all the entities in a confirmed theory receive such support. In short, holism blocks the withdrawal of the provisional support supplied by naturalism. And that gives us the first premise of the Quine/Putnam indispensability argument.

In the next chapter I will continue the defence of Quinean naturalism. In particular, I will demonstrate some of the serious problems that face the Eleatic Principle. (As we have seen in this chapter, this principle is the central tenet of Quinean naturalism's main rival: the causal version of naturalism.) There are three aspects to this strategy; two relevant to the task of motivating Quinean naturalism and one that's more general. First, discrediting the Eleatic Principle undermines much of the plausibility of the causal version of naturalism, leaving Quinean naturalism looking all the more attractive. Second, I will show that when one considers how the Eleatic Principle might be defended against some of the objections I raise for it, one is drawn toward a position not unlike Quinean naturalism. This again helps fortify the Quinean position. Finally, if the Eleatic Principle were tenable, it would present a quite general problem for mathematical realists, independent of the brand of naturalism subscribed to. Undermining the Eleatic Principle thus clears the way for Platonist philosophies of mathematics.

3

The Eleatic Principle

In the last chapter I identified the crucial difference between Quinean naturalism and Armstrong's causal version of naturalism. The difference, I argued, is that the latter subscribes to the Eleatic Principle, while the former does not. In this chapter I will continue the defence of Quinean naturalism. In particular, I will highlight some problems faced by the Eleatic Principle and suggest that Quinean naturalism is better equipped to deal with these difficulties.

The Eleatic Principle or causal criterion, you will recall, is a causal test that entities must pass in order to gain admission to some philosophers' ontologies. This principle justifies belief in only those entities to which causal power can be attributed, that is, to those entities that can bring about changes in the world. The idea of such a test is rather important in modern ontology, since it is neither without intuitive appeal nor without influential supporters. Its sympathisers (if not supporters) have included David Armstrong (1978, Vol. 2, p. 5), Brian Ellis (1990, p. 22), and Hartry Field (1989, p. 68), to name but a few.

Clearly this principle requires some justification. In this chapter I will look at the arguments that have been put forward for such a principle and suggest some problems for each of these. Of course in such a survey I cannot provide a decisive refutation of the principle, but I do hope to show that, despite its intuitive appeal, the Eleatic Principle's main justifications either look ad hoc or do not justify a version of the principle that delivers the intuitively correct results about some fairly uncontroversial cases. This is not an entirely negative result though. Once we look at the shortcomings of the motivations for the Eleatic Principle, a more general principle suggests itself. This more general principle looks very much like Quine's thesis that we are ontologically committed to all and only the entities that are indispensable to our current best scientific theories.

Recall that I am not intending to enter into the realist/anti-realist debate in this work. I assume realism is true but that some criterion is needed to distinguish real entities from fictional ones. It should also be stressed that any criterion put forward for such a task is a *criterion of acceptance*. As Keith Campbell puts it (emphasis in original):

> This search for a criterion for the real must be understood as a search for a criterion *for us to count something as real*; it will be a principle to apply in determining whether to accord that status, given our current stage of epistemic development. There need not be, and probably cannot be, any critical mark of the real itself; the real is what is, period. (Campbell, 1994, p. 28)

Also a word or two about a precise formulation of the Eleatic Principle. Graham Oddie (1982) attacks the Eleatic Principle by systematically questioning a number of specific formulations of it. Although he does this with considerable success, I will not follow him down this path since my criticisms of the Eleatic Principle, for the most part, will not depend on any particular formulation. What I take issue with is the motivation for *any* formulation of the principle. In any case, it may be useful to specify a particular version, just by way of example. I suggest the following is as good as any (although I stress that nothing hangs on this particular statement of the principle):[1]

Principle 1 (The Eleatic Principle) *An entity is to be counted as real if and only if it is capable of participating in causal processes.*

In one direction principle 1 is reasonably uncontroversial. Most realists agree that causal activity is a *sufficient* condition for an entity to be counted as real.[2] (It is more controversial that entities *capable* of participating in causal processes ought to be counted as real [as Principle 1 asserts], for it might be argued that Pegasus would pass such a test.) The important question, for our purposes, is whether causal activity is a *necessary* condition for an entity to be counted as real. In what follows I will argue that it is not.

[1] At this stage I will remain vague about what participation in a causal process amounts to. I will discuss this matter further in section 3.4.

[2] Ian Hacking (1983) argues for this claim by applying his interventionist test: Those entities that can be manipulated as tools in scientific investigations, as opposed to those that are merely tested for, are to be granted real status.

3.1 The Inductive Argument

The first motivation for a causal criterion I will consider is an inductive argument from uncontroversial cases of real entities. We start by noting that there are some fairly widely held intuitions (among realists at least) about roughly where the demarcation between the real and the instrumental should be. It should include physical objects, including theoretical entities, perhaps fields and hence waves as disturbances in these fields, among the real entities, but should not include (concrete) possible worlds[3] and frictionless planes. We thus see an initial motivation for the causal criterion: All the things that we intuitively think of as real are the sorts of things that participate in causal processes (in this world at least), whereas those that we intuitively think of as unreal do not participate in such processes.

Thus motivated, the Eleatic Principle is an inductive hypothesis about the way the world is. We look at the things in the world that are uncontroversially real, such as tables and chairs, and notice that they are all causally active. Then, by induction, we conclude that all real entities are causally active. At first glance a causal criterion thus motivated looks as if it is purely descriptive—it lacks the normative force that a criterion of acceptance, such as principle 1, ought to have. This defect, however, is easily rectified by appeal to naturalism (of almost any variety), which does make substantial normative claims about what we ought to believe. For example, if all evidence suggested that all real entities are causally active (a purely descriptive claim), then naturalism commands us to *believe* that all real entities are causally active (the corresponding normative claim). It is important to bear in mind that something like naturalism is required to get from purely descriptive theses to normative ones. In any case, since we are presupposing naturalism of some form or other, I will not dwell upon the descriptive–normative distinction.

There is a more substantive objection to this motivation though, which is that there are many other properties that the uncontroversially real entities share and that the uncontroversially unreal entities lack. All the uncontroversially real entities are spatio-temporally located, for instance. Indeed, there is argument on this very issue between protagonists of a strictly causal criterion, such as Brian Ellis, and the likes of Hartry Field, who require causal efficacy *or* spatio-temporal location. The reality of space-time points hangs crucially on whether it's causal efficacy or spatio-temporal location that is the important property. Similarly, we could opt for the property of having a positive rest mass as the mark of the real and

[3] Even David Lewis (1986b, p. 135) grants that modal realism is counterintuitive, which is all I'm claiming here.

again the demarcation would be different.

Worse still, it seems that such an inductive argument is going to depend on what our set of uncontroversially real entities is taken to be. For instance, if we decide to be fairly cautious about selecting the members of this set, and only admit medium-sized solid objects, we might conclude that all real entities are coloured. In light of these considerations, how do we come to decide to pin our hopes on causal efficacy and not on other properties?

One possible answer to this question is to argue that it is by virtue of an entity's causal efficacy that we have epistemic access to it, whereas other properties don't force themselves on us like this.[4] In any case, it seems fairly natural to suppose that we have epistemic access to any entity that we are to count as real. I am thus suggesting that it looks as though it is epistemic access that is in fact the prime motivation for the Eleatic Principle and not the inductive argument at all. At the very least, the inductive argument needs supplementing and this epistemic argument seems to fit the bill.

3.2 The Epistemic Argument

The epistemic considerations are perhaps the most common motivation for the Eleatic Principle. The argument is simply that even if there were causally idle entities, we would have no reason to believe in them, since their causal idleness would ensure that they didn't causally interact with us.[5]

A little consideration, however, reveals that this motivation is also somewhat thin, since the Eleatic Principle, depending on exactly how we formulate it, will require either that entities are causally active or that they are causally active with humans. The latter alternative, I suggest, looks far too anthropocentric to warrant serious consideration, for surely we ought to believe in stars and planets outside our own light cone[6] even though they are not causally interactive with us. To deny the existence of such entities is to effectively believe that the earth is the centre of the universe. This leaves only the "causally active (not necessarily with us)" formulation. On this reading, though, the epistemological motivation is lost altogether, for there may be many perfectly legitimate real entities involved in causal

[4] Recall that we seek "a criterion for us to count something as real" (Campbell, 1994, p. 28), so we may well admit that there are other properties that all the uncontroversially real entities share, but if we don't (or can't) have epistemic access to these entities by virtue of these other properties, then we can hardly use such properties as criteria of existence.

[5] This argument is, of course, due to Paul Benacerraf (1973).

[6] That part of the universe close enough to us so that there has been sufficient time since the Big Bang for light to reach us from it.

networks, but because they are not causally interactive with us, they suffer the same epistemic worries as causally idle entities. Again stars, planets, and so on outside our light cone are the prime examples here. That is, the set of entities whose existence may be motivated by this epistemic concern form a proper subset of the set of causally active entities.

In fact, the Eleatic Principle motivated by epistemic concerns seems to suffer all the same worries that the causal theory of knowledge suffers.[7] In particular, it seems that we have no reason to believe in future objects (whether causally active or not), and even universal empirical facts obtained by induction are likely to be problematic. Colin Cheyne (1998), who defends an epistemically motivated causal criterion, suggests the rather promising move to *kinds* of entities in order to overcome some of these problems. He suggests the following principle:[8]

Principle 2 *We cannot know that F's exist unless our belief in their existence is caused by at least one event in which an F participates.*

He argues for this principle based mainly on evidence from scientific practice—what it takes to convince scientists of the existence of a new type of entity. For example, he cites the discovery of the planet Neptune as evidence for principle 2. The existence of Neptune was predicted in 1845–6 jointly by Leverrier and Adams based on Newtonian gravitational theory and anomalies in the orbit of Uranus. It seems that the scientists of the time were unwilling to acknowledge the existence of Neptune until Galle first directly observed the planet in 1846. Cheyne claims that the moral to be drawn from such episodes in the history of science is "interacting is knowing" (1998, p. 40).

First note that this example is not entirely appropriate, since we certainly had causal contact with Neptune prior to Galle's visual contact, as Cheyne admits—we had indirect causal contact via its disturbance on the orbit of Uranus. Furthermore, a new planet may not qualify as a new kind of entity since it is of the same kind (namely, planet) as the earth.[9] In any case, leaving these points aside, there is another moral to be drawn from this and the other of Cheyne's examples. The moral I draw is "don't settle

[7] See Steiner (1975) for some of these.

[8] Cheyne offers this principle as a "first attempt" (p. 38) and later refines it to meet some objections. The objections I have to the principle, though, are not deflected by his later modifications, so I shall be content to deal with his first statement of the principle. See also Beall (2001) for related criticisms of Cheyne's principle.

[9] This point illustrates a difficulty with the move to kinds of entities—the difficulty of deciding whether a new entity is of a different kind to other entities already accepted. I will not pursue this point though, because I think Cheyne's approach has other, more serious problems.

for indirect evidence if you can do better." Clearly in the Neptune example direct visual evidence was better than orbital disturbances of Uranus, and scientists sought this better evidence because it was possible. Contrast this with the announcement by Wolszczan and Frail in 1992 of one of the first discoveries of planets outside our own solar system. These planets were detected because of the effects they were having on a nearby pulsar, PSR1257 + 12. Here visual contact was out of the question due to their distance from earth, and yet it seems there was no reservation on behalf of the discoverers about making full-blooded existence claims:

> [W]e demonstrate that ... the pulsar is orbited by two or more planet-sized bodies. The planets detected so far have masses of at least 2.8 $M_\oplus$ and 3.4 $M_\oplus$, where $M_\oplus$ is the mass of the Earth. Their respective distances from the pulsar are .47AU and .36AU, and they move in almost circular orbits with periods of 98.2 and 66.6 days. (Wolszczan and Frail, 1992, pp. 145–147)

Perhaps a better example is that of the discovery of the element germanium. In 1871 there had been no (known) causal contact with this element, and, in fact, causal contact wouldn't come until Winckler isolated the metal in 1887. However, because of the "gap" in Mendeleeff's periodic table corresponding to the position of germanium, much was known of its chemical behaviour. Cheyne claims that

> [i]f [Mendeleeff] believed, prior to 1887, in the existence of germanium, that belief, although true, would not count as knowledge. It could only be a lucky guess, unless it was actually caused, in an appropriate way, by events in which germanium atoms participated. (Cheyne, 1998, p. 36)

Even if I were to agree that if Mendeleeff believed in the existence of germanium prior to 1887, it would not count as knowledge, it seems extremely harsh to call such a belief "a lucky guess." After all, it's not as though he would have had *no* reason to believe that germanium existed, for there was surely reason to believe that something ought to fill the relevant gap in the periodic table. (Perhaps because of some argument from past predictive success of the table or some appeal to symmetry.) It's not a lucky guess in the same sense that a lottery winner guesses the winning numbers. I claim that, at the very least, Mendeleeff would have had *reason to believe* in germanium prior to 1887. This is all we are interested in here. If you accept this claim, then this example, while perhaps not a counterexample to the causal theory of knowledge, looks like a counterexample to an epistemically motivated Eleatic Principle (such as one motivated by principle 2), since Mendeleeff had justified belief in a novel substance without

the causal contact that principle 2 requires.[10]

It seems clear that appeal to the causal theory of knowledge (which is, after all, what is at the bottom of the epistemic justification), and all its notorious difficulties, is not the right approach for a justification of a causal restriction on ontological commitment. For one thing, an Eleatic Principle thus justified leaves out too many uncontroversially real entities (stars and planets outside our light cone), and second, even if one were to accept the causal theory of knowledge, there is no reason to insist that such acceptance implies that causal contact is necessary for *justified belief*, as the last example illustrates.[11] Furthermore, Cheyne's move to *kinds* of entities does not save the Eleatic Principle from these objections.[12]

3.3 The Argument from Causal Explanation

In this and the next section I will address what I take to be the most important argument for the Eleatic Principle. This argument, in its most compelling form at least, is due to David Armstrong. Armstrong has defended the Eleatic Principle in various places. For example, he proposes the following dilemma. "Are these [abstract] entities capable of *acting upon particulars*, or are they not," he asks (1978, Vol. 1, p. 128). He then raises difficulties for the first horn of the dilemma, since typically causation involves one change bringing about another, and yet here we have unchanging abstract entities presumably bringing about changes by "some sort of steady, unchanging, pressure" (1978, Vol. 1, p. 129). He concludes that "[s]uch a notion is perhaps barely possible, but it is impossible to see how such alleged causal operation could ever be identified" (1978, Vol. 1, p. 129). The other horn of the dilemma is simply that "Occam's razor ... enjoins us not to postulate them" presumably because causally idle entities have no role to play in science (1978, Vol. 1, p. 130). As Graham Oddie puts it (on behalf of Armstrong):

[10] Of course, germanium is *capable* of participating in causal processes, so the Eleatic Principle as I set it out in principle 1 would not rule out belief in germanium prior to 1887. This just serves to highlight the gulf between the Eleatic Principle and an epistemic motivation for it.

[11] It might even be reasonable to argue that many of our basic justified beliefs require causal contact, but this does not mean that inferential beliefs require causal contact. For example, if I am justified in believing P (because I have had causal contact with the truthmaker of P, say) and I am justified in believing that P implies Q, then surely I am justified in believing Q, whether or not I have causal contact with Q's truthmaker. I am indebted to Peter Forrest and John Bigelow for this point.

[12] In fairness to Cheyne, though, he is interested in a causal criterion of *existence*, whereas I am interested in a causal criterion of *justified belief*.

> They are 'causally idle' and hence 'idle'. Respectable entities work for their living, and there is no social security in Armstrong's universe. (Oddie, 1982, pp. 285–286)

Armstrong has more to say about this second horn of the dilemma:

> To postulate entities which lie beyond our world of space and time is, in general, to make a speculative, uncertain, postulation. The postulation may perhaps be defended if it can be presented as *explaining* some or all of the spatio-temporal phenomena. But if the entities postulated lie beyond our world, and in addition have no causal or nomic connections with it, then the postulation has no explanatory value. Hence (a further step of course) we ought to deny the existence of such entities. (Armstrong, 1989, pp. 7–8)

Here Armstrong explicitly cites (and emphasises) the role entities play in explanations as the key to justifying the Eleatic Principle.[13] In a way this is not at all surprising, since explanation often plays a key role in justifying scientific realism.[14] It would seem only fitting that explanation should also provide some guide as to the extent of that realism.

While I agree with Armstrong's rejection of the atypical causal action of the first horn of the dilemma, I don't think that the postulation of causally idle entities has no explanatory value. If this were true, then all genuine explanations in science would have to make essential reference to causally active entities. That is, all scientific explanations would be fully causal explanations, but this is not the case.[15] There are many instances of causally idle entities playing important explanatory roles in scientific theories and I will give a few examples of such cases.

Before I proceed to the examples of explanations in science featuring non-causal entities, I should mention that I do not wish to presuppose any particular model of explanation,[16] nor, for that matter, any particular model of causation. For obvious reasons I wish to be as broad-minded as

[13] He also explicitly appeals to explanation in Armstrong (1980a, pp. 154–155).

[14] For instance, J. J. C. Smart's (1963, p. 39) cosmic coincidence argument which, very crudely, is that the best explanation for the world behaving *as if* there were theoretical entities, is that the entities in question actually exist (if not it would be a remarkable coincidence).

[15] In the discussion following a presentation of material from this chapter, David Armstrong denied that his argument for the Eleatic Principle relies on all explanation being causal (despite the textual evidence I present here). Be that as it may, the argument I present in this section, whether or not it is the argument Armstrong intended, has considerable plausibility and is worthy of careful attention. For convenience I will continue to attribute the argument (as I read it) to Armstrong.

[16] Although I am inclined toward Philip Kitcher's (1981) account of explanation as unification.

possible about both causation and explanation. I will assume only that an explanation must be enlightening—it must make the phenomena being explained *less mysterious*. If it does not fulfil this minimal requirement, I suggest that it has no right to be called "an explanation" at all. As for the nature of causation, here I need only appeal to some widely held intuitions about what sorts of entities are taken to be causally idle (for example, geometric entities and functions). Beyond this, I leave the nature of causation completely open.

I should also explain why I won't be appealing to a couple of obvious examples of non-causal explanations. The first is mathematical explanations of mathematical facts. For instance, when seeking an explanation of Gauss's *Theorema Egregium*[17] one does not find a causal one. Perhaps this goes without saying. Anyway, I will not be discussing such examples because, although they are undoubtedly non-causal explanations, they are not explanations of *events*, and it is for these that it is sometimes claimed that no non-causal explanations may be found.[18]

The second example of non-causal explanations that I won't be considering comprises the Einstein-Podolsky-Rosen cases in quantum mechanics. It might seem a little odd to leave these aside, since many would argue that these offer some of the best candidates for non-causal explanations. My reason for leaving these cases aside, however, is that I think they are very difficult to interpret. It is not clear whether the explanations are non-causal, or as some suggest, reverse-causal.[19] Such cases cast significant doubt on traditional notions of causal explanations. The examples I present, however, I take to be less controversial and hence more damaging to the view that all explanations are causal.

3.3.1 The Bending of Light

We are told that the path of a beam of light is bent in the vicinity of a massive object; the more massive the object, the greater the bending. This result was first observed in 1919 by comparing the position of a star when its light passed near the sun (during a solar eclipse) with its usual position. What is the explanation for this bending?

The preferred explanation, offered by general relativity, is geometric. It's not that something *causes* the light to deviate from its usual path; it's

[17] This remarkable theorem states that the Gaussian curvature of a surface is intrinsic (McCleary, 1994, p. 148).

[18] David Lewis (1986a, p. 221), for one, claims that there are no non-causal explanations of (particular) events.

[19] See, for example, Price (1996; 1999) and Dowe (1997) for interesting discussions of what amount to reverse-causal accounts of the Einstein-Podolsky-Rosen cases.

simply that light travels along space-time geodesics[20] and that the curvature of space-time is greater around massive objects. Typically the defender of causal explanations will point out that it is the mass of the object that *causes* the curvature of space-time, and so there *is* an underlying causal explanation after all. There are two fairly serious problems with this reply though. The first is the difficulty of spelling out, in a causally acceptably way, how it is that mass brings about the curvature of space-time. After all, it can't be that there is an exchange of energy or momentum between the object and space-time, as some accounts of causation require. As I have already said, I wish to remain noncommittal with respect to the details of causation, but it seems that any account that permits mass to *cause* the curvature of space-time is unintuitive to say the least.

There is undoubtedly covariance between mass and curvature, but all covariance need not be cashed out in terms of causation.[21] For example, the angle sum of a triangle covaries with the shape of the space in which it is embedded, but one is not inclined to say that the angle sum of a triangle *causes* the shape of the relevant space. It seems to me that the case of mass and the shape of space-time is similar to this. Another way of looking at this difficulty is by asking the question: Why is it not the case that the curvature of space-time causes the mass? Simple covariance doesn't guarantee that one of the factors causes the other.

The second problem for this line of argument is that there are solutions to the Einstein equation for empty space-times in which the curvature of space-time is not identically zero. These are the non-Minkowski vacuum solutions (Peat, 1992, p. 17).[22] Thus, we see that, at the very least, mass cannot be the *only* cause of curvature. What then is causing the curvature in the vacuum solutions case? There is nothing *to* cause it! It simply looks as though the curvature is uncaused. Why then insist on a causal explanation of the curvature in the universes *with* mass? I suggest that

[20] In fact, it's probably better to say (tenselessly) that light *lies* along space-time geodesics.

[21] Physicists, I'm told, are more inclined to think of curvature as a *manifestation* of mass.

[22] The situation is somewhat complicated though, since the positive mass theorem of general relativity states, in effect, that such solutions must have a singularity, without which the space-time would be flat. This is assuming the ADM (after R. Arnowitt, S. Deser, and C. W. Misner) conception of mass. (This is a *global* conception of mass.) Adopting the stress-energy tensor conception of mass (which is a *local* conception of mass), however, non-singular, non-Minkowski, vacuum solutions are possible. For example, the analytic extension of the Schwarzchild metric (see d'Inverno (1992, pp. 219–221)) through the singularity has non-zero ADM mass but the stress-energy mass is everywhere zero. What is more, this space-time is non-singular and non-flat. I am indebted to Robert Bartnik and Matthew Spillane for their help with this point.

there is no reason at all and we ought to simply accept the geometric explanation for the bending of light.

3.3.2 Antipodal Weather Patterns

We discover that at some time t_0 there are two antipodal points p_1 and p_2 on the earth's surface with exactly the same temperature and barometric pressure. What is the explanation for this coincidence?

Notice that there are really two coincidences to be explained here: (1) Why are there *any* such antipodal points? and (2) Why p_1 and p_2 in particular? The first explanation I will offer is a causal explanation (i.e., featuring only causally active entities) and it addresses the second question. This explanation, presumably, will trace the causal history of the current weather patterns, to arbitrary fine detail if necessary, to account for the weather patterns at p_1 and p_2. In particular, the temperature and pressure readings at p_1 and p_2 at time t_0 will be accounted for. Notice that an explanation such as this does not explain why p_1 and p_2 have the *same* temperature and barometric pressure, just why each has the particular temperature and pressure that they have, and that these *happen* to be the same. Thus, an important part of the original phenomenon is left unexplained.

This case looks similar to that of explaining why there were 11 fatalities on New South Wales roads over the 1995 Easter break. The causal story will give the causal history of each fatality but will not explain why, in particular, there were 11 fatalities. This does not seem like such a deficiency in the road fatalities case, since it seems as though there is nothing (significant) left to be explained above and beyond what the causal story tells us. The case of the antipodal weather conditions, though, is entirely different.

The difference is due to a theorem of algebraic topology that states that for any time t there are antipodal points on the surface of the earth that simultaneously have the same temperature and barometric pressure.[23] This theorem, or more correctly the proof of this theorem, provides the missing part of the causal explanation. It guarantees that there will be two such antipodal points at any time, and, furthermore, the explanation makes explicit appeal to non-causal entities such as continuous functions and spheres.

[23] This theorem is a corollary of the Borsuk-Ulam theorem, combined with some minor structural assumptions (i.e., that the earth is topologically equivalent to a sphere and that temperature and pressure change continuously across its surface) (Kosniowski, 1980, pp. 157–159).

Notice, though, that this explanation also has its limitations—it does not explain why it is p_1 and p_2 in particular that have the same temperature and pressure. So we see that for a complete explanation of the phenomenon in this example, we require both causal and non-causal elements in the explanation.

3.3.3 The FitzGerald-Lorentz Contraction

The special theory of relativity tells us, amongst other things, that a body in motion, relative to some inertial reference frame $\mathcal{F}$, suffers a FitzGerald-Lorentz contraction. This is a reduction in the length of the body in the direction of motion, as measured by an observer stationary with respect to $\mathcal{F}$. What is the explanation for this contraction?

Minkowski's great contribution to relativity was in offering an elegant explanation for the Lorentz transformations (including the FitzGerald-Lorentz contraction). This explanation appeals to the now familiar concept of space-time, that is, a three-plus-one-dimensional manifold, which consists of three spatial dimensions and one temporal dimension. Minkowski realised that one of the key assumptions of special relativity, the constancy of the speed of light, could be formalised as the satisfaction of the equation:

$$(\Delta x_1)^2 + (\Delta x_2)^2 + (\Delta x_3)^2 - c^2(\Delta t)^2 = 0 \tag{3.1}$$

in any inertial frame. Here x_1, x_2, and x_3 are the spatial coordinates, t is the temporal coordinate, and c is a constant (the speed of light in a vacuum). Minkowski then introduces the imaginary time coordinate

$$x_4 = ict$$

where as usual $i = \sqrt{-1}$. So (3.1) becomes:

$$(\Delta x_1)^2 + (\Delta x_2)^2 + (\Delta x_3)^2 + (\Delta x_4)^2 = 0 \tag{3.2}$$

and (3.2) will be satisfied in every inertial frame if the quantity

$$\sigma^2 = (\Delta x_1)^2 + (\Delta x_2)^2 + (\Delta x_3)^2 + (\Delta x_4)^2 \tag{3.3}$$

is invariant under Lorentz transformation. This, says Einstein, "shows that the Lorentz transformation so defined is identical with the translational and rotational transformations of Euclidean geometry, if we disregard the number of dimensions and the relations of reality" (Einstein, 1967, p. 31). That is, the FitzGerald-Lorentz contraction is nothing more mysterious than the apparent shortening of an object in one dimension when a new set of axes are chosen, inclined at some angle to the old. This latter thesis

is the invariance of length under translation and rotation and is expressed mathematically as the invariance of the quantity

$$s^2 = (\Delta x_1)^2 + (\Delta x_2)^2 + (\Delta x_3)^2 \tag{3.4}$$

under linear transformations with determinate $|1|$ (i.e., the transformations are neither contractions nor expansions).

The explanation for the FitzGerald-Lorentz contraction is seen very clearly when one realises that the quantity s^2 in equation (3.4) is not invariant under Lorentz transformation in Minkowski space (although it *is* under rotation and translation in $\mathbb{R}^3$, as we have seen). The relevant invariant in Minkowski space is σ^2, as given by equation (3.3). I also stress the obvious here: this is a purely geometric explanation of the contraction, featuring such non-causal entities as the Minkowski metric and other geometric properties of Minkowski space.[24]

3.4 Causal Relevance

In this section I investigate another reply that supporters of the causal criterion are liable to make. This is to deny the causal idleness of the entities in examples such as those presented in the last section. One plausible way this can be done is to claim that the entities in the explanation are *causally relevant* but not *causally efficacious*. Frank Jackson and Philip Pettit (1990) give a good account of this approach. Although Jackson and Pettit don't specifically put the notion of causal relevance to work salvaging the causal criterion, nonetheless, their program could be used for this purpose.

Consider a case of trying to fit a square peg of side length ℓ into a round hole of diameter ℓ. Clearly it will not go. The first reason is non-causal: because of the squareness of the peg (and the roundness of the hole). The second is causal: the resistance offered by the overlapping portion of the peg. Furthermore, it seems that someone in possession of the squareness explanation knows more than someone who knows only the overlapping explanation. Jackson and Pettit suggest, and I agree with them here, that although the abstract property of squareness did not *cause* the overlapping, nor did it combine with the overlapping to produce the blocking, it is certainly true that the squareness was efficacious only if the overlapping was. They conclude that the abstract property of squareness is not causally efficacious (at least in this example). There is a sense, though, in which it

24 I am indebted to Jack Smart for drawing my attention to this examples such as this in both discussion and in his article (1990).

is not causally irrelevant either. It is not irrelevant in the way in which, say, the colour of the peg is. On this, Jackson and Pettit have the following to say:

> Although not efficacious itself, the abstract property was such that its realization ensured that there was an efficacious property in the offing. (Jackson and Pettit, 1990, p. 116)

That is, the property of squareness *programs* for the efficacious property of overlapping portions.

While I think there is much to be said for the causal relevance approach, in the end it won't save the causal criterion, for as I see it there are two serious difficulties facing this approach. First, I don't think that this defence will work for all explanations making use of non-causal entities. It will work only for those in which a fully causal explanation (i.e., one in which *all* the entities in question are causally efficacious) is on offer as well as the non-causal one, or where there are non-causal elements in a largely causal explanation. Thus, this strategy won't work for the FitzGerald-Lorentz contraction case where only one explanation is on offer and it is non-causal. Second, although this strategy enables supporters of the causal criterion to classify many apparently non-causal entities as causal, this is done at a fairly high price: significant blurring of the distinction between the causal and the non-causal. This blurring, if serious enough, is just the sort of thing that antagonists of the causal criterion would welcome. After all, if the property of squareness can enter into causal explanations, albeit in a subsidiary role (i.e., as causally relevant rather than causally efficacious), it seems that the causal requirement lets in too much. I am more inclined to admit that causally idle entities can have explanatory power than to fiddle with the definition of "causal" in this way.

One final move is left open to the supporter of the claim that only causally active entities can have explanatory power, and this is to argue that the geometry of space-time, for instance, while not being causally efficacious, nor programming for causally efficacious properties, may predetermine the range of possibilities. Space-time is thus seen as a *structuring cause* in Dretske's language. This move will allow the supporter of the causal criterion to classify the remaining recalcitrant explanations I've presented as fully causal explanations. But now I think that the difference between such a position and my own is entirely terminological. After all, what is the difference between holding (a) that there are causally idle entities with explanatory power and (b) that only causally active entities have explanatory power but that some of those entities might be structures that are not directly involved in causal chains? I suggest that whether one classifies such structures as causal or not, the important point is the recognition

of the importance of such structures in scientific explanations.

Where does this leave us then? Either there are causally idle entities with explanatory power, such as the geometry of space-time, or only causally active entities have explanatory power but they may include structural elements such as the geometry of space-time and programming properties such as continuous temperature distribution functions. Clearly it is the former conclusion I have been arguing for, but the latter will do as well. If the causal criterion is motivated by a notion of causally active entity that must include geometric properties, continuous functions, and the like, then it is ill equipped to make the demarcation required of it. Geometric properties and mathematical entities are just the sorts of entities the causal criterion is usually thought to eliminate. Perhaps this is not a terribly damaging argument against the causal criterion. After all, you could just bite the bullet and accept that the causal criterion does not rule out mathematical and geometric properties as is commonly thought. But surely mathematical and geometric entities are paradigm cases of non-causal entities. It looks as though the causal criterion is preserved in name only.

3.5 Rejecting Inference to the Best Explanation

In the previous two sections I discussed, at some length, the motivation for the Eleatic Principle that rested ultimately on the claim that only causally active entities can have explanatory power. It might be useful at this stage to make explicit one further assumption that defenders of the Eleatic Principle have thus far accepted. They have all accepted the view that we have ontological commitment to the entities in our best scientific explanations. That is, the defenders of the Eleatic Principle we've met so far accept inference to the best explanation. But now another defence of the Eleatic Principle presents itself. This defence is to accept that there are non-causal entities with explanatory power, but to reject inference to the best explanation in its most general form. This position has been defended by Nancy Cartwright (1983). Cartwright argues for inference to the most likely cause instead of the more general inference to the best explanation. Brian Ellis argues for a similar position. Ellis accepts that science makes extensive use of non-causal explanation, but he argues that only fully causal explanations carry ontological commitment.[25] If some restriction on inference

[25] Of course there are other issues on which Cartwright and Ellis have substantial disagreement. For instance, Cartwright is anti-realist about (most) of our best scientific

to the best explanation to causal explanations can be sustained, then the Eleatic Principle is justified trivially. I consider Cartwright's and Ellis's arguments for such a restriction in this section.

3.5.1 Ellis's Argument

Ellis is a scientific realist and, like many other realists, is so largely because of Smart's cosmic coincidence argument (which I mentioned in section 3.3). There is one difference, though. Ellis does not accept inference to the best explanation as Smart does. Ellis claims that "[o]ntological commitment can derive only from causal process explanations" (1990, p. 22). The latter is enough for a restricted version of Smart's argument to go through. The resulting realism is restricted to causally active entities:

> The ontology does not admit abstract entities like propositions and sets, unless these can somehow be reduced to entities of other kinds. For such entities have no causes or effects, have no location in space or time, and cannot influence any causal processes. It is argued that while such entities may have a role in model theoretic explanations, acceptance of such explanations carries no ontological commitments; only the acceptance of causal explanations carries any such commitment to the entities involved. The entities occurring in our model theories should generally be regarded as fictions. (Ellis, 1990, p. 5)

His reason for restricting inference to the best explanation in this way is apparent once we distinguish between two quite different types of scientific explanation. The first is the causal explanation, which, on Ellis's account of causation, will typically involve a story about exchanges of energy between physical entities. The second type of explanation is what Ellis calls *model theoretic explanations*. These typically idealise away from real situations and they are used as backgrounds for causal explanations. For example, Newton's first law provides the background for a causal explanation of why some moving object comes to rest. These model theoretic explanations typically feature such obviously fictional entities as frictionless planes, non-turbulent, laminar flow and inertial reference frames, so we should not accord real existence to the entities that feature in such explanations.

While I agree that this argument presents good reason to be suspicious of entities in such model theoretic explanations, it says nothing of abstract entities that feature in causal explanations. For example, in a fully causal account of a billiard ball collision (i.e., with frictional forces, etc.), we will

theories, whereas Ellis is not.

find reference to vectors.[26] Ellis acknowledges as much in the following passage:

> The main argument for realism about theoretical entities is also, apparently, an argument for the existence of forces, fields, numbers, sets, spatio-temporal relationships, possible worlds, and many other kinds of things. (Ellis, 1990, pp. 60–61)

But he has another reason for insisting that abstract entities have no real existence:

> The basic reason for resisting abstract particulars is that the world we can know about would be the same whether or not they existed. (Ellis, 1990, p. 79)

The key phrase here is "the world we can *know* about." Clearly some causal theory of knowledge is alluded to here, for otherwise the statement is patently false.[27] So in the end Ellis's restriction of the application of inference to the best explanation to causal process explanations will not provide a justification for the Eleatic Principle (by his own admission—see the second last quotation), so we are back to the epistemic justification that I discussed and dismissed in section 3.2.

3.5.2 Cartwright's Argument

Nancy Cartwright is also a realist of sorts. She is a realist about theoretical entities but not about scientific theories. She sums up her view rather nicely in the following passage:

> I believe in theoretical entities. But not in theoretical laws. Often when I have tried to explain my views on theoretical laws, I have met with a standard realist response: 'How *could* a law explain if it weren't true?' Van Fraassen and Duhem teach us to retort, 'How could it explain if it *were* true?' What is it about explanation that guarantees truth? I think there is no plausible answer to this question when one law explains another. But when we reason about theoretical entities the situation is different. The reasoning is causal, and to accept the explanation is to admit the cause. (Cartwright, 1983, p. 99)

[26] Recall that it is this feature of mathematics that Putnam so forcefully argued for in his realist days. He argued that mathematical entities feature indispensably in *the very same explanations* that lead realists to believe in theoretical entities. See Putnam (1971) and section 1.2.2.

[27] Elsewhere (1990, p. 7) Ellis is more explicit about his endorsement of a causal theory of knowledge.

In her rejection of inference to the best explanation she aligns herself more with anti-realists such as Bas van Fraassen[28] but accepts theoretical entities that feature in causal explanations for the same sorts of reasons as Ian Hacking.[29] So whereas Ellis is a realist who rejects inference to the best explanation in its most general sense, it's perhaps more appropriate to see Cartwright as an anti-realist who accepts inference to the most likely cause. Hacking's arguments ensure that causal activity is a sufficient condition for ontological commitment, whereas general anti-realist considerations ensure that it is also a necessary condition. Once put this way it is clear, I think, that for me to reply to Cartwright would involve entering into the realist/anti-realist debate, which I said at the outset I was not going to tackle in this work. Nevertheless, I feel obliged to say something in reply to Cartwright, but before I do this I need to clear up an ambiguity in Cartwright's position.

Recall that Cartwright admits theoretical entities that are causes of some phenomena which require explaining. The ambiguity revolves around what constitutes an event (or phenomenon) in need of explanation. For example, suppose that all entities that are causally active are the cause of some event or other. Then Cartwright's inference to the most likely cause may warrant belief in these entities as the most likely causes of their respective events. Notice that nothing here ensures that the entities in question are causally active with us. On the other hand, one could argue that what she takes to be an event in need of explanation must be an event which *we know about*. That seems uncontroversial enough, but now depending on how we spell out the "we know about" claim, it looks as though Cartwright cannot admit causally active entities that are not causally active with us.

As it turns out I think that it is the latter position that Cartwright intends. This can be best seen by considering a case where there is some event, e, which we have no causal contact with and asking the question: What reason does Cartwright have to believe that such an event occurred? One reason would be if we directly observe e or observe a result of e, but this would mean that we have causal contact with e, and this is ruled out by construction. The important question is whether, for Cartwright, there can be any other way of knowing about e. It seems not, given what she says about inference to the best explanation. Recall that inference to the best explanation is not an admissible inference for Cartwright, so it can't be that e explains some other event or phenomena, unless of course e is

[28] For details of van Fraassen's rejection of inference to the best explanation, see van Fraassen (1980).

[29] See footnote 2 in this chapter or Hacking (1983) for further details.

the cause of that event or phenomena, and this is also ruled out by construction. What other reason can we have to postulate a causally isolated entity on Cartwright's account? I can think of none, so I must conclude that Cartwright is indeed committed to only those entities that are causally active with us.[30] I will now use this reading of Cartwright's position on some examples I have considered previously. These examples will show how Cartwright's position seems unable to give the intuitively correct result in what are fairly uncontroversial cases.

The first example is that of the stars and planets outside of our light cone. These are theoretical entities that are not the cause of anything (that we can observe), so it seems on Cartwright's account they ought not be granted "real" status. As I've mentioned previously, in relation to the epistemically motivated Eleatic Principle (cf. section 3.2), this seems like the wrong answer. Note, however, that I am not claiming that Cartwright is committed to this view, just that to avoid this conclusion will require some additional argument. In the absence of such argument, the undesirable conclusion does seem to follow from inference to the most likely cause alone.

The other example I'd like to reconsider is the case of belief in the existence of germanium prior to 1887. Recall that I assume that there had been no known causal contact with germanium until 1887, but that it was postulated in 1871 on the basis of a "gap" in Mendeleeff's periodic table. It seems Cartwright must deny any good reason to believe in germanium at that time since its causal isolation guarantees that it couldn't have been the cause of anything (that we knew about) and the only appropriate inference she allows is inference to the most likely cause. Again, I trust that your intuitions suggest that this is the wrong answer.

Much more could be said about Cartwright's ontology and, in particular, her inference to the most likely cause, but in the end I find examples such as these give us good reason to suspect that her project cannot be used to support a causal constraint on existence claims.

[30] Cartwright has pointed out to me that she could appeal to some other form of inference besides inference to the best explanation to justify belief in causally isolated events and entities. Nothing in her position rules out such a move. Such a move, however, is not open to defenders of the Eleatic Principle interested in restricting the admissible, relevant inferences to inference to the most likely cause. Such a move clearly undermines the Eleatic Principle by (presumably) allowing belief in entities without any restriction on their causal histories.

3.6 The Content of Scientific Theories

Both Jody Azzouni (1997a) and Mark Balaguer (1996b; 1998) have been critical of Quinean naturalism. They have suggested (in different ways) that it is a mistake to read our ontological commitments simply from what the quantifiers of our best scientific theories range over. Rather, we should pay attention to what our best scientific theories are about and count only those entities as real. Azzouni and Balaguer argue that our best scientific theories are not about mathematical entities, for instance, but about causally efficacious entities such as electrons, stars, and such. They thus suggest that only the latter causally active entities should be admitted to our ontology.

Although Azzouni's and Balaguer's intuitions are similar—it's the content of the theories that's important, not the domains of quantification—their arguments are different enough to warrant separate consideration. In each case I'll focus attention on the parts of their arguments that might be used to support the Eleatic Principle. I begin with Balaguer.

3.6.1 Balaguer's Fictionalism

Balaguer's position,[31] in a nutshell, is that "the nominalistic content of empirical science is all empirical science is really 'trying to say' about the world" (Balaguer, 1998, p. 141). We thus need not have ontological commitment to anything other than the entities that are part of this nominalistic content—that is, the causally efficacious entities.

Balaguer begins by pointing out that entities that lack causal powers could not make a difference to the way the *physical* world is. He thus defends the view that the content of our scientific theories can be separated into nominalistic and Platonistic components and that the nominalistic content (i.e., the purely physical facts described by such theories) is true (or mostly true), while the Platonistic content (i.e., the abstract, mathematical facts described by such theories) is fictional (1998, p. 131).

To illustrates his position, he considers a scientific claim such as

(A) The physical system S is forty degrees Celsius.

Balaguer argues that while (A) does assert that a certain relation holds

[31] I should point out that Balaguer does not ultimately accept the position that I'll be attributing to him in this section. (He adopts a rather anti-metaphysical stance on such questions as the existence of mathematical entities (1998, pp. 151–179).) He has, however, argued rather persuasively for the position I'll be discussing. I think it is fair to say that he has considerable sympathy with the position and he certainly sees it as a position worthy of serious attention.

between S and the number 40, it is reasonable to maintain that since the number 40 is causally inert, the truth of (A) depends on purely nominalistic facts about S and purely Platonistic facts about the natural numbers; Balaguer argues that these two sets of facts hold or don't hold independently of one another. Thus, again it is reasonable to maintain that facts of the one sort obtain, whereas facts of the other sort do not (i.e., the nominalistic content of (A) is true whereas its Platonistic content is fictional). Balaguer claims that we can take this view of the whole of empirical science: empirical science is not true (because there are no abstract objects), but the nominalistic content of empirical science is true (because, as he puts it, "the physical world holds up *its end* of the 'empirical-science bargain'" (1998, p. 134).

Now it is clear that if Balaguer is right about all this, then we have a rather interesting argument for the Eleatic Principle. There are problems, though, with Balaguer's position. For instance, consider Balaguer's argument for the claim that empirical science does not confirm the existence of mathematical objects:

> Empirical science *knows*, so to speak, that mathematical objects are causally inert. That is, it does not assign any causal role to any mathematical entity. Thus, it seems that empirical science *predicts* that the behavior of the physical world is not dependent in any way upon the existence of mathematical objects. But this suggests that what empirical science says about the physical world—that is, its complete picture of the physical world—could be true even if there aren't any mathematical objects. (Balaguer, 1998, p. 133)

It's not at all clear, however, that the physical universe cannot depend upon causally inert (or at least causally isolated) entities. After all, physicists posit causally isolated universes (i.e., universes with no causal influence on *this* universe) in order to explain why our universe is fine-tuned for carbon-based life (Barrow and Tipler, 1986). It seems, then, that certain features of the physical universe (namely, its "fine-tuning") may be explained by appeal to causally isolated entities (i.e., other universes) and thus, in some sense, the physical universe may indeed be said to depend upon causally isolated entities.[32]

In a review of Balaguer (1998), Ed Zalta and I raised a somewhat vague concern about Balaguer's nominalist project (Colyvan and Zalta, 1999). We claimed that Balaguer's driving intuition here—"the nominalistic content of empirical science is all empirical science is really 'trying to say' about

[32] In fact, Balaguer seems to be assuming that all explanation is causal. I argued against this position in section 3.3.

the world"—is little more than an intuition in favour of nominalism. We suggested that

> at least part of the business of science is to describe reality. To suppose that reality can be described by the nominalistic content of scientific theories is something akin to begging the question against the Platonist. (Colyvan and Zalta, 1999, p. 344)

Let me now try to flesh out this concern.

It seems to me that there is a very interesting ambiguity in the claim that the nominalistic content is all that science is really trying to say about the world. The ambiguity is appreciated if we consider the following example. Suppose we wish to tell the story of why my hand won't pass through a solid object such as a wall. Now in one sense this story is simply about the wall and my hand. But in another sense, it's (somewhat surprisingly) about quite a bit more, because the explanation of why my hand won't pass through the wall involves a story about the electro-repulsive forces of the electrons in both the wall and my hand. Of course, all we're trying to do is talk about walls and hands, but in doing so we are forced to discuss electrons. The confusion is between what the theory is supposed to be describing or explaining and what resources it requires to do the describing or explaining. Now I think it is clear that there *is* a sense in which the explanation of why my hand will not pass through the wall is simply about hands and walls. This, however, does little to convince us that we have no reason to believe in electrons.[33]

Now the ambiguity that I see in Balaguer's intuition that the nominalistic content is all science is really trying to say is between the following two readings: (1) what science is ultimately trying to describe or explain is simply the nominalistic content and (2) the scientific descriptions and explanations ultimately will not need to appeal to anything other than the nominalistic content. Now Balaguer might well be correct that (1) is true, but it is clear that (2) does not follow from (1). Moreover, my hand, wall, and electron example shows that ontological commitments hang on (2) not (1). If, as I suggest, Balaguer is conflating (1) and (2), his intuition is indeed something akin to begging the question against Platonism. For (2) supposes that our ultimate descriptions and explanations of the world will not require anything other than nominalistic resources. In any case, this is not an intuition that I share.

[33] I'm restricting my attention to scientific realists here. Of course, anti-realists such as van Fraassen do reject electrons for reasons along these lines. Balaguer, however, is trying to show that one can be a nominalist and a scientific realist.

3.6.2 Azzouni's Challenge

Let's turn our attention to Azzouni. Although Jody Azzouni has criticised (the letter of) Quinean naturalism in a number of places,[34] I'll focus on Azzouni (1997a) where he comes closest to defending a version of the Eleatic Principle. Azzouni suggests that when we consider a problem such as the movement of two (nearly) homogeneous and (nearly) perfectly spherical masses connected by a spring and tossed up into a gravitational field, we find three quite distinct sorts of objects quantified over. The first are the spheres themselves; the second are items like centres of mass; the third are mathematical quantities (such as the product of the two masses). Now Azzouni suggests that the first category is what the theory is about. The second, although physically significant, is not real and the third category, according to Azzouni, is not even physically significant. He argues that it is only the two spheres that we interact with, not centres of mass and certainly not mathematical entities. The point is that:

> there is in place an implicit causal story about the sorts of objects we interact with. Even when quite sophisticated instruments are involved ..., still, in the physical description of how the instrument does what it does, we find implicit and explicit causal relations described. When such a story is missing, as it invariably is with physically significant items which are not physically real (such as centres of mass) or with merely mathematical objects (such as various functions of physical quantities that we could define), we find no causal interactions with the item (either between it and us, or between it and other physical items). (Azzouni, 1997a, p. 200)

So the suggestion is that scientific theories are about the physically real and these we can causally interact with. Moreover, it is only the physically real that we are ontological committed to.

The problem with this suggestion, however, should be quite familiar by now: There are physically real objects outside our light cone. As we've already seen, we do not causally interact with these, and so they fail the causal test. Azzouni admits that being an item with which we can causally interact is a sufficient but not a necessary condition on being physically real. But this is enough to suggest that Azzouni's proposal is of little use as a defence of an Eleatic Principle. But Azzouni makes a very interesting suggestion. He suggests that causal power might provide a *defeasible* criterion:

[34] See, for example, Azzouni (1994; 1997b; 1998; 2000).

> Notice that if something is physically real, but out of our causal reach, the physicist will usually have a story for why this is so (e.g., it's too far away). Such a story is unnecessary, and never given, when it comes to centres of mass or purely numerical items. (Azzouni, 1997a, p. 210)

As it turns out, I have a great deal of sympathy with this suggestion. The suggestion is in effect that the Eleatic Principle is not the final arbiter, but a rule of thumb that can be overturned when it gives the wrong results. We know when it gives the wrong results by considering the scientific theory as a whole. Moreover, when the Eleatic Principle gives the wrong result, total theory should provide a story about why we got the wrong result. I will pursue this suggestion further in the next section, where I will argue that this and similar suggestions are very much in the spirit of Quinean naturalism. Of course, there is room for disagreement about when entities with which we have no causal interactions are to be admitted into our ontology and what the story concerning such entities should look like,[35] but for the moment, it is sufficient to simply note that Azzouni's suggestion of a defeasible Eleatic Principle marks a serious departure from the causal version of naturalism under discussion.[36]

3.7 The Moral

In this section I will show that there is a positive message to be gleaned from the lack of support for the Eleatic Principle. In section 3.1 we saw how the inductive argument for the principle is in itself inadequate but leads rather naturally to epistemic considerations. These, in turn, suffer from the same sorts of objections as the causal theory of knowledge (which is, after all, what underlies the epistemic motivation). In particular, we saw how this motivation yields some undesirable consequences in some fairly straightforward cases; for instance, that we ought not believe in stars and planets

[35] For example, Azzouni suggests that there is an important difference between mathematical entities and items outside our light cone. The latter, he suggests, are admitted despite failing the causal test because an explanation for their failure is available. Although mathematical entities also fail, no excuses are offered on their behalf. But it seems that a very good story is available: *Mathematical entities do not have causal powers*. Moreover, that this is so well known explains why in this case "such a story is unnecessary, and never given" (Azzouni, 1997a, p. 210).

[36] I don't want to give the impression that Azzouni sets out and fails to defend the Eleatic Principle. In the article in question (Azzouni, 1997a), at least, he simply explores the cogency of a non-Quinean naturalism (not unlike the causal version we've been considering). Moreover, he does this without providing an argument for the position in question.

outside of our own light cone. Such a position is surely unpalatable to even the staunchest defenders of the principle. In section 3.3 I examined what I take to be the most important argument for the Eleatic Principle, due to David Armstrong. Armstrong claims that causally idle entities have no explanatory value and hence should not be considered to be real. This argument was seen to rely on the assumption that only causally active entities have explanatory power, and this was shown to be a mistake. Finally, in section 3.5 we saw that motivating the Eleatic Principle by appeal to a restriction of inference to the best explanation to causal explanation also faces serious problems.

The moral to be drawn from all this is best illustrated by considering the deficiencies of the arguments for the Eleatic Principle and how they might be overcome. For instance, the epistemic justification failed because it ruled out entities such as germanium prior to 1887 and stars and planets outside our light cone, which seem reasonable to believe in. There is a way of getting around these problem cases though, and that is to appeal to some sort of 'rounding out' principle. Thus, a defender of the causal test may argue that even though an epistemically motivated Eleatic Principle rules out stars and planets outside our light cone, such entities may, nevertheless, be included in our ontology on the basis of this rounding out principle.[37] But now one wonders what the purpose of the Eleatic Principle is, if it is so easily overruled when it gives the wrong answer. Why not try to be more explicit about what the rounding out principle amounts to and just appeal to it in the first place?

It seems this rounding out principle, if it allows belief in stars and planets outside our own light cone, is going to do so for reasons of symmetry or, more generally, for reasons of theoretic virtue. That is, the astronomical theory that posits stars and planets outside our own light cone is a better theory than its counterpart that does not posit such entities. (For one thing, the latter theory is stuck with the problem of explaining why the earth is apparently at the centre of the universe.) Now notice that most entities that are causally active are likely to be needed in (causal) explanations of certain phenomena.[38] Thus, using this rounding out principle, at least as I've outlined it here, will effectively subsume one direction of the Eleatic Principle—causally active entities will gain admission to our ontology because they are needed for our scientific theories. This principle looks more promising than the Eleatic Principle, since it seems to avoid the

[37] For instance, in private conversation Alan Musgrave has suggested this. Note also the similarity between this suggestion and Azzouni's suggestion (in section 3.6.2) of a defeasible Eleatic Principle.

[38] For example, the stars and planets inside our own light cone will be needed to explain why we see points of light when we look up at the night sky.

more obvious pitfalls of the epistemically motivated Eleatic Principle and yet, because in many cases both principles yield the same results, it retains much of the Eleatic Principle's intuitive appeal. In effect, I'm suggesting that the Eleatic Principle may be a good ontological "rule of thumb" but may not be the final arbiter on such matters.[39]

Alternatively, let's cast our minds back to the Armstrong argument for the Eleatic Principle. This argument placed the responsibility for questions of ontology firmly upon the explanatory power of the entities in the theories in which they occur. My point of disagreement was that acausal entities also have explanatory power, so we should not confine our attention to causally active entities. I agree that explanations featuring causally active entities are often very good explanations, and perhaps we ought to seek such explanations when appropriate. It's just that it's pointless to restrict our attention to these explanations, since they do not exhaust the acceptable explanations of science. There is, however, an important insight in Armstrong's argument. This insight is the move to using explanatory power as a guide to deciding questions of ontology. If we heed this advice, but do not restrict the acceptable explanations of science to those featuring causally active entities, we find ourselves once again with a very general principle for deciding ontology, not unlike the rounding out principle we arrived at in the last paragraph. In both cases the principle places the responsibility for ontology on theoretic virtue (symmetry in the former case and explanatory power in the latter), and in both cases the principle goes well beyond the scope of the Eleatic Principle.

This more general principle now begins to look very much like Quine's ontic thesis. So, by responding to the difficulties faced by the Eleatic Principle, we are lead to something very much like Quinean naturalism. Now of course this does not mean that Quinean naturalism is the *only* way to avoid these difficulties, but it does at least suggest that the Quinean position is better equipped to embark on the difficult task of deciding which entities we allow entry to our ontology.

[39] Indeed, there is some evidence to suggest that even some of the principle's supporters may agree with me here. For instance, in the following passage David Armstrong suggests that the causal criterion is a heuristic device for isolating those entities for which it is hoped that a reductive analysis may be given:

> The argument from lack of causal power is simply intended as a reason for thinking that the research programme [of giving an account of the truth conditions of numbers, classes, propositions etc. in terms of particulars and their properties and relations] is a promising one. (Armstrong, 1978, Vol. 2, p. 5)

In other words, he suspects that causally idle entities are dispensable to science but admits that the real work is the required reduction of such entities to more respectable entities.

3.8 Recapitulation

So far I have discussed the Quine/Putnam indispensability argument and shown how it depends on the Quinean doctrines of naturalism and holism. I then defended the Quinean reading of naturalism against causal versions of naturalism. The latter explicitly or implicitly appeal to a causal test in order to decide which entities are legitimately deserving of real status, with only the causally active entities able to pass this test. In this chapter I argued against the use of such a test.

It is important to see the problems associated with such a test for at least two reasons. First, assuming that it is desirable to hold some version of naturalism, discrediting the causal version raises the stocks in its main rival: Quinean naturalism. The second reason is that in considering the problems with the Eleatic Principle and how they might be rectified, we found that one is drawn in a very natural way to Quinean indispensability theory. None of this counts as a positive argument for the Quinean position; indeed, as is often the case with such basic doctrinal beliefs, it is difficult to produce persuasive positive arguments. The best one can hope for is that the doctrine in question fare well in cost–benefit comparisons with the main contenders. I hope you agree that the preceding chapters have shown that the Quinean position, if not the best approach to matters of ontology, is at least highly plausible.

Much more could be said concerning the Quinean backdrop, particularly with regard to naturalism, but instead I shall move on to consider some explicit attacks on the indispensability argument. In discussing these attacks, however, issues concerning the Quinean backdrop will arise again. This will hopefully bring into sharper focus the doctrines of naturalism and holism and underline their importance to the whole debate.