this communal function. We find practitioners employing ideas which they by their own admission do not believe. And we find them adopting ideas which clearly lack persuasive evidence. None of this makes any sense if we view the adoption of a scientific idea as the adoption of a belief about nature. If, instead, we view the adoption of a scientific idea as the adoption of a *technique* used to *study* nature, we are able to fit a lot more of what researchers do into a coherent picture of knowledge production. In adopting a technique, we do not ask whether the technique is true; techniques are not the sorts of things that can be true. In adopting a technique, we do routinely demand something like evidence – but not evidence of its truth. Rather, we seek evidence of its efficacy. There are better or worse techniques, or techniques which are more or less useful.

Viewed from this perspective, an idea becomes scientific knowledge when members of the relevant community adopt it as part of a general community-wide approach to the study of nature. Prior to that, an idea is, well, just an idea. Peer-reviewed publication of an idea is normally the *beginning* of a process which may or may not eventuate in that idea's

[9] Following Gettier (1963), it is now customary to include a fourth condition to the effect that conditions (1)–(3) could not easily have been otherwise. This reflects Socrates' emphasis on stability in the *Meno*.

becoming knowledge. A peer-reviewed research publication functions as an invitation to members of the scientific community to consider an idea as a candidate for community uptake and, thus, for scientific knowledge. Even if we think that only scientific ideas that are true can be published in peer-reviewed journals, time will tell whether a certain true idea becomes knowledge. Because of the essentially *social* nature of scientific knowledge, each idea – even the *true* ones – must go through a community-level process of determining whether it will serve the communal function performed by scientific knowledge.

Although it has taken a very long time, this thesis – which has been obvious to historians of science over the last century – is gradually beginning to gain a substantial foothold within the philosophy of science. But wherein lies its significance for the question of whether the humanities generate knowledge? I have devoted considerable space to introducing this socially-oriented conception of scientific knowledge for two reasons. First, I believe that it offers a plausible and accessible model for a certain kind of knowledge that is held specifically by research *communities*, regardless of the focus of their research. It thus offers a guide with which we can assess questions related to the research output of humanities disciplines. Humanists produce lots and lots of ideas. I think it is an open question whether, given the nature of contemporary humanities research communities, any of these ideas can be placed on a trajectory that might eventuate in knowledge.

Another advantage relates to the way in which this conception of scientific knowledge is inoculated against some of the frailties that afflict efforts to apply individualistic, proposition-oriented conceptions of knowledge, *mutatis mutandis*, to the context of scientific inquiry. In particular, it avoids the difficulties we encounter when we insist that an idea needs to be true in order to qualify as knowledge. In avoiding this complication, moreover, we remove the temptation to insist on a deal-breaking disanalogy between scientific and humanistic knowledge, based on the supposition that scientific knowledge takes the form of truths about nature, whereas humanistic knowledge takes the form of … well, something else altogether. I agree that humanistic knowledge, were it to exist, would normally take a form other than that of truths about nature. My primary aim in this book is to explore the question of whether that alternative form can be understood as closely related to the form that we now associate with scientific knowledge, as well as the question of whether a mismatch between these two domains can help us understand something deeper about the peculiar nature of humanistic knowledge.

Thirdly, the community-centered conception of scientific knowledge allows us to connect the idea of knowledge in the humanities with some of the alternatives to knowledge we explored above. For, the process by which ideas become knowledge in science is akin – perhaps, even identical – to the process by which ideas come to constitute items of deeper understanding. Upon initial completion, the idea is submitted to a select group of peer referees. This unrepresentative sample of scholars eventually makes a determination as to whether the idea is fit for declaration to the broader scholarly community, that is, whether it is fit for publication. As with knowledge, the scholarly community's reaction to the idea will determine whether deeper understanding has been achieved. If the idea comes to form part of the framework through which members of the community develop their own ideas – as Darwin's and Einstein's ideas eventually did – then we can assert with confidence that the publication "expresses deepened understanding which some individual has acquired." However, if the idea is rejected, or ignored, it is hard to see upon what basis a verdict of "deepened understanding" could rest. If, after publication, the relevant scientific community rejects the idea as confused or false, or just pays no attention to it, then to nevertheless insist that the idea constitutes deepened understanding is to imply that the community's assent is immaterial to whether understanding is achieved. But if that is the case – if the community's uptake is genuinely irrelevant – then why go through the charade of peer review and publication? Scientific journals do not exist simply for practitioners to express their thoughts and feelings. That is what diaries are for. Scientific journals are intended to function as venues for the proposal and evaluation of ideas at the community level, to see whether those ideas have properties that satisfy the community's norms for adoption. If we embrace a conception of scholarly understanding – or of knowledge – that obviates the need for community-level adjudication, we erase any meaningful distinction between an individual's sense of intellectual satisfaction, on the one hand, and genuine insight.

The notion that community-level adjudication is essential to the acquisition of scientific understanding can be clearly discerned in the development of Darwin's theory of natural selection. Darwin hit upon the basic idea of natural selection in the September of 1838, after reading Thomas Malthus's *An Essay on the Principle of Population*. We know this because he recorded the event in a notebook he was keeping at that time (Ospovat 1981, 61). Darwin published a basic outline of the theory in 1858 in a short paper, with the more carefully elaborated version appearing in 1859 in *On the Origin of Species*. That book, in turn, would go through five subsequent

https://doi.org/10.1017/9781009067508.003 Published online by Cambridge University Press

editions, the sixth containing roughly 50% new material as compared with the first. The question we now want to answer is, at what point did someone – *any*one – come to understand how populations of organisms are modified through natural selection, "the preservation of favorable variations and the destruction of injurious variations"?

Part of the difficulty in answering this question lies in the substantive revisions that Darwin would make to the theory between 1838 and 1859, such as his shift from a belief in perfect adaptedness to one of adaptedness relative to conspecifics (Ospovat 1981, Chapter 3). That is probably the most significant update Darwin made to the idea of natural selection *per se*, although other core components of the Darwinian picture – such as how natural selection produces entirely *new* species – would also undergo fundamental rethinking (Kohn 2008). So, if we believe it was Darwin who first understood that populations are modified over time by their environments, that achievement could have occurred as late as 1859ish, when *On the Origin of Species* was completed.

There are a variety of features related to the book's reception that appear to undercut this conclusion. There isn't time to survey the community's responses,[10] but two reactions bear particular notice, given their relevance to the acquisition of understanding. One, which a number of critics seem to have had, was that the projection from domestic breeding practices to a principle governing the modification of wild populations was not valid. If that inference was generally regarded by practitioners as invalid (as indeed it appears to have been), it is difficult to support the claim that Darwin achieved understanding of how populations are modified prior to the publication of the *Origin*. His argument rested on an inferential step that violated accepted norms of scientific inference at the time (Hull 2003). For all anyone had reason to think, he "may be as woefully wrong as Humphrey Belcher, who believed the time was ripe for a cheese cauldron."[11]

The violation of accepted inferential norms exemplifies some of the difficulties we face, because it provides an illustration of a clear sense in which the very meaning of notions such as *knowledge* and *understanding* depend on the existence of community-wide conventions concerned with how to properly conduct inquiry. Despite the intuitive pull of the idea that Darwin achieved genuine insight after reading Malthus (more so in the ensuing two decades), we have to resist this temptation if the scientific

[10] See Hull 1973; Engels and Glick 2008; Ruse 1979.
[11] Rowling 2005, Chapter 10.

https://doi.org/10.1017/9781009067508.003 Published online by Cambridge University Press

community's adjudicatory role is to carry any weight. If we think it should carry weight, then we need to at least suspend judgment with respect to the question of whether between 1838 and 1859 Darwin acquired a deeper understanding of how populations of organisms are modified.

Another famous reaction to Darwin's theory provides, I think, a more decisive verdict on the question of what Darwin understood. In 1865, Scottish engineer (and, according to Wikipedia, inventor of the cable car) Fleeming Jenkin, published a review of the *Origin* which was to haunt Darwin. In his review, Jenkin showed that Darwin's suggested mechanism for evolutionary change was fundamentally incompatible with his views about the nature of biological inheritance. In particular, what Jenkin showed was that if one assumes that (1) biological traits are inherited through a "blending" process – rather than in discrete units like genes – and that (2) the variations that provide environmental advantages tend to be "minute," as Darwin claimed, then favorable variations will inevitably be diluted over time and will cease to make a difference to the properties of organisms. In short, the process of evolutionary change that Darwin envisioned was simply not possible. If his image of evolutionary change was known to be self-contradictory, it is far from clear that that image nevertheless amounted to some kind of understanding at the time. The problem that Jenkin raised for natural selection was deceptively difficult, as evinced by the fact that the conceptual waters remained cloudy for roughly sixty years after Jenkin's initial criticism. It was not until the 1920s that the muddle surrounding natural selection and the nature of biological inheritance was sorted out. It took three truly brilliant mathematicians – Haldane, Fisher, and Wright – who, during that process, created modern statistics. Up to that time, the community of naturalists was largely skeptical as to whether Darwin's theory described a coherent process of evolutionary change, and they were right to be so (Gayon 1998). Even after that, it still took another twenty years or so for the majority of practitioners to embrace the idea that, not only was natural selection *possible*, in fact it appeared to be the dominant evolutionary force – the "paramount power," as Darwin had called it (Brush 2009).

I have taken care to describe this historical vignette in some detail because it offers a clear and compelling illustration of the *processual* nature of scientific understanding, a process that in this case begins with a flash of insight in 1838 and ends with a series of publications between 1918 and 1924. The publication of Darwin's views was just one, relatively early and underdeveloped, node in this process. In a certain sense, though, so were the publications of Haldane, Fisher, and Wright; our picture of natural selection is still evolving in fundamental ways.[12] Most importantly, I want this episode to exemplify the way in which the historical trajectory of the process of understanding is dictated in large part by the community of practitioners. This process *could* have stopped with the publication of the *Origin*. But the cogency and publicity of Jenkin's criticisms in particular ensured that that would not happen; Darwin's idea could not and *did* not gain widespread acceptance until the community was convinced that those criticisms had been resolved. And yet, neither do we want to say that Darwin acquired *no* understanding in 1838. After all, it's not for nothing that we call it the *Darwinian* theory of natural selection. But this is as it should be. Given the way in which significant ideas are shaped, refined, reconsidered, and ultimately either transformed or rejected, often over several generations of thinkers, we should expect confusion to arise when we attempt to specify a particular moment or individual out of which understanding emerges. Characteristically, *there is no such moment or individual*. There are communities that undergo identifiable periods of relative equilibrium – not *stasis*, but periods in which members of a field partake of a shared framework for inquiry that encompasses everything from what constitutes an important research question to how specific bits of the natural world behave. "What we already understand" is determined by the elements of this shared framework. "Cultivating understanding," then, is often a matter of drawing out less salient, perhaps tacit, features of that framework and submitting them to the community for consideration.

Rather than test the reader's patience with yet further ways in which the large-scale cognitive aims of the humanities – knowledge, understanding, interpretation, and so forth – parallel those of the natural sciences, I propose to adopt as a working hypothesis that both domains participate in the general cultural form picked about by the term *disciplinary inquiry*. While the many manifestations of this form are as different as the disciplines that serve as its hosts, they exhibit a pleasing generality with respect to the intellectual goals that rational inquiry is perceived to satisfy. It is also, I think, a predictable one. From the moment we are born, we are in search of patterns that can give us insight into a broader system. Although we use these insights to an enormous variety of ends, from the purely cognitive to the doggedly practical, what unites their application is the hope that we can stand upon them to extend our grasp of experience even further.

[12] See, for example, Orr 2005. Also Wallace 1991, who chronicles the history of genetic load, a theoretical problem in population genetics that is still not resolved.

https://doi.org/10.1017/9781009067508.003 Published online by Cambridge University Press

There is now overwhelming historical evidence that this practice thrives especially well in communities with very specific properties, modern science being a relatively recent and powerful instantiation of those properties. As I have detailed above, the community's contribution to rational inquiry is indeed so distinctive that the meaning of terms such as *knowledge* and *understanding* have come to be defined partly in terms of communities when used in connection with disciplinary inquiry. It is strange that it has taken us so long to make this dimension of inquiry explicit in our thinking about the nature of knowledge and its cognitive associates. Looking back, it does seem that this was perhaps so obvious to scholars of the premodern era that it could scarcely bear mention. Modern science is in many ways a very late outgrowth of this ancient and widely embraced scholarly tradition. There is a sense in which, rather than being the first of its kind, modern science was a bizarrely delayed application of community-centered principles of intellectual growth to the domain of nature.

Although our understanding of community-driven rational inquiry continues to develop, many of its components have been subject to able scrutiny over the past half-century or so, such that we can now fit them into a reasonably coherent picture of how, in general, intellectual communities advance. And although this picture has been pieced together largely through the historical study of the natural sciences, it is important to appreciate that it need not have been. Had we instead derived our understanding of disciplinary inquiry from ancient intellectual communities (to whom natural science would have appeared as perverse and inadequate in many ways), we would, I believe, have developed essentially the same lens through which we now view the growth of scientific knowledge. It is time that we turned this lens toward the humanities.