

7 Experimental Realism Reconsidered

How Inference to the Most Likely Cause Might Be Sound¹

Mauricio Suárez

INTRODUCTION

In her first book, *How the Laws of Physics Lie*, Nancy Cartwright provided a wide-ranging critique of a realist attitude to any explanatory scientific theory. Cartwright argued, roughly, that the explanatory power of a theory is at odds with its descriptive accuracy. The greater the “covering power” of a theory the more idealised and further from the truth the theory will be. Cartwright promoted instead the position known as “entity realism” or “experimental realism”, according to which it may be possible to have a justified belief in the existence of some unobservable entity postulated by science, independently of any justification for our current best theory about that entity. Experimental realism thus achieves a combination of common sense realism about some unobservable entities with a principled nonrealism about theories.

By contrast Cartwright’s last book *The Dappled World*, is perfectly happy to accept a robust form of scientific realism about theories. As she puts it herself: ‘Nowadays I think I was deluded about the enemy; it is not realism but fundamentalism that we need to combat’ (Cartwright 1999: 23). Fundamentalism entails that there is only one true set of laws about the world; antifundamentalism on the other hand allows a large number of scientific theories, postulating alternative sets of laws, to be true at once, each of them in their particular domain. Cartwright’s current view is “anomalous dappling”, according to which different laws govern different patches of the world, but no law may govern some patches at all (Cartwright 2002).² “Anomalous dappling” allows us in principle to take a full-blown realist attitude to many more than just one empirically adequate theory, as long as they don’t contradict each other, thus yielding the promiscuous or patchwork realism that is in accordance with the metaphysics of the disunified world.³

Cartwright’s quote suggests that she might be happy to accept that if one remains unconvinced about realism about theories, one would not be particularly inclined to defend either antifundamentalism or fundamentalism,

since both presuppose realism. The fundamentalist claims that all regions of the world are law-governed, and moreover by the same laws; the antifundamentalist, such as Cartwright's "anomalous dappler", claims that only some regions are law-governed, and not necessarily by the same laws. Both presuppose that science aims, through its laws, to represent the way the world really is. So their dispute about whether there is one law that subsumes all phenomena is also a substantive ontological dispute about what the world is really like.

By contrast, an antirealist, or nonrealist, will find that there is no substantive ontological issue at stake. He or she will find no offence in the search for the system that best organises and economises our thought or even in supposing that there is one system that does it best. For the existence of such system does not show anything about what the world is really like but only, at best, about how we conceptualise it. To put it in a nutshell, only from the perspective of realism about theories can fundamentalism be the "enemy"; and only from that perspective can anomalous dappling be defended. From the perspective of nonrealism, both fundamentalism and anomalous dappling are metaphysical views underdetermined and not required by the practice of science or by an abductive inference to the best explanation of that practice. From this perspective, it is just as mistaken to draw metaphysical lessons from scientists' failure to find a unified system of laws that fits all phenomena as it is to draw them from their success in finding it. Fundamentalism and anomalous dappling appear to be equally unwarranted and unnecessary.⁴

This suggests that the choice that we are presented with is not really exhaustive or, more precisely, that we need not share its presuppositions. I was surprised to find Cartwright essentially conceding this point at the Konstanz conference, where she presented anomalous dappling not so much as a metaphysics of its own but as an attempt to break away from the dominance of fundamentalist metaphysics. It then seems to me that the anomalous dappler and her fundamentalist opponent share the mistaken assumption that theoretical physics and philosophy of physics have so far been totally dominated by fundamentalist metaphysics (see for instance Hoefer 2003). At least the work of those physicists and philosophers of physics that I am most familiar with (e.g., Bohr, Schrödinger; Hempel, Reichenbach, Fine, Van Fraassen) presupposes neither fundamentalism nor anomalous dappling. I don't see, for instance, how one attacks, or defends, philosophy of physics as a discipline by attacking, or defending, fundamentalism.

Cartwright would probably reply that the antirealist is still "deluded about the enemy". She has advanced several arguments against fundamentalism and in favour of anomalous dappling. These are detailed and intricate arguments, which require careful and thorough philosophical analysis. I will not attempt to evaluate them here. My aim in this chapter is the more modest one of discounting one possible motivation for Cartwright's present belief that she was "deluded about the enemy". One way to motivate the dilemma between anomalous dappling and fundamentalism is of a negative

sort: These are the only alternatives to understand science; no other alternatives will work.

More specifically, one reason that may lead (and which may have led Cartwright) to the theoretical realism that underpins the dilemma of *The Dappled World* is dissatisfaction with the experimental realism of *How the Laws of Physics Lie*. There have been many papers directly criticising experimental realism since 1983. The objections fall into three different kinds, depending on whether they charge experimental realism with (i) inadequacy, (ii) incoherence, or (iii) implausibility. Did Cartwright abandon experimental realism because these critics convinced her? Is that the reason why she changed the focus of her criticism? I would like to argue that there is no sound argument, even of such a negative kind, in favour of Dappled World metaphysics. There was no need to abandon the antirealism about theories of *How the Laws of Physics Lie* in the first place: The arguments against “experimental realism” are inconclusive, and a version of the position is defensible.⁵

The experimental “realism” that is rendered plausible by the arguments in this paper is distinct from what Cartwright intended in the first place, and seems to me to be incompatible with her recent *dappled world* metaphysics. In fact, I do not think of it as a realism at all. A more appropriate name for my views would be “the experimental attitude”, as among all contemporary epistemological views it is closer to the “neither realism-nor-antirealism” of Arthur Fine’s NOA. In the remainder of this paper I try to sketch out what this position would amount to if forced to express it as an epistemic thesis—as I believe Cartwright should have done in 1983.

EXPERIMENTAL REALISM: COMMON-SENSE EPISTEMOLOGY OR FANCY METAPHYSICS?

It is widely believed that, when first introduced by Hacking and Cartwright, experimental realism was conceived primarily as a metaphysical thesis about what kind of entities are, or can be, real. Hacking’s manipulability criterion, in particular, is often taken to express a condition on what should count as real. For example, Margaret Morrison writes:

Hacking contrasts the metaphysical questions concerning scientific realism with those that deal with rationality, the epistemological questions. The former raise issues such as, Are those entities postulated by physics theories real?, What is true of those entities?, What is truth?, and so on . . . In arguing for entity realism Hacking takes himself to be addressing only the metaphysical questions. (Morrison 1990: 1)

I shall argue that Morrison’s description of Hacking’s position as *only* metaphysics cannot be correct. Certainly Hacking’s slogan ‘if you can spray

them then they are real' lends itself to this interpretation, and indeed Hacking has explicitly defended a metaphysical claim on behalf of experimental realism:

Reality has to do with causation and our notions of reality are formed from our ability to change the world . . . We shall count as real what we can use to intervene in the world to affect something else, or what the world can use to affect us. (Hacking 1983: 146)

However, in addition to this metaphysical claim about what should be counted as real, Hacking also makes an explicit epistemic claim:

The *best kinds of evidence* for the reality of a postulated or inferred entity is that we can begin to measure or otherwise understand its causal powers. The *best evidence*, in turn, that we can have this kind of understanding is that we can set out, from scratch, to build machines that will work fairly reliably, taking advantage of this causal nexus. (Hacking 1984: 170, my emphasis)

The metaphysical claim then aims to establish a conceptual link between reality and manipulation. The epistemic claim asserts that manipulation provides particularly robust warrant for our ontological commitments. Which one is primary? Suppose that the metaphysical claim was primary. This might be enough to substantiate Morrison's description of Hacking's experimental realism as a piece of metaphysics. Hacking sometimes writes as if the epistemic claim is required only to bring out the practical consequences of the metaphysical one. Or as he puts it, the metaphysics would be "idle" without the epistemology (Hacking 1983: 28); it is because of the tight conceptual link between reality and manipulation that our best evidence for an entity's existence is our manipulating it.

However, Hacking accepts that manipulation is not a necessary condition on reality. There could be real entities out there that we would never be able to manipulate: black holes and gravitational lenses are possible examples. Our failing to manipulate them does not necessarily mean that these entities are unreal; it simply precludes us from having grounds to justify their existence (Hacking 1989). Manipulation is then meant, if anything, as a sufficient condition on reality. Hacking is not defending the conceptual equivalence of what is real and what can be manipulated, but rather, it seems, that manipulation is one important hallmark of reality. We may enunciate Hacking's metaphysical claim as follows:

Metaphysical Experimental Realism (MER)

Manipulation is a sufficient condition on reality: x is real if x can be manipulated.

A fair amount of criticism has been devoted to this metaphysical claim. In particular the critics have argued that if manipulation is a hallmark of reality, it is hard to see how we could classify or describe some type of entity as “real” independently of the theory that describes its causal powers and our possible manipulations of them.⁶ Would it not be incoherent to classify entities as “real” that we have no theoretical description of?

I think that there is something right about this incoherence objection as applied to (MER), and I will return to it in due course; but I’ll argue that it can only be an objection to the metaphysical version of experimental realism. I noted that Hacking also makes an epistemic claim on behalf of experimental realism, and I want to suggest that experimental realism must be understood as primarily making this epistemic claim:

Epistemic Experimental Realism (EER)

Manipulation is a necessary and sufficient condition on causal warrant: Our belief that x exists acquires this special kind of warrant if and only if we believe that we manipulate x .

EER is consistent with many passages in both Hacking’s and Cartwright’s original papers. I already quoted a passage from Hacking’s *Representing and Intervening* to this effect. Let me quote a couple from Cartwright’s *How the Laws of Physics Lie*, where I have also emphasised the phrases with undeniable epistemic content:

Causal reasoning provides good grounds for our beliefs in theoretical entities. (Cartwright 1983: 6)

I agree with Hacking that when we can manipulate our theoretical entities in fine and detailed ways to intervene in other processes, then we have *the best evidence possible* for our claims about what they can and cannot do; and theoretical entities that have been *warranted* by well-tested causal claims like that are *seldom* discarded in the progress of science. (Cartwright 1983: 98)

The view that I want to defend manipulation is indeed taken as an indication, or symptom of reality, but not a certain one; for it is not part of the notion of warranted belief that warrant be infallible and the corresponding belief always true. Hence our taking ourselves to manipulate x cannot be, on this view, a sufficient condition on x ’s reality. MER does not follow from EER.⁷

This chapter constitutes a first step in an argument to the effect that experimental realism needs to make no metaphysical commitments at all and is in particular not committed to MER.⁸ In other words, I want to turn the presumed primacy of MER on its head, in order to defend experimental realism as *only* epistemology. Our belief in the existence of x acquires a

special sort of warrant when we come to convince ourselves that we manipulate x ; and it is precisely this fact about our epistemic practice that grounds the secondary claim that manipulation is a good indicator of reality; a good guide—not an infallible one.

Experimental realism then needs to establish that EER is true by (i) elucidating the notion of causal warrant and (ii) showing that manipulation affords it. We find some clues for (i) and a partial but essentially sound defence of (ii) in Cartwright's arguments in favour of inference to the most likely cause.

INFERENCE TO THE MOST LIKELY CAUSE (IMLC)

In Chapter 4 of *How the Laws of Physics Lie*, Cartwright argues that inference to the most likely cause (IMLC) is a success term: a putative causal explanation of a phenomenon is only a genuine explanation if the cause is real.⁹ By contrast, inference to the best theoretical explanation, or explanation by subsumption under a theory is, according to Cartwright, not a success term: A theory may provide a good explanation of some phenomenon, regardless of its truth-value. So our acceptance of a theoretical explanation of a phenomenon *qua* explanation is not in itself a reason to believe in the explanation; but our acceptance of a causal explanation is a reason to believe in the existence of the cause or causes cited. As an illustration, Cartwright gives the following everyday example:

My newly planted lemon tree is sick, the leaves yellowing and dropping off. I finally explain this by saying that water has accumulated in the base of the planter: the water is the cause of the disease. I drill a hole in the base of the oak barrel where the lemon tree lives, and foul water flows out. That was the cause. Before I had drilled the hole, I could still give the explanation and to give that explanation was to present the supposed cause, the water. There must be such water for the explanation to be correct. An explanation of an effect by a cause has an existential component, not just an optional extra ingredient.

(Cartwright 1983: 91)

Cartwright takes the reality of the cause to be an intrinsic characteristic of the causal explanation: pointing to a nonexistent cause cannot explain anything. A cause can only constitute a genuine explanation if it actually exists. By contrast, Cartwright follows Duhem and Van Fraassen in arguing that theoretical explanation is not a success term. Providing a satisfactory explanation of a phenomenon by subsuming it under a theory gives no reason to believe that the theory is true. The theory can be explanatory without being true. This is because theoretical explanation does not in general

meet the requirement of nonredundancy, which is met in the case of causal explanation:¹⁰

We can infer the truth of an explanation only if there are no alternatives that account in an equally satisfactory way for the phenomena. In physics nowadays . . . there is redundancy of theoretical treatment, but not of causal account. (Cartwright 1983: 76)

In addition, Cartwright distinguishes between two different senses of “theoretical explanation”. First, there is the explanation of a phenomenon by showing that the law that describes the phenomenon is a special case of a theoretical law. Hempel’s Deductive Nomological account of explanation is a case in hand. Inference to the best D-N explanation, admits Cartwright, gives *some* reason to believe in the theoretical law. This reason is not conclusive, though, since the phenomenological law could have been derived from an alternative theoretical law. We are never in a position to rule out such alternative law, hence the requirement of nonredundancy is not generally met. For instance, suppose that we are able to deductively derive the exact future positions of most planets from the laws of Newtonian mechanics plus facts about the present positions and forces. This provides some explanation of the motions of the planets in the solar system; and the success of this explanation in turn provides some evidence in favour of Newton’s theory; but it does not conclusively show that Newtonian mechanics is true, since the motions might also be derived from another set of laws (as indeed is the case with Einstein’s theory).

Cartwright distinguishes D-N explanations carefully from the type of theoretical explanation that Duhem and Van Fraassen discuss. A theory according to Duhem is an abstract system whose aim is to summarize and classify logically a group of experimental laws. The requirement of nonredundancy fails maximally here: We can be certain that alternative summaries and classifications of experimental laws always exist; ours is just one that we find convenient for our own purposes. Hence there is no reason at all, conclusive or otherwise, to expect a theory, in this Duhemian sense, to be true.

What underlies this discussion is a difference between two types of inference to the best theoretical explanation (IBTE): those IBTEs that are completely unwarranted (inferences to the best Duhemian theoretical explanation) and those that are warranted but only mildly, on the supposition that no alternative superior explanations are available (inferences to the best Hempelian theoretical explanation). The former type of IBTE transmit no warrant to their conclusions; while the latter transmit what we may call theoretical warrant. By contrast, according to Cartwright, the requirement of nonredundancy is always met in the case of causal explanation. This is the explanation of a phenomenon by direct appeal to its cause. Since this type of explanation is only successful to the extent that the cause is real,

inference to the most likely cause (IMLC) is strongly warranted. I will refer to this type of warrant, which accrues from successful causal explanation, as causal warrant.

We have thereby established a distinction between three types of inference and the corresponding type of warrant that they are able to transmit to their conclusions. An inference to the best Duhemian theoretical explanation of a phenomenon is fully unwarranted, one to the best Hempelian theoretical explanation is theoretically warranted, while an inference to the most likely cause of a phenomenon is causally warranted. What is left to establish is the sense in which causal warrant is stronger than theoretical warrant.

Cartwright claims that causal explanations obey the requirement of non-redundancy because experimental testing and manipulation of the cause under controlled laboratory conditions allows us to establish the most likely cause of a phenomenon. Only then can we say that we have provided a causal explanation for it. Hence success in causally explaining a phenomenon by citing some entity *x* and its causal properties gives us the most conclusive reason that we may ever have to believe in *x*'s existence. Cartwright writes:

If God tells you that the rotting of the roots is the cause of the yellowing of the leaves, or that the ionisation produced by the negative charge explains the track in the cloud chamber, then you do have reason, conclusive reason, to believe that there is water in the tub and that there is an electron in the chamber. (Cartwright 1983: 93)

It is tempting to interpret “conclusive reason” in this passage as “infallible reason”, in a sense that would defeat a Cartesian sceptic: The existence of causal explanations would entail the existence of real entities and their causal properties; and this in turn would entail the existence of an external world. IMLC would then become the best proof of metaphysical realism. This interpretation is surely mistaken, however, even if it is unfortunately suggested by Cartwright's above appeal to divine revelation and by the contrast she draws between inference to the best (theoretical) explanation and the Cartesian *cogito ergo sum* argument (Cartwright 1983: 89). But it would be a mistake to suggest that IMLC has the global or radical scepticism-defeating character of Descartes' argument in favour of the existence of the external world. For that would make experimental realism susceptible to the dazzling battery of arguments in favour of global scepticism and would thus turn it into a thoroughly untenable position. In addition, it would be attempting to provide much more than is needed to defend a qualified form of selective realism. Experimental realism should not be required to show that an external world exists, but rather—on the assumption that there is such a world—that our beliefs in the unobservable entities of science are no less warranted than our beliefs in the objects of our ordinary experience.

Is there a coherent interpretation of the term “conclusive” that could make Cartwright’s statement true? I believe there is. It turns on the fact that the warrant afforded by the inference to the most likely cause of a phenomenon is more robust than the warrant that inference to the best theoretical explanation could ever provide. “Conclusive” is thus to be understood as a relative term: Causal warrant is conclusive in comparison with the warrant provided by an IBTE. No existential commitment derived from an IMLC can be defeated by any amount of theoretical warrant to the contrary. If we believe that we have manipulated an entity, in the way required for an IMLC, in order to causally explain a phenomenon, then no theoretical explanation of that phenomenon, no matter how empirically successful the theory, ought to lead us to withdraw the causal commitment.

Once I come to be convinced that I have manipulated the water in the basis of the lemon tree planter, in such a way as to establish it to my satisfaction as the cause of the yellowing, then I would not give up my belief that water causes the yellowing (and it would be epistemically irresponsible for me to do so) on account of any alternative, purely theoretical or speculative, explanation. The only defeater of causal warrant in favour of the existence of x is causal warrant of the same strength against x . I would only abandon the belief that there is water that causes the yellowing if I obtained causal warrant of the same strength in favour of a different substance causing the yellowing—forcing me to conclude that I was wrong to believe that I was manipulating water in the first instance.

To sum up my improved version of Cartwright’s argument in three compact theses:

- (i) Duhemian theoretical explanation is not a success term—in the sense that a false theory T may provide a satisfactory explanation of a phenomenon. But causal explanation is a success term—if the cause is not real there is no genuine explanation.
- (ii) From the fact that a theory T explains a phenomenon y , we cannot infer that T is true. But from the fact that x (probably) causally explains y we may infer that x is (probably) the real cause of y .
- (iii) We can accept a theoretical explanation, *qua* explanation, even if we do not believe that the theory is true. But we cannot accept a causal explanation, *qua* explanation, unless we believe that the cause is real.

Many of the objections to experimental realism, which I review in the next section, suppose that (i), (ii), and (iii) are the epistemological consequences of MER. Consequently, the inference to the cause in (ii) would have to be certain, or infallible. For in order to establish that x causally explains we need to have manipulated x , and if x has been manipulated then, by MER, it surely is real. Yet, (i), (ii), and (iii) can be argued for directly, without presupposing MER; and thus without presupposing that an inference to the most probable cause is certain, or infallible.

INADEQUACY, INCOHERENCE, IMPLAUSIBILITY

The objections to experimental realism fall roughly into three types. I will discuss them in order of what I take to be their increasing importance.

Inadequacy

The first objection to experimental realism is that it provides an inadequate picture of the actual aims and particular objectives of scientific research and inquiry. Such an artificially mixed combination of realism about entities and antirealism about theories does not drive any particular scientific inquiry, and it does not accurately describe science, including the most experimental, nontheoretical branches of applied science. Resnik, for instance, writes:

Experimenters do not operate without genuine scientific theories and laws about the phenomena they investigate: The gulf between experiment and theory is not nearly as large as Hacking supposes. . . . A person running experiments with a particle accelerator may not be aware of the latest developments in theoretical physics, but he (or she) is likely to be familiar with most of the commonly accepted background theories in physics, including some theories about the particles he (or she) is studying.

(Resnik 1994: 410)

Let us suppose Resnik is correct about the intertwining, in practice, of theoretical and practical knowledge. It seems plausible that scientists use both theoretical and experimental knowledge in their work; and of course many of their beliefs are infused by theory. The experimental realist need not deny any of this, though. He or she need only claim that theoretical and experimental knowledge have some distinct cognitive and epistemic functions but not that all their functions are distinct or fully separate. Resnik's objection will draw out slightly different replies depending on whether experimental realism is understood as primarily defending MER or EER, but the core of the reply will be the same in either case: The origin and content of scientists' beliefs and knowledge is irrelevant to experimental realism. MER claims that scientists' ontology is ultimately derived from their phenomenological knowledge and the results of their laboratory manipulations. EER claims that experimental laboratory practice provides the strongest form of warrant for scientist's existential commitments. Each of these views is fully compatible with scientists' possessing a mixture of theoretical and experimental knowledge.

The issue concerning experimental realism is a more fine-grained question about the specific role that these two types of knowledge play in actual scientific practice. For MER the question is what type of knowledge scientists

ultimately base their ontologies on; for EER it is what warrants their existential commitments; but both can accept that theoretical and experimental knowledge is always deeply intertwined in practice.

Incoherence

The second objection is that experimental realism is incoherent. It is not possible to coherently separate an unobservable entity from the theory that describes it, because the only concept that we possess of an unobservable scientific entity is the one given to us by the theoretical description that we happen to accept. We may have both a theoretical and a perceptual concept of an observable entity, such as the Konstanz train station; but about an unobservable entity, we can have no perceptual concept, only theoretical.

Hacking is clear that in order to describe our causal interaction with an unobservable entity we need only appeal to a set of phenomenological “home truths” about the entity and its properties; we need not believe the full theoretical description of it. Suppose we infer the reality of an entity x on the basis of our interaction with it, as described by some such set of home truths. The thorny question for the experimental realist is this: what is it that we are inferring to, when we infer to the reality of an entity on the grounds that we can manipulate it? Whatever it is, it is not that entity that we take ourselves to be inferring, since the properties that we must ascribe to the entity in order to manipulate it are typically only a subset of the full set of properties that informs our theoretical concept.

For instance, the entities that we must suppose are real because we manipulate them in the electron microscope are not quite electrons as we understand them: They are particles (call them *flectrons*) that have some of the properties of electrons but not all of them. (MER) would then not allow us to claim that electrons are real but only that flectrons are. And in each distinct manipulation of “electrons”, each circumscribing our causal interaction to a different subset of their properties, we would *really* be manipulating different entities: “flectrons” this time, “plectrons” the next, and so on. Yet, this is deeply counterintuitive, if not plainly wrong. Scientists do not take themselves to be confronting different particles when they carry out a scattering experiment on electrons as opposed to operate an electron microscope. In both cases they take themselves to be confronting electrons.

This is a powerful objection against the metaphysical version of experimental realism because, on its most natural reading, MER sanctions only those inferences to the existence of the particular properties that are actively being manipulated. Since it’s only the properties of flectrons that scientists manipulate in an electron microscope, it should only be flectrons that scientists are entitled to presume are real, and so on. Hence the incoherence charge shows that MER comes into conflict with the actual ontological commitments of scientists.

In addition, the incoherence objection also undermines the following cognitive claim:

Belief Possession Experimental Realism (BPER)

Manipulation is a necessary and sufficient condition on possessing existential beliefs: A subject *S* can have the belief that *x* exists if and only if *S* manipulates *x*.

For suppose BPER is true; the incoherence charge then shows that scientists rarely ever have—warranted or unwarranted—beliefs in any theoretical entities, as they rarely ever manipulate all of an entity’s properties at once. But this would conflict with scientific practice once again. For example, Margaret Morrison has provided the details of two case studies, the cloud chamber and charmed quarks, which show that manipulation can fail to induce the appropriate beliefs and even ‘can occur in a context where there are no firmly held beliefs about the entities being manipulated’ (Morrison 1990: 6). These two case studies show empirically that BPER is false in general.

However, I argue that experimental realism is not committed to BPER or MER but only to EER. On this epistemic version of experimental realism the “home truths” about an entity *x*, which we need to believe in for our inference that *x* is real to be causally warranted, need not in any way exhaust our concept of *x*. EER does not entail MER, so it does not entail that we can only infer those properties of *x* that we can manipulate or interact with. It entails instead that these are the properties that best ground our inference that *x* exists.¹¹

This is just the commonsense view that we apply in our ordinary life. I infer certain properties of the city of Konstanz (for instance those pertaining to the relative positions of its station with respect to the Hotel Barbarossa and the University, and the average time intervals to walk or travel by bus between them, which Stephan Hartmann accurately described in his instructions sheet on how to get there) because I have manipulated and interacted with them in order to find my way around. But those “home truths” neither exhaust the city of Konstanz nor my concept of it. There are many other interesting properties of Konstanz, some of them observable (i.e. its Town Hall building), some of them not (i.e. its founding year, the number of its inhabitants, or its geographical borders) that I can suppose are real, and indeed believe to be real, but because I’ve had no opportunity to manipulate them, I have no causal warrant for them.

Neither does EER entail BPER: Manipulation is neither a sufficient nor a necessary condition on having, or acquiring, the belief that *x* exists. The source of our belief in some unobservable entity may in no way be related to the grounds that warrant that belief. Theory is undoubtedly, for most of us, the source of our belief in most unobservable scientific entities, including electrons. Most of us learn about electrons from our high-school classes on

electromagnetism and particularly Maxwell's theory. This is a fact about our psychology. It is also possible that this is indeed the most educationally efficient and appropriate way to acquire such a belief. That would be a fact about pedagogy. Experimental realism, in the epistemic version that I wish to defend, makes the distinct and additional claim that our belief in electrons possesses a special sort of warrant, causal warrant; we can be particularly confident in our belief in electrons because we are confident that we routinely manipulate electrons, or their causal properties, in experimental conditions. The fact that this belief has this type of warrant pertains to epistemology and is *prima facie* independent of the psychological and pedagogical facts.

Most of us will only know that electrons have been manipulated by description; only a few of us have direct acquaintance with the operations required to correctly use an electron microscope. No matter. According to EER it is only because we believe that we can manipulate electrons that we have a causally warranted belief in their existence: The methods employed in acquiring the beliefs that there are electrons and that they can be manipulated are irrelevant to the epistemic, warranting-transferring relation that holds between them. Again this is the common sense view: I may have learnt about the facilities at Hotel Barbarossa from a theoretical description; but my corresponding beliefs acquire causal warrant only to the extent that I have manipulated those facilities. The acquisition of the belief need not have the same source, nor follow the same route, as the acquisition of the causal warrant.

To sum up, experimental realism is not a metaphysical thesis about what is real, nor is it a psychological thesis about the source or origin of our beliefs in the entities postulated by science; it is an epistemological thesis about the grounds that warrant those beliefs. The incoherence charge does not undermine *this* thesis.

Implausibility

The third, and in my view the sharpest, objection to experimental realism accepts that the position is coherent and that it does not conflict with scientific practice, but argues that it provides an implausible epistemology for science. Although one could come to have a warranted belief in the existence of an entity on the basis of laboratory manipulations, without believing in the truth of any particular theory about it, and although scientists often do so, it cannot be an epistemological principle that they *ought* to do so.

The objection has been pursued in an interesting paper by Christopher Hitchcock, which addresses directly Nancy Cartwright's argument for IMLC (Hitchcock 1992). In this paper Hitchcock presents two challenges for the experimental realist, which pertinently track what in my view are the two stages in the overall argument in favour of experimental realism. Hitchcock's first challenge is that antirealists about theoretical entities may also

be able to accept causal explanations without thereby committing themselves to the reality of the cause.¹² In other words, Hitchcock questions that causal explanation is really a success term. ‘What is special’, he asks, ‘about the role of explanation such that causal stories filling that role, and not some other, must be believed if accepted?’ (Hitchcock 1992: 174).

The core of Hitchcock’s first challenge is a couple of examples of putative causal explanations where the cause is most definitely not real. I will discuss only one of the examples, as they are argumentatively identical. He considers an explanation of the two-slit experiment that is sometimes offered in quantum mechanics textbooks. Electrons are fired through a screen with two slits *A* and *B* on it and are then detected in a further screen (see Figure 3). The pattern of detections of particles in the faraway screen does not correspond to the sum of the patterns registered when the experiment is repeated with slit *A* closed and *B* open (Figure 7.1), and slit *A* open and *B* closed (Figure 7.2). This is true even when only one electron at a time is passed through the slit.

Interestingly the pattern of Figure 7.3 is destroyed as soon as a measurement is made to detect which slit the electron actually goes through. Hitchcock considers the following possible explanation of this fact. A detection process of the electron in either slit will ultimately consist in bouncing a photon off the electron as it passes through the slit. This will impart momentum on the electron, which will affect its trajectory, thereby destroying the interference pattern. This story, argues Hitchcock, seems perfectly explanatory and is a description of causal processes. We must thereby, by IMLC, infer that there is such photon-electron interaction taking place.

However, as Hitchcock quickly points out, this explanation is unacceptable because it ‘contradicts almost every interpretation of quantum mechanics, and as such would not be believed to be true by any but the most




Figure 7.1




Figure 7.2

stubborn believer in hidden variables’ (Hitchcock 1992: 171). According to quantum theory electrons do not have classical, continuous trajectories; but according to the causal story above, the entity putatively responsible for the observed interference pattern is precisely taken to be the electron’s trajectory. Hitchcock concludes that IMLC cannot provide the sort of warrant that Cartwright takes it to: A causal story may be acceptable, as an explanation, even if the cause is most certainly not real.

However the defender of EER will reply that Hitchcock’s argument is flawed, and cannot refute EER, because there is no causal explanation in the first place. In a causal explanation it is not the causal *story* that does the explaining but the causes themselves. The problem with the causal story above is that it presupposes an account of the interference pattern that we lack causal warrant for—and arguably have some causal warrant against. We are invoking the photon-electron interaction in order to explain not the interference pattern but rather its disappearance when we detect the electron




Figure 7.3

in the first screen. So the causal story presupposes that were we *not* to detect the electron's passage in the first screen, the electron's trajectory would have been causally responsible for the interference pattern. In other words we are presupposing that the interference pattern in a two-slit experiment is causally explained by the electrons' trajectories. And this "explanation" of the interference pattern is not causal—if it is an explanation at all—since it "would not be believed to be true by any but the most stubborn believer in hidden variables". And this is not so on simply interpretational or theoretical grounds but on the grounds of the experimental evidence against the existence of classical trajectories in quantum mechanics, in the form of all kinds of interferometry experiments.¹³

The defender of causal explanation takes causal explanation to be a success term. So if we don't believe in the "causes" appealed to in the story then the explanation the story offers—if any—cannot be said to be causal. The causal story that we are told about photon-electron interactions can only be accepted as an explanation in the sense that theoretical explanations can be; that is, we are given a theoretical account, and we are invited to deduce the phenomenon from it. But it has already been established that there is no truth-requirement on theoretical explanation. The fact that the theoretical account employs causal vocabulary is irrelevant. For a causal explanation of a phenomenon does not merely subsume a phenomenon under a theory that uses causal vocabulary: A causal explanation of a phenomenon such as the breakdown of the two-slit pattern must cite the actual causes of that phenomenon. And this is patently not the case in Hitchcock's example, as he himself acknowledges.

WHAT IT TAKES TO MAKE IMLC SOUND

In the epistemological version of experimental realism that I am defending, successful causal explanation provides warrant for the existence of the cause cited. But similarly, following Cartwright, some theoretical explanations provide warrant for the truth of the theories involved. The key to experimental realism, I am suggesting, is to distinguish carefully two types of warrant, which we may refer to as causal and theoretical.¹⁴ The claim is then that causal warrant, i.e. warranted inference to the most likely cause, is conclusive (not "infallible") in the sense that only causal warrant to the contrary would force us to withdraw the existential commitment. In other words the commitment won't be defeated by an alternative theoretical explanation that dispenses with the entities so warranted.

By contrast, even the most warranted inference to the truth of a theoretical explanation would *ipso facto* be defeated by causal warrant to a cause whose existence contradicts the theory. Hence I am suggesting that inference to the most probable cause provides a type of warrant that is uniquely strong in that it can only be defeated by warrant of the same type.

Thus the full argument in favour of experimental realism has not one but two stages. The first stage has been provided by Cartwright's claim, already reviewed, that causal explanation, unlike theoretical explanation in general, is a success term. When "causal" and "theoretical" are understood properly this argument is, in my view, sound. Causal explanation has a truth-requirement built in: A cause can genuinely explain only to the extent that it is real, and we can only take it to explain if we believe it to be real. I have argued that Hitchcock's putative counterexamples are not actually such; and I know no other convincing counterexamples to Cartwright's claim so far.

The second stage of the argument is equally important but was, if anything, left implicit in Cartwright's original writings. It is the claim that an inference to the existence of an explanatory cause is more robustly warranted than an inference to the truth of an explanatory theory. Note that this claim, although not unrelated, is distinct from the truth-requirement one. Hitchcock's second challenge directly addresses this fact. (In what follows I change Hitchcock's terminology in order to distinguish a causal explanation from a theoretical explanation that employs causal vocabulary. I assume throughout that capital letters P , Q refer to sentences in some language, which may include causal terms—so " P causally explains Q " is a sentence in some theory that uses causal language. I reserve noncapital letters p , q to refer to entities, their properties, or phenomenological facts directly—so " p causally explains q " is not a theory but a statement of fact).¹⁵

Suppose that causal explanation *is* a success term; and suppose that I know that some phenomenon q occurs, and I find out, through careful manipulation in laboratory conditions, that p is the most likely cause of q . I am then invited by IMLC to infer that p is real. Hence " p is real" is the conclusion of an argument that has, as premises (i) q occurs and (ii) " p causally explains q ". But, asks Hitchcock, since causal explanation is a success term, do I not need to believe in the existence of p in order to accept that ' p causally explains q ' in the first place? And if so, in what way is the inference to p in this argument providing me with any warrant in p 's existence that I did not already have?

The challenge is interesting and to the point, but it can be answered—and in an illuminating way. Let me first discount a trivially off-the-mark interpretation of Hitchcock's second challenge. Hitchcock is not criticising experimental realism for supposing that we must simultaneously believe that p is real and that p causally explains q ; that is, he is not merely pointing out that in order to believe that p causally explains q , we need to already have the belief that p is real. This at best would argue against BPER, which I have already discounted as a misinterpretation of experimental realism. The fact that my belief " p causally explains q " necessarily presupposes my belief " p is real" is part and parcel of what it means for causal explanation to be a success term, so EER cannot be in the business of denying it.

The potential problem that Hitchcock is pointing to here is deeper, and quite general. The question is: How can we take the argument above—which

appeals to a specific explanation of a particular phenomenon q by means of some putative cause p —to provide warrant for my (antecedent) belief that p is real? What are the conditions for a deductive argument to transmit warrant from its premises to its conclusion? Is it not clear that any such argument (where the truth of the conclusion must be presupposed in order to believe in the truth of the premises) would fail to transmit warrant and fail to provide me with a new reason to back up my belief in the conclusion?

Crispin Wright and Martin Davies have for some time been studying the mechanisms that underlie loss of warrant transmission in a deductively valid argument (Davies 1998; Wright 2000, 2003). Their particular target is McKinsey's argument for the incompatibility of externalism about mental content and privileged access to one's own mental contents. Roughly, McKinsey tried to show that these two premises together entail a priori knowledge of natural kinds, which he took to provide a reductio refutation of externalism (McKinsey 1991). More specifically, the argument is as follows: (i) I believe that water is wet; (ii). If I believe that water is wet then I belong to a community of speakers that has had contact with water; hence (iii) my community has had contact with water. Since (supposing privileged first person access) I know (i) a priori, and (supposing strong content externalism) I know (ii) a priori; it follows that I can know (iii), that there is water in my environment, a priori.

Wright and Davies defend the compatibility thesis against McKinsey's argument. They suggest that McKinsey's argument may be valid, but not cogent, in the sense that it fails to transmit warrant to its conclusion. So it offers us no reason to abandon externalism—or first-person access: We do not, in following the argument, acquire any warrant or justification for the conclusion. Wright and Davies claim, roughly, that the only evidence that I may ever possess in favour of the first premise in McKinsey's argument ("I believe that water is wet") is an introspective experience of the content of my mental state, which entails, on the strong conception of externalism that grounds the second premise, that I can know a priori that there is water in my environment. So the evidence for the first two premises jointly entail the conclusion; and McKinsey's argument, although valid, and possibly sound, does not transmit warrant.

The method is rather general, and can be—now quite precisely—summed up as follows: If the conclusion of an argument is a necessary presupposition for the evidence that we actually have to hand for its premises then the argument is not capable (for us) of transmitting warrant from the premises to the conclusion (even if the conclusion as well as the premises is true—and even if we correctly believe them all to be true!) A deductively valid argument with true premises will not warrant its conclusion if the only evidence that we possess in favour of the premises would not be evidence had the conclusion of the argument been false.¹⁶

With this in mind, let us now turn to Hitchcock's second challenge. Is the argument in question—(i) q is real, (ii) p causally explains q ; therefore (iii) p

is real—warrant transmitting? Is the conclusion (iii) a necessary presupposition of the evidence that we possess in favour of (i) and (ii)?

The reality of p is not a presupposition of any evidence we may possess for q —otherwise it would be impossible to empirically establish a phenomenon without thereby also establishing the reality of distinct, apparently unrelated, unobservable entities: Any empirical evidence in favour of, say, electrical conductivity would ipso facto establish the reality of electrons, without any need for further experimenting or reasoning. So the question is whether the reality of p is required for the evidence that we need to have to hand in order to accept premise (ii) to be counted as evidence. I do not believe this is generally the case, and hence I do not believe that IMLC fails, in general, to be warrant-transmitting. Let me explain why. There are two features that typically distinguish causal explanation and provide evidence that some explanatory claim is causal: lack of redundancy and a material mode formulation. I argue that these features provide evidence for “ p causally explains q ”, regardless of whether p is real.

Nonredundancy

The nonredundancy requirement is met to a much larger degree by IMLC than by IBTE. Scientists establish which putative cause is nonredundant through controlled intervention and manipulation in laboratory conditions—and only then have they got reason to believe that the putative cause is genuinely responsible for the phenomenon.

By contrast, we have much poorer and controversial methods to establish which among all possible empirically adequate theoretical explanations of a phenomenon is the most probable one. (In the case of Duhemian theoretical explanation, we have no methods at all.) As is well-known, opinions on this matter differ enormously, both among philosophers and among practitioners: Is the most probable theoretical explanation the simplest one, the most ontologically parsimonious, the most familiar, the one that preserves the greatest amount of structure from previous theories, the one that explains a greater number of independent phenomena, etc, etc, etc. As a consequence there is much more redundancy, in the form of underdetermination, in the case of theoretical explanation. Lack of redundancy in an explanation is typical evidence in favour of it being a *causal* explanation.

In other words, we intervene and control variables in situations where we expect p to be operating in order to rule out redundancy in the explanation of q , which is in turn evidence that p causally explains q ; and this provides warrant—by IMLC means—that p is real. The question is whether lack of redundancy on its own entails that p is real. If it does then Hitchcock is right, and IMLC cannot transmit warrant.

EER entails that we would only revise our causally warranted existential commitments through further experimentation that shows some other cause is more probable. This is an extremely rare occurrence, but it is not

impossible. We arguably once had causal warrant for phlogiston but no longer do. The explanation of combustion is nowadays to be found in the interaction of oxygen with flammable materials—and we have acquired plenty of causal warrant in favour of oxygen and its role in combustion. So what we have here, arguably, is a case of causal warrant for the existence of an entity (phlogiston), and its role in combustion, that has been overturned by causal warrant for another entity (oxygen). We were convinced that we were able to manipulate phlogiston, and on that basis discarded any competing explanation of combustion; but we have since learnt that what we are actually able to manipulate is oxygen, which we have shown by experimental means to be involved in our present-day nonredundant explanation of combustion.¹⁷

This is in illuminating contrast with the case of the electromagnetic ether. It was generally accepted that it was not possible to manipulate or causally interact with the ether—even those few who thought that the ether might be manipulated were unable to convince themselves or others to have manipulated it.¹⁸ So it would be wrong to say that we once had causal warrant for the existence of the ether. At best we had some theoretical warrant; and this was lost when we abandoned ether theories in favour of Einstein's relativity theory.

I am thus suggesting that manipulation provides only a particularly robust kind of warrant—causal warrant. We can never be certain that we are in fact manipulating *p*; at best we can be certain that we believe that we are manipulating *p*. This belief allows us to establish, by means of intervention in laboratory conditions, that *p* is nonredundant as a causal explanation of *q*, and this nonredundancy is one type of evidence that we need to possess in order to accept (ii) that “*p* causally explains *q*”. But it should be obvious that both nonredundancy and the required belief in the manipulability of *p* are at best fallible evidence for (iii) “*p* is real” and do not, separately or jointly, logically entail it. (The fact that someone believes that he or she has interacted with aliens does not entail that there are aliens!)

Now the question is: does nonredundancy cease to be evidence in favour of the causal character of the explanation of *q* by means of *p* were *p* not real? Note that the question is not: “Can *p* be the causal explanation of *q* if *p* is not real?”, to which we already know we must give a negative answer. Rather the question is: “Would the nonredundancy of explanation—established by what we take to be manipulation of *p* under experimental conditions—cease to constitute (defeasible) evidence in favour of the causal character of the explanation of *q* by *p* were *p* not real?”

Suppose that Priestley did establish to his own satisfaction, by the experimental means of intervention and (what he took to be) manipulation of phlogiston under laboratory conditions, that the only explanation of combustion involves the presence of phlogiston. Does all this painstaking and careful experimental and laboratory work not amount to (defeasible) evidence for Priestley in favour of the claim “phlogiston-release is the causal

explanation of combustion?” It is hard to see what else could count as evidence in favour of such a claim. The reality or otherwise of phlogiston makes no difference whatever to the “evidential” character of the evidence in favour of the claim.¹⁹ Priestley carried out every single experiment, intervention, and manipulation that he could have been expected to carry out in order to establish such a fact experimentally; and (we may suppose) he reported his experimental activity and results in full honesty. He was led by his prior belief in phlogiston to interpret all his experimental manipulations as providing grounds for the nonredundant role of phlogiston in the explanation of combustion.

Of course such evidence was defeasible and in fact turned out to be defeated; and we no longer interpret his work in those terms. But if (iii) is not presupposed by the character of the evidence for (ii) and (i), then IMLC, unlike McKinsey’s argument, is in general capable of transmitting warrant. We do learn something after all when we infer from “*p* causally explains *q*” to “*p* is real”; what we learn is not that *p* is real—we already believed this—but that we have as good a reason as we could have to believe that *p* is real.

Hitchcock’s second challenge then fails to refute EER. What it does undermine in fact is a competitor view, defended by neither Cartwright (1983) nor Hacking (1983), which we may refer to as the internalist version of “metaphysical entity realism”:

Internalist MER: *x* is real if we believe that we manipulate *x*.

This principle must strike everyone as obviously too strong, for reasons already pointed out: the mere belief that we have interacted with aliens does not make them real. But suppose it was true, then (iii) “*p* is real” would follow from the evidence that we need to possess in order to accept (ii) “*p* causally explains *q*”, namely that we believe that we manipulate *p*; and IMLC could not transmit warrant, for the putative evidence in favour of (ii) would not be genuine evidence were the conclusion (iii) false. Hitchcock’s second challenge thus refutes MER, and it provides yet one more argument in favour of understanding experimental realism as EER *only*.

Material Inference

One aspect of experimental realism that is rarely mentioned is Hacking’s and Cartwright’s insistence on the importance of semantic descent from the formal mode to the material mode. They themselves have not made clear what precise role this distinction plays in the argument in favour of experimental realism. The following conjecture is, I think, plausible: Causal warrant can accrue to the conclusion of an inference entirely carried out in the material mode; theoretical warrant on the other hand is always the result of an inference in the formal mode. Although the conclusions of such inferences

are, as Carnap argued, intertranslatable, the vehicle of the inference, and the corresponding strength of the warrant transmitted to the conclusion, differ (Carnap 1935).

The difference might be best explained by means of the following example. Suppose that we would like to explain the phenomenon that metals dilate in the presence of heat. We could:

1. Formulate the corresponding phenomenological law and state it in a “protocol sentence”.
2. State formally the solid-state physics theoretical treatment of metals, including the formal hypothesis that heat makes molecules vibrate with higher energy and thus forces them move further apart from each other.
3. Deduce from this theory together with the required boundary and initial conditions, the “protocol sentence” in 1.
4. Infer, by IBTE, the truth of the theory including the molecular assumption.
5. Infer by semantic descent the reality of highly energetic molecules in a solid.

Or, alternatively, we could:

1. Formulate the phenomenological law to be explained.
2. Assume that molecules vibrate with higher energy in heat, thus move further apart from each other.
3. Causally explain the law by appeal to the assumption (i.e. describe the experiments that show that no other cause of the expansion is as likely by manipulating the molecules of different samples in order to vary their energy and then checking whether the heat of the solid covaries accordingly).
4. Infer directly, by IMLC, the reality of highly energetic molecules in a solid.

The former inference is a formal inference to the best theoretical explanation, whereas the latter is a material inference to a most likely cause. The latter type of inference is more robust—it contains fewer steps where it may go wrong. In particular we need not worry about how appropriate or fair the translation is into the formal mode description in the first place.

But can the formal mode description not be applied to the causally explanatory argument directly? Yes, indeed: material-mode speech is not always required for causes; we can refer directly to the causes or refer to our statements and theories about those causes. The statement in material mode that ‘magnetic fields can cause electrons to deflect’ is translatable into (although not synonymous with) the statement: ‘according to Maxwell’s theory, “magnetic fields” are correlated with deflections in the trajectories

of “electrons”.’ What is relevant here is not how to couch the explanation but what the actual relata of the explanatory relation are: the distinction between genuine causal explanations that refer directly to the causes, and causal “sounding” theoretical explanations, or causal *stories*, is crucial in countering Hitchcock’s first challenge to experimental realism.

However, no similar options are available for theoretical explanations. We cannot describe theories in the material mode, by definition, on pain of incurring a use/mention distinction fallacy. I cannot describe Newton’s principles or Schroedinger’s equation in a mode of speech that does not allow me to refer to sentences, theories, and language but only to real entities and their properties—unless I turn these principles and equations into the reality that they are aiming to describe. In other words, material mode speech is another hallmark of causal explanation, as only causal explanations can be cast entirely in this mode. The fact that an attempt at an explanation is given entirely in the material mode—that the relevant manipulations of the cause are presented or pointed to and not merely theoretically described—is fallible evidence that the attempted explanation is causal.

To conclude, the belief that we manipulate p turns out to be essentially involved in both types of evidence that we may possess for “ p causally explains q ”. The question then arises as to whether our having this belief, on its own, entails “ p is real”. For if so, IMLC transmits no warrant and cannot provide us with any reason to believe that “ p is real”. But as a matter of fact, according to EER, the belief that we manipulate p does not entail that p is real. So there is no real reason to expect failures of warrant transmission in an IMLC: This type of inference transmits causal warrant, thus providing us with a new and particularly strong reason to back up our belief in the existence of the cause.

AN EXPERIMENTAL REALISM?

Material inference to the most likely cause is the norm in ordinary abductive reasoning; for examples one need go no further than one’s own kitchen appliances or car mechanics. Whenever solving a problem with the normal functioning of our most familiar tools and appliances, we manipulate possible causes, provide evidence for causal explanations for the machines’ malfunctions, and thereby infer most likely causes. Experimental realism, in the epistemic form that I defend it here (EER), brings the epistemology of science in line with our everyday epistemology. What I have called causal warrant is not special to science; it is precisely the kind of warrant typical of successful inquiries in everyday life. If scientific realism is characterised as the view that our beliefs in the unobservable entities postulated by science are in principle as warranted (or unwarranted) as the beliefs in the objects of ordinary life, then EER is just enough on its own to furnish a kind of realist epistemology.

However, this is a very limited and modest realism, and it is questionable whether it deserves the honorific, much sought-after title.²⁰ EER says that we have a stronger type of reason to believe in car engine pistons, washing machine filters, and electrons in electron microscopes than we do to believe in quarks and quasars. But it does not say we have no reason at all to believe in quarks and quasars, nor does it say that the existence of pistons, filters, and electrons is beyond any possible doubt.

Hence a commitment to EER is rather minimal, and it is hard to see how anyone, regardless of his or her additional epistemological commitments, would disagree with it. Even Van Fraassen's constructive empiricist could accept EER, although he would naturally add an independent principle to further privilege warrant in observable entities—a principle that the experimental realist will not accept. This suggests that EER is a good candidate for part of that elusive “core” that Arthur Fine argued realism and antirealism share (Fine 1997). Fine's Natural Ontological Attitude (NOA) was explicitly designed to capture this “common core”, so I offer EER as a good candidate for a (part of) NOA.²¹

NOTES

1. This paper was written in 2002. Since then my views on this topic have sharpened considerably, in ways that do not always line up with Cartwright's advice in her response. For instance, I would now be explicit that inference to the most likely cause (IMLC) is not defensible as a method of causal *discovery*. It can only be used as a method to warrant the ordinary presuppositions of laboratory work, whether or not they play any causal role. So I would no longer use the phrase “causal warrant”, but “experimental warrant” instead. The ensuing shifts in terminology and argument are unfortunately too subtle to record in detail in the proofs (January 2008), and must await a paper of their own. In the meantime I emphasize these points in two brief footnotes in the final part of the paper.
2. The term “anomalous dappling” is due to Peter Lipton (Lipton 2002).
3. *The Dappled World* claims that theories postulating radically different laws and ontologies for different domains or at different levels of complexity may be simultaneously true; it does not follow from this that inconsistent theories may simultaneously be true in a way that would require a revision of classical logic.
4. Lipton's 2002 article was brought to my attention in the last stages of writing this chapter; it also argues for agnosticism, rather than atheism, about fundamentalism. Teller (this volume) essentially concurs with this agnosticism—but his arguments are orthogonal to those I present here.
5. I have occasionally tried to pull Nancy back to some of her earlier views. We even coauthored a brief paper (Cartwright et al. 1994) sketching a form of methodological instrumentalism about theories. This is an instrumentalism that, as I understand it, does not require any metaphysical backup, whether antifundamentalist or otherwise (see also Suárez, 1999).
6. This is a line of criticism adopted by Dorato (1988), Morrison (1990), and Resnik (1994). See also Elsamahi (1998), Gross (1990), and Reiner and Pierson (1995).

7. It is an interesting question (which I can not pursue here) whether the implication would be restored if EER were given an externalist twist, as follows. Externalist EER: our belief that x exists acquires causal warrant if we actually manipulate x . Manipulation is arguably a success term, so causal warrant would then become nondefeasible, and MER would seem to be implied by Externalist EER. There are however two important caveats. First, Externalist EEP stipulates that manipulation is a sufficient not a necessary condition on causal warrant—so there might be other sources of causal warrant, which would explain how it is sometimes defeasible. And second, even if there were no other sources of causal warrant (i.e. even if actual manipulation was a necessary as well as sufficient condition), one could still accept (Externalist EER) and insist that causal warrant is defeasible while denying that manipulation is in actual fact a success term. In either case (MER) would not follow from (Externalist EER).
8. Although of course, (MER) and (EER) may both happen to be true. My point is that experimental realism, in the epistemic version I wish to defend, neither requires nor provides grounds for (MER).
9. Cartwright's original term is "most *probable* cause". I prefer "most *likely*", because, strictly speaking, probabilities are only defined over statements, theories, hypotheses, or events but not entities or their properties.
10. It will become clear that I do not think the requirement is met exactly as worded here—but I do agree that there is a significant difference in the *degree* of redundancy in each case.
11. In correspondence, Peter Lipton points out that an analogy with descriptivist semantics illustrates vividly why the incoherence objection is off the mark: 'LSE philosopher whose name ends in "right" refers to all of Nancy Cartwright and her properties, even though it only mentions one most trivial property'. I agree: The flectron objection trades in a sense on ignoring the distinction between reference and mention.
12. This challenge was first raised by Fine (1991).
13. Neutron interferometry experiments are good instances. A noteworthy attempt to resist the weight of evidence is of course Bohmian mechanics, which notoriously restores well-defined trajectories by reinterpreting the evidence as a consequence of the radical nonlocal character of the quantum potential or guiding wave in configuration space. Brown et al. (1995) provide a critical analysis.
14. Nowadays I would distinguish "theoretical" and "experimental" warrant. (Note added to proofs, January 2008).
15. Hitchcock assumes that all inferences are carried out in the formal mode, since he assumes that the only way to infer that p is real from the fact that p causally explains q is to infer the sentence P : " p is real" from the sentences Q : " q is real" and the theory T : " P causally explains Q ". That is, he assumes incorrectly that the relata of causal explanations are, as is the case for theoretical explanation, sentences or theories. Steve Clarke, in his otherwise cogent defence of Cartwright's argument, also turns IMPC into a subspecies of IBTE, i.e. precisely those IBTE that do not suffer from redundancy (Clarke 2001). The above considerations suggest that this move upwards in the semantic ladder is not without consequences and in fact already gives half the game away against experimental realism.
16. In other words, an argument does not transmit warrant if the evidence for the premises would lose its character as evidence were the conclusion false. And indeed, the introspective experience of my own mental states would remain the same even if there were no external world, but it could no longer be taken as evidence that the water in my environment is wet.

17. Musgrave (1976) is a good summary of the complicated history of the overturn of phlogiston theories by Lavoisier's theory.
18. Warwick describes the wonderful case of Joseph Trouton's failed efforts to build a perpetuum mobile machine out of the earth's interaction with the ether (Warwick 1995).
19. It of course makes a crucial difference to the actual truth-value of the claim. My point is that it makes no difference to the fact that Priestley's manipulations were rightly taken by him as evidence that the phlogiston explanation was nonredundant. We now have collected much stronger evidence in favour of oxygen, which makes Priestley's own explanation redundant—but he could not have anticipated this.
20. I would nowadays emphasize this point even more strongly. The view defended in this paper is not intended to provide an inductive method of discovery for unobservable entities. What I call *experimental warrant* can only provide support for our antecedently held existential commitments. The aim of this paper's project is to start deflating Hacking's and Cartwright's views, by extracting the "realism" out of the "experimental realism" (Note added to proofs, January 2008).
21. I would like to thank the participants at the Konstanz conference (December 2002) and the students at my doctoral course at Complutense (2002–2003) for their reactions; and Hasok Chang, Christopher Hitchcock, Carl Hoefer, and Peter Lipton for their detailed comments and helpful suggestions. This chapter has been supported by projects BFF2002-01552 and HUM2005-07187-C03-01 of the Spanish Ministry of Education and Science.

REFERENCES

- Brown, H. et al. (1995) 'Bohm particles and their detection in the light of neutron interferometry', *Foundations of Physics*, 25: 329–347.
- Carnap, R. (1935) *Philosophy and Logical Syntax*, Bristol: Thoemmes Press.
- Cartwright, N. (1983) *How the Laws of Physics Lie*, Oxford: Oxford University Press.
- . Cartwright, N. (1999) *The Dappled World*, Cambridge: Cambridge University Press.
- . (2002) 'Summary' and 'Reply', *Philosophical Books*, XLIII: 241–243, 271–278.
- Cartwright, N. et al. (1995) 'The tool box of science', in *Theories and Models in Scientific Processes*, W. Herfel et al. (eds) Poznan Studies in the Philosophy of Science and the Humanities, Rodopi, 44: 137–149.
- Clarke, S. (2001) 'Defensible territory for entity realism', *British Journal for the Philosophy of Science*, 52: 701–722.
- Davies, M. (1998) 'Externalism, architecturalism, and epistemic warrant', in C. Wright et al. (eds) *Knowing Our Minds*, Oxford: Oxford University Press.
- Dorato, M. (1988) 'The world of worms and the quest for reality', *Dialectica*, 42: 3.
- Elsamahi, M. (1996) 'Could theoretical entities save realism?', *PSA 1994*, pp. 173–180.
- Fine, A. (1997) *The Shaky Game: Einstein, Realism and the Quantum Theory*, 2nd edn, Chicago: University of Chicago Press.
- . (1991) 'Piecemeal realism', *Philosophical Studies*, 61: 79–96.
- Gross, A. (1990) 'Reinventing certainty: The significance of Ian Hacking's realism', *PSA 1990*, 1: 421–431.

- Hacking, I. (1983), *Representing and Intervening*, Cambridge: Cambridge University Press.
- . (1984) ‘Experimentation and scientific realism’, in J. Leplin (ed.) *Scientific Realism*, Berkeley: University of California Press.
- . (1989) ‘Extragalactic reality: The case of gravitational lensing’, *Philosophy of Science* 56: 578.
- Hitchcock, C. (1992) ‘Causal explanation and scientific realism’, *Erkenntnis*, 37: 11–178.
- Hofer, C. (2003) ‘For fundamentalism’, *Philosophy of Science, PSA 2002*, 70.5, pp. 1401–1412.
- Lipton, P. (2002) ‘The reach of the law’, *Philosophical Books*, XLIII: 254–260.
- McKinsey, M. (1991) ‘Anti-individualism and privileged access’, *Analysis*, 51: 9–16.
- Morrison, M. (1990) ‘Theory, intervention and realism’, *Synthese*, 82: 1–22.
- Musgrave, A. (1976) ‘Why did oxygen supplant phlogiston: Research programmes in the chemical revolution’, in C. Howson (ed.) *Method and Appraisal in the Physical Sciences*, Cambridge: Cambridge University Press.
- Resnik, D. (1994) ‘Hacking’s experimental realism’, *Canadian Journal of Philosophy*, 24: 395–412.
- Reiner, R., and R. Pierson. (1995) ‘Hacking’s experimental realism: An untenable middle ground’, *Philosophy of Science*, 62: 60–69.
- Warwick, A. (1995) ‘The sturdy protestants of science: Larmor, Trouton and the motion of the earth through the ether”, in J. Z. Buchwald (ed.) *Scientific Practice: Theories and Stories of Doing Physics*, Chicago: University of Chicago Press.
- Wright, C. (2000) ‘Cogency and question-begging: Some reflections on McKinsey’s paradox and Putnam’s proof”, in *Philosophical Issues*, 10: 140–163.
- . (2003) ‘Some reflections on the acquisition of warrant by inference’, in S. Nuccetelli (ed.), *New Essays on Semantic Externalism and Self-Knowledge*, MIT Press, pp. 57–78.