INTERNATIONAL STUDIES IN THE PHILOSOPHY OF SCIENCE,
VOL. 17, NO. 3, 2003

Manipulative success and the unreal

AXEL GELFERT
Department of History and Philosophy of Science, University of Cambridge, Cambridge, UK

Abstract *In its original form due to Ian Hacking, entity realism postulates a criterion of manipulative success which replaces explanatory virtue as the criterion of justified scientific belief. The article analyses the foundations on which this postulate rests and identifies the conditions on which one can derive a form of entity realism from it. It then develops in detail an extensive class of counterexamples, drawing on the notion of quasi-particles in condensed matter physics. While the phenomena associated with quasi-particles pass the entity realist's criterion of manipulative success, quasi-particles themselves are illusions, and can be seen to be so even on the basis of the largely non-theoretical "home truths" that one must be ready to admit as background knowledge. Hence, Hacking's entity realism is shown to be incoherent.*

1. Introduction

Of the great number of philosophical positions in the debate about scientific realism, entity realism, in its original form due to Ian Hacking, has the catchiest battle cry: "If you can spray them, then they are real." The slogan is catchy because it seems a truism. Surely, whatever it is that we succeed in using or manipulating, it must be real. Or must it?

Following a brief analysis of the slogan, I shall sketch, in Section 2, the two standard lines of criticism to entity realism.[1] The first usually criticizes the scope of entity realism as too narrow, whereas the second questions the alleged primacy of experimental success over explanatory success. After this sketch, I shall lay the groundwork, in Sections 3 and 4, for a novel criticism of entity realism, namely that it is *too permissive*—that it would sanction belief in things that don't exist. In a nutshell, Hacking's entity-realist proposal can be regarded as an attempt to arrive at generalizations without theorizing. Continued manipulative success, so it is claimed, takes the place of explanatory success as the criterion for when one has hit on an alleged general feature of reality. In his appeal to our being causally connected to entities, the criterion of manipulative success builds on the causal theory of reference. However, as I shall attempt to show, for this to lead to a convincing form of realism about entities it is required that "real" entities satisfy determinate identity conditions: instances of token entities must be identifiable as causal relata in order to serve as a starting point for any experimental generalization to entity types. Where this is not the case there lurks the danger of reference failure.

This is precisely what occurs in a class of counterexamples to entity realism, to be described in Sections 4–6. In these cases, far from being too narrow, the criteria of entity

ISSN 0269-8595 print/ISSN 1469-9281 online/03/030245-19 © 2003 Inter-University Foundation
DOI: 10.1080/0269859032000169451

realism turn out to be too permissive: an entity-realist interpretation of science would end up, strangely enough, with more kinds of (supposedly) "real" entities than science itself. In particular, Sections 5 and 6 develop in some detail the notion of *quasi-particles*, which is deployed in many areas of condensed matter physics, and which, as I shall argue, subverts the entity realist's criterion of manipulative success in a novel way, by demonstrating that manipulative success is not sufficient as a guide to the reality of entities. Quasi-particles are no proper kinds of entities at all—they are merely collective effects of (an indeterminate number of) "real" entities, and they must be acknowledged as *illusory entities* even from a non- or anti-theoretical position such as entity realism.[2] Taken in conjunction with previously proposed counterexamples (to be summarized briefly in the next section) where the criterion is *too narrow*, our finding that the criterion is at the same time *too permissive* puts considerable pressure on the proponent of the criterion to demonstrate that, despite these limits, his position is still on defensible ground.

2. "If you can spray them, then they are real"

In *Representing and Intervening* (1983), Hacking recalls vividly how he witnessed a Stanford University experiment for the detection of fractional charges that convinced him of the reality of electrons and positrons. The experiment was based on Millikan's old idea that small charges can be detected by observing the movement of a macroscopic superconducting metal sphere in an electric field. The Stanford experiment required neutralizing any initial surplus charges present on the sphere. This charge neutralization was achieved by transmitting electrons and positrons onto the sphere. During this process of "spraying", the sphere's behaviour in a magnetic field changed—much like stripping oil droplets of an electron in the Millikan experiment altered their behaviour in a static electric field. The success in *using* electrons and positrons, thereby manipulating the behaviour of the sphere, should, according to Hacking, suffice to convince us of the reality of electrons and positrons: *"So far as I'm concerned, if you can spray them, then they are real"* (Hacking, 1983, p. 23). We come to acknowledge the existence of an entity (e.g. electron or positron) not by making it the primary object of observation, but rather by using it as a tool: "When we use entities as tools, as instruments of inquiry, we are entitled to regard them as real" (Hacking, 1989, p. 578). Electrons and positrons may, at some point, have been merely theoretical entities within a physical theory. But once they can be used "to manipulate other parts of nature in a systematic way," they have "ceased to be something hypothetical, something inferred" (Hacking, 1983, p. 262).

Previous criticism of Hacking's proposal has moved in two directions. Some authors, for example, Dudley Shapere, have questioned the range of applicability of the criterion of manipulative success, pointing out that "Hacking tends to equate (1) anything we cannot interfere with with … (2) anything we cannot 'use"' (Shapere, 1993, p. 146). Because of this, many non-manipulative uses are discounted by entity realism, and so are many well-established theoretical entities in astronomy, which—despite there being good reason to believe they are real—cannot be used for manipulation. Gravitational lenses are case in point. They appear to have striking observable effects on the light that reaches the telescope, but Hacking, on the basis of his criterion, says he is "very disinclined to say that we can observe the lens system", where "the lens system" refers to "an instance of a conjectured kind of entity, a gravitational lens" (Hacking, 1989, pp. 561 and 563). Shapere objects further that instances of "direct" observation, for example of solar neutrinos in large-scale underground detectors, are merely "legislated away" as "passive" and "noninterfering" (Shapere, 1993, p. 148).[3] In other words, Shapere urges us to acknowledge that there are many scientific entities that fall through the entity realist's net, but that we have good reason to believe in nonetheless. The criterion of manipulative success, so the conclusion goes, is too strict, and the resulting class of acceptable entities *too narrow*.

The second main criticism of Hacking's entity realism concerns the relation between (theoretical) explanation and (experimental) manipulation. Not only are many experiments heavily theory laden, but, as David Resnik argues:

> one cannot rationally claim to use a theoretical entity as a tool of inquiry without some evidence, or justification.... Hence, if one regards an entity as a tool of inquiry, one must also claim that its place in the world's causal structure explains some phenomena. (Resnik, 1994, p. 404)

There are two aspects to this objection. First, the contrast between theory and experiment may not be as great as the entity realist thinks. Recognizing instances of manipulative success may already require a substantial amount of theory: the behaviour of superconducting spheres in a magnetic field is hardly trivial, and knowing which variables are relevant and need to be monitored requires substantial background knowledge. Scientists, in their publications and other accounts of their work, frequently stress that experiments do not take place in a theoretical vacuum, that there is an interplay between manipulation and its place within theoretical frameworks. Scientists, one might say, are too clever to rely on bare manipulation, stripped of all explanatory, or otherwise theoretically informed, intuitions. Hacking, however, is adamant in his opposition towards explanatory inferences:

> Once upon a time the best reason for thinking that there are electrons might have been success in explanation.... Luckily we no longer have to pretend to infer from explanatory success (i.e. from what makes our minds feel good). (Hacking, 1983, p. 271)

Given this refusal to let theoretical and explanatory considerations dominate our picture of what is real, any proposed counterexample to entity realism must, in a sense, first beat the entity realist at his own game, before making the further claim that theory *ought to* constrain the criterion of manipulative success. We shall later, in Sections 5 and 6, encounter a class of counterexamples that achieves just that.

The second aspect of Resnik's objection concerns the question of when we are justified in claiming to use a theoretical entity as a tool of inquiry. The slogan "If you can spray them, then they are real", it turns out, is only superficially tautological. In an actual physical experiment one can at best hope to achieve individual instances of manipulation, whereas the reality of a theoretical entity turns on the existence, or non-existence, of a type, or kind, of entity. Only if the successful acts of manipulation are indeed due to token entities of the same kind, can we hope to correctly pick out a feature of the real world. "Experimental generalization" to entity types can succeed only in cases where what we spray are causally potent token entities that can be used as tools for experimental manipulation.

One could further argue that without some act of inference to the best explanation, observed instances of manipulation cannot count at all as evidence for our "access to unobservable entities". Even if, as Richard Reiner and Robert Pierson argue, such individual "experimenter's entities" are conceded—that is, the presence of unobservable causal factors in a given case—it would only be by another inference to the best explanation that one could establish warranted belief in the existence of exactly one corresponding kind of theoretical entity, “rather than two kinds or a thousand kinds” (Reiner & Pierson, 1995, p. 66). At both levels, failure is possible: one may fail to achieve prima facie manipulation in the first place, or one may wrongly ascribe the manipulative success to a kind that does not exist.

Entity realism is sometimes construed as believing in the existence of entities without believing in (any particular set of) statements about those entities. But this is a somewhat unfair and simplistic view. What entity realism denies is not that statements about entities can ever be true, and known to be true; rather, it denies that one can ever hope to be justified in believing in the truth of *theories* about these entities. Entity realism is anti-theory, not anti-truth. Steve Clarke puts this nicely:

> Entity realists have no objection to low-level generalisations about entities; in fact they rely on the viability of these to give sense to the descriptions of entities that they wish to endorse. (Clarke, 2001, p. 704)

Other low-level generalizations are required as background knowledge for setting up an experiment in the first place. Hacking (1983, p. 265) calls these “home truths”—truths that are implicit in the working knowledge of a skilled scientist, have stood the test of time, or may have themselves passed the criterion of manipulative success. They are a diverse lot and, as Margaret Morrison describes it, “do not constitute anything like the kinds of complex frameworks that are normally taken to be definitive of a theory” (Morrison, 1990, p. 1).

Truth, then, enters the entity realist’s picture in two important ways: first, via home truths, which are themselves well established and are a prerequisite for setting up an experiment, and second, when it comes to judging an act of manipulation successful or not. Entities have the role of truth-makers—not for theories, but for *judgements* concerning the success of acts of interfering in other parts of nature. This role of entities as truth-makers also explains why manipulation, at least at first sight, seems a useful guide to the reality of entities, for, in David Lewis’s words, “truth is supervenient on what things are and which perfectly natural properties and relations they instantiate” (Lewis, 1992, p. 216). And this is exactly what entity realism holds manipulation can probe.

3. Manipulative success and the possibility of reference failure

The criterion of manipulative success appeals to a causal theory of reference to bridge the gap between instances of token entities and types (or kinds) of entities: a dubbing event is successful if it picks out entities that are instantiations of the same kind. Such an “initial baptism” (Kripke, 1972, p. 96) may occur in different ways and is bound to be fallible: the beliefs we have about the things we refer to will typically be subject to later revision. However, if a dubbing event is successful, then it will give rise to a corresponding kind term, in the sense of “*that kind of thing*, where the kind can be identified by paradigmatic instances” (Kripke, 1972, p. 122).

Not all dubbing procedures are equally congenial to Hacking’s proposal. If, for example, one were to pick out an entity, such as the neutrino, by describing its place within an elaborate theoretical framework, this would not satisfy the entity realist. After all, the theoretical content scientists assign to terms such as *neutrino*, *phlogiston*, or *gene* in the light of later theories frequently turns out to be false. For this reason, it may seem unpromising to turn to theoretical descriptions as the arbiter of the reality of scientific entities. And indeed, the entity realist believes that we can do better than this. What he purports to demonstrate is that we can commit to the existence of an entity without adopting a body of theoretical expert knowledge—that we can stick to the referent throughout theoretical re-interpretations.

For the traditional scientific realist, the difficulty with theory changes lies in exactly such retrospective reference assignments. How can one secure reference across theory changes, when the referent is determined by the theoretical statements we can derive about it from our best current theory? As Putnam suggests, some form of "principle of benefit of the doubt" will have to come into play when dealing with retrospective reference assignments on the level of theory (Putnam, 1984, p. 143). The entity realist, by contrast, claims to have a more straightforward solution to the problem of the continuity of reference. What is important, he claims, is not the question of whether, or how, we can successfully interpret past theoretical expressions in terms of present theoretical expressions, but rather whether we are still dealing with the same sorts of things. This, the entity realist assures us, will be the case as long as continued manipulative success indicates that the initial dubbing procedure has indeed picked out token entities that are representative of an entity type that is itself a feature of the real world.

Manipulation, as the entity realist conceives of it, is to be regarded as one of the alternative "other possibilities of initial baptism" that Kripke speculated could stand alongside ostension and description (Kripke, 1972, p. 97). Its importance, however, goes beyond the initial phase of introducing new entities, in that manipulative success comes in degrees, can be improved in terms of precision, and can be extended in scope: "The more we can understand some of the causal powers of electrons, the more we can build devices that achieve well-understood effects in other parts of nature" (Hacking, 1983, p. 262). In other words, entity realism holds that manipulation succeeds in tracking entities better than other procedures such as theoretical description.

However, there seems to be good reason for caution. The philosophical literature is replete with examples of how descriptive theories of reference can lead to dubbing failures. When the police in Victorian London used "Jack the Ripper" to refer to whoever committed all (or most of) the notorious murders, nothing ruled out the possibility that there never was any such person. Perhaps a contamination of water supplies with psychoactive drugs had triggered murderous instincts in a large number of people, each of whom then committed exactly one of the murders. Under such circumstances, the term "Jack the Ripper" would have no referent, and dubbing by description would have failed. What guarantees that there cannot be similar failures in the case of manipulation, too? Common intuitions about the role of description in establishing reference have long since been proven wrong, and nothing guarantees that manipulation can be saved from sharing this fate. In fact, in Sections 5 and 6, we will encounter a whole class of just such dubbing failures based on illusory instances of manipulation.

4. Prerequisites for the criterion of manipulative success—and how to subvert it

For the purposes of scientific realism, to say that electrons are real is to say that they are *of a kind*. If one adopts Kripke's account of rigid designation, then the semantic function of natural kind terms consists in picking out the same referents in every possible world in which they have any reference at all. Essentialism—the view that for any given natural kind there are "properties which nothing can lack and still be of the kind" (Mellor, 1977, p. 299)—would achieve such "trans-world" reference, but it is not an option for the entity realist with his strong anti-theoretical leanings, and indeed Hacking openly distances himself from it (Hacking, 1983, p. 82). Essentialism typically requires an amount of theory that is unacceptable to the entity realist. Putnam, for example, when it comes to characterizing the rigidity of natural kind terms such as "water", explicates it as "whatever bears a certain equivalence relation … to the piece of liquid referred to as 'this' *in the actual world*" (Putnam, 1973, p. 708), where the equivalence relation is to be defined in theoretical terms by natural science (Cassam, 1986, p. 87). To the entity realist, such an appeal to a theoretical relation must seem suspect; it would go against his conviction that manipulative success, rather than any explanatory or theoretical desideratum, is the criterion for the reality of entities.

Given that entity realism is concerned with our knowledge of what is real, it must insist on criteria that tell us when we have good reason to believe that there actually are entities of a certain kind, not that, once their existence is established, there is a sense in which the corresponding theoretical term rigidly designates the same kind of entities in other possible worlds as well. While rigid designation is something the causal theory of reference can deliver, an answer to the question of whether we actually are empirically right about what there is in the world lies outside the causal theory's reach. It might not be inconsistent to stipulate that there is a difference in relative metaphysical status of manipulable, as opposed to merely theoretical, entities, but in order to recognize instances of manipulation, which is surely something the entity realist requires, it is essential that we have epistemic access to the world. Hence, manipulative success as an epistemic criterion is subject to the same constraints as other ways of acquiring knowledge about the world, and is a largely contingent matter. In short, it is far from clear that manipulative success by itself can bear the weight the entity realist assigns it.

What is indeed presupposed by entity realism, at least for natural (manipulable) kinds, is a commitment to the existence of a workable "mapping relation" between instances of entities and kinds of entities. This is so because, for the entity realist, the kind structure of the world as we know it is logically supervenient upon the instances of our interaction with it. If one could not re-identify instances of manipulation in a systematic way, or distinguish successful cases of manipulation (or observation) from failures in the first place, one could never hope to probe this relation between instances and kinds.

One might prefer to regard manipulability as "a mark of the real" (Nola, 2002, p. 5)—which some entities may, and others may fail to exhibit—rather than an infallible criterion, since in some cases, for example when it comes to gravitational lenses, one might simply want to withhold one's judgement. As Robert Nola puts it, "[m]anipulation is a success term (one cannot manipulate the unreal), and is a sufficient condition for realism; but it is not a necessary condition" (ibid.).[4] It would indeed be an achievement if entity realism succeeded in formulating such a sufficient condition, for it would mean that, though the criterion might miss certain kinds of entities, at least it could not be wrong about those cases where it gives a positive verdict. Alas, manipulative success alone does not suffice, as our discussion (from Section 6 onwards) of a large class of counterexamples will demonstrate. These will be cases of seemingly successful manipulation, which, despite their non-accidental nature, systematically miss their target—simply because the target is, in a sense, a mirage. The illusory nature of the target and the success in terms of experimental reproducibility can be explained only by appealing to a level of theory that goes beyond the entity realist's home truths. In other words, manipulative success is not only not sufficient, but it can even be misleading.

Having pointed out, at the beginning of this section, some of the metaphysical-*cum*-

epistemological prerequisites for using manipulative success as a criterion for realism about entities, let us consider how one can hope to subvert this criterion. We may ask: what are the desiderata for any class of "subversive" counterexamples to the criterion of manipulative success? First, if one wants to beat the entity realist at his own game, one should avoid postulating an explanatorily more fundamental level of analysis, as this will leave the entity realist quite unimpressed. As an example consider those realists about the quantum field, who would have us dispose of particles altogether and interpret them instead as excitations of a more basic real entity, namely the quantum field (Teller, 1995). On this account, particles—electrons for example—appear to evaporate into an ethereal notion of an all-pervasive substrate (the "quantum field"), which, so one might argue, calls into question the validity of "electron talk", and hence also of manipulative success using electrons. Would such a line of argument convince an entity realist to abandon his position? Hardly. For one, it is quite obvious that the quantum field picture of what particles are, is warranted only by a bold theoretical inference that the entity realist is unwilling to accept in the first place. Furthermore, to the extent that one may hold the quantum field to underlie the causal processes observed (and exploited) in the world, it does so via particles—at least on the level of experimental interaction with manipulable properties. Even for the most liberal of entity realists, the quantum field is, at best, one way of many to resolve the metaphysical underdetermination of quantum entities—and one resounding ad nauseam with explanatory overtones. Hence, introducing a higher-level theoretical redescription as in the case of the quantum field does not succeed in subverting the criterion of manipulative success. This is not to say that the entity realist's refusal to let theoretical considerations enter the picture should not be criticized (as indeed was done in Section 2), but merely that it strengthens the case of the critic, if he can grant the entity realist his anti-theoretical sentiments and still succeed in pointing out the incoherence of his position.

The second of our desiderata for proposed counterexamples is related to the first and concerns the way in which the entity realist's "experimental inference" from manipulative success (using token entities) to the reality of types of entities is supposed to occur. If the counterexample aims at demonstrating the possibility of reference failure despite apparent manipulative success, then this, too, should be made clear not on a highly theoretical level of analysis or formalism, but rather on the very level entity realism itself takes to be fundamental: namely that of experimentation and manipulation. More specifically, if the prospective counterexample involves the claim that, in the particular example at hand, manipulative success provides a wrong indication as to *what* is being used as a tool, the characterization of the mistake should not involve *more* theory than the interpretation one attempts to refute. Which propositions count as acceptable "home truths" and which don't, will of course depend on the level one starts from, and for illustrative purposes at any rate it should be admissible to draw on theoretical concepts—if they succeed in bringing out the experimental content more clearly.

5. What do we spray, when we spray a particle?

In view of entity realism's commitment to the primacy of the experiment, and of Hacking's sympathies for "small science" in particular, let us turn to an example that combines the two; namely, the transistor effect discovered by John Bardeen, William Brattain, and William Shockley in 1947.[5]

For a basic understanding of the transistor effect, nothing more is needed than a familiarity with some of the home truths concerning metals and semiconductors. The notion of home truth, here as before, should not be taken to require detailed theoretical accounts of how the behaviour of substances can be explained in terms of their specific microstructure; rather, a "home truth" must be of a kind that has itself passed the test of manipulability or, alternatively, is so basic that questions concerning its validity simply do not arise. Indeed, the properties and phenomena of semiconductors had been known, and experimentally exploited, for over a century since the 1830s before a qualitative explanation in terms of quantum theory became available through the work by Alan Wilson on energy bands in solids in 1931. Amongst the most widely discussed phenomena associated with semiconductor physics was the rectifying effect that occurred at a metal–semiconductor contact. When a voltage is applied to such a device (or, more cautiously, "proto-device"), the flow of current in one direction (the so-called forward current) is considerably higher than that in the reverse (so-called backwards) direction with opposite polarity. It does not require a high level of theoretical sophistication to conceive of the possibility of an arrangement of metal–semiconductor interfaces that would allow one to control the flow of current by manipulating the properties of one of the interfaces. One might, for example, think of changing the geometry of the interface, thereby choking the flow of charges through the contact.[6] Or one might influence the behaviour at the interface by applying a voltage across it. This idea, in the form of Shockley's field-effect amplifier, was in fact the historical starting point for the transistor. Such a device would exploit the fact that if the internal contact-potential field in a rectifier could produce a space-charge layer at the interface between a metal and a semiconductor, then an externally applied electric field, too, should create such a barrier region.

I shall, for illustrative purposes, summarize two quite different theoretical models that, at the time, were used in order to explain the processes on the metal–semiconductor interface. It must be emphasized that the entity realist cannot help himself to explanations of this kind, at least not for deciding which kinds of entities are involved in these processes. The two models do, however, provide a helpful background for the class of counterexamples to be discussed in Section 6.

Both models draw to some extent on the band theory of metals (Wilson, 1931), semiconductors and insulators, which states that electrons in a solid cannot occupy a continuum of energy levels. Instead they are confined to certain energy intervals, called "bands", which are separated by a gap. In this picture, semiconductors and insulators differ only in that the separation between the last band fully occupied by electrons (i.e. the valence band) and the first empty band to which electrons must be excited in order to contribute to conduction, is smaller for semiconductors than it is for insulators. It requires much more energy than is available at room temperature in an insulator to excite electrons into conducting states where they are mobile and can produce a current. In semiconductors room temperature is sufficient to excite a non-negligible number of electrons into such states. Metals, finally, are qualitatively different in that the highest occupied band is only partially filled, so that an infinitesimal energy already suffices to excite electrons into conducting states.

The quantum theory of solids also predicts that in addition to thermal excitation of electrons, impurities and imperfections within a crystal can contribute to conduction via donor and acceptor states. Such donor and acceptor states lead to excess electrons and spatially extended "holes" with a net positive charge, both of which can propagate through a crystal, thus carrying a current. Depending on which process is dominant,

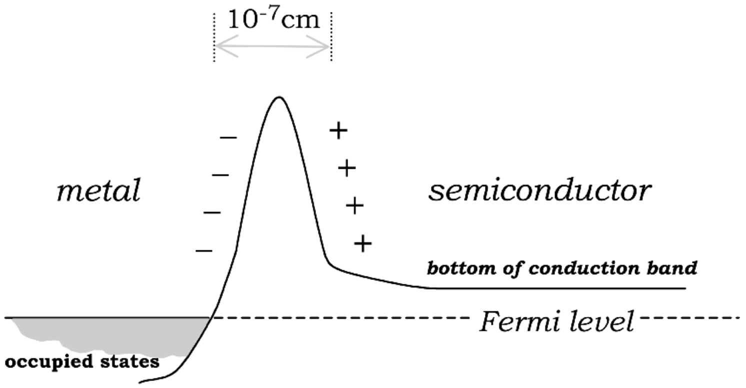

Figure 1. *The Wilson model. Electrons in the metal (left) are separated from the semiconductor side (right) by a symmetrical potential hump (after Hoddeson, 1981).*

either electrons or holes can be majority carriers, with the other one being the minority carrier.

Wilson himself proposed a model of rectification at metal–semiconductor contacts, in which the interface is pictured as a symmetrical potential hump separating the electrons in the metal from those in the semiconductor (Figure 1). Electrons can pass from the metal to the semiconductor (and vice versa) only via quantum-mechanical tunnelling. In order for this to happen, the hump must be extremely narrow ($\sim 10^{-7}$ cm). The rectifying effect arises as follows: applying a negative voltage on one side of the contact raises the energy level of the electrons there, thus increasing the probability of their participating in quantum-mechanical tunnelling to the other side. As far more electrons are available on the metal side, the current from the metal to the semiconductor should be much greater than that in the opposite direction. Experiments, however, showed the opposite behaviour; and when later experiments established a much greater width of the potential hump, the assumption that tunnelling was the transfer mechanism had to be dropped altogether.

An alternative model for the rectifying junction, the Mott–Schottky model, also makes use of this band theory of conduction in metals and semiconductors. However, the process of charge transport is treated in terms of classical mechanics and thermodynamics. In this model, the transfer mechanism is not based on quantum effects such as quantum-mechanical tunnelling through a barrier, but is based instead on thermal excitation over a barrier of potential. The directional preference that manifests itself in the rectifying behaviour is pictured as a consequence of a built-in asymmetry of the barrier: it is much steeper on one side than on the other (Figure 2). Electrons on the semiconductor side see a reduced, or, depending on the polarity, increased barrier, which they must overcome by thermal excitation. While the assumption of the asymmetry of the barrier is rather ad hoc, the Mott–Schottky model both explained the direction of rectification and gave a plausible estimate of the width of the barrier. Explanatory virtue, whether ad hoc or not, is of course nothing that impresses the entity realist. What counts for him is the experimental success in manipulating the behaviour of well-established other items of experimentation. Suffice it to say, therefore, that the Mott–Schottky model provided a heuristic agenda for the research that led to the transistor.

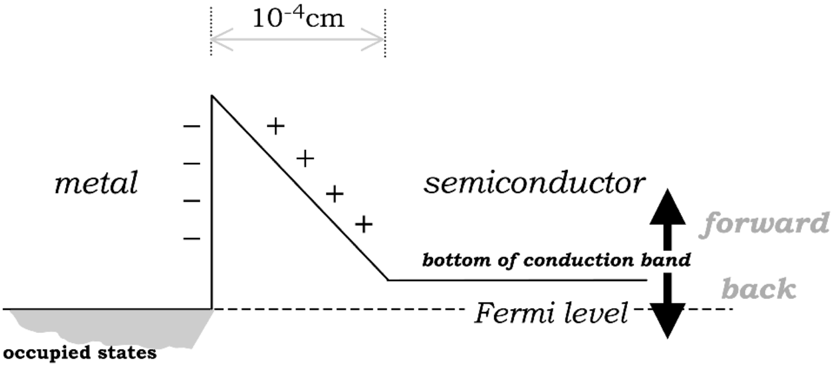

Figure 2. *The Mott–Schottky model. Metal and semiconductor sides are separated by an asymmetrical potential barrier. Depending on the polarity of the external field (forward/back), the effective barrier is decreased or increased (after Hoddeson, 1981, modified).*

The experimental attempts at constructing a field-effect amplifier, led by Brattain and Bardeen, underwent various changes and modifications—the kind of "debugging" procedure Hacking agrees is typical of experimental research: "We spend a lot of time building prototypes that do not work" (Hacking, 1982, p. 76). Various geometries were tried out, both with and without immersion in liquid electrolytes and dielectrics, and different possibilities to connect the input and output circuits with the semiconductor were explored. At Bardeen's suggestion, sharp point contacts were used, and when one of the two contacts, originally supposed to be arranged symmetrically with respect to an oxide film (which had formed as a result of etching the surface), was misplaced by accident, the device finally showed substantial amplification—the effect needed for developing functioning commercial applications.

Contrary to expectations, however, the desired effect showed up only when the polarity was opposite to that in the proposed set-up. This did not preclude *using* the contact in a systematic fashion, so as to manipulate the flow of current and amplify it significantly. The new type of behaviour—the transistor effect—was just as systematic and reliably manipulable as in the original experiments with metal–semiconductor junctions, but the effect went, so to speak, "the wrong way". The experiment behaved as though not electrons had been injected via the point contact, but particles with an opposite charge. The manipulative success was undeniable: the flow of a current, the amplification in proportion to the base voltage—all this is surely as suggestive as spraying a particle. Had Bardeen and Brattain stumbled upon a new entity? No, they hadn't. One of the home truths in this case is that the only real entities that move freely in a metal or semiconductor, or through an interface between the two, are electrons. But, as suggested in our earlier discussion of charge transport, in an ensemble of electrons there can be holes: areas of charge depletion that appear to function as charge carriers—bubbles, if you like, in a sea of electrons. Their positive charge is due, in an entirely derivative way, to the positively charged "seabed" of positive ions in the crystal lattice "shining through". Bardeen and Brattain's puzzling result was that positively charged holes, instead of the expected negatively charged electrons, had been injected into the sample due to the asymmetry arising from the misplaced point contact. "Holes have positive charge"—this is nothing more than a *façon de parler*, a way of talking about holes as entities. We already know—without having to resort to an act of inference to the best explanation, but merely by virtue of the home truths that were established before anyone ever set out to invent the transistor—that the origin of the positive charge lies in the positive ion cores that make up the crystal lattice of the metal (or semiconductor). However, as we shall see in the next section, a satisfactory explanation of the particle-like properties of holes—which is needed in order to resolve this puzzle of seemingly finding new entities where, as is clear from the home truths account, there cannot be any—requires an appeal to higher-level theory, which goes beyond the level of home truths and, hence, is unavailable to the entity realist.

Before turning to the physics of (electron) holes, let us briefly discuss some general metaphysical aspects of holes and absences. Holes, in the broadest sense, were (re-) discovered as a topic in metaphysics by David and Stephanie Lewis in a delightful dialogue they published in 1970, in which it is argued (at least by one of the fictitious discussants) that absences, quite generally, cannot be anything like things and, hence, cannot instantiate natural properties and relations the way things do (Lewis & Lewis, 1970).[7] But why not bite the bullet and accept absences, perhaps not as things, but as the relata of causal-manipulative interaction? Is that not all the entity realist needs? The difficulty with such a project of "reifying" absences is that it is not compatible with entity realism. The problem is not about whether or not the entity realist should be committed to reductionism. Atoms are entities, and are to be considered real by the entity realist's standards, just as much as their main constituents, electrons and nucleons, are, as well as the things they are part of—perhaps a strand of DNA in a cancer cell. By contrast, absences cannot be described, except superficially, in terms of any such relation of containment or "emergence". We may be able to superficially

> cook up ersatz absences to serve as relata of causal relation—though surely they will seem to be the wrong relata, since we don't really think of these ersatz absences as having the same effects (or causes) as the absences they stand in for. (Lewis, forthcoming)

Even those who, contra Lewis, defend absences and holes against attempts to "deontologize" them, concede that "unlike an actual entity or property" they can only be "causally *relevant* but not causally *operative*" (Martin, 1996, p. 64). This I take to be just another way of saying that holes are at best *fake* entities—though "sprayable" ones, as the transistor example demonstrates.

For "manipulation" to be a success term, entities must be truth-makers for positive existential statements. Absences—and a whole class of more complex "quasi-entities" to be discussed in the next section—do not satisfy this condition. Even if one is a non-reductionist about holes, as, for example, C. B. Martin is, one can at best hold that they are either truth-makers for negative existentials or false-makers for positive existentials (Martin, 1996, p. 58).

Lewis, in his discussion in "Void and Object", is right to say that "[t]he best response is to concede that a void is nothing at all, and that a lesser absence [e.g. a hole] is nothing relevant at all, and therefore cannot furnish causal relata" (Lewis, forthcoming). *Relata*, however, are what the entity realist needs, once he has discarded belief in *relations* as described by scientific theory. In order to be true to his basic commitments, one might say, an entity realist, too, had better keep his eye on the doughnut, not on the hole.[8]

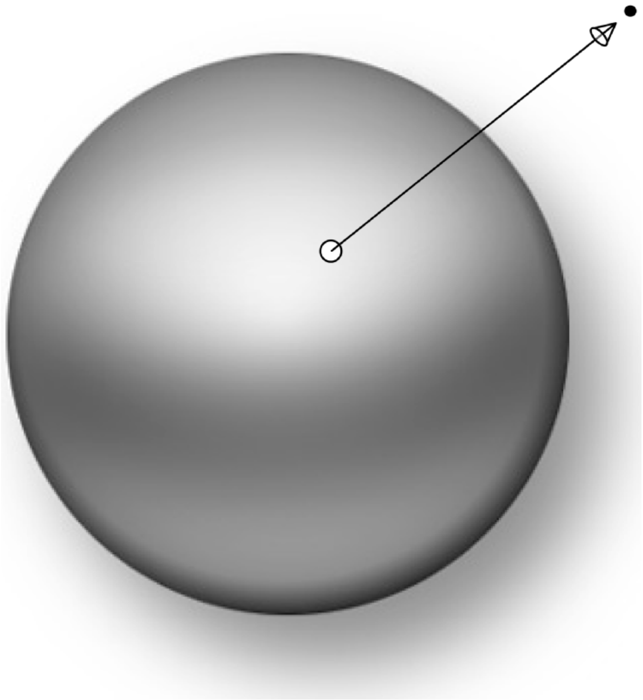

Figure 3. *Removing an electron leaves behind an electron hole in the Fermi sphere.*

6. Quasi-particles: spraying the unreal

The example of electron holes is by no means an anomaly. In fact, it is representative of a much larger class of physical phenomena to do with what physicists call *quasi-particles*. Before giving further examples of quasi-particles and how they can be used for manipulative purposes like their real counterparts—particles—let me illustrate some common features of quasi-particles by pursuing the example of electron holes a little further. Again, I shall draw on some amount of theory, and again this mainly serves illustrative purposes.

As mentioned earlier, electrons in solids occupy energy states within certain energy bands, and the filling level of the highest energy band will determine the conductive behaviour of the solid. The reason why bands "fill up" at all, and not all electrons settle in the same single-particle state with minimal energy, lies in the Pauli exclusion principle. It states that no more than two electrons occupy the same energy level. Hence, two electrons occupy the energy level with lowest single-particle energy; the next two will have to put up with a state of slightly higher kinetic energy, and so on, until all the 10^{23} or so particles are accommodated. The picture that is often invoked is that of a sphere filled with electrons, where the distance from the centre of the sphere is a measure of an electron's kinetic energy.[9] The ground state, then, is a sphere with a sharply defined surface, all electrons inside, and no electron outside. If one were to produce a hole by taking away an electron, it would leave a blank in the sphere—it would appear like a permanent positively charged particle with infinite lifetime (see Figure 3).

In the presence of interactions and fluctuations, when the system is at finite (non-zero) temperature, the sphere will be deformed and the sharp surface will become blurred due to thermal excitations that scatter electrons to states outside the sphere. Some electrons will now be outside the original sphere, leaving space for other particles on the inside of the sphere. If an electron were sprayed—excited to a high energy level, perhaps so high that it could escape the solid—it could recombine again by falling back into a lower energy level. In effect, excited electrons can lose some of their energy through interactions and recombine with positively charged holes. Manipulation of electrons, excitation to (and recombination from) high energy levels happens relative to the backdrop of the Fermi sphere. The collective behaviour of the electrons in a sharply defined Fermi sphere at zero temperature is causally inert—"what is important physically is the behavior of the lower excited states relative to the ground state" (Anderson, 1992, p. 99).

As the transistor example shows, these excitations are indistinguishable in terms of their manipulability from the actual particles. This is captured by the concept of *quasi-particles* in physics, which are nothing but these collective excitations of a many-particle system. What is important for the question whether or not quasi-particles should be regarded as entities in their own right, is the fact that their causal powers are real—but even on the comparatively low level of home truths acceptable to the entity realist, one knows that there are no causal powers other than those of the ensemble of $\sim 10^{23}$ electrons. Solids do not contain two kinds of mobile charge-carrying particles, holes and electrons; they *only* contain electrons. That electrons in a sample can produce a combined effect that makes it appear as if there were a separate kind of charge-carrying entity, ought, by the entity realist's own standards, not affect this home truth. If one were to grant quasi-particles the same degree of reality as electrons, one would violate the very intuitions that lie at the heart of entity realism; namely, that there is a set of basic substantive entities that have priority over composite or derivative phenomena. A proliferation of entities would evolve into what one might call *inflationary realism*—after all, there are various other quasi-particles that one would have to include along with holes.

Excitons are another example. Sometimes described in textbooks as deriving from a pairing of electrons and holes, this is in fact slightly misleading. For there is nothing more composite about excitons than there is about electron holes. Excitons do not consist of a pair of two different entities, they are collective effects of the many-electron system as a whole. Yet they display, in experiments and in their causal interaction with external perturbations of the system, precisely the features associated with bound particles. For example, when luminescence is induced, the formation of excitons shifts the spectral line to lower energy—something that can be accounted for in a single-particle picture by attributing a binding energy to the quasi-particle (Reynolds & Collins, 1981, p. 42). Hence, excitons can be used for spectroscopic purposes, and spectroscopic experiments can be performed on them, in much the same way as on atoms. It is perhaps not surprising, then, to come across a title such as *Excitons: Their Properties and Uses* in the research literature on quasi-particles (Reynolds & Collins, 1981). This rhetoric, of course, is not meant to make a statement about the reality of excitons as causally operative entities; all it indicates is "what it feels like" for a researcher to use excitons: you can spray them, and analyse their spectrum, *but they ain't real.*

This leads over to an aspect of quasi-particles that was hinted at in previous sections, namely that explaining new quasi-particles as collective excitations goes beyond the means available to the entity realist, in that it requires an appeal to theory

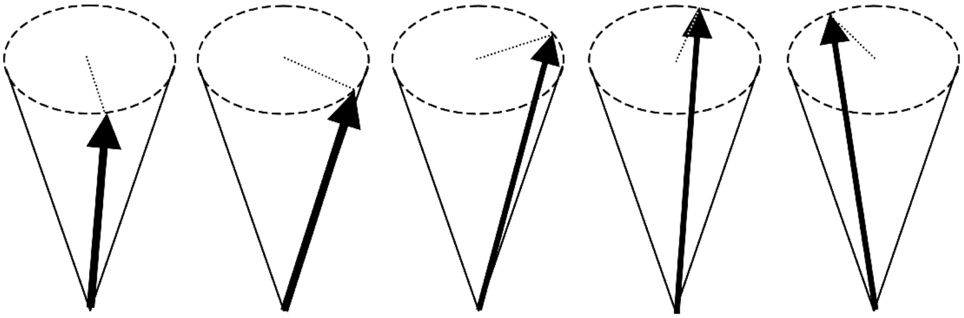

Figure 4. *Spin waves can be pictured as collective deviations from perfect ferromagnetic alignment of individual spins.*

beyond the home-truth level. This is best illustrated by looking at the way scientists themselves introduce different types of quasi-particles. The experimental observation of particle-like phenomena is only a first step in a series of progressively more theoretical moves. First it is recognized that the phenomena are indeed particle-*like*. After that, the range over which such particle-like behaviour is displayed in the many-body system, is experimentally tested or theoretically estimated, usually by drawing on some kind of established technique or theoretical model. Finally a mapping is constructed such that the original system (as described by its many-body Hamiltonian) is mapped onto a mathematically simpler theoretical description in terms of the newly introduced "fictitious" quasi-particles. One way of describing this procedure is to say that "the concept of elementary excitations is a way of *linearizing* the equations of the system about the true ground state rather than about some independent particle approximation" (Anderson, 1992, p. 102).[10] The properties attributed to quasi-particles, such as the effective mass of electron holes or the lifetime of excitons, turn out to be functions of the total many-body system. It is important to note that whereas the properties of quasi-particles can be exploited experimentally, the composite nature of their collective dependency is opaque to experimental methods. The experimenter who sprays a quasi-particle is, in a sense, blind not to manipulative success generally but to the nature of what he uses as a tool.

A particularly striking example of quasi-particles are *spin waves*. They involve excitations of a system consisting of a number of spins (e.g. a system of N electrons, each of which carries an intrinsic spin of 1/2) around the ground state. Consider a ferromagnetic ground state of minimal energy, where all the spins point in the same direction. A spin wave can then be visualized as a collective deviation from this perfectly ordered ferromagnetic state, which is distributed in a wave-like form over the whole system and reduces the total magnetization of the spin system by one unit. Loosely speaking, instead of flipping the spin of one particle ("reversing one arrow"), the spins ("arrows") of all particles begin to rotate like tops (see Figure 4). As a result, each of the spins deviates slightly from its maximum alignment—namely, in such a way that the z-component is reduced by $1/N$ units, giving the same total effect as the complete reversal of the spin of a single particle. The quasi-particle associated with the excitation of a spin wave that reduces the total magnetization of the sample by one unit is called a *magnon*.

In the magnon case the causal property of "spin" is detached from individual electrons and is associated instead with a new quasi-particle. However, in his discussion of Uhlenbeck and Goudsmit's original experiment, which demonstrated that a stream of electrons separates under the influence of an inhomogeneous magnetic field, Hacking claims (in accordance with his criterion of manipulative success) that an inference to a property such as "electron spin" as the best explanation for the experimental result is not sufficient to establish that *individual* electrons indeed have this intrinsic property: "The clincher is when we can put a spin on the electrons, polarize them and get them thereby to scatter in slightly different positions" (Hacking, 1983, p. 274). Do quasi-particles pass this test?

They do indeed. It is possible to selectively "put a spin on" quasi-particles, as experiments with epitaxial thin-film heterostructures confirm. This gives rise to a quasi-particle current that can be injected—*sprayed*, if you like—across barriers between films of different materials. Wei *et al.* (1999) refer to this process as "spin-polarized quasiparticle injection". Spin-polarized currents of this kind can be further exploited to interfere with macroscopic structures, such as magnetic domains in a metal or semiconductor. By employing a point-contact geometry similar to that of the early transistors, Myers *et al.* (1999) injected spin-polarized current densities into a thin-film structure and were able to "demonstrate how spin-transfer can be used to controllably reverse magnetic domains", thereby succeeding in "manipulating nanomagnets with a spin-polarized current" (Myers, 2000).

Other researchers have taken the transistor idea further and have implemented transistor-like devices that can amplify a current passing through them by a process called "quasi-particle trapping" (Booth, 1987): superconducting quasi-particles ("Cooper pairs") in the injector electrodes are transformed into another kind of quasi-particles ("excitation waves"), which in turn can push "electrons through the aluminium to a top electrode, where they are collected as a current" (Ball, 2000), thereby leading to a measurable net gain. The idea of a "quatratran" (Pepe *et al.*, 2000), a transistor entirely based on exploiting the properties of quasi-particles, symbolizes the degree of manipulative success that physicists attribute to quasi-particles—even though they are well aware of the fact that all they really deal with are collective phenomena of electrons and crystal lattices.

The examples given in this section demonstrate that quasi-particles are typically both manipulable and exploitable for purposes of interfering in other parts of nature. On the entity realist's criterion of manipulative success, they should therefore be considered real in the same sense that atoms or electrons are real. On the other hand, even on the non-theoretical body of home truths that the entity realist is forced to admit, we know that solids consist of crystals formed by atoms and of electrons travelling through the crystal, rather than of a plethora of emergent quasi-particles. This knowledge does not involve a theoretical inference; rather it represents a preservation of home truths. Obviously, there is a contradiction here between the inference from manipulative success and the established corpus of home truths, and given that home truths provide the very backdrop against which instances of manipulative success are to be evaluated, the contradiction cannot easily be brushed aside as irrelevant. Furthermore, it points to an asymmetry that follows from the denial of theoretical inferences. While it is true that no advanced theoretical considerations enter into the elementary finding that all causal effects in a solid are ultimately due to its constituents (of which there are only atoms and electrons), only theory can explain how quasi-particle behaviour in a solid comes about and what makes its occurrence stable. The entity realist, thus, is faced with the dilemma of either following the criterion of manipulative success and ending up with a permissive form of inflationary realism that violates basic home truths, or trying to build on these home truths while depriving himself of the very theoretical tools necessary for "explaining away" quasi-particles as collective many-body excitations.

Could the entity realist perhaps, in a desperate final attempt to save his position, claim that it is possible to bite the bullet and accept quasi-particles as new entities, on a par with the more entrenched entities such as electrons? I think not, as the entity realist would find himself in a very awkward position, for two reasons. First, since quasi-particles cannot exist independently of electrons, they would have to be defined as in some way parasitic on sets of electrons. However, no one set of electrons could be identified as the material basis of a particular quasi-particle. (This is precisely the reason why quasi-particles can be "sprayed" from one material sample to another without physically transforming one *into* the other!) Hence, the only solution would be to define as the referent the total ensemble of $\sim 10^{23}$ electrons. But this would be the very ensemble of particles that can support other quasi-particles (e.g. magnons in addition to excitons) as well. Hence, there would be a built-in mismatch between actual and intended referent—a mismatch that would be unavoidable in principle.

Second, when we speak of entities as "real", we want to say something criterial about their existence, not merely that in adopting a particular jargon we refer to an amorphous *je ne sais quoi*. If we were to call anything that we are ever causally in touch with an "entity", this would undermine the alleged privileged status of manipulation as ultimate arbiter over what should and should not count as real. After all, "observation", too, could be construed in such a way as to be synonymous with "whatever causally affects us".

In short, if the entity realist were to bite the bullet on quasi-particles, he would have to give up either on entities as we know them, or on being a realist. Neither seems to be a viable option.

7. Conclusion

I began my analysis of Hacking's proposal by pointing out that in order for it to be a convincing form of scientific realism, it must bridge the gap between encountering individual entities and claiming that they are *of a kind*. The statement "electrons are real" is more than a claim about past successful experiments: it is a claim about what the world is like. A key motivation for entity realism lies in the causal theory of reference. A dubbing event, which in the entity-realist picture is to be equated with an act of manipulation, is believed to fix the reference of an entity type, thereby allowing one to trace entities throughout theory changes. Typically, a dubbing event will be singular and (largely) theory-free; it involves an individual instantiation rather than a descriptive or theoretical account of an entity type. However, for a dubbing event to "stick" and in fact designate a "real" entity, it is required that instances be individuatable in the first place, so that one can reproducibly pick out the same referent over time. In the case of quasi-particles, manipulative success fails as a criterion for identifying the right dubbing conditions: when an experimenter "sprays" a quasi-particle, she exploits perfectly real causal properties—but not the causal properties of a new entity. Rather, the causal properties she exploits are those of a large (and possibly varying) indeterminate number of electrons behaving in a collective fashion. In other words, quasi-particles lack the determinate identity conditions required for making a dubbing event stick, since their causal properties are spread out over the total, and entirely contingent number of electrons in a sample. Hence, the criterion of manipulative success fails—not because the act of manipulation is *vacuous*, but because it is *promiscuous*.

Some critics of entity realism, such as Shapere, have argued that manipulative success is too strict a criterion in that it excludes entities such as gravitational lenses, for which there seems to be plenty of observational evidence. Other critics have called for a more liberal attitude towards alternative criteria of realism—one that gives proper heed to theoretical and explanatory considerations, which, in one way or another, inform most of scientific practice. What the present article shows is that, far from staying on the safe side and admitting only the best-established entities as real, the criterion of manipulative success in fact entails an inflationary form of realism with respect to a plethora of quasi-particles and collective effects. Not only would this violate the minimalist intuitions of entity realism; as discussed at the end of Section 6, it would also contradict the proposed relation between experimental research and home truths: if, on the basis of elementary experiments, it is clear that solids are made up of ions and electrons, how would one account for the sudden explosion of the number of entities?

It appears, then, that there is no easy way out for the proponent of the criterion of manipulative success. As is demonstrated by the quite general class of examples discussed in this article, "manipulation" is not a primitive notion and cannot bear the metaphysical-*cum*-epistemological weight attributed to it. Nancy Cartwright once struck a cautious note when she proposed a modification in emphasis of Hacking's slogan: "*When* you can spray them, then they are real" (Cartwright, 1999, p. 34). In view of our discussion there seems to be good reason for adding another caveat: if you don't know what you are spraying, you cannot tell whether it is real or not.

Acknowledgements

For helpful discussions and constructive remarks on earlier drafts, I would like to thank Peter Lipton (Cambridge) and Anjan Chakravartty (Toronto). Parts of this article were presented at the 28th Annual Philosophy of Science Conference at the Inter-University Centre in Dubrovnik in April 2002. I am grateful to those present for feedback and criticism, especially Robert Nola (Auckland) and Mauricio Suárez (Bristol). I would also like to thank two anonymous referees of this journal for their constructive comments.

Notes

1. Unless otherwise noted, in what follows, "entity realism" refers to Ian Hacking's proposal as summarized in this and the next section. As has recently been pointed out, there may well be defensible ground for other forms of entity realism (Clarke, 2001). The relation between the present argument and Nancy Cartwright's quite different form of entity realism (Cartwright, 1983) will be the topic of a future article.
2. That the number of "real" entities involved in the creation of quasi-particles is indeterminate is due to correlation effects. The dependency between quasi-particles (as collective excitations) and particles (as "constituents") is one of correlation rather than containment. Both aspects can be successfully explained only in quantum theoretical terms and indicate that quasi-particles are importantly, and relevantly, different from ordinary composites.
3. Shapere, in fact, goes further than this and argues that by means of the detection of neutrinos, one can observe, in a straightforward way, processes at the core of the sun, where the neutrinos originate (Shapere, 1982).
4. I would like to thank Robert Nola and Mauricio Suárez for a helpful discussion regarding this point.
5. For an historical account see Hoddeson (1981).
6. An uncharitable reading of entity realism might construe the first half of this sentence as a counterexample to Hacking: surely there are many external constraints, such as the geometry of a sample, that can be "used" for manipulation. Similar cases could be made for "manipulating" fictitious quantities such as

averages—for example, the centre of mass of a system. I grant Hacking that such claims miss the point somewhat. Conversely, however, when I speak of "the flow of charges" in the second half of this statement, I take this not to be unduly theoretical talk. That conduction involves electrons I take to be a home truth.

7. See also Casati and Varzi (1994).
8. This is Martin's caricature of Lewis's position.
9. This sphere, called Fermi sphere, is to be taken literally only in reciprocal (i.e. momentum, or wave-vector) space, not in physical space; hence, electron holes are not even properly localizable as "ordinary" absences would be.
10. For the purposes of this article, the terms "quasi-particle", "collective excitation" and "elementary excitation" can be used largely interchangeably. Hence the term "elementary excitation" in the quote from Anderson (1992) can be read as "collective excitation" without any relevant change in meaning.

References

ANDERSON, P.W. (1992) *Concepts in Solids: Lectures on the Theory of Solids* (Menlo Park, CA, Addison-Wesley).

BALL, P. (2000) Cool new gadget, *Nature Science Update*, online edition at http://www.nature.com/nsu/000803/000803–2.html (28 July).

BOOTH, N.E. (1987) Quasi-particle trapping and the quasi-particle multiplier, *Applied Physics Letters*, 50, pp. 293–295.

CARTWRIGHT, N. (1983) *How the Laws of Physics Lie* (Oxford, Oxford University Press).

CARTWRIGHT, N. (1999) *The Dappled World: A Study of the Boundaries of Science* (Cambridge, Cambridge University Press).

CASATI, R. & VARZI, A.C. (1994) *Holes and Other Superficialities* (Cambridge, MA, MIT Press).

CASSAM, Q. (1986) Science and essence, *Philosophy*, 61, pp. 95–107.

CLARKE, S. (2001) Defensible territory for entity realism, *British Journal for the Philosophy of Science*, 52, pp. 701–722.

HACKING, I. (1982) Experimentation and scientific realism, *Philosophical Topics*, 13, pp. 71–87.

HACKING, I. (1983) *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science* (Cambridge, Cambridge University Press).

HACKING, I. (1989) Extragalactic reality: the case of gravitational lensing, *Philosophy of Science*, 56, pp. 555–581.

HODDESON, L. (1981) The discovery of the point-contact transistor, *Historical Studies in the Physical Sciences*, 12, pp. 41–76.

KRIPKE, S. (1972) *Naming and Necessity* (Cambridge, MA, Harvard University Press).

LEWIS, D. (1992) Critical notice of D.M. Armstrong, *A Combinatorial Theory of Possibility, Australasian Journal of Philosophy*, 70, pp. 211–224.

LEWIS, D. (forthcoming) Void and object, in: J. COLLINS, N. HALL & L. PAUL (Eds) *Causation and Counterfactuals* (Cambridge, MA, MIT Press).

LEWIS, D. & LEWIS, S. (1970) Holes, *Australasian Journal of Philosophy*, 48, pp. 206–212.

MARTIN, C.B. (1996) How it is: entities, absences, and voids, *Australasian Journal of Philosophy*, 74, pp. 57–65.

MELLOR, D.H. (1977) Natural kinds, *British Journal for the Philosophy of Science*, 28, pp. 299–312.

MORRISON, M. (1990) Theory, intervention and realism, *Synthese*, 82, pp. 1–22.

MYERS, E.B. (2000) Manipulating nanomagnets with a spin-polarized current, Presentation C7.001, March Meeting of the American Physical Society, Minneapolis, MN.

MYERS, E.B., RALPH, D.C., KATINE, J.A., LOUIE, R.N. & BUHRMAN, R.A. (1999) Current-induced switching of domains in magnetic multilayer devices, *Science*, 285, pp. 867–870.

NOLA, ROBERT (2002) Realism through manipulation, and by hypothesis, in: S.P. CLARKE & T.D. LYONS (Eds) *Recent Themes in the Philosophy of Science* (Dordrecht, Kluwer), pp. 1–23.

PEPE, G.P., AMMENDOLA, G., PELUSO, G., BARONE, A., PARLATO, L., ESPOSITO, E., MONACO, R. & BOOTH, N.E. (2000) Superconducting device with transistor-like properties including large current amplification, *Applied Physics Letters*, 77, pp. 447–449.

PUTNAM, H. (1973) Meaning and reference, *Journal of Philosophy*, 70, pp. 699–711.

PUTNAM, H. (1984) What is realism?, in: J. LEPLIN (Ed.) *Scientific Realism* (Berkeley, University of California Press), pp. 140–153.

REINER, R. & PIERSON, R. (1995) Hacking's experimental realism: an untenable middle ground, *Philosophy of Science*, 62, pp. 60–69.

RESNIK, D.B. (1994) Hacking's experimental realism, *Canadian Journal of Philosophy*, 24, pp. 395–412.

Reynolds, D.C. & Collins, T.C. (1981) *Excitons: Their Properties and Uses* (New York, Academic Press).

Shapere, D. (1982) The concept of observation in science and philosophy, *Philosophy of Science*, 49, pp. 485–525.

Shapere, D. (1993) Discussion: astronomy and antirealism, *Philosophy of Science*, 60, pp. 134–150.

Teller, P. (1995) *An Interpretive Introduction to Quantum Field Theory* (Princeton, NJ, Princeton University Press).

Wei, J.Y.T., Yeh, N.-C., Fu, C.C. & Vasquez, R.P. (1999) Tunneling spectroscopy study of spin-polarized quasiparticle injection effects in cuprate/manganite heterostructures, *Journal of Applied Physics*, 85, pp. 5350–5352.

Wilson, A.H. (1931) The theory of electronic semiconductors. Parts 1 and 2, *Proceedings of the Royal Society of London A*, 133, pp. 458–491; 134, pp. 277–287.

Note on contributor

Axel Gelfert is in the Department of History and Philosophy of Science, University of Cambridge. Before turning to philosophy of science, he completed a degree in theoretical physics at the Humboldt University, Berlin. He has published articles on many-body physics, including a survey of the theory of phase transitions for the *Journal of Physics (Condensed Matter)*, on the history and epistemology of risk, and on his current area of research, the epistemology of testimony. *Correspondence*: Department of History and Philosophy of Science, University of Cambridge, Free School Lane, Cambridge CB2 3RH, UK. E-mail: axel@gelfert.net