CHAPTER I

Introduction

In the nearly perfect movie, *Office Space*, some business consultants are brought in to a software company to trim a little fat. As part of that process, employees are called to meet with the consultants, to more or less interview for the job they already have. During one of these meetings, a guy named Tom with a job that clearly did not need to exist becomes irate, spluttering about this and that in a desperate plea for relevance before his inevitable exposure as a superfluous, and thus expendable, accoutrement.

The world is full of Toms. We hope that we are not among them. Sometimes this is because we do not want to lose our jobs. Tom's protests in *Office Space* were focused on preserving his livelihood. Often enough, though, we simply lament the prospect that other people might not value what we contribute … and that they might be right to do so. We lament this prospect, not because we think we deserve their respect or recognition, but because we don't want to waste our time. Moment by moment, I feel the crushing threat of wasted time bearing down on me. I hear the echoes of Tom's futile, self-serving outburst as I try to tell myself that what I do is important.

The book you are reading is the result of an extended meditation on the question of whether the humanities produce knowledge. And I do mean *extended*. For half of my life, I have tried to work out for myself what it is exactly I do with the time I spend reading, writing, and talking to academics across a wide range of disciplines. When I describe this pattern to people outside of academia, they often burst out laughing; this frequently happens during interactions with natural scientists as well. My own children are pretty much convinced that I do nothing all day long. The way in which I and most other humanists work simply does not comport with their image of what knowledge production looks like. And the fact that I'm literally wearing a tweed jacket as I write this probably doesn't help, either.

https://doi.org/10.1017/9781009067508.002 Published online by Cambridge University Press

Their image of knowledge production comes from the natural sciences. They have a clear sense of what it means to produce knowledge because research in the natural sciences has produced concrete results that we can use to make stuff. Thanks to the natural sciences, my iPhone can accurately predict local temperature changes in ten-minute increments. Thanks to the natural sciences, we have developed increasingly powerful and efficient energy alternatives to fossil fuels. For many of us, the awe-inspiring concrete achievements of modern science have made the difference between life and death, between joy and grief. My son, a type 1 diabetic, has an insulin pump that can predict high blood sugar and provide compensatory insulin so that he can basically lead a normal life. The world owes the natural sciences a humbling debt for the development of the COVID-19 vaccine. In the soaring words of Richard Dawkins, "Planes fly. Cars drive… It works, bitches."[1]

To complement these tangible and often personally meaningful results, we have also been treated to a highly refined schematic representation of *how* they have been achieved – the Scientific Method: observation, hypothesis, experiment, analysis, and conclusion. Boom! Knowledge accomplished. In addition to being easy to grasp, there is something very intuitive about the process encapsulated here. Its essence is to think and to check, with some bespoke add-ons for the purposes of ensuring that the "checking" part is worthwhile. Thinking and checking are reflexes of the mind. The Scientific Method gives voice to those mental impulses, trains them a bit, and crowns them as the Royal Road to knowledge land. Hence, it is not merely that science works. It is that we understand why it works. Indeed, we understand why it *should* work. By disciplining and extending certain built-in habits of the mind, the Scientific Method allows us to learn about the world in a way that asks very little from us in the form of fundamentally novel behavior. Perhaps the Scientific Method is simply a natural stage in the development of humanity's relentless quest to acquire knowledge. Perhaps it is the ultimate stage.

It has long been fashionable in many circles to deride the notion of the Scientific Method – not just the canonical version I laid out above but the very idea that there is an identifiable method according to which scientists structure their investigations. There is some justice in these criticisms. Yet it would not be difficult to accurately model a great many actual scientific investigations as if they were following this template, at least at a very general level. That is to say, there is something right about

[1] www.youtube.com/watch?v=0OtFSDKrq88. Last accessed March 2, 2022.

https://doi.org/10.1017/9781009067508.002 Published online by Cambridge University Press

how the schematic Scientific Method portrays the nature and the spirit of scientific investigation. Were we to show that schematic representation to the likes of Galileo, Newton, Boyle, and countless researchers who lived before the emergence of the concept of *the Scientific Method*, they would doubtless instantly recognize the essential process it depicts, as would any practicing scientist of today who had never heard of the Scientific Method. That means something. It suggests that there is a gross pattern to scientific inquiry, and that scientists have seen themselves in that schematic depiction of science precisely because it accentuates the salient features of real scientific investigation.

The problem, I think, is that this schematic depiction elides an enormous amount of detail, detail which, were it to be widely appreciated, would almost certainly effect "a decisive transformation in the image of science by which we are now possessed" (Kuhn 1962, 1). Indeed, the volume of stuff it does not talk about is so copious that I actually need to divide it into two categories. First, each of the model's components – observation, hypothesis, etc. – represents an incredibly complex and confusing phenomenon. Thus, to say that, for example, observation is part of the Scientific Method is to make a claim that appears to be as clear as day but is in fact very poorly understood (see, e.g., Hacking 1981; Daston 2008). For instance, Nasim (2013) provides compelling evidence that the practice of drawing nebulae in nineteenth-century astronomy was *part* of the observation process, not simply an *aid* to observation. Rudwick (1972) shows how early attempts to draw fossils forced Renaissance naturalists to focus on representing certain features, without any knowledge of whether those features were biologically or taxonomically significant. What is more, each of these components "has a history," as they say; that is, each of them becomes part of the Scientific Method through a complex historical process, a process which might have gone in another direction. It took a long time for the use of hypotheses to be viewed as an acceptable way to do science. The eighteenth-century physicist Georges-Louis Le Sage endured widespread resistance to his use of hypotheses in trying to understand the nature of gravity (Laudan 1981). Charles Darwin was roundly criticized for the *Origin*'s liberal use of hypotheses (Hull 2003). *Darwin*. In addition, the form that many of these components take in our time would have been unrecognizable to our predecessors. To a contemporary scientist, "analyzing data" often just means running it through Stata or some other off-the-shelf statistical package. Newton never used Stata, because Newton didn't have a computer. Also, he didn't have statistics. In sum, to the extent that the schematic depiction of science is accurate, it is woefully underspecified, and that underspecification significantly affects how we understand scientific knowledge and its place in human culture. In particular, it suppresses the massive role of human judgment in the development of scientific knowledge.

That brings us to the second category. The traditional portrayal of the Scientific Method leaves out entirely what is in my opinion truly distinctive about scientific inquiry – namely, its social dimension. Observation, hypothesis, experiment, analysis, and conclusion – none of it amounts to much in the hands of a single individual, even when perfectly executed. Scientific knowledge is not the conclusion one draws from his properly carried-out scientific methodology. It is a social sausage-making process, during which groups of similarly inclined but often antagonistic scientific charcutiers select the finest cuts from the world's prize-winning pigs, applying a prescription-strength dewormer before surgically slicing them to bits and packing them in the highly elastic but appropriately constrictive intestinal casing of publication, so they can be smoothly digested by members of the community. This social dimension is what distinguishes the *idea* of a single person from the *knowledge* of a community. But the Scientific Method doesn't even mention this. It invites us to conceive of scientific knowledge as the output of a single Cartesian investigator. I have no idea why, and I won't speculate. Yet, here again, we see the exclusion of a defining feature of scientific inquiry that just happens to center around the necessity of human judgment. Indeed, even to say it is a *defining* feature of inquiry credits the social dimension with too small of a role. I'm sensing a pattern here….

We have concocted a story about how scientific knowledge is acquired that is very satisfying but not very accurate, which is actually pretty ironic if you think about it. That story, because of its generality and its plausibility, has proven to be as seductive as it is resilient. And as the eminent historian of science, John Heilbron has observed, "you do not have to be right to make a revolution. You have to have a plausible and comprehensive programme" (Heilbron 2013, 15). Now, if it were just another innocent seduction, that would be one thing. This dalliance, however, has spawned all manner of poisoned fruit. For, the story of the Scientific Method has precipitated a revolution not only in our conception of scientific knowledge but in the very idea of knowledge itself. In doing so, it has severely undermined our ability to recognize and appreciate other forms of knowledge. Before I present an overview of the specific alternative that interests me, I want to spend some time giving due respect to the forgivable slide from respect for science to disrespect for the humanities.

The philosopher of physics Harvey Brown once remarked that anyone who is not mystified by the law of inertia has not properly understood it (Brown 2005, 15). Something similar must be held in relation to the successes of modern science. Looking closely at the intellectual and experimental triumphs of modern science, particularly since the seventeenth century, we are presented with countless manifestations of human genius, creativity, and intellectual fortitude that are truly humbling. For reasons that are understandable but not altogether satisfying, these deeply human struggles are hidden from view by certain constraints imposed on the various forms of science communication – including education and journalism, the vehicles through which most of us make contact with modern science. The actual historical details of a scientific investigation – the undaunted, often spiritually taxing journey from some inchoate sense of confusion to (say) a mathematically precise expression of the behavior of unobservable particles – accentuate the human features of scientific inquiry. They reveal the scientist to be nothing more (or less) than an ordinary person with a dogged commitment to developing some understanding of nature with which she can be satisfied, if only for a moment. As we trace her struggles, we can relate to her passion, her frustration, and her desperate groping for a lifeline by which she can, with effort, momentarily breach the surface of a problem before being plunged back into the depths of befuddlement where light fails to penetrate. We sympathize with her temptation to cut corners and to be too easily convinced by signs of promise. We admire her disciplined refusal to take the easy way out, just as we cringe at the less-than-admirable depths to which she sometimes sinks to outdo her competitors. When afflicted with setbacks, we feel her disappointment. We celebrate her triumphs. In scientific investigation, we find a microcosm of the cognitive and emotional tumult that is human life.

And yet, the intellectual content of this recognizably human struggle is often something that relatively few of us can appreciate. What are Einstein's field equations? What form do solutions to those equations take, and what do they even mean? What are gravitational waves? Why is our ability to detect them significant? How does the "spike protein" facilitate our efforts to inoculate people against COVID-19? And on and on and on. While most of us have encountered these and other terms from the natural sciences, far more (including me) have at most a tenuous grasp of their meanings. Our ability to reflect on them, to turn them over in our minds, to get a feel for them, to understand their implications is even more poorly grounded. We lack the sensibilities that give rise to the mathematician's aesthetic appreciation of the Macdonald equation, described by the polymath Freeman Dyson as "the most beautiful thing that I ever discovered."[2] The simplicity of Schrödinger's equation is lost on us; it was not lost on his fellow physicist Werner Heisenberg. Although none of us finds any difficulty grasping the meaning of the symbol H_2O, the ease with which we can relate to it masks the intimidating cognitive demands of the conceptual and experimental morass that impeded its formulation over much of the nineteenth century (Rocke 1984). We understand that Darwin's theory of natural selection explains adaptation, yet how many of us are tempted by the meaningless suspicion that the theory of evolution is "just a theory"? Our scientific literacy (whatever that means) might be up to par. However, when confronted with demands that require more than a surface understanding, those of us without advanced scientific training or mathematical training – people like me – are essentially alienated from the practitioner's capacity to admire the relentless torrent of cognitive victories which characterizes the history of science and mathematics.

Thus, even if are able to see something of ourselves in the emotional tumult reflected in the finer historical details of scientific inquiry, our capacity to appreciate the significance of scientific achievement is typically quite limited. In this way, we join a venerable lineage of cognoscenti who, not for lack of intelligence, are unable to take part in the pleasure of developments made outside the range of our experience. Archimedes could have fit any living mathematician in his hip pocket. But he would have found our modern approach to measuring the area under a curve to be quite incomprehensible, despite the fact that it would have spared him much exhaustion. Cut off from a community of practitioners due to time, space, or specialization, he would lack the background of accumulated knowledge required to relate to these developments on more than a superficial level (though I suspect he'd get up to speed fairly quickly). To be perfectly forthcoming about my own limitations, following any discussion that goes beyond high school mathematics can quickly go from taxing to hopeless. It can sometimes take several uninterrupted days of concerted effort for me to reconstruct (for teaching purposes) some historically important geometrical proof. If you find yourself unable to list from memory Einstein's field equations, welcome to the club. If you cannot see what Freeman Dyson sees in the Macdonald equation, you're not alone. If the Schrödinger equation strikes you as no less simple than anything else in quantum mechanics, I feel your pain. Our paths diverged long ago from those who went on to the kind of training that results in a connoisseur's esteem for scientific achievement.

[2] www.ias.edu/ideas/2015/dyson-concinnitas. Last accessed March 5, 2022.

The corollary to all this is that those with training in the natural sciences and in mathematics often have a keen sense of the significance of these triumphs as well as, albeit somewhat less frequently, the struggles that went into achieving them. And although the outcomes can be exhilarating for practitioners, most science is extremely tedious and boring. No psychologically normal human would elect to endure for a span of a few minutes the sorts of things on which many scientists spend the best years of their graduate students' lives. The philosopher Robert Paul Wolff reflected on this gulf in an anecdote recalling the occasion of his meeting the entomologist and popular science author, Edward O. Wilson:

> We met in Wilson's office in the Museum. After the usual greetings, he showed me the centerpiece of the office, a large table on which, under a Plexiglas dome, was a bustling, complex ant colony. Wilson banged the side of the table, which set the ants scurrying, and as they poured out of the anthill he pointed out the soldier ants, worker ants, and so forth. I didn't have much in the way of conversation. What can you say about an anthill, after all? So, casting about for something to say, I mused aloud, "I wonder how many ants there are in the entire colony." "Fifteen thousand," Wilson replied. "How can you be sure?" I asked. "I counted them," he said.
>
> There are moments in life when the scales fall from your eyes and you suddenly see clearly something that has hitherto been obscured from view. This was one of those moments. I had from time to time reflected on how different the workaday lives are of people in different corners of the Academy, even though we all call ourselves "Professor." Here was E. O. Wilson, the creator of Sociobiology, who thought nothing at all about counting fifteen thousand ants. Had anyone asked me to figure out the number of ants in an anthill, the farthest I would have gone was watching eight or ten walk by and then guesstimating the rest.
>
> To be sure, philosophers sometimes descend to the level of the particular. But our tendency is to go in somewhat the opposite direction. Confronted with the real world, the reflex reaction of philosophers is to ask about possible worlds. It was clear to me that although we were both professors and authors, Wilson and I led lives so utterly different that no real mutual understanding was likely. It was also clear that however much the world might think of Wilson as the tendentious, controversial author of Sociobiology, his real interest was in those ants.[3]

This is not to deny "the pleasure of finding things out," as the great physicist Richard Feynman put it.[4] That pleasure, however, tends to come in

[3] https://robertpaulwolff.blogspot.com/2013/09/what-have-i-been-reading.html.

[4] In needs to be borne in mind that, by any standard, Feynman was an absolute genius; he was such a genius that the title of his biography is literally just *Genius* (Gleick 1992). Perhaps Feynman was not well-acquainted with the humdrum of science, given his use of "the Feynman Method," which his

the moment of insight, not the years-long slog through data, grant applications, disappointment, grant applications, and more data. The ability to achieve anything of real scientific value – to gain access to the delights of discovery – requires a degree of disciplined meticulousness and persistence that most people simply do not have. As the founder of atomic theory, John Dalton expressed it, "If I have succeeded better than many who surround me, it has been chiefly, nay, I may say almost solely from unwavering assiduity." Dalton himself amassed more than 200,000 entries in his meteorological diary (Pennock 2019, 153. See Strevens 2020 for many similar examples). Indeed, novice scientists are often stunned to discover that actual scientific research is not the blazing state of perpetual ecstasy that it is often depicted as being.

Caught between the mind-numbing monotony of actual scientific research and the general inaccessibility of scientific pleasures for most people, we have constructed a handful of convenient yet intuitively plausible distortions of the nature of scientific inquiry that we use to bring scientific knowledge into the consciousness of the nonscientific public. As I write, I am sitting in an auditorium with about 300 little kids (my daughter says more like 150), watching a man in an affected red bowtie and white lab coat boisterously pour liquid nitrogen on stuff to shock and amaze. His grand finale is dropping Cheetos into a bowl of liquid nitrogen and then feeding them to volunteers, to the delight of the entire audience. Professor McInquiry and his Dazzling Display of Wizzbangery are both parts of a tried-and-true method of inviting – or better, *luring* – unsuspecting marks into the scientific adventure through awe-inspiring manipulations of nature's hidden properties. Lying just beneath the surface, these marvels are ready-made vehicles for bringing the pleasures of scientific knowledge to a public that is ill-equipped to appreciate the more recondite versions of excitement to which the practicing scientist hopes to gain access. Because we can all partake in the spectacle of a rehearsed demonstration, and follow along with the explanation of how cold things get smaller and hot things get bigger, it becomes possible to cultivate an appreciation for the majesty of scientific knowledge without any background knowledge, without any particular talent, and without any particular predilection. Professor McInquiry was a hit with audience members of all ages. Wrapping things up, the man in charge of the event tellingly quipped, "He's also available for birthdays and bar mitzvahs."

colleague Murray Gell-Mann characterized in the following way: "Dick's method is this. You write down the problem. You think very hard. (He shuts his eyes and presses his knuckles parodically to his forehead.) Then you write down the answer" (Gleick 1992, 315).

https://doi.org/10.1017/9781009067508.002 Published online by Cambridge University Press

A parallel phenomenon exists in science journalism and in the popularized accounts of science that are published mostly by trade presses. Much of the way science is depicted in these venues is shaped by constraints like that of being interesting to "an old lady on her way to the grocery store," as one science journalist described it to me. Her tongue-in-cheek way of framing this constraint was intended to emphasize that science stories need an audience. They are not aimed at practicing scientists. They must be stories that speak to people from a panoply of backgrounds. In order to do that, you need to either (1) describe something totally amazing or (2) contort the history of discovery into a thrilling yet tidy narrative arc: the low rumble of complication, the promising early efforts, the setback about two-thirds of the way through the story, and the eventual triumph. Each of these strategies is designed to exploit well-known human tendencies – respectively, the love of novelty and the love of a good story. Potent in their own way, such tactics can do for anyone's morning commute or weekend reading what Professor McInquiry can do for your birthday party. They put scientific knowledge into a form that generates an immediate response in pretty much anybody.

The combined force of these publicly accessible distortions is, I think, incredibly powerful. We have perpetuated a base caricature of how scientific knowledge is produced in the form of the Scientific Method. We carefully select the most titillating and mystifying bits of that knowledge for public consumption. And we emphasize the (often illusory) astonishing potential applications of this knowledge. To add to this, there is a vast plenitude of actual applications that are quite legitimately astonishing. Scientific inquiry led to a COVID-19 vaccine in significantly less than a year. Scientific inquiry led to the transistor and the computer. Scientific inquiry has made it possible for my own child to live an essentially normal life through effortless management of his diabetes. Even if we acknowledge the genuinely horrific downstream consequences of some inquiry, we all know that we owe science big time. In this curious melange of legitimate admiration and ill-gotten awe, the idea that scientific knowledge is exemplary of knowledge itself is not a very hard sell. Literary criticism cannot cure diabetes. Philosophy cannot send something to the moon, not even those chimps that were rumored to have come back super-intelligent.[5] Art history cannot reprogram cells to search and destroy. The study of poetry in ancient Rome will not lead to insights into any of the Millennium problems in mathematics. We know that science can or will do these things.

[5] *The Simpsons*, Season 5, Episode 15.

https://doi.org/10.1017/9781009067508.002 Published online by Cambridge University Press

And we have a convincing explanation – the Scientific Method – of exactly how science does them. In short, we have a clearly articulated model for understanding how truly stunning triumphs of human ingenuity occur. They occur through *scientific* inquiry.

By contrast, the output of humanistic inquiry, however enjoyable it may be, generally lacks the power to astonish, particularly when it comes to the uninitiated. It consists almost entirely of words on pages – increasingly, on digital pages through which we rapidly scroll. There are no blinding flashes of light. There are no iPhones. There are no high-res three-dimensional computer-generated animations of the processes described by humanists; and if there were, humanists would probably object to them for some reason or other. There is no humanist counterpart to Professor McInquiry. Go ahead, try inviting an art historian to your seven-year-old's birthday party. What a raucous hoot that will be. Unless you can get Anthony Grafton to let the kids ride his book wheel, odds are that this is going to be a very disappointing and embarrassing birthday for your son or daughter.[6] While the humanities are just as capable as the natural sciences at producing that profound sense of wonder and appreciation among connoisseurs, we simply cannot compete with the power of the natural sciences to provide high-potency hits of dopamine to the novices.

But there is more to it than that, and here humanists would appear to bear total responsibility. This second asymmetry concerns the ability of nonspecialists to grapple with the nature, process, and significance of humanistic inquiry. In the natural sciences, the Scientific Method gives nonspecialists the sense that the wonders of science are not just magic. Yes, that sense is based on a way of representing scientific inquiry that is not faithful to loads of significant detail. But there's a word for that: *idealization*. Idealizations are valuable precisely because they afford a kind of cognitive grasp that does not get bogged down in certain particulars. That sense of understanding the process behind the magic, I have argued, has partly fueled the reputation of the natural sciences as having cornered the knowledge market. It makes sense of why science can do all these amazing things. If you take that sense of understanding away, all that is left is the Dazzling Display of Whizzbangery and the technological marvels. As early campaigners on behalf of science well knew, in the absence of an accessible narrative that could capture the process by which these marvels are achieved, the natural sciences were likely to remain in their stature as an

[6] www.princeton.edu/~paw/archive_new/PAW06-07/11-0404/features_grafton.html. Last accessed March 7, 2022.

https://doi.org/10.1017/9781009067508.002 Published online by Cambridge University Press

entertaining cabinet of curiosities (Laudan 1993). Thus, before the middle of the nineteenth century, long before the natural sciences could claim responsibility for world-changing devices, popular narratives of scientific progress had solidified the reputation of the natural sciences as knowledge generators *par excellence*. But we do not even have an *idealization* of humanistic inquiry, let alone a fine-grained model that faithfully captures the details. To put it bluntly, *there is no humanities equivalent of the Scientific Method*. There is no model, distorted or otherwise, of how humanistic knowledge is produced. There is no account of how the output of humanistic scholarly endeavor qualifies as knowledge of any kind. This is a serious problem.

To be sure, there is no shortage of books on the *value* of the humanities. Written largely in response to the declining appreciation for the humanities, they tend to be more concerned with the content and effect of ideas that have emerged through the tradition of humanistic thought, rather than with the process by which those ideas are refined and accepted (see, e.g., Nussbaum 2010; Small 2013). While there is much in these efforts to admire, what this approach does not address is the fact that not all ideas are created equal, and that the Scientific Method is now understood to be the all-purpose tool by which we are able to distinguish the ideas that we ought to take seriously from the ideas with which we need not bother. The Scientific Method is how we gain access to *truth*. The profound influence of philosophical and literary ideas, while not to be denied, does not tell us that they are true, and it does not establish them as knowledge. Aristotle's physics was enormously influential; still wrong. How do we know? The Scientific Method. Galileo disproved the Aristotelian idea that heavier objects fall faster by performing an *experiment* involving spheres of different masses dropped from the Leaning Tower of Pisa. Case Closed.[7] Now, where is the experimental evidence for the value of democracy? What predictive successes does John Rawls' theory of justice as fairness have under its belt? How do we even test the hypothesis that love is all you need? We do not need to reject the beauty of any of the results of humanistic inquiry in order to acknowledge the straightforward sense in which those results clearly fail to qualify as knowledge. The seventeenth century gave us modern science. Modern science gave us the tools to distinguish truth from really nice-sounding ideas. But a rogue group of sentimental bookworms chose to ignore those developments and continue playing with words like Aristotle did. Like Thomas Hobbes did. Like Galen did. The humanities today are the contemporary descendants of that group. They are not on the winning side of the history of knowledge.

[7] The veracity of this report, first made by Galileo's disciple and biographer Viviani, is widely disputed.

The Enlightenment produced an epistemological vacuum so completely devoid of substance that something as appallingly inaccurate as the Scientific Method had no trouble filling it. Since that time, natural science has continued to shore up its claim to epistemic hegemony through persistent efforts to refine and promulgate its theory of knowledge and by producing certain results that everyone can and should appreciate. Rather than develop an alternative model of knowledge, though, we have watched as more and more territory falls under the ostensible purview of natural science. Consequently, we now live in a time where the humanities are viewed as being constitutionally incapable of producing knowledge because they are not the natural sciences, and in which surveys conducted by evolutionary psychologists are perceived as having more to teach us about love than *Romeo and Juliet*. (Which romantic love measurement framework do *you* find most compelling – the Passionate Love Scale, the Triangular Love Scale, or the Love Attitudes Scale? I just can't make up my mind, although they say the Passionate Love Scale is really "only valid in people who are in a romantic relationship with their loved one," which makes total sense if you really think about it.)[8] Would that Shakespeare had lived in a time awash with such wonders as these. Perhaps then he might have produced something of real value.

1.1 Toward an Epistemology of Disciplinary Knowledge

The mid-twentieth century produced a perspective on scientific knowledge which sort of runs parallel to the Scientific Method and which was almost entirely neglected by theorists of knowledge. I take this alternative tradition to be principally the progeny of four highly original thinkers, who sort into two natural pairings. The first pair: Ludwik Fleck and Thomas Kuhn. The second: Michael Polanyi and Noam Chomsky. Each of these mavericks worked as an empirical scientist before training their sites on questions of knowledge. I'm going to briefly state their respective contributions to this alternative picture of scientific knowledge, before explaining how it is relevant to the problem of knowledge in the humanities.

In *The Genesis and Development of a Scientific Fact* (1935), Fleck, a bacteriologist and physician, shows how the knowledge basis upon which any scientific inquiry rests is an accumulation of ideas that have survived the scrutiny of a scientific community. Scientific facts are those propositions that are accepted by a scientific community as facts as a result of

[8] Bode 2021, 4.

https://doi.org/10.1017/9781009067508.002 Published online by Cambridge University Press

this process. Doubtless there are facts that no one knows about. There is a fact about how many dust particles adorn my computer screen. But that sense of *fact* is significantly different from the one employed in the phrase, "scientific fact." To say that some proposition is a *scientific* fact is to say something about a perspective at which a scientific community has arrived after proper scrutiny of some assertion of fact.

The community norms that determine whether a mode of scrutiny is proper, and whether the subject of that scrutiny is warranted, emerge from reflection on the disciplinary import of certain exemplars. So argued Thomas Kuhn, a physicist trained under the Nobel Prize winner John Van Vleck, in an elegant model of a specific historical pattern that many sciences appear to exhibit, described in his famous book, *The Structure of Scientific Revolutions* (1962). Kuhn developed a picture according to which exemplars are used in various ways to govern the process by which certain facts are selected for scrutiny, as well to as to govern the nature of scrutiny itself. In his model, the norms followed by practitioners in the course of scientific inquiry are acquired through training and experience, rather than through explicit instruction as to what the norms are and why they matter. The intensity of this process affects perception, cognition, and language, shaping practitioners into the communicative communities that engage with each other on a variety of levels to produce scientific facts in the sense outlined by Fleck.

Together, Kuhn and Fleck give us a picture of what real scientific knowledge looks like, because they give us a framework for seeing how that knowledge develops at the level of a disciplinary community. Any conception of scientific knowledge that fails to give due consideration to the centrality of the disciplinary community cannot be held to be credible. This is the principal reason why the Scientific Method is so inexcusably inadequate. Without the community component, the Scientific Method is ultimately just another appeal to someone's "lived experience."

The Chomsky-Polanyi pairing is focused on the recognition of a variety of knowledge that was radically out-of-step with the prevailing trends in twentieth-century epistemology. While somewhat idiosyncratic, Polanyi's (1958) *Personal Knowledge* provides an exhaustive catalog of instances of the kind of knowledge that people have but cannot articulate, which he called "tacit knowledge." Many of the instances on which he focuses were inspired by his training as a scientist and his long and distinguished research career in physical chemistry. Perhaps unsurprisingly, these instances of knowledge that derive from the context of scientific practice align precisely with the sorts of considerations which Kuhn portrayed as the principal determinants of the direction of scientific research. Scientists know things that they cannot or do not articulate, things that guide the development of scientific knowledge in profound and explicable ways.

No one who reads Polanyi's book and considers his staggering wealth of examples can fail to come away with the conclusion that propositional knowledge is but a subspecies of human knowledge more generally, and possibly not an overwhelmingly important one at that. What is more, this alternative variety – *tacit knowledge* – does not appear to result from the application of the Scientific Method, or of any explicit "method" at all. Rather, it travels via the same channels as our acquisition of skills and of cultural norms. It requires no explicit instruction. Indeed, attempts to explicate necessarily fall short and systematically mislead. In sum, Polanyi's impressive survey grounds Kuhn's generalizations about the kind of knowledge that scientists derive from exemplars, knowledge which shapes the nature of scientific practice within a community of practitioners.

The last contributor, Noam Chomsky, is the well-known founder of modern linguistics (also known for other stuff). Beginning in the late 1950s, Chomsky revived the ancient practice of systematically probing the value judgments of people who had achieved mastery over a certain body of knowledge, producing a model of that knowledge by constructing a careful mapping of its outer edges (Chomsky 1957).[9] As his principal focus was language, he took an interest in the kinds of utterances that respondents found intuitively unacceptable, even if they could not explain what it was about those utterances that displeased them so. From those objections, Chomskyan linguists formed a kind of negative image of the content of the norms governing speakers' use of language, norms which the speakers themselves could not articulate and of which they in general showed no signs of awareness. These norms are precisely the sort of phenomenon that Polanyi described as tacit knowledge, and Chomsky developed their extraction into a precise experimental research program.

Together, these four distinct components combine to give us a framework for understanding what I'll be calling *disciplinary knowledge.* Disciplinary knowledge consists of knowledge of the norms governing value judgments in a discipline. It is acquired through exemplification rather than explicit instruction. It is typically tacit. And it can be/has been/is studied through deliberate attempts to violate those norms. The humanities have a distinguished history of systematically exploring these tacit norms, one that dates back to long before the Renaissance. Indeed, such explorations

[9] See Sprouse 2020; Ludlow 2013, eps. chapters 1, 3–4.

helped to define humanistic inquiry itself. The humanities, I will argue, produce disciplinary knowledge of human experience.

Humanistic knowledge is deceptively difficult to obtain. The deceptive part lies in the fact that, unlike the exotic formulae and methods of the natural sciences, the humanities overwhelmingly involve reading, writing, listening, and looking. What could be simpler? One needs no special training to look or to listen. Reading and writing are acquired often well before the age of five. What these utterly quotidian practices disguise are the significant differences between a concerto buff and someone with a trained ear, or between a Renaissance art buff and an art historian. These people hear and see things that I do not hear, because for the expert, looking and listening are more than just using your eyes and your ears. They involve cultivated capacities to notice and examine specific features from among the multitude. Of course, the existence of such capacities is the most familiar thing in the world to the practicing scientist; they are quite literally indispensable. They are what make him an expert. They are what enable him to separate the telling phenomena from the surrounding cacophony of nature. Distracted by the incomprehensibility of equations or the high-octane antics of Professor McInquiry, we forget how much of scientific expertise comes from reading and writing, looking, and listening. Or, for chemists before the twentieth century, tasting and smelling (yes, many of them had severe brain damage). We do not consider that dimension of his expertise to be part of scientific knowledge. That needs to change.

I develop a model of disciplinary knowledge over the next four chapters. Chapter 2 concentrates on clearing away a lot of epistemological brush so that we can get a better, more realistic picture of scientific knowledge, one which foregrounds its fundamentally social and disciplinary nature. In Chapter 3, I peel away that surface layer to reveal the role of consensus in the development of disciplinary knowledge. Chapter 4 explains the way in which exemplars factor into the production of consensus. And Chapter 5 describes the cross-disciplinary practice of studying norms by attempting to violate them. This model is as at home in the natural sciences as it is in the humanities, which is just as it should be. For, the model rests on features of inquiry that are common to every group of practitioners that forms a genuine research community. Not every group does, just as the group of people who recognize the superiority of *Office Space* do not thereby constitute a culture (a cult, maybe). These special groups – the disciplines – share more than just a passion. They share a whole world of norms that are specific to their narrow research focus. The sum of these norms is the corpus of disciplinary knowledge. Chapter 6 moves away from concerns specific to practitioners of the humanities to look at how disciplinary knowledge can be extended beyond the context of research to touch the lives of everyone.

Chapter 7 marks an inflection point in our study of the humanities, where we begin the painful process of looking at some of the humanities' current weaknesses and trying to understand what's gone wrong. The model of disciplinary knowledge comes in handy here as well, showing how some of the contemporary threats to humanistic knowledge described in Chapters 8 and 9 can be understood as deviations from the model.