8

Experimentation and Scientific Realism

Ian Hacking

Experimental physics provides the strongest evidence for scientific realism. Entities that in principle cannot be observed are regularly manipulated to produce new phenomena and to investigate other aspects of nature. They are tools, instruments not for thinking but for doing.

The philosopher's standard "theoretical entity" is the electron. I will illustrate how electrons have become experimental entities, or experimenter's entities. In the early stages of our discovery of an entity, we may test hypotheses about it. Then it is merely a hypothetical entity. Much later, if we come to understand some of its causal powers and use it to build devices that achieve well-understood effects in other parts of nature, then it assumes quite a different status.

Discussions about scientific realism or antirealism usually talk about theories, explanation, and prediction. Debates at that level are necessarily inconclusive. Only at the level of experimental practice is scientific realism unavoidable—but this realism is not about theories and truth. The experimentalist need only be a realist about the entities used as tools.

A PLEA FOR EXPERIMENTS

No field in the philosophy of science is more systematically neglected than experiment. Our grade school teachers may have told us that scientific method is experimental method, but histories of science have become histories of theory. Experiments, the philosophers say, are of value only when they test theory. Experimental work, they imply, has no life of its own. So we lack even a terminology to describe the many varied roles of experiment. Nor has this one-sidedness done theory any good, for radically different types of theory are used to think about the same physical phenomenon (e.g., the magneto-optical effect). The philosophers of theory have not noticed this and so misreport even theoretical enquiry.

Different sciences at different times exhibit different relationships between "theory" and "experiment." One chief role of experiment is the creation of phenomena. Experimenters bring into being phenomena that do not naturally exist in a pure state. These phenomena are the touchstones of physics, the keys to nature, and the source of much modern technology. Many are what physicists after the 1870s began to call "effects": the photoelectric effect, the Compton effect, and so forth.[1] A recent high-energy extension of the creation of phenomena is the creation of "events," to use the jargon of the trade. Most of the phenomena, effects, and events created by the experimenter are like plutonium: they do not exist in nature except possibly on vanishingly rare occasions.[2]

In this paper I leave aside questions of methodology, history, taxonomy, and the purpose of experiment in natural science. I turn to the purely philosophical issue of scientific realism. Simply call it "realism" for short. There are two basic kinds: realism about entities and realism about theories. There is no agreement on the precise definition of either. Realism about theories says that we try to form true theories about the world, about the inner constitution of matter and about the outer reaches of space. This realism gets its bite from optimism: we think we can do well in this project and have already had partial success. Realism about entities—and I include processes, states, waves, currents, interactions, fields, black holes, and the like among entities—asserts the existence of at least some of the entities that are the stock in trade of physics.[3]

The two realisms may seem identical. If you believe a theory, do you not believe in the existence of the entities it speaks about? If you believe in some entities, must you not describe them in some theoretical way that you accept? This seeming identity is illusory. *The vast majority of experimental physicists are realists about entities but not about theories.* Some are, no doubt, realists about theories too, but that is less central to their concerns.

Experimenters are often realists about the entities that they investigate, but they do not have to be so. R. A. Millikan probably had few qualms about the reality of electrons when he set out to measure their charge. But he could have been skeptical about what he would find until he found it. He could even have remained skeptical. Perhaps there is a least unit of electric charge, but there is no particle or object with exactly that unit of charge. Experimenting on an entity does not commit you to believing that it exists. Only manipulating an entity, in order to experiment on something else, need do that.

Moreover, it is not even that you use electrons to experiment on something else that makes it impossible to doubt electrons. Understanding some causal properties of electrons, you guess how to build a very ingenious, complex device that enables you to line up the electrons the way you want, in order to see what will happen to something else. Once you have the right experimental idea, you know in advance roughly how to try to build the device, because you know that this is the way to get the electrons to behave in such and such a way. Electrons are no longer ways of organizing our thoughts or saving the phenomena that have been observed. They are now ways of creating phenomena in some other domain of nature. Electrons are tools.

There is an important experimental contrast between realism about entities and realism about theories. Suppose we say that the latter is belief that science aims at true theories. Few experimenters will deny that. Only philosophers doubt it. Aiming at the truth is, however, something about the indefinite future. Aiming a beam of electrons is using present electrons. Aiming a finely tuned laser at a particular atom in order to knock off a certain electron to produce an ion is aiming at present electrons. There is, in contrast, no present set of theories that one has to believe in. If realism about theories is a doctrine about the aims of science, it is a doctrine laden with certain kinds of values. If realism about entities is a matter of aiming electrons next week or aiming at other electrons the week after, it is a doctrine much more neutral between values. The way in which experimenters are scientific realists about entities is entirely different from ways in which they might be realists about theories.

This shows up when we turn from ideal theories to present ones. Various properties are confidently ascribed to electrons, but most of the confident properties are expressed in numerous different theories or models about which an experimenter can be rather agnostic. Even people in a team, who work on different parts of the same large experiment, may hold different and mutually incompatible accounts of electrons. That is because different parts of the experiment will make different uses of electrons. Models good for calculations on one aspect of electrons will be poor for others. Occasionally, a team actually has to select a member with a quite different theoretical perspective simply to get someone who can solve those experimental problems. You may choose someone with a foreign training, and whose talk is well-nigh incommensurable with yours, just to get people who can produce the effects you want.

But might there not be a common core of theory, the intersection of everybody in the group, which is the theory of the electron to which all the experimenters are realistically committed? I would say common lore, *not* common core. There are a lot of theories, models, approximations, pictures, formalisms, methods, and so forth involving electrons, but there is no reason to suppose that the intersection of these is a theory at all. Nor is there any reason to think that there is such a thing as "the most powerful nontrivial *theory* contained in the intersection of all the theories in which this or that member of a team has been trained to believe." Even if there are a lot of shared beliefs, there is no reason to suppose they form anything worth calling a theory. Naturally, teams tend to be formed from like-minded people at the same institute, so there is usually some real shared theoretical basis to their work. That is a sociological fact, not a foundation for scientific realism.

I recognize that many a scientific realism concerning theories is a doctrine not about the present but about what we might achieve, or possibly an ideal at which we aim. So to say that there is no present theory does not count against the optimistic aim. The point is that such scientific realism about theories has to adopt the Peircean principles of faith, hope, and charity. Scientific realism about entities needs no such virtues. It arises from what we can do at present. To understand this, we must look in some detail at what it is like to build a device that makes the electrons sit up and behave.

OUR DEBT TO HILARY PUTNAM

It was once the accepted wisdom that a word such as 'electron' gets its meaning from its place in a network of sentences that state theoretical laws. Hence arose the infamous problems of incommensurability and theory change. For if a theory is modified, how could a word such as 'electron' go on meaning the same? How could different theories about electrons be compared, since the very word 'electron' would differ in meaning from theory to theory?

Putnam saved us from such questions by inventing a referential model of meaning. He says that meaning is a vector, refreshingly like a dictionary entry. First comes the syntactic marker (part of speech); next the semantic marker (general category of thing signified by the word); then the stereotype (clichés about the natural kind, standard examples of its use, and present-day associations. The stereotype is subject to change as opinions about the kind are modified). Finally, there is the actual referent of the word, the very stuff, or thing, it denotes if it denotes anything. (Evidently dictionaries cannot include this in their entry, but pictorial dictionaries do their best by inserting illustrations whenever possible.)[4]

Putnam thought we can often guess at entities that we do not literally point to. Our initial guesses may be jejune or inept, and not every naming of an invisible thing or stuff pans out. But when it does, and we frame better and better ideas, then Putnam says that, although the stereotype changes, we refer to the same kind of thing or stuff all along. We and Dalton alike spoke about the same stuff when we spoke of (inorganic) acids. J. J. Thomson, H. A. Lorentz, Bohr, and Millikan were, with their different theories and observations, speculating about the same kind of thing, the electron.

There is plenty of unimportant vagueness about when an entity has been successfully "dubbed," as Putnam puts it. 'Electron' is the name suggested by G. Johnstone Stoney in 1891 as the name for a natural unit of electricity. He had drawn attention to this unit in 1874. The name was then applied to the subatomic particles of negative charge, which J. J. Thomson, in 1897, showed cathodes rays consist of. Was Johnstone Stoney referring to the electron? Putnam's account does not require an unequivocal answer. Standard physics books say that Thomson discovered the electron. For once I might back theory and say that Lorentz beat him to it. Thomson called his electrons 'corpuscles', the subatomic particles of electric charge. Evidently, the name does not matter much. Thomson's most notable achievement was to measure the mass of the electron. He did this by a rough (though quite good) guess at *e*, and by making an excellent determination of *e/m*, showing that *m* is about 1/1800 the mass of the hydrogen atom. Hence it is natural to say that Lorentz merely postulated the existence of a particle of negative charge, while Thomson, determining its mass, showed that there is some such real stuff beaming off a hot cathode.

The stereotype of the electron has regularly changed, and we have at least two largely incompatible stereotypes, the electron as cloud and the electron as particle. One fundamental enrichment of the idea came in the 1920s. Electrons, it was found, have angular momentum, or "spin." Experimental work by O. Stern and W. Gerlach first indicated this, and then S. Goudsmit and G. E. Uhlenbeck provided the theoretical understanding of it in 1925. Whatever we think, Johnstone Stoney, Lorentz, Bohr, Thomson, and Goudsmit were all finding out more about the same kind of thing, the electron.

We need not accept the fine points of Putnam's account of reference in order to thank him for giving us a new way to talk about meaning. Serious discussion of inferred entities need no longer lock us into pseudo-problems of incommensurability and theory change. Twenty-five years ago the experimenter who believed that electrons exist, without giving much credence to any set of laws about electrons, would have been dismissed as philosophically incoherent. Now we realize it was the philosophy that was wrong, not the experimenter. My own relationship to Putnam's account of meaning is like the experimenter's relationship to a theory. I do not literally believe Putnam, but I am happy to employ his account as an alternative to the unpalatable account in fashion some time ago.

Putnam's philosophy is always in flux. His account of reference was intended to bolster scientific realism. But now, at the time of this writing (July 1981), he rejects any "metaphysical realism" but allows "internal realism."[5] The internal realist acts, in practical affairs, as if the entities occurring in his working theories did in fact exist. However, the direction of Putnam's metaphysical antirealism is no longer scientific. It is not peculiarly about natural science. It is about chairs and livers too. He thinks that the world does not naturally break up into our classifications. He calls himself a transcendental idealist. I call him a transcendental nominalist. I use the word 'nominalist' in the old-fashioned way, not meaning opposition to "abstract entities" like sets, but meaning the doctrine that there is no nonmental classification in nature that exists over and above our own human system of naming.

There might be two kinds of internal realist, the instrumentalist about science and the scientific realist. The former is, in practical affairs where he uses his present scheme of concepts, a realist about livers and chairs but thinks that electrons are only mental constructs. The latter thinks that livers, chairs, and electrons are probably all in the same boat, that is, real at least within the present system of classification. I take Putnam to be an internal scientific realist rather than an internal instrumentalist. The fact that either doctrine is compatible with transcendental nominalism and internal realism shows that our question of scientific realism is almost entirely independent of Putnam's internal realism.

INTERFERING

Francis Bacon, the first and almost last philosopher of experiments, knew it well: the experimenter sets out "to twist the lion's tail." Experimentation is interference in the course of nature; "nature under constraint and vexed; that is to say, when by art and the hand of man she is forced out of her natural state, and squeezed and moulded."[6] The experimenter is convinced of the reality of entities, some of whose causal properties are sufficiently well understood that they can be used to interfere *elsewhere* in nature. One is impressed by entities that one can use to test conjectures about other, more hypothetical entities. In my example, one is sure of the electrons that are used to investigate weak neutral currents and neutral bosons. This should not be news, for why else are we (nonskeptics) sure of the reality of even macroscopic objects, but because of what we do with them, what we do to them, and what they do to us?

Interference and interaction are the stuff of reality. This is true, for example, at the borderline of observability. Too often philosophers imagine that microscopes carry conviction because they help us see better. But that is only part of the story. On the contrary, what counts is what we can do to a specimen under a microscope, and what we can see ourselves doing. We stain the specimen, slice it, inject it, irradiate it, fix it. We examine it using different kinds of microscopes that employ optical systems that rely on almost totally unrelated facts about light. Microscopes carry conviction because of the great array of interactions and interferences that are possible. When we see something that turns out to be unstable under such play, we call it an artifact and say it is not real.[7]

Likewise, as we move down in scale to the truly unseeable, it is our power to use unobservable entities that makes us believe they are there. Yet, I blush over these words 'see' and 'observe'. Philosophers and physicists often use these words in different ways. Philosophers tend to treat opacity to visible light as the touchstone of reality, so that anything that cannot be touched or seen with the naked eye is called a theoretical or inferred entity. Physicists, in contrast, cheerfully talk of observing the very entities that philosophers say are not observable. For example, the fermions are those fundamental constituents of matter such as electron neutrinos and deuterons and, perhaps, the notorious quarks. All are standard philosophers' "unobservable" entities. C. Y. Prescott, the initiator of the experiment described below, said in a recent lecture, that "of these fermions, only the t quark is yet unseen. The failure to observe $t\bar{t}$ states in e^+e^- annihilation at PETRA remains a puzzle."[8] Thus, the physicist distinguishes among the philosophers' "unobservable" entities, noting which have been observed and which not. Dudley Shapere has just published a valuable study of this fact.[9] In his example, neutrinos are used to see the interior of a star. He has ample quotations such as "neutrinos present the only way of directly observing" the very hot core of a star.

John Dewey would have said that fascination with seeing-with-the-naked-eye is part of the spectator theory of knowledge that has bedeviled philosophy from earliest times. But I do not think Plato or Locke or anyone before the nineteenth century was as obsessed with the sheer opacity of objects as we have been since. My own obsession with a technology

that manipulates objects is, of course, a twentieth-century counterpart to positivism and phenomenology. Its proper rebuttal is not a restriction to a narrower domain of reality, namely, to what can be positivistically seen with the eye, but an extension to other modes by which people can extend their consciousness.

MAKING

Even if experimenters are realists about entities, it does not follow that they are right. Perhaps it is a matter of psychology: maybe the very skills that make for a great experimenter go with a certain cast of mind which objectifies whatever it thinks about. Yet this will not do. The experimenter cheerfully regards neutral bosons as merely hypothetical entities, while electrons are real. What is the difference?

There are an enormous number of ways in which to make instruments that rely on the causal properties of electrons in order to produce desired effects of unsurpassed precision. I shall illustrate this. The argument—it could be called the 'experimental argument for realism'—is not that we infer the reality of electrons from our success. We do not make the instruments and then infer the reality of the electrons, as when we test a hypothesis, and then believe it because it passed the test. That gets the time-order wrong. By now we design apparatus relying on a modest number of home truths about electrons, in order to produce some other phenomenon that we wish to investigate.

That may sound as if we believe in the electrons because we predict how our apparatus will behave. That too is misleading. We have a number of general ideas about how to prepare polarized electrons, say. We spend a lot of time building prototypes that do not work. We get rid of innumerable bugs. Often we have to give up and try another approach. Debugging is not a matter of theoretically explaining or predicting what is going wrong. It is partly a matter of getting rid of "noise" in the apparatus. "Noise" often means all the events that are not understood by any theory. The instrument must be able to isolate, physically, the properties of the entities that we wish to use, and damp down all the other effects that might get in our way. *We are completely convinced of the reality of electrons when we regularly set to build—and often enough succeed in building—new kinds of device that use various well understood causal properties of electrons to interfere in other more hypothetical parts of nature.*

It is not possible to grasp this without an example. Familiar historical examples have usually become encrusted by false theory-oriented philosophy or history, so I will take something new. This is a polarizing electron gun whose acronym is PEGGY II. In 1978, it was used in a fundamental experiment that attracted attention even in *The New York Times*. In the next section I describe the point of making PEGGY II. To do that, I have to tell some new physics. You may omit reading this and read only the engineering section that follows. Yet it must be of interest to know the rather easy-to-understand significance of the main experimental results, namely, that parity is not conserved in scattering of polarized electrons from deuterium, and that, more generally, parity is violated in weak neutral-current interactions.[10]

PARITY AND WEAK NEUTRAL CURRENTS

There are four fundamental forces in nature, not necessarily distinct. Gravity and electromagnetism are familiar. Then there are the strong and weak forces (the fulfillment of Newton's program, in the *Optics*, which taught that all nature would be understood by the interaction of particles with various forces that were effective in attraction or repulsion over various different distances, i.e., with different rates of extinction).

Strong forces are 100 times stronger than electromagnetism but act only over a miniscule distance, at most the diameter of a proton. Strong forces act on "hadrons," which include protons, neutrons, and more recent particles, but not electrons or any other members of the class of particles called "leptons."

The weak forces are only 1/10,000 times as strong as electromagnetism, and act over a distance 100 times greater than strong forces. But they act on both hadrons and leptons, including electrons. The most familiar example of a weak force may be radioactivity.

The theory that motivates such speculation is quantum electrodynamics. It is incredibly successful, yielding many predictions better than one part in a million, truly a miracle in experimental physics. It applies over distances ranging from diameters of the earth to 1/100 the diameter of the proton. This theory supposes that all the forces are "carried" by some sort of particle: photons do the job in electromagnetism. We hypothesize "gravitons" for gravity.

In the case of interactions involving weak forces, there are charged currents. We postulate that particles called "bosons" carry these weak forces.[11] For charged currents, the bosons may be either positive or negative. In the 1970s, there arose the possibility that there could be weak "neutral" currents in which no charge is carried or exchanged. By sheer analogy with the vindicated parts of quantum electrodynamics, neutral bosons were postulated as the carriers in weak neutral interactions.

The most famous discovery of recent high-energy physics is the failure of the conservation of parity. Contrary to the expectations of many physicists and philosophers, including Kant,[12] nature makes an absolute distinction between right-handedness and left-handedness. Apparently, this happens only in weak interactions.

What we mean by right- or left-handed in nature has an element of convention. I remarked that electrons have spin. Imagine your right hand wrapped around a spinning particle with the fingers pointing in the direction of spin. Then your thumb is said to point in the direction of the spin vector. If such particles are traveling in a beam, consider the relation between the spin vector and the beam. If all the particles have their spin vector in the same direction as the beam, they have right-handed (linear) polarization, while if the spin vector is opposite to the beam direction, they have left-handed (linear) polarization.

The original discovery of parity violation showed that one kind of product of a particle decay, a so-called muon neutrino, exists only in left-handed polarization and never in right-handed polarization.

Parity violations have been found for weak *charged* interactions. What about weak *neutral* currents? The remarkable Weinberg-Salam model for the four kinds of force was proposed independently by Stephen Weinberg in 1967 and A. Salam in 1968. It implies a minute violation of parity in weak neutral interactions. Given that the model is sheer speculation, its success has been amazing, even awe-inspiring. So it seemed worthwhile to try out the predicted failure of parity for weak neutral interactions. That would teach us more about those weak forces that act over so minute a distance.

The prediction is: slightly more left-handed polarized electrons hitting certain targets will scatter, than right-handed electrons. Slightly more! The difference in relative frequency of the two kinds of scattering is 1 part in 10,000, comparable to a difference in probability between 0.50005 and 0.49995. Suppose one used the standard equipment available at the Stanford Linear Accelerator Center in the early 1970s, generating 120 pulses per second, each pulse providing one electron event. Then you would have to run the entire SLAC beam for twenty-seven years in order to detect so small a difference in relative frequency. Considering that one uses the same beam for lots of experiments simultaneously, by letting different experiments use different pulses, and considering that no equipment remains stable for even a month, let alone twenty-seven years, such an experiment is impossible. You need enormously more electrons coming off in each pulse. We need between 1000 and 10,000 more electrons per pulse than was once possible. The first attempt used an instrument now called PEGGY I. It had, in essence, a high-class version of J. J. Thomson's hot cathode. Some lithium was heated and electrons were boiled off. PEGGY II uses quite different principles.

PEGGY II

The basic idea began when C. Y. Prescott noticed (by chance!) an article in an optics magazine about a crystalline substance called gallium arsenide. GaAs has a curious property; when it is struck by circularly polarized light of the right frequencies, it emits lots of linearly polarized electrons. There is a good, rough and ready quantum understanding of why this happens, and why half the emitted electrons will be polarized, three-fourths of these polarized in one directon and one-fourth polarized in the other.

PEGGY II uses this fact, plus the fact that GaAs emits lots of electrons owing to features of its crystal structure. Then comes some engineering—it takes work to liberate an electron from a surface. We know that painting a surface with the right stuff helps. In this case, a thin layer of cesium and oxygen is applied to the crystal. Moreover, the less air pressure around the crystal, the more electrons will escape for a given amount of work. So the bombardment takes place in a good vacuum at the temperature of liquid nitrogen.

We need the right source of light. A laser with bursts of red light (7100 Ångstroms) is trained on the crystal. The light first goes through an ordinary polarizer, a very old-fashioned prism of calcite, or Iceland spar[13]—this gives linearly polarized light. We want circularly polarized light to hit the crystal, so the polarized laser beam now goes through a cunning device called a Pockel's cell, which electrically turns linearly polarized photons into circularly polarized ones. Being electric, it acts as a very fast switch. The direction of circular polarization depends on the direction of current in the cell. Hence, the direction of polarization can be varied randomly. This is important, for we are trying to detect a minute asymmetry between right- and left-handed polarization. Randomizing helps us guard against any systematic "drift" in the equipment.[14] The randomization is generated by a radioactive decay device, and a computer records the direction of polarization for each pulse.

A circularly polarized pulse hits the GaAs crystal, resulting in a pulse of linearly polarized electrons. A beam of such pulses is maneuvered by magnets into the accelerator for the next bit of the experiment. It passes through a device that checks on a proportion of polarization along the way. The remainder of the experiment requires other devices and detectors of comparable ingenuity, but let us stop at PEGGY II.

BUGS

Short descriptions make it all sound too easy; therefore, let us pause to reflect on debugging. Many of the bugs are never understood. They are eliminated by trial and error. Let me illustrate three different kinds of bugs: (1) the essential technical limitations that, in the end, have to be factored into the analysis of error; (2) simpler mechanical defects you never think of until they are forced on you, and (3) hunches about what might go wrong.

Here are three examples of bugs:

1. Laser beams are not as constant as science fiction teaches, and there is always an irremediable amount of "jitter" in the beam over any stretch of time.

2. At a more humdrum level, the electrons from the GaAs crystal are back-scattered and go back along the same channel as the laser beam used to hit the crystal. Most of them are then deflected magnetically. But some get reflected from the laser apparatus and get back into the system. So you have to eliminate these new ambient electrons. This is done by crude mechanical means, making them focus just off the crystal and, thus, wander away.

3. Good experimenters guard against the absurd. Suppose that dust particles on an experimental surface lie down flat when a polarized pulse hits it, and then stand on their heads when hit by a pulse polarized in the opposite direction. Might that have a systematic effect, given that we are detecting a minute asymmetry? One of the team thought of this in the middle of the night and came down next morning frantically using anti-dust spray. They kept that up for a month, just in case.[15]

RESULTS

Some 10^{11} events were needed to obtain a result that could be recognized above systematic and statistical error. Although the idea of systematic error presents interesting conceptual problems, it seems to be unknown to philosophers. There were systematic uncertainties in the detection of right- and left-handed polarization, there was some jitter, and there were other problems about the parameters of the two kinds of beam. These errors were analyzed and linearly added to the statistical error. To a student of statistical inference, this is real seat-of-the-pants analysis with no rationale whatsoever. Be that as it may, thanks to PEGGY II the number of events was big enough to give a result that convinced the entire physics community.[16] Left-handed polarized electrons were scattered from deuterium slightly more frequently than right-handed electrons. This was the first convincing example of parity-violation in a weak neutral current interaction.

COMMENT

The making of PEGGY II was fairly nontheoretical. Nobody worked out in advance the polarizing properties of GaAs—that was found by a chance encounter with an unrelated experimental investigation. Although elementary quantum theory of crystals explains the polarization effect, it does not explain the properties of the actual crystal used. No one has got a real crystal to polarize more than 37 percent of the electrons, although in principle 50 percent should be polarized.

Likewise, although we have a general picture of why layers of cesium and oxygen will "produce negative electron affinity," that is, make it easier for electrons to escape, we have no quantitative understanding of why this increases efficiency to a score of 37 percent.

Nor was there any guarantee that the bits and pieces would fit together. To give an even more current illustration, future experimental work, briefly described later in this paper, makes us want even more electrons per pulse than PEGGY II can give. When the aforementioned parity experiment was reported in *The New York Times*, a group at Bell Laboratories read the newspaper and saw what was going on. They had been constructing a crystal lattice for totally unrelated purposes. It uses layers of GaAs and a related aluminum compound. The structure of this lattice leads one to expect that virtually all the electrons emitted would be polarized. As a consequence, we might be able to double the efficiency of PEGGY II. But, at present, that nice idea has problems. The new lattice should also be coated in work-reducing paint. The cesium-oxygen compound is applied at high temperature. Hence the aluminum tends to ooze into the neighboring layer of GaAs, and the pretty artificial lattice becomes a bit uneven, limiting its fine polarized-electron-emitting properties.[17] So perhaps this will never work. Prescott is simultaneously reviving a souped up new thermionic cathode to try to get more electrons. Theory would not have told us that PEGGY II would beat out thermionic

PEGGY I. Nor can it tell if some thermionic PEGGY III will beat out PEGGY II.

Note also that the Bell people did not need to know a lot of weak neutral current theory to send along their sample lattice. They just read *The New York Times.*

MORAL

Once upon a time, it made good sense to doubt that there were electrons. Even after Thomson had measured the mass of his corpuscles, and Millikan their charge, doubt could have made sense. We needed to be sure that Millikan was measuring the same entity as Thomson. Thus, more theoretical elaboration was needed, and the idea had to be fed into many other phenomena. Solid state physics, the atom, and superconductivity all had to play their part.

Once upon a time, the best reason for thinking that there are electrons might have been success in explanation. Lorentz explained the Faraday effect with his electron theory. But the ability to explain carries little warrant of truth. Even from the time of J. J. Thomson, it was the measurements that weighed in, more than the explanations. Explanations, however, did help. Some people might have had to believe in electrons because the postulation of their existence could explain a wide variety of phenomena. Luckily, we no longer have to pretend to infer from explanatory success (i.e., from what makes our minds feel good). Prescott and the team from the SLAC do not explain phenomena with electrons. They know how to use them. Nobody in his right mind thinks that electrons "really" are just little spinning orbs about which you could, with a small enough hand, wrap your fingers and find the direction of spin along your thumb. There is, instead, a family of causal properties in terms of which gifted experimenters describe and deploy electrons in order to investigate something else, for example, weak neutral currents and neutral bosons. We know an enormous amount about the behavior of electrons. It is equally important to know what does *not* matter to electrons. Thus, we know that bending a polarized electron beam in magnetic coils does not affect polarization in any significant way. We have hunches, too strong to ignore although too trivial to test independently: for example, dust might dance under changes of directions of polarization. Those hunches are based on a hard-won sense of the kinds of things electrons are. (It does not matter at all to this hunch whether electrons are clouds or waves or particles.)

WHEN HYPOTHETICAL ENTITIES BECOME REAL

Note the complete contrast between electrons and neutral bosons. Nobody can yet manipulate a bunch of neutral bosons, if there are any. Even weak neutral currents are only just emerging from the mists of hypothesis. By 1980, a sufficient range of convincing experiments had made them the object of investigation. When might they lose their hypothetical status and become commonplace reality like electrons?—when we use them to investigate something else.

I mentioned the desire to make a better electron gun than PEGGY II. Why? Because we now "know" that parity is violated in weak neutral interactions. Perhaps by an even more grotesque statistical analysis than that involved in the parity experiment, we can isolate just the weak interactions. For example, we have a lot of interactions, including electromagnetic ones, which we can censor in various ways. If we could also statistically pick out a class of weak interactions, as precisely those where parity is not conserved, then we would possibly be on the road to quite deep investigations of matter and antimatter. To do the statistics, however, one needs even more electrons per pulse than PEGGY II could hope to generate. If such a project were to succeed, we should then be beginning to use weak neutral currents as a manipulable tool for looking at something else. The next step toward a realism about such currents would have been made.

The message is general and could be extracted from almost any branch of physics. I mentioned earlier how Dudley Shapere has recently used "observation" of the sun's hot core to illustrate how physicists employ the concept of observation. They collect neutrinos from the sun in an enormous disused underground mine that has been filled with old cleaning fluid (i.e., carbon tetrachloride). We would know a lot about the inside of the sun if we knew how many solar neutrinos arrive on the earth. So these are captured in the cleaning fluid. A few neutrinos will form a new radioactive nucleus (the number that do this can be counted). Although, in this study, the extent of neutrino manipulation is much less than electron manipulation in the PEGGY II experiment, we are nevertheless plainly using neutrinos to investigate something else. Yet not many years ago, neutrinos were about as hypothetical as an entity could get. After 1946 it was realized that when mesons disintegrate giving off, among other things, highly energized electrons, one needed an extra nonionizing particle to conserve momentum and energy. At that time this postulated "neutrino" was thoroughly hypothetical, but now it is routinely used to examine other things.

CHANGING TIMES

Although realisms and antirealisms are part of the philosophy of science well back into Greek prehistory, our present versions mostly descend from debates at the end of the nineteenth century about atomism. Antirealism about atoms was partly a matter of physics; the energeticists thought energy was at the bottom of everything, not tiny bits of matter. It also was connected with the positivism of Comte, Mach, K. Pearson, and even J. S. Mill. Mill's young associate Alexander Bain states the point in a characteristic way, apt for 1870:

> Some hypotheses consist of assumptions as to the minute structure and operation of bodies. From the nature of the case these assumptions can never be proved by direct means. Their merit is their suitability to express phenomena. They are Representative Fictions.[18]

"All assertions as to the ultimate structure of the particles of matter," continues Bain, "are and ever must be hypothetical. . . . The kinetic theory of heat serves an important intellectual function." But we cannot hold it to be a true description of the world. It is a representative fiction.

Bain was surely right a century ago, when assumptions about the minute structure of matter could not be proved. The only proof could be indirect, namely, that hypotheses seemed to provide some explanation and helped make good predictions. Such inferences, however, need never produce conviction in the philosopher inclined to instrumentalism or some other brand of idealism.

Indeed, the situation is quite similar to seventeenth-century epistemology. At that time, knowledge was thought of as correct representation. But then one could never get outside the representations to be sure that they corresponded to the world. Every test of a representation is just another representation. "Nothing is so much like an idea as an idea," said Bishop Berkeley. To attempt to argue to scientific realism at the level of theory, testing, explanation, predictive success, convergence of theories, and so forth is to be locked into a world of representations. No wonder that scientific antirealism is so permanently in the race. It is a variant on "the spectator theory of knowledge."

Scientists, as opposed to philosophers, did, in general, become realists about atoms by 1910. Despite the changing climate, some antirealist variety of instrumentalism or fictionalism remained a strong philosophical alternative in 1910 and in 1930. That is what the history of philosophy teaches us. The lesson is: think about practice, not theory. Antirealism about atoms was very sensible when Bain wrote a century ago. Antireal-

12. But excluding Leibniz, who "knew" there had to be some real, natural difference between right- and left-handedness.

13. Iceland spar is an elegant example of how experimental phenomena persist even while theories about them undergo revolutions. Mariners brought calcite from Iceland to Scandinavia. Erasmus Bartholinus experimented with it and wrote it up in 1609. When you look through these beautiful crystals you see double, thanks to the so-called ordinary and extraordinary rays. Calcite is a natural polarizer. It was our entry to polarized light which for three hundred years was the chief route to improved theoretical and experimental understanding of light and then electromagnetism. The use of calcite in PEGGY II is a happy reminder of a great tradition.

14. It also turns GaAs, a 3/4 to 1/4 left-hand/right-hand polarizer, into a 50-50 polarizer.

15. I owe these examples to conversation with Roger Miller of SLAC.

16. The concept of a "convincing experiment" is fundamental. Peter Gallison has done important work on this idea, studying European and American experiments on weak neutral currents conducted during the 1970s.

17. I owe this information to Charles Sinclair of SLAC.

18. Alexander Bain, *Logic, Deductive and Inductive* (London and New York, 1870), 362.

9

Explanation and Realism

Clark Glymour

I

One way to argue to a theory is to show that it provides a good explanation of a body of phenomena and, indeed, that it provides a better explanation than does any available alternative theory. This pattern of argument is not bounded by time or by subject matter. One can find such arguments in sociology, in psychometrics, in chemistry, and in astronomy, in the time of Copernicus and in the most recent scientific journals. The goodness of explanations is a ubiquitous criterion; in every scientific subject it forms one of the principal standards by which we decide what to believe. The ambition of philosophy of science is, or ought to be, to obtain from the literature of the sciences a plausible and precise theory of scientific reasoning and argument: a theory that will abstract general patterns from the concreta of debates over genes and spectra and fields and delinquency. A philosophical understanding of science should, therefore, give us an account of what explanations are and of why they are valued but, most important, it should also provide us with clear and plausible criteria for comparing the goodness of explanations. One half of the subject of this essay concerns a fragment of such a theory. I will try to describe, without gratuitous formality, some features that generate clear and powerful criteria for judging the goodness of explanations. That is half of my subject; the remainder concerns what such criteria determine about what we ought to believe. Both halves are prompted by Bas van Fraassen's recent and delightful book, *The Scientific Image*.