

SCIENCE, TRUTH, AND
DEMOCRACY



Philip Kitcher

OXFORD
UNIVERSITY PRESS

FIVE



Mapping Reality

SOME SCIENTISTS DO NOT SPEND THEIR LIVES formulating equations or Seven inscribing sentences. The products of their labor are pictures, or diagrams—or maps. In both geology and genetics, map-making is a respected activity, but I want to clarify the picture of the sciences I have been developing by looking at the core field, the academically rather unfashionable discipline of cartography.¹

The history of map-making illustrates the modest realism with which I began. Consider some of the maps of our planet offered by the geographers of the past, maps of the entire globe. Later maps appear superior to earlier ones in two major respects. First, they include entities that were previously omitted, the New World and Australasia being the most striking examples. Second, their depictions of the spatial relations among the entities commonly represented are more accurate; the margins of the various countries follow actual coastlines more closely. We make these judgments without believing that *any* of the maps ever produced is completely accurate, even while admitting the possibility that earlier maps might occasionally deliver a more accurate representation of some local features, and that the kind of convergence we appreciate visually need not be monotonic.

Past explorers might have had reason to think their maps were accurate because they were able to use them to navigate successfully (although there were, of course, any number of failures). The maps that superseded their charts have

1. Philosophers have sometimes turned to maps for comparisons with science. One particularly fruitful exploration, unjustly neglected in subsequent discussions, is Stephen Toulmin's treatment of the topic in his *Philosophy of Science: An Introduction* (New York: Harper and Row, 1953).

enabled their descendants to achieve success more systematically across a far wider range of voyages. In retrospect, we can also see why the earlier expeditions prospered to the extent that they did. Some features of the older maps are preserved in their later counterparts because, even though the old cartographers were not entirely right (and sometimes, of course, they were wildly wrong), they managed to achieve approximate accuracy about particular aspects of the globe. When they planned journeys that depended just on these features, they were able more frequently to reach their destinations.

Cartography displays in a particular instance just the type of progress and just the relation between success and accuracy that the modest realism of chapter 2 claimed for the sciences generally. Yet, when we think more carefully about the variety of maps—even the variety of maps of the earth's surface—we recognize complications. Map-makers are invariably selective. They introduce conventional elements, and, in consequence, standards of accuracy vary. The most obvious convention in the case of global maps concerns the way in which the three-dimensional relations are projected on a two-dimensional surface. It would be foolish to protest that a map of the globe using the Mercator projection is inaccurate because it makes the area of Greenland appear much greater than it is (relative, say, to South America)—just as foolish as denouncing a map for its uniformly pink coloration of the British Commonwealth. Associated with any map there are conventions that determine which aspects of the visual image are to be taken seriously.

Expanding our purview to embrace a variety of maps, we understand how maps designed for different purposes pick out different entities within a region or depict those entities rather differently. One map of a California resort region may display major roads, sports facilities, restaurants and services surrounded by a sparsely marked expanse of green, grey, and blue; another, designed for the serious backpacker, may show the roads only as conduits to the wilderness, while lavishing detail on the courses of streams, the sinuosity of the tree-line, the contour lines, and the trails. The shape of the same stretch of coast may be drawn differently—because of decisions about the coastal margin—in maps intended for the casual yachtsman, the holiday-maker in search of secluded bathing spots, the marine biologist, the geologist concerned with the fault structure, and the urban planner. What counts as an omission or an inaccurate spatial representation depends on the conventions associated with the kinds of maps, and, in their turn, those conventions are in place because of the needs of the potential users.

It might appear there is a limit to the variety of maps, some minimal set of conventions common to all, and some common standard of accuracy stemming from these conventions that every map is required to meet. So, for example, one might propose that if a map represents as collinear geographical features that do not lie on a line, then the map is *ipso facto* inaccurate. The proposal is mistaken.

Consider the map that figured in the scenarios of the second chapter. If practical success in navigating is to serve as our test of accuracy, then the map of the London Underground must count as accurate—for it figures in the successful activity of tens of thousands of people each day. But the map lies, in portraying as collinear places that do not fall on a straight line through space: the stations on the Central Line share a common Line but no common line.

These elementary reflections can help us to a more precise account of the conventionality and the accuracy of maps. Think of a map as a visual display coupled to a set of conventions. The set of conventions divides into two parts, the *intended content* and the *reading conventions*. The intended content of the map consists of the region and the types of entities and properties that the map intends to portray. The reading conventions link items in the visual display to those entities and also specify which features of the display do not correspond to any aspect of nature. In some instances, they will divide lines in the visual display into meaningful units and specify how these units correspond to parts of nature. Other conventions stipulate which spatial relations in the visual display correlate with spatial relations in the world. Yet others say that an aspect of the visual display is to be ignored. (Thus, old-fashioned maps of the globe, on which the British Commonwealth is colored pink, do not carry the information that Canada is uniformly pink.)

Consider again the map of the London Underground. The domain of the map (the region identified in the intended content) is London, and the objects of interest include the various Tube stations and the railway lines connecting them. The intended content identifies properties that are of interest, such as the relational property of two different lines being connected by a walkway (more colloquially, the map picks out stations where people can make connections). The spatial relations of special concern include being connected by the same railway line and being adjacent along the same rail line—but not being collinear in physical space nor any metrical relations among points depicted.

The reading conventions for this map connect dots and lines with entities in London: the dot marked “Clapham Common” is linked to a station in South London, the horizontal red line stands for the Central line, and so forth. Those conventions also tell the competent reader which parts of the display shouldn’t be taken seriously, not just the coloration but also the ordinary conventions about points of the compass (the map doesn’t inform us that Notting Hill Gate is exactly west of Oxford Circus).

I belabor the relatively obvious in order to show how to make sense of the notion of accuracy. Suppose that, in the visual display, two elements, identified by the reading conventions as bearers of meaning, stand in a particular spatial relation. According to the conventions, there are items in the domain of the map that correspond to those elements and there is a real-world counterpart of the spatial relation. The map is accurate, in this respect, if the two items from

the domain stand in the counterpart of the displayed relation. For the London Underground, part of the map's accuracy consists in its depiction of Victoria and Green Park as connected, without intermediates, by a blue line, and, in light of the reading conventions, this signifies that Victoria and Green Park are adjacent stations on the Victoria Line—as indeed they are. The map is also accurate in showing a large black circle at Notting Hill Gate, signifying that it is possible to change there between the Central Line and the Circle Line. So it goes. On the account I have offered, the map of the Underground is not *approximately* accurate. It is exact.

This may be unusual in that this particular map is thoroughly discrete and finite. More typical maps have reading conventions identifying a large infinite number of signifying elements. Consider, for example, an outline of Manhattan island, accompanied just by a pointer indicating North. Any connected segment of the line can be regarded as a meaningful element, and, given the reading conventions, we can appraise the accuracy of that element in terms of the conformity between the shape and orientation of the line and the shape and orientation of the pertinent portion of the island's margin (which is, of course, itself conventionally fixed). In this way, the map is equivalent to a truly enormous number of claims about spatial relations (continuum many): a picture is not worth a thousand words, but rather a staggering infinity of sentences. Further, although the map says many things that are incorrect, it also expresses an infinite number of true statements, for there are infinitely many truths of the form "A is within φ of being θ from due North of B," where A, B are places on the Manhattan shoreline and θ , φ are angular measurements. Like most scientific theories, the map, taken as a whole, is false, even though it contains large amounts of truth.

There is no problem, then, in talking of the accuracy of maps and hailing some maps as accurate in some (even all) respects. Our realist commitments are, however, perfectly compatible with recognizing the fact that human interests change and, in consequence, maps are drawn with very different reading conventions. We now have little use for the Tudor maps that showed sheepherding trails and the boundaries of manors. Nor do we expect that our maps of the globe will display the significance of Christ's passion. The reading conventions for many older maps are very different from those of the present, but the change should not surprise us. Reading conventions identify the ways of dividing the spatial domain that are of interest to the map-maker, and those conventions depend on the goals and the institutions of the society in which the map is to be used.

There is no unique correct way for a map of the globe, or of some smaller region, to draw boundaries. Sometimes part of the purpose is to recognize political divisions, to demarcate nations, states, counties, administrative districts

from one another. Sometimes we are interested in streets, roads, and highways. On other occasions the significant boundaries show the movements of migratory animals, the zones of common climate, the distribution of plant species, the topography of a mountain range, the extent of public and private lands, the successive inundations of a flood, the distribution of areas with specified population densities, incidence of disease, or availability of minerals. Which features are crucial as landmarks in drawing these alternative divisions varies from case to case: the botanist studying the distribution of arctic flora may have no use for any background markers at all, the hiker's needs may be met by a topographical map that indicates no more than access roads.

I argued in the last chapter for the analogous thesis about the sciences generally. Our ways of dividing up the world into things and kinds of things depend on our capacities and interests. The history of map-making extends the point by showing how cartographical conventions and divisions evolve in response to changing human purposes. Maps lose their place in human lives as the projects they once served are superseded, and those maps then retain value only for the historian. The intricate relations among Tudor manors and the minutiae of eighteenth-century waterways have no purchase on contemporary actions, and we can easily envisage that our descendants may find maps that reflect our administrative distinctions—and even the much-used map of the Underground—equally beside the point.

The map-makers' task is to produce maps that are pertinent to the enterprises and interests of their societies. By the same token, I suggest, the aim of the sciences is to address the issues that are significant for people at a particular stage in the evolution of human culture. Languages are fashioned to draw those distinctions that are most helpful in carrying out the lines of investigation those people want to pursue. As the history of cartography reveals a succession of maps with very different reading conventions, so too the history of the sciences generally should disclose a succession of languages framed, often imperfectly, to the pursuit of inquiries that appear, at the time, most important.

At this point we confront again an issue initially joined in the last chapter. Surely there is a grand scientific project constitutive of inquiry at any time, in any place, independent of culture, social institutions, or mutable human concerns? Whatever language, or compendium of languages, is apt for this large purpose will mark out privileged divisions in nature. It will identify the *real* natural kinds, the *genuine* objects and properties. Scientific inquiry aims to learn this language and to enunciate in it the basic truths about nature.

Although this conception of an overarching aim for inquiry has been influential in most discussions of the sciences, I am skeptical.² My skepticism surfaced in the last chapter, and it will be articulated further in the next. For the

2. I was not always so. See *The Advancement of Science*, especially chap. 4.

moment, however, I want to motivate it by drawing on the special case of maps.

Imagine a philosopher of cartography devoted to the idea that map-making has context-independent goals. A natural way to present the philosopher's thesis would be to suppose that the goal of cartography is the construction of an *ideal atlas*. The maps actually produced in human history are a selection of sheets from this atlas (to the extent, that is, that they are accurate). Of course the selections actually made are informed by interests that vary from group to group, epoch to epoch, but behind the contingent choices stands the inclusive ideal. Individual maps are significant because they belong to a hypothetical compendium, towards which we aim but which we shall never achieve.

What exactly would this ideal atlas be? It's implausible to suppose that it contains a single map that reveals all the spatial relations on our planet, for, as Lewis Carroll pointed out long ago (in *Sylvie and Bruno*), nothing could substitute for that except the terrain itself. Rather, the view must be that there are certain fundamental kinds of maps from which all spatial information can be generated, and that they collectively provide a unified presentation of the wide diversity of kinds of knowledge drawn from our actual ventures in cartography (and, presumably, projects we might have undertaken). But now a new question arises. What kinds of maps furnish information to advance any conceivable human project? Simply surveying the vast diversity of maps produced in actual human history—from which I have drawn a tiny sample—exposes the difficulty of reducing the atlas to any manageable compendium. Moreover, for any atlas we can envisage, we can easily conceive of projects that would require maps of different types—for example the project of recognizing the distribution of copies of the atlas itself.

There is no good reason to believe in the ideal atlas. A much more straightforward approach to the variety of maps, prefigured in my earlier expression of skepticism, is to relativize the notion of cartographical significance to communities, seeing some kinds of decisions about what to represent and how to represent it as the results of central aspects of those communities' ways of life. We would abandon the idea that cartography is governed by a context-independent goal. Perhaps we should lose similar baggage in thinking about the sciences generally.

The analogy between cartography and science invites a further step. Current ventures in map-making often carry the traces of past endeavors. To understand why present maps take the forms they do, we need to recognize the ways in which past projects of map-making have led to modifications of the part of nature that is mapped. Cartography generates a counterpart of the thesis that classifications may play a causal role in the reshaping of reality.

Consider a straightforward example in which map-making has contributed

to the alteration of the physical environment and the development of various pieces of technology. Backpacking Californians use topographical maps to explore the wilderness of the High Sierras. Older maps (together with guidebooks, lightweight camping equipment, and so forth) made it possible for more people to experience the beauty and solitude of the mountains. In consequence, certain lakeshores became degraded from over-camping, the foraging habits of animals (most obviously bears) were disturbed, and hikers seeking solitude were pressed to explore higher altitudes. Forest rangers have responded by marking more clearly (and sometimes widening) trails once viewed as "cross-country routes" for mountain climbers. At the same time, as backpackers ascend above the tree line, pursued by bears who are ever more adept at liberating their food, new kinds of backpacking technology have been introduced, such as the "bear boxes" installed at some wilderness sites, or the plastic canisters that hikers can carry to protect food. The maps of today show more detail for the more remote elevations than did their ancestors, as well as recording the changes caused by human activity, and we can envisage that the maps of the future may display information about food storage that would have been irrelevant a generation ago. This simple example shows quite clearly how the full story of why one set of conventions is chosen must include the past choices of map-makers and the projects their maps made possible, for those maps and projects influence the desires of later map-users, the resources available to them, and even the character of the terrain that they will explore.

So it is, I suggest, with the sciences generally. Like maps, scientific theories and hypotheses must be true or accurate (or, at least, approximately true or roughly accurate) to be good. But there is more to goodness in both instances. Beyond the necessary condition is a requirement of significance that cannot be understood in terms of some projected ideal—completed science, a Theory of Everything, or an ideal atlas. Recognizing that the ideal atlas is a myth, I hope to have provoked concerns about the analogue for inquiry generally. A rival vision proposes that what counts as significant science must be understood in the context of a particular group with particular practical interests and with a particular history. It further suggests that just as maps can play a causal role in reshaping the terrain that later cartographers will depict, so too the world to which scientists of one epoch respond may be partially produced by the scientific endeavors of the past—not in any strange metaphysical sense but in the most mundane ways.

I've offered a motivational analogy, not an argument. Fans of the traditional idea that there is a context-independent aim of inquiry could accept everything I have said about maps, their accuracy, their conventionality, and the sources of their significance and also argue that map-making is intimately bound up with practical projects, so that there are no implications for the significance of pieces of theoretical science. To claim that maps are invariably drawn for specific prac-

tical ends would be overstatement—historians often employ maps to advance our understanding of the past—but I shall not try to move beyond my motivational exercise to extend the argument in this direction. The analogy helps to frame the main issue that will occupy us: If science is indeed different, what is the genuine counterpart of the admittedly fictitious ideal atlas? Can we provide an account of the goals of inquiry, a specification of what constitutes significant science, that will apply across all historical contexts, independently of the evolving interests of human beings? Let us see.



Scientific Significance

THE SCIENCE STORY OF 1997 centered on a Scottish sheep. In an article widely discussed in scientific journals and in the broader press, Ian Wilmut, a researcher at an agricultural station near Edinburgh, reported the birth of Dolly as a result of nuclear transplantation. Wilmut had extracted the nucleus from an egg taken from one ewe, replaced it with the nucleus from an udder cell of another ewe, implanted the resultant egg, and allowed gestation to proceed. His report was accepted as correct at the time, and, despite challenges, it has been upheld since.¹ But why all the fuss? Why is it a significant piece of science to show that a mammal can be born in this way?

An obvious part of the answer recognizes the practical importance of Wilmut's work. This doesn't lie in the absurd fantasies about cloning people that have been widely touted in newspapers and popular magazines. Consonant with the character of his position, the primary practical import of Wilmut's achievement lies in its opening up the possibility of breeding domestic animals with desired characteristics (resistance to common diseases, preferred musculature, and so forth). At the junction of animal husbandry and medicine, researchers also envisage the possibility of modifying the genome of a future nuclear donor, inserting alleles to direct the production of a useful drug, and obtaining a flock of clones whose milk would be laced with this substance. Farther from such everyday benefits is the potential use of nuclear transplantation to generate genetically identical mammals (probably mice) with some particu-

1. The challenges focused on the possibility that Dolly did not share the nuclear genetic material from the adult female Wilmut assigned as her "nuclear mother." Any such doubts have now been resolved by DNA sequencing. See D. Ashworth et al., "DNA Microsatellite Analysis of Dolly," *Nature*, 394, 1998, 329.

larly interesting genotype, one associated, say, with the analogue of a recalcitrant human disease, so that the effects of this genotype can be systematically explored.

This last implication bestows on his work an indirect import for inquiries in theoretical biology. But there's a more straightforward connection, one noted in the forthright appraisal of the commentary that accompanied Wilmut's original article: "The results are of profound significance."² The judgment is supported in the commentary's opening sentences: "A hoary old question that has interested developmental biologists for years has been the continuity of the genome during animal development. Put another way—do growth, differentiation and development of the embryo involve irreversible modifications to the genome in somatic cells?" Since the 1950s, developmental biologists have understood that virtually all the cells of an organism contain the same complement of nuclear genes. At the same time, it's clear that cells differentiate into types with very different properties—muscle cells, blood cells of various types, liver cells, neurons, and so forth. Given our understanding of interactions between nuclear genes and other molecules in the cell (particularly proteins), we infer that different genes must be switched on and off in the different cell types. What we know of the mechanisms of gene regulation has shown that inactivation of genes sometimes occurs through the binding of molecules that block the "reading" enzymes from the pertinent parts of the DNA. So the question arises: Is there some modification of the DNA (perhaps an indissoluble coating with proteins) that permanently prevents differentiated cells from expressing some genes?

The previous failure of attempts to transplant nuclei from adult vertebrate cells suggested that the answer to this question is "Yes." Wilmut's approach to the problem focused on the different stages of the cell cycle. He starved the cells from which the nuclei were taken, forcing them into the rest phase. His success in nuclear transplantation can be understood in one of two ways: either as the result of synchrony between the cell cycle of the donor from which the nucleus is drawn and the recipient, the enucleated oocyte into which it is inserted; or because the biochemistry of the rest phase makes the DNA of the transplanted nucleus susceptible to reprogramming (perhaps by weakening the bonds between DNA and the protein "blockers"). Whichever of these explanations is correct, the answer to the old question is now more complex. Differentiated cells do modify the genome in ways that make some genes inaccessible at some stages of the cell cycle, but at the rest phase they aren't so modified. This answer is a contribution to the project of understanding the dynamics of differentiation, and

2. The original article is I. Wilmut et al., "Viable Offspring Derived from Fetal and Adult Mammalian Cells," *Nature*, 385, 27 Feb. 1997, 810–813. The commentary is Colin Stewart, "An Udder Way of Making Lambs," *Nature*, 385, 27 Feb. 1997, 769–771.

thus of development as a whole, and it also points towards further projects (What is the exact character of the modification, and how is it undone?).

This example, or any of a thousand like it, can help us see the shortcomings of traditional ideas about the aims of the sciences. Nobody should be beguiled by the idea that the aim of inquiry is merely to discover truth, for, as numerous philosophers have recognized, there are vast numbers of true statements it would be utterly pointless to ascertain. The sciences are surely directed at finding significant truths. But what exactly are these?

One possible answer makes significance explicitly relative—the significant truths for a person are just those the knowledge of which would increase the chance she would attain her practical goals. Or you could try to avoid relativization by focusing on truths that would be pertinent to anyone's projects—the significant truths are those the knowledge of which would increase anyone's chance of attaining practical goals.

Neither of these is at all plausible as a full account of scientific significance, and the deficiency isn't just a result of the fact that both are obviously rough and preliminary. Linking significance to practical projects ignores areas of inquiry in which the results have little bearing on everyday concerns, fields like cosmology and paleontology. Moreover, even truths that do facilitate practical projects often derive significance from a different quarter. Surely the principles of thermodynamics would be worth knowing whether or not they helped us to build pumps and engines (and thereby attain further goals). Besides the notion of practical significance, captured perhaps in a preliminary way by the rough definitions given above, we need a conception of "theoretical" or "epistemic" significance that will mark out those truths the knowledge of which is intrinsically valuable.

Prominent efforts to understand the notion of epistemic significance, embodied in the writings of philosophers during the last three centuries and in the rhetoric of public paeans to scientific inquiry, attempt to show that inquiry is directed towards discovering a particular kind of truth, a kind scientists seek at all times, whatever practical projects they (or their contemporaries) may favor. The disciplines we pick out as sciences count as part of *science* because they aim at, and sometimes deliver, truths of this special kind, and they can be distinguished from technology precisely because the latter is focused on the practical. An allegedly context-independent notion of epistemic significance insulates science, or "basic science" at least, from social and moral values, by claiming that the achievement of epistemically significant truth is valuable in principle—even though, in actuality, that value might be compromised by ways in which the recognition of some truths would generate unfortunate consequences. Because I believe no such conception can be found, I take moral and social values to be intrinsic to the practice of the sciences.

My argument for this view doesn't depend on abandoning the idea that the

sciences yield truth about nature or on giving up the ideal of objectivity; chapters 2 and 3 distinguish my position from one of the unacceptable images with which we began. Instead, I think there's a serious problem with traditional ideas about scientific significance, more specifically about epistemic significance. A gallery of bad pictures has held us captive. Dolly, the Scottish sheep, can help us recognize what is wrong with the traditional views and can point us in a better direction.

Traditional approaches suppose the notion of epistemic significance has nothing to do with us and our ephemeral practical concerns, and everything to do with the structure of the world. There are various ways to try to articulate the point. It is easy to start with muddy metaphors: some questions are "on Nature's agenda"; inquiry aims to discover "how Nature works." Personification is, however, hardly pellucid.

From the early modern period to the present, scientists and philosophers have tried to do better by invoking one of a family of interlinked concepts. So there arise the following well-known proposals:

The (epistemic) aim of science is to achieve objective understanding through the provision of explanations.

The (epistemic) aim of science is to identify the laws of nature.

The (epistemic) aim of science is to arrive at a unified picture of nature.

The (epistemic) aim of science is to discover the fundamental causal processes at work in nature.

Many thinkers have accepted more than one of these theses because they have recognized conceptual connections among the crucial terms ("explanation consists in subsuming the phenomena under laws," "explanation consists in identifying the fundamental causal processes," "laws of nature are those generalizations that figure in a unified account of the phenomena," "the fundamental causes are those described by the most general principles of a unified account of nature," and so forth). As we shall see later, it is not entirely clear that any of these grand conceptions will enable us to understand the hoopla about Dolly, but, for the moment, let us suspend worries that the particularities of her birth fail to live up to the large advertisements.

For our purposes, the important issue is not one that has figured in the large majority of worries about how we might come to apply the difficult concepts of law, cause, and explanation, but the question of which, if any, of the formulations of the last paragraph might identify an epistemic aim of science whose value could be convincingly defended. With respect to any of these projected achievements, it's appropriate for us to inquire, "What would be so valuable about gaining that?" There are immediate difficulties with the last two formu-

lations. A unified picture of the world isn't something that wears its worth on its face—the question of why unity is so wonderful remains open. The commendation of causal knowledge does a bit better. Such knowledge plainly facilitates intervention in the world. Practical concerns are, however, not pertinent when we're out to fathom epistemic significance, and, when we bar them, there is again an open question about why knowledge of fundamental causal processes should be valuable.

Turn next to the suggestion that science is the search for natural laws, a proposal underlying many influential discussions of the sciences in the last three centuries. For the present, let's grant that a satisfactory account of natural law can be given, one that will distinguish genuine laws from accidental regularities. We can still ask why it's valuable to identify true statements with these special features.

Some major figures from the history of modern science would have answered this question by supposing that talk of laws is more than a bad pun: laws of nature are prescribed to the Creation by the ultimate sovereign, the Creator, and the world must conform to them. To seek the laws of nature is thus to reveal the divine rulebook, and to rejoice in the wisdom and beneficence of God. Copernicus, Kepler, Descartes, Boyle, and Newton all sounded the theme, and Newton's theological justification of his physics in a letter to Richard Bentley is typical: "When I wrote my treatise about our system, I had an eye upon such principles as might work with considering men, for the belief of a Deity; and nothing can rejoice me more than to find it useful for that purpose."³ Similar ideas of a divine lawmaker whose statutes, once revealed, will inspire our admiration, resound throughout the eighteenth century and into the nineteenth (especially in Britain, where they are prominent in the Bridgewater treatises).

So here's *an* answer to the question I have posed. Knowledge of God ought to be our highest concern; disclosure of God's laws will promote this knowledge, thereby enabling us to "think God's thoughts after him"; what can be a more worthy goal than that? I doubt, however, that this will seem particularly persuasive. In light of our increasing knowledge of the history of the cosmos and of life on our planet, even committed theists are unlikely to feel that divine wisdom and beneficence are manifest in the creation: it all seems a curiously roundabout, baroque, inelegant, wasteful, and savage way of doing things.

In the twentieth century (and even in the nineteenth) the dominant articulation of the view that science aims to disclose the laws of nature has been thoroughly secular. Of course, even if there is no Creator and no divine rulebook, the universe might still be organized *as if* there had been a Creator with a rule-

3. Newton, *Opera*, vol. 4, 429. Quoted in E. A. Burtt, *Metaphysical Foundations of Modern Physical Science* (London: Routledge and Kegan Paul, 1924), 285. Burtt's work also contains excellent examples of similar ideas in the works of Copernicus, Kepler, Descartes, and Boyle.

book, but the secular surrogate loses the immediacy of the explanation of epistemic value. Recognizing the rules of organization might assist us with practical projects, but these concerns are irrelevant when we are trying to fathom epistemic significance.

The best way to develop the traditional approach is to appeal, at this point, to the idea of objective understanding and its correlate, objective explanation. Some truths are significant because they enable us to explain nature. Now in one very obvious sense, explanation is an activity provoked by actual or anticipated questions that arise in particular contexts, and explanations are directed at satisfying an envisaged audience that poses these questions. No defender of objective explanation should question this elementary point. Rather, what is claimed is that objective understanding consists in recognizing special relationships among the facts or events, so the criterion of success for an explanatory episode is the generation of this recognition, not any subjective satisfaction that the person given the explanation may feel. The sciences supply an *explanatory store* from which information or arguments can be drawn and adapted in the particular contexts in which understanding is sought and explanations given, and this store contains fully specific delineations of the relationships on which understanding depends. To put the point in its simplest terms, there is something which science supplies that provides an all-purpose basis for the practice of giving explanations: whatever interests people may have, whatever feelings of satisfaction or puzzlement they harbor, there is a set of relationships that, ideally, science presents to us, and the presentation brings a distinctive epistemic benefit.

One thought about those relationships is that they are revealed by showing how individual occurrences and states of affairs fall under general principles—explanation is a matter of subsumption under laws. Another thought is that they are recognized by seeing how the diverse phenomena of the natural world are integrated within a unified picture. Yet a third is that the relationships are causal, and that we appreciate them when we can identify the fundamental causes at work. But an important part of the view must be that the store is somehow systematic. It will fail as an all-purpose explanatory device if it is simply a long list of potential explanations, one for each context in which the desire for understanding might arise. Were that to be so, there would be no basis for a distinction between the epistemically significant and the epistemically insignificant—for every truth about the world would surely figure somewhere on the list, in the quite particular explanation that accounted for it, and, almost certainly, in giving explanations of particularities that flow from it.

The traditional search for a context-independent conception of epistemic significance is thus committed to the idea of a systematic organization of the truths about nature from which objective explanations may be drawn. I now

want to scrutinize this commitment, and it will be useful to begin with a view that presented it most forthrightly.

The Unity-of-Science movement drew inspiration from examples in which particular scientific achievements were exhibited as derivative from others: Galileo's law of free fall and Kepler's laws of planetary motion could be viewed as consequences of Newton's gravitational theory; the laws of ideal gases could be incorporated first into the kinetic theory of heat and subsequently into statistical mechanics; thanks to the recognition that chemical bonds involve transfer or sharing of electrons, ascriptions of valence properties and, in consequence, laws of chemical combination, could be derived ("in principle") from atomic physics, ultimately perhaps from the basic equations of quantum mechanics. Extrapolating from these instances, it was proposed that all laws of chemistry could be derived from principles of physics, that all laws of biology could be derived from principles of physics and chemistry, that all laws of psychology could be derived from principles of biology (most notably neurobiology), chemistry, and physics, and so forth. Proponents of the unity of science understood quite clearly that the different sciences used special vocabularies, so that the envisaged derivations would depend on coordinative definitions that would link these vocabularies, and they pointed to the kinetic-theoretic identification of temperature as mean molecular kinetic energy as an example of the kinds of definitions that would ultimately be provided.

Attractive as the view may be, it has suffered from scrutiny of crucial junctions, most particularly those between the physical sciences and biology and between biology and psychology. Major difficulties have emerged. First, the successes achieved in the motivating examples seem to depend on the fact that the theories *reduced* (those exhibited as consequences of more fundamental parts of science) were individual laws or small collections of laws. Nobody has even the faintest idea what it would be to present biology or psychology as a cluster of laws. There are serious doubts concerning whether these sciences contain *any* genuine laws, and it is uncontroversial that there are highly significant parts of them that are not simply collections of laws: Dolly points to no general law of ovine (let alone mammalian) development.

Second, both biology and psychology seem to employ concepts that are not definable in the terms of the sciences proposed as reducing them. Defenders of the autonomy of psychology have pointed out how unlikely it is that there is a single characterization in terms of physics-plus-chemistry-plus-biology of the psychological state of thinking about the Unity-of-Science view (say), for the neural realizations and the underlying physicochemical conditions are very likely to vary from person to person. The situation is even clearer in the case of genetics. Nobody currently knows how to achieve a specification of the concept

of gene in physicochemical terms: more pedantically, it is hard to see how to complete the schematic sentence, "x is a gene if and only if x is . . ." by filling the blank with a structural description that will enable us to apply laws of physics and chemistry to derive conclusions about the behavior of genes. To be sure, an important necessary condition on genes is that they be segments of DNA or RNA; but of course there are lots of segments of DNA and RNA (most of them, in fact) that are not genes. The task is thus to identify the property that distinguishes the right segments of nucleic acid from the wrong ones. Contemporary efforts to discern the genes in reams of sequence data would be greatly aided if some such structural description were available, but, as molecular geneticists know all too well, the best they can do is to look for "Open Reading Frames"—stretches of DNA that show a relatively long interval between a codon for the initiation of transcription and a stop codon. (Computer searches pick out the ORFs, and investigators then use functional criteria to decide if new ORFs are genuine genes.) In addition, it's sometimes plausible to count regulatory "regions" as regulatory genes, to consider other nontranscribed regions as genes that have lost crucial parts of the regulatory apparatus, to see disconnected regions as parts of the "same gene," to identify overlapping genes or even genes within genes. Molecular genetics tells us at a staggering rate about the chemical structures of *individual* genes, but fails to provide a *general* specification. Matters are even worse when we move away from genetics and consider such important biological notions as *cell*, *organism*, *Drosophila melanogaster*, *species*, and *predator*.

The actual deliverances of the sciences accord rather badly with the Unity-of-Science view. But I now want to press a deeper point, scrutinizing the commitment to the provision of understanding through incorporation within a single overarching framework. Consider a fundamental Mendelian regularity which I'll formulate roughly as follows: genes sufficiently far apart on the same chromosome or on different chromosomes assort independently. How do we explain why this regularity obtains? The unity of science view sees us as gaining "objective understanding" from a physico-chemical specification of the notions of gene and chromosome and a derivation employing laws of physics and chemistry to yield the result about independent assortment. But this is quite unpersuasive. Rather we understand the regularity—as objectively as you please—by recognizing that the transmission of genes to gametes is a process of pairing and separation. Homologous chromosomes are brought together at meiosis, they may exchange some genetic material (hence the qualification about being "sufficiently far apart"—segments too close have a higher probability of being transmitted together), and one member of each pair is then passed on to a gamete. We assimilate the transmission of genes to all kinds of other processes which involve bringing together pairs of similar things and selecting one from each pair. What's crucial is the form of these processes, not the

material out of which the things are made. The regularity about genes would hold so long as they could sustain processes of this form, and, if that condition were met, it wouldn't matter if genes were segments of nucleic acids, proteins, or chunks of Swiss cheese.

To reinforce the point, consider the regularity discovered by Dr. John Arbuthnot in the early eighteenth century. Scrutinizing the record of births in London during the previous 82 years, Arbuthnot found that in each year a preponderance of the children born had been boys; in his terms, each year was a "male year." Why does this regularity hold? Proponents of the Unity-of-Science view can offer a recipe for the explanation, although they can't give the details. Start with the first year (1623); elaborate the physicochemical details of the first copulation-followed-by-pregnancy showing how it resulted in a child of a particular sex; continue in the same fashion for each pertinent pregnancy; add up the totals for male births and female births and compute the difference. It has now been shown why the first year was "male"; continue for all subsequent years.

Even if we had this "explanation" to hand, and could assimilate the details, it would still not advance our understanding. For it would not show that Arbuthnot's regularity was anything more than a gigantic coincidence. By contrast, we can already give a satisfying explanation by appealing to an insight of R. A. Fisher. In considering sex ratios from an evolutionary point of view, Fisher recognized that, in a population in which sex ratios depart from 1:1 at sexual maturity, there will be a selective advantage to a tendency to produce the underrepresented sex. It is easy to show from this that there should be a stable evolutionary equilibrium at which the sex ratio at sexual maturity is 1:1. In any species in which one sex is more vulnerable to early mortality than the other, this equilibrium will correspond to a state in which the sex ratio at birth is skewed in favor of the more vulnerable sex. Applying the analysis to our own species, in which boys are more likely than girls to die before reaching puberty, we find that the birth sex ratio ought to be 1.04:1 in favor of males—which is what Arbuthnot and his successors have observed. We now understand why, for a large population, all years are overwhelmingly likely to be male.

I've been opposing two commitments of the Unity-of-Science view, the claim that the sciences can be hierarchically unified, and the view that integration within a single unified framework is the essence of objective understanding. It would be natural to respond to the arguments by proposing that, while they may doom a particular way of articulating the traditional view that epistemic significance attaches to those truths that can figure in an explanatory system, that is not the approach that the tradition should have preferred. But this is too sanguine. Invoking an ideal of objective understanding, based on a single unified framework of laws, was no arbitrary extension of the basic commitment, but an elaboration of the idea that science supplies a structure that is a

Red
do
ab
ex

resource for "objective understanding," whatever our contingent interests. Appealing to the Unity of Science specifies the character of this systematic, all-purpose structure. If the Unity-of-Science view fails, we need a substitute.

That, you might suppose, is easy enough. An obvious suggestion is that the discovery of laws (or the identification of causal processes) really does advance our understanding, although not in the way that the Unity-of-Science view suggests. Instead of a single system within which all "objective" explanations are subsumed, we proceed piecemeal, gaining understanding of nature by recognizing the laws of nature that govern the phenomena or the causal processes at work in the phenomena. This seems a salutary development, but it invites an obvious question. What is meant by "the phenomena"? Either the traditionalist intends, in accordance with the original motivation of science's agenda as set by nature, to think of providing resources for explaining *all* phenomena, or what is at issue is the explanation of the phenomena *that we find in need of explanation*, a vision that brings our contingent and evolving interests into the picture. What must be shown then, is how to reconcile the idea of *some sort of system* of laws or causes with the considerations that doomed the Unity-of-Science view.

Waiving concerns about the omnipresence of laws (that will occupy us shortly), the critique of the ideal of unified science displayed areas of inquiry that classified overlapping parts of nature in distinctive ways and offered their own (locally) unifying frameworks in terms of these classificatory schemes. Genetics, for example, focuses on DNA molecules but groups them in ways that do not map neatly onto physicochemical classifications and approaches the transmission of genes in ways that connect with (for example) processes of pairing and separation. So any complete system of laws of nature will consist of a patchwork of *locally unified pieces*, sciences with their own schemes of classification, their own favored causal processes, and their own systematic ways of treating a cluster of phenomena.⁴ When we think about scientific inquiry as responding to a relatively narrow range of explanatory projects, to wit the kinds of questions we find worth posing, there's little harm in conceiving of this type of patchwork. But when we drop the reference to ourselves and our concerns, I see no reason to think there's any manageable system at all. To put the worry bluntly, why should we suppose that the number of classificatory schemes and unified treatments for all nature's phenomena is *finite*? The Unity-of-Science view had a simple answer to that question, since it proposed that the classificatory schemes of all the sciences were, ultimately, one, but once we've admitted plurality there's no reason for thinking we can stop. To revert to the motivating analogy of the

4. Nancy Cartwright thinks of the sciences as offering us a patchwork of laws. I agree with the point about the patchwork but believe that she places too much emphasis on the notion of natural law—if only in reacting against it. See her book *The Dappled World* (Cambridge: Cambridge University Press, 1998).

last chapter, the Unity-of-Science view made it look as though there was a fundamental set of maps from which any map we might care to use could be constructed, and so gave content to the conception of the ideal atlas. Once we abandon that view, it looks as though all that may remain is a collection of charts that may proliferate indefinitely with our changing interests.

At this point, a defender of context-independent goals for inquiry can reply that there is an as yet unformulated notion of objective understanding that will serve. We may not see yet how to divide the class of true statements about the world into those that are epistemically significant and those that aren't, but this is an important research project for philosophy of science. Let me explain why I am skeptical.

Those who seek context-independent goals for inquiry should admit that the explanation-seeking questions people pose take many different forms: we ask, "How?," "What?," "How is it possible?," and, of course, "Why?" The search for a philosophical notion of objective explanation has focused on Why-questions, conceding, tacitly or explicitly, that the topics of such questions reflect changeable human interests. But it has presupposed that there's a certain kind of information or argument that ought to be supplied in response to any explanation-seeking Why-question. The idea, then, has been that if we identify the explanatory store with the collection of all complete answers to Why-questions whose topics are true, there will be some propositions that pervade these answers, and these propositions are the epistemically significant truths. More exactly, given any Why-question whose topic proposition is true, there will be a special relation—the *relevance* relation—that holds between the topic and the objective complete answer. This relevance relation is independent of time and context, and whatever topics interest people, inspiring them to ask "Why?," the objective answers (the things that bear the relevance relation to the topic propositions) will always contain members of a set of true statements, the epistemically significant truths. So, for example, when explanation is taken to consist in showing how phenomena fall under laws of nature, any objective answer will contain some statement of law, and the laws will be selected as the epistemically significant truths.⁵

Plainly it's a disaster for this approach if all sorts of humdrum truths turn out to be epistemically significant. Now for typical mundane truths, statements about the contents of my cluttered desktop for example, there are everyday explanations in which those truths figure; my failure to find some pieces of paper

5. For a careful and precise account of the ways in which explanations work in everyday contexts, see Bas van Fraassen, *The Scientific Image* (Oxford: Oxford University Press), chap. 5; the proposal to see explanation as lawlike subsumption is most extensively developed by C. G. Hempel in the title essay of *Aspects of Scientific Explanation* (New York: Free Press, 1965).

when I'm looking elsewhere is explained by noting that the things I sought are buried in a pile on my desk. So it looks as though any truth, however banal, will occur somewhere in the explanatory store, unless we are offered a filter that lets just the "pervasive" truths enter the class of the epistemically significant.

We can now see why the approach in terms of unified science was so promising, for it imagined that ideally complete explanations use the same fundamental principles again and again. Without recourse to the Unity-of-Science view, we have to look for some context-independent relevance relation that will generate the right filter. One natural suggestion is that explanations are given by furnishing causal information that bears on the topic. Now in everyday explanations the kinds of causal factors that people provide are heavily context-dependent: the lawyer, highway engineer, automobile mechanic, and psychologist may offer quite different accounts of why the Princess of Wales had a fatal accident. Perhaps, however, we should view the diverse appeals to causal factors as context-dependent selections from the *complete* cause, supposing that there's a context-independent relation between the topic (specifying the circumstances of Princess Diana's death) and a complete description of its causal antecedents.

But we can completely specify the causal factors that produced an effect at any given time prior to the effect, so that focusing on a particular time already involves a further selection. To avoid context-dependence, one must invoke the idea of a complete causal history, an imagined account that shows how the effect described by the topic occurred as a consequence of events in the remote past. The view, then, must be that the objective answers describe some vast causal history, and that these serve as a store for ordinary explanations by permitting selections that are attuned to the interests of the intended audiences.

This idea is vulnerable to two difficulties. First, recall one of the considerations that doomed the Unity-of-Science view. Neither the sequence of "male" years in London nor the independent assortment of genes is understood by grinding out the full causal details: a narrative drawn from the deep recesses of the past would fail to offer the type of information sought. We might honor the idea that here too explanation consists in the provision of causal information, but only by recognizing that it is a different type of causal information, one not captured in the allegedly ideal and complete causal history. We understand sex-ratios by seeing the state we wish to understand as an equilibrium and identifying the factors that maintain the equilibrium—in a sense a causal account, but one that doesn't relate effects to completely specified antecedent causes.

The second problem arises when we consider the goal at which the account of explanation aims. We worried earlier that all sorts of mundane truths would figure somewhere in the explanatory store. But with respect to virtually any truth, however humdrum, we can devise a complete causal narrative in which

that truth plays an essential role. So how is the filtering to be done? Perhaps by supposing that the epistemically significant truths occur in *every* complete causal narrative. Given the difficulties of the Unity-of-Science view, it looks as though ideal explanations will describe causal processes that occur at different levels, but perhaps there are *some* truths common to all members of the explanatory store, for example descriptions of very early stages of the universe. To broaden the epistemically significant beyond early cosmology and particle physics, one might suggest that epistemically significant truths are those that occur in a very large number of the complete causal narratives. How do you count? Given the continuity of time, it looks as though any statement that occurs in a complete causal narrative figures in an infinite number of such narratives (indeed, continuum many), for it will be an essential part of the "objective explanations" of all those statements that describe "downstream" states and events. In terms of numerical frequency, all truths are on a par.

The enterprise thus strikes me as hopeless. For there's a general problem. Everyday explanations seem quite varied in their offerings of causal information (and maybe in other types of information as well). To pick out a context-independent relevance relation that covers all this diversity requires one to portray individual acts of explanation as selecting from much vaster entities, ideal explanations that have many constituent propositions. Just about any truth will turn up in the resultant store. Because of the deficiencies of the Unity-of-Science view, the statements one would like to pick out as epistemically significant will not be all-pervasive. So there will be no simple solution to the filtering problem. The only option seems to be to resort to counting, and this fails because the classes to be compared are all infinite.

We can free ourselves from this bind by developing a different approach to "objective explanation." Given a topic that is of interest to us and a relevance relation, the objective explanation is whatever complex of truths stands in the appropriate relation to the topic. Just as the topics of interest to us evolve in the growth of the sciences, so too with the relevance relations. Perhaps many of these are broadly causal (although we do recognize other types in seeking mathematical explanations, for example). Even the most cursory survey of our practices reveals the heterogeneity of the relevance relations that pertain to our questions: sometimes we are interested in triggering events, sometimes with enduring features that are taken to constitute the "natures" of the things under study, sometimes with the intentions of agents, sometimes with conditions that maintain an equilibrium, sometimes with factors that are to the advantage of an organism. Frequently, relevance relations reflect our interest in the covariation of properties we find salient or in factors that we can manipulate and control. Objective explanation goes on in the sciences, then, but only against the background of our questions and our interests. The most we can expect from a theory of explanation is some understanding of how these questions and interests

shift as our inquiries, and the complex environments in which they occur, evolve.⁶

I now want to approach the issue in a different fashion, considering the ways in which judgments of significance are made in the everyday practice of the sciences. Whether one turns to the specialized journals of particular subdisciplines (*Physical Review*, *Cell*) or to the general journals in which publication is most difficult (*Science*, *Nature*), it's overwhelmingly obvious that new laws are very hard to find. Prominent articles tell us about the distribution of minerals in particular parts of the earth's crust, about the relative sizes of australopithecine skulls, about the sequences of the genomes of bacteria, worms, and flies, and, of course, about that celebrated sheep. Why should any of these studies be hailed as significant?

In all such instances, we can tell the same sort of story I summarized in the case of Dolly. There are broad questions we find interesting—What were our hominid ancestors like? How do single-celled organisms regulate their metabolism?—and we can see the findings as advancing the project of answering them. Often there are practical problems—of understanding earthquake zones or combating Lyme disease—on which the research bears. Indeed, in many cases (though not in all) epistemic and practical interests are interwoven.

Defenders of science as the search for laws and objective explanations have an obvious strategy for responding to these examples. The goals of a vast and ambitious enterprise are not necessarily revealed in everyday activities, and a myopic focus on the brushwork will not reveal the splendor of the picture. So, they might contend, the very particular investigations I report are significant because they are small contributions to attaining the types of epistemic significance that the tradition celebrates: the laws and the descriptions of fundamental causes will emerge from them, perhaps in one of those rare articles published every few decades in which some fortunate scientist, standing on the shoulders of a pyramid of under-laborers, displays what the entire venture has been aiming towards all along.

There's a valuable point here, in the recognition that significance accrues to work that would strike many outsiders as arcane, because of the advancement of a much larger project. But I think the response errs in misunderstanding the

6. In an earlier essay, "Van Fraassen on Explanation," *Journal of Philosophy*, 84, 1987, 315–330, Wesley Salmon and I framed issues about explanation in terms of a choice between giving an "objective" account of explanation and a pluralism in which "anything goes." We overlooked an important possibility, that variable interests might promote different topics and different relevance relations. Not every relation counts as a relevance relation, but one can think of a family of relevance relations, bound together by loose resemblances, indefinitely extendable and coevolving with the history of inquiries and the social ventures they serve. This may have been van Fraassen's own position.

interrelations among pieces of scientific work, the channels along which scientific significance flows. It thinks in terms of a Baconian hierarchy: the contents of *Nature* and *Science* are pieces of information that will be systematized into general laws, and, ultimately, into overarching theories—significance runs (drips?) from the envisaged theoretical top to the mundane accomplishments at the bottom. The demise of the Unity-of-Science view ought to make us suspicious about parts of this, but the examples I have cited, with my brief explanations of how they integrate into larger projects, should inspire a quite different concern: is there any reason to think that significance flows from the general (or the “causally fundamental”) to the particular, rather than having its source in very specific concerns about particular types of properties of entities that matter to us (the crust of our planet, our ancestors, bacteria that are human pathogens)? Rather, the connections that confer significance seem to radiate in many different directions, so that a map of an area of inquiry that reveals how its claims and projects earn their significance might look more like a tangled skein than a hierarchy. An elaboration of this view will show, I believe, how it does much greater justice to the way in which scientific significance attaches to the work scientists actually do.

Back to Dolly. In all epochs, and in all cultures, people have been struck by the fundamental phenomenon of development, the unfolding of characteristics from initially tiny fragments of organic material, with preservation of species-typical traits. This phenomenon must have been evident to those who first domesticated plants and animals (the collection and planting of seeds bears witness to an appreciation of it), and some aspects of it were apparent even in the case of our own species. Conceiving of reproduction as a process in which something is transmitted from parents to offspring, investigators since antiquity have struggled to learn what this something is and how it interacts with the rest of nature (particularly the bits of the world with which it comes into contact) to produce a new organism. The first question—“What is the hereditary material?”—obviously fuels the large questions of genetics, from Mendel’s observations to the present. The second bifurcates into two kinds of issue: the description of the processes that lead from the first stages of the nascent organism to its adult form, and the fathoming of causal processes of particular interest to us. As we learn that the first stage of the new organism is a zygote (a fertilized egg) and that development proceeds via cell division, we pick out some causal processes for attention. Thus we ask, “How are genes activated and inactivated?” (because we think of the hereditary material as playing a role in guiding the organism towards a species-typical phenotype), “How do the major cellular movements that lay down the organism’s body plan take place?” (because we see, for example, that all vertebrates have a common form we can trace to changes that take place in the formation of the notochord), “How do cells of different types dif-

ferentiate?" (because we recognize physiological, and subsequently biochemical, differences among the cells from different bodily systems), and so forth. The last of these questions forms the backdrop to Wilmut's work, in just the ways I indicated earlier.

So we have found our way from a natural preoccupation of virtually all people to the birth of a lamb in a Scottish research station. We could trace similar paths from the initiating question for developmental biology to many other inquiries—efforts to map and sequence the genome of the soil amoeba *Dictyostelium discoideum*, attempts to breed mutant zebrafish of particular types, computer programs that try to simulate the growth of a chick forelimb, biochemical assays of tissues in *Drosophila*, mathematical analyses of snail shell patterns. Instead of automatically assuming that these efforts, to which highly intelligent and extensively trained people devote large portions of their lives, are directed at some enterprise of great generality, we can actually explore why they have come to be at the focus of inquiry, recognizing the affiliations to practical projects as well as to large questions that naturally excite human curiosity.

As with the example of map-making considered in the last chapter, where we saw that the kinds of maps we construct are shaped by evolving interests, so too the questions we take to be significant and the endeavors we pursue in attempts to answer them co-evolve with all sorts of practical projects. Fields of science are associated with structures I shall call significance graphs that embody the ways in which their constituent research projects obtain significance. A significance graph is constructed by drawing a directed graph with arrows linking expressions, some of which formulate questions that workers in the field address, others that encapsulate claims they make, yet others that refer to pieces of equipment, techniques, or parts of the natural world (figures 1, 2). The significance graph reveals how to explain the significance of various items—where 'item' is an all-purpose term for questions, answers, hypotheses, apparatus, methods, and so forth. One would account for the significance of the item to which the arrow points in terms of the significance of the item from which it comes. Arrows thus display the inheritance of scientific significance.

In talking of the "explanation of significance," I intend to make explicit what workers in the field know at the time to which the net is indexed. The commentary on Dolly told the broader public what researchers typically take for granted. As a field grows, however, the character of the significance graph changes so that later explanations of significance are quite different from those that would have been given earlier. Moreover, the significance of an item may well be overdetermined, and researchers with different interests may give priority to alternative linkages. Some prize Dolly because of the possibilities she represents for improving livestock, others because she contributes to our understanding of cellular differentiation. We can adopt a *field-centered* perspective on significance graphs, one that shows how significance is inherited within a par-

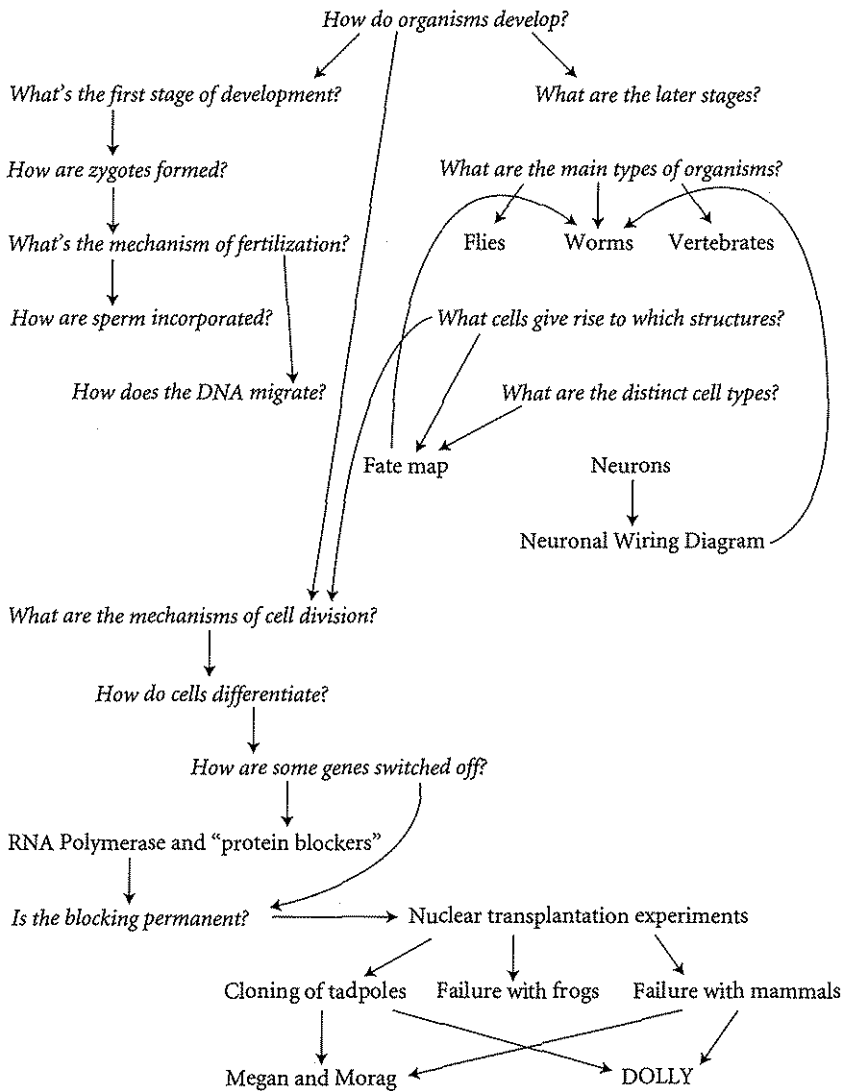


Figure 1. An extremely partial depiction of the significance graph for developmental biology.

ticular area of research, viewing Dolly solely within the purview of developmental biology (say). Or we can take an *item-centered* perspective, looking at all the ways in which a particular node in the significance graph, the one designating Dolly for example, gains significance for science. The perspectives are compatible and valuable for different purposes:

One principal difference between thinking in terms of significance graphs and the more traditional conception of the sciences as seeking laws is a far more

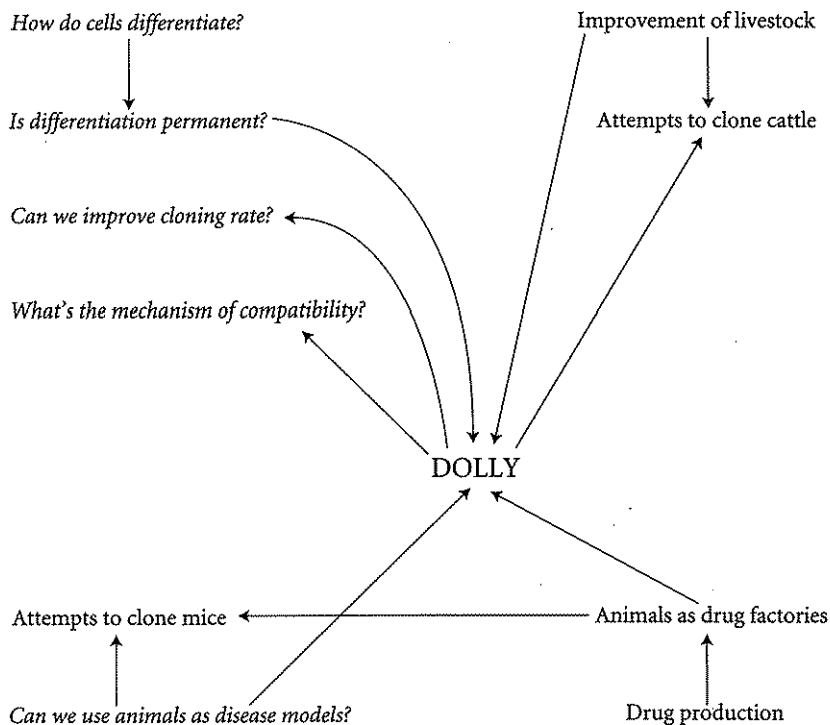


Figure 2. A Dolly-centered view of inquiry.

pragmatic approach to generality. Large generalizations are good where we can find them, for they enable us to fit phenomena into a broader framework, to explain more, to predict more, and maybe to intervene more successfully. There may be some areas of scientific inquiry where we can achieve precise, accurate, generalizations of large scope—the parts of physics that have inspired many philosophers may be like this—but the most common aspect of *la condition scientifique* seems to be that we have to make compromises among generality, precision, and accuracy.⁷ Hence significance graphs do not embody the idea that significance (or epistemic significance) is always a matter of achieving, or pointing to, universal laws.

What, then, is the ultimate source of epistemic significance? The answer is, I think, commonplace and disappointing to those who expect a grand theory that will invest the sciences with overriding importance. Recall our explanation of Dolly's significance. It began from the idea that we wanted to understand

7. The point was formulated beautifully by Richard Levins in *Evolution in Changing Environments* (Princeton: Princeton University Press, 1966). For a deep skepticism about the possibility of finding precise generalizations even in physics, see Cartwright, *The Dappled World*.

how an organism's characteristics unfold from a tiny piece of organic material. If someone asked why we want to understand that—or why we want to know why the heavenly bodies move as they do, or why we are interested in the evolution of our hominid ancestors—it would be hard to say very much. We expect other people to see the point of such questions, and we describe those who don't as lacking in "natural curiosity." Partly as the result of our having the capacities we do, partly because of the cultures in which we develop, some aspects of nature strike us as particularly salient or surprising. In consequence we pose broad questions, and epistemic significance flows into the sciences from these. ✕

Human beings vary, of course, with respect to the ways in which they express surprise and curiosity. Some are disposed to ask more, others less. Typically, we respond to the diversity with tolerance, explaining some of the variation in terms of differences in cultural or educational context. But tolerance has its limits, and we do count some of our fellows as pathological, either because they obsess about trifles or because they are completely dull. In claiming that the sciences ultimately obtain their epistemic significance from the broad questions that express natural human curiosity, I am drawing on this practice of limited tolerance, on our conception of "healthy curiosity" and the commonplace thought that most of us, given minimal explanation, would find interesting the global questions that stand at the peripheries of significance graphs.

Significance graphs evolve. As information accumulates, new connections are forged, new practical projects are designed, new questions emerge as tractable ways of pursuing inquiries already established. Precisely because the same entities sometimes serve as resources both for important practical ventures and for theoretical work, epistemic significance today may bear the traces of yesterday's practical significance. Because some instruments, techniques, sites, or model organisms become embedded in the significance graphs of different fields, so that researchers know how to use them, the evolution often shows a kind of inertia. Alternative choices made earlier would have led to a different development of the field, so that, in quite particular ways, the development of the sciences is thoroughly contingent.⁸

As our inquiries evolve and different phenomena become salient for us, we introduce new classifications, dividing up the world in novel ways. Sometimes we arrive at new views about which entities are single objects (think of the recognition that lichens consist of symbiotic pairs of organisms); more frequently, we group old entities together in different ways. Our understanding of objects and kinds of things evolves with our significance graphs, driven by changing expressions of natural curiosity and by our practical needs. Further-

8. For articulation of the point in the context of experimental physics, see Peter Galison, *Image and Logic* (Chicago: University of Chicago Press, 1997). The dominance of certain organisms—for example, *Drosophila*—in contemporary biology is extremely obvious.

more, the contingent decisions made today modify the pressures under which the graphs evolve, for the phenomena into which we inquire sometimes change in response to our activities. One obvious example comes from twentieth-century biomedical research, in which one generation's struggle against what it views as the most devastating diseases can open up niches for disease vectors, thereby allowing the evolution of new pathogens. If pessimistic forecasts about infectious diseases are even partly justified, the medical problems our descendants are likely to confront will be caused by microorganisms whose existence depends on our ancestors' choices.

I have tried to outline a view of scientific significance that is very different, both in its character and its consequences, from that which has dominated traditional reflections on inquiry. If I am right, then the analogy of the last chapter has been vindicated. Like maps, scientific theories—or, better, significance graphs—reflect the concerns of the age. There is no ideal atlas, no compendium of laws or “objective explanations” at which inquiry aims. Further, the challenges of the present, theoretical and practical, and even the world to be mapped or understood, are shaped by the decisions made in the past. The trail of history lies over all.