A Critique of Pure Existence

Reynier Pet

MA Thesis

Supervisor: Victor Gijsbers

History and Philosophy of the Sciences

Institute for Philosophy, Leiden University

July 11, 2020

Contents

Introduction

> I suggest that everything which possesses any power of any kind, either to produce a change in anything of any nature or to be affected even in the least degree by the slightest cause, though it be only on one occasion, has real existence. For I set up as a definition which defines being, that it is nothing else than power.
>
> —Plato, *Sophist*

What does it mean to exist? There is perhaps no other question as pervasive in philosophy as this. In fact, describing the nature of reality, which we might commonly know as the endeavor of metaphysics, is perhaps the philosopher's most favorite pastime, and has been since the ontological philosophies of the ancient Greeks, if not before. It may seem then that we're setting ourselves a most ambitious goal when we put forward that the answer to this question can be formulated as a simple principle: that something exists if *and only if* it is causally connected to me.

Nevertheless, this will be the central thesis of this paper. And while it may legitimately be understood as a specific formulation of David Armstrong's Eleatic Principle, it will be important to distinguish ours from it, since we shall see that how such a principle is specifically formulated will play a crucial role in its defense, and formulations of the Eleatic Principle have varied significantly (Colyvan 1998, 314; Cowling 2014, 247; Oddie 1982, 286). For this purpose we shall, where relevant, refer to our formulation as simply the Principle.

For the same reason we will also find that defending a principle will require a somewhat unorthodox approach. A principle, as we will come to understand it here, cannot be defended in the same way as a normal thesis. This is because they are things that determine rather than things that are to be determined, like how principles of mathematics determine mathematical truth. Thus, rather than a simple refutation or affirmation based on

logic or empirical evidence, it will require a specific process explaining how principles can apply absolutely, yet be adopted arbitrarily. The reasons for this will be addressed at the start of Chapter 1, and expanded on throughout the paper where relevant, as we shall see that the totality of this explanation is deeply tied to the defense itself.

In Chapter 1, relying on a Kantian understanding of principles, we will focus first on demonstrating that a principle can apply absolutely, even though this might not be immediately apparent, and that this is because the implications of such a principle may apply necessarily in the same sense that synthetic a priori judgments do. The rest of the chapter is then dedicated to analyzing the exact meaning of the principle, i.e. what meaningful truth it conveys in the same sense that a synthetic a priori judgment conveys a meaningful truth.

In Chapter 2, we consider the question of how it might be possible for there to be discussions about causal principles of existence such as the Eleatic Principle, and by extension the Principle, if this principle is supposed to apply absolutely. Particularly of note will be the fact that alternative principles can and have been suggested. We will demonstrate that these alternatives all in fact can only make sense if they, by one way or another, implicitly imply the truth of the Principle. Of particular note will be Mark Colyvan's suggestion of inference to the best explanation (Colyvan 1998, 326–30), which will require us to revisit our Kantian approach the principles to conclude that his suggestion cannot serve as an alternative to the Principle, as that would rest on a confusion regarding what these principles are really about.

Finally, in Chapter 3, having established how the Principle can be arbitrary in its formulation, yet have its truth apply absolutely, we can turn to arguing in favor of our specific formulation. This takes a form that one might expect to be part of the defense of a principle: a demonstration of the merits of adopting it. These are embodied in the concrete consequences of applying this abstract principle, which we determine by examining the

specific types of situations within which these consequences manifest. The final and most important merit of the Principle, however, will relate to a more general advantage gained by the fact that it is, as we'll put it, true to form, arguing that adopting any other principle would result in insurmountable problems that are the consequence of a disingenuous representation of the nature of existence.

Chapter 1

The Principle

What is a Principle?

If we're going to consider the truth of a principle, we need to start off by asking ourselves an important question: how can a principle even be defended? Defending (or refuting) a normal claim, such as "the sun is roughly eight light-minutes away from the earth" or "murder is bad," is fairly straightforward: empirical facts or logical deductions will support or refute the truth of such a thesis. But a principle is neither empirical nor logical. It doesn't lend itself to such a possibility. Because it is *principal,* it goes before anything else. It comes first. It is an absolute rule that is not something to be determined, but rather is itself something that *determines*. A principle of ethics determines how to distinguish right from wrong, a principle of logic determines how to distinguish valid from invalid arguments.

And this seems to pose something of a problem: does the truth of a principle not simply hinge on whether or not one chooses to adopt it? Aren't claims defended *based on* principles, and, as such, does that not mean that the truth of our thesis, that something exists if and only if it is causally connected to me, is arbitrary? To a degree this is true, and this would seem to suggest that our only means of defending this principle would be to demonstrate its merit; that it's superior to the other (equally arbitrary) options. And, again, to a degree this is certainly true. And we will eventually see our defense return to that point. But there is more to the story. Consider, for example, that this would apply just as much to mathematical principles as it would to ours, yet most of us wouldn't be inclined to argue that mathematical principles such as that (for all natural numbers), if $x = y$, then $y = x$ (the third axiom of Peano), are arbitrary. This raises a new question about these principles: how aren't they arbitrary then?

To find an answer to this question, perhaps the best place to look is Immanuel Kant's *Critique of Pure Reason*. This is not only because there he provides what must have been at least a very compelling answer, but also because the work itself is actually dedicated to establishing principles, the principles of pure reason to be precise. Depending on the nature of his answer, we may wish to emulate his method.

So what is his answer? To answer this in full, we would likely have to recount the entire *Critique* from scratch, but for our purposes the short answer will suffice: because synthetic a priori judgments are possible. Now, we don't need to completely adopt Kant's views (if only to avoid having to go off on a tangent defending the entirety of Kant's metaphysics). All we need is to know what it is about synthetic a priori judgments that would make it possible for principles to not be arbitrary, and this answer can be found in what he considers to be a typical example of such a judgment, conveniently for us also taken from mathematics: 7 + 5 = 12 (Kant 1998, 144).[1] Simply put, what Kant points out using this example is that one can fully grasp the concepts of "five" and "seven," yet still not (yet) know that five plus seven equals twelve. Whether this specific formulation is a fair one, one might wish to debate, but part of this observation is a fact that is difficult to deny: some act is performed when we add five and seven together to make twelve, a calculation. When we have both five and seven, we already have all the information necessary to conclude that we have twelve, but before we perform the calculation, we lack the knowledge of it. We don't know seven plus five equals twelve until after we've added seven and five together.

Another way to put this, perhaps somewhat closer to how Kant would have put it, might be to say that we might know the meaning of five plus seven, and we might know the meaning of twelve, and still not (yet) know that they mean the same thing, that five plus seven equals twelve. Yet another way might be to say that we may "know" five plus seven

[1] B 15-16

equals twelve without (truly) knowing it. In this sense one might say it's about awareness, about the difference between "knowing" and "knowing that." This is where the matter becomes relevant for trying to defend a principle. What this implies, is that they aren't arbitrary because they're already known, we simply don't yet know that we know them; we aren't necessarily aware of them. Thus when we say that, as a principle, something exists if and only if it is causally connected to me, we're essentially saying that an awareness of the existence of something amounts to an awareness of our causal connection to that thing.

This awareness isn't simply achieved by understanding this specifically. We've shown how a Principle can be non-arbitrary (because it's already known), yet something to be defended (because we're not yet aware of it). This applies as much to our phrasing of the Principle as it does to our experience of existence. It's one thing to say the words, it's another thing to actually know what you're saying. So the same line of reasoning applies to the use of language, and not only to a potentially skeptical listener, but also to those who "know" what they are saying when we say that something exists if and only if it is causally connected to them. Even to them, to us, it may still be unclear what that, as one might say, *amounts* to. We have what one might call an intuitive understanding of our words, just as with any phrase, argument, or story one might have uttered in the moment, but have yet to realize, to understand, what we've just said, in the same sense that we might have just summed up five and seven, but have yet to understand that that means we have made twelve. We have to perform the calculation. What we need to do is consider what we have just said, what each element of the Principle must reasonably mean (understanding it as a principle), and add it all together. Continuing in the spirit of Kant, one might consider that our current task can be summarized as an attempt to answer the question: how is it possible that, as a principle, something exists if and only if it is causally connected to me?

Why "Me"?

First and foremost, it should be established that there is an obvious intuitive appeal to the Principle. By involving causality in the story about the nature of existence, we are provided with an excellent, clear means of excluding from the list of possibly existing things the most pesky of metaphysical (supposed) entities: things that cannot interact with us. In fact, it is for this very reason that David Armstrong defended the Eleatic Principle, stating:

> The argument is this. A spatio-temporal realm of particulars certainly exists (it includes our bodies). Whether anything else exists is controversial. If any entities outside this realm are postulated, but it is stipulated further that they have no manner of causal action upon the particulars in this realm, then there is no compelling reason to postulate them. Occam's razor then enjoins us not to postulate them. (Armstrong 1978, 1:130)

It is perhaps for this same reason that there is some significance to referring to a causal connection to "me." It may seem a rather awkward formulation, but it emphasizes an important consideration regarding the relevance of causal interaction, the same consideration that Armstrong himself brings up in the following passages (Armstrong 1978, 1:131–32): causal interaction is a precondition for the possibility of detectability, arguably of knowability. It is an appeal to the idea that to suggest the possibility of (fundamentally) undetectable entities is being epistemologically irresponsible. The implication is that such a thing must be a contradiction, based on arguments of the kind such as Hilary Putnam's Brains in a Vat (Putnam 1981), the thought being that one can't refer to undetectable (unknowable) entities, just as much as one can't refer to things outside the simulation of the brain in a vat if the brain's concepts have no link to the external world. Fundamentally, regardless of the specific argument behind it, the idea is that there is an inescapable connection between detectability and "be-ability," between the epistemology and ontology.[2]

[2] Obviously this has a Kantian ring to it, but the essential quality of the argument should appeal as much to an empiricist (particularly a Berkeleyan) or a phenomenologist.

There is, however, another significance to the specific reference to ‘me’ in the Principle. This is because, besides emphasizing this connection between epistemology and ontology, it has an important implication. Earlier, we considered that this type of formulation might be awkward, and the reason for this is that who “me” is referring to, depends on who is uttering the phrase. One could see this as a potential weakness, because the reference point changes depending on who is uttering the phrase. However, that is precisely the point. The fact that the reference point is relative to the person uttering it, is a statement in and of itself. And this statement implies two possibilities, since we haven’t yet established whether these causal connections can manifest as one thing being able to affect another while remaining unaffected by it, making them unidirectional, or whether it must be the case that affecting something means being affected by it, making them bidirectional. In the former case, that would mean you being causally connected to something doesn’t necessarily imply it being connected to you. Thus either these connections can be unidirectional, and what exists is relative to the observer (i.e. “me”), or they must be bidirectional, and what exists is always the same, regardless of the observer. So now we must consider the question: which is it?

Direction

When considering the possible directionality of the connections between entities, there are two major considerations. The first is based simply on how causation would work; if causal connections are bidirectional by nature, then obviously the connection can only be bidirectional. The second consideration is how much is needed to satisfy the conditions of existence. If a unidirectional connection suffices, then, as a principal condition for existence, there should be nothing more to the directional nature of the connection.

This brings up a problem, as the appeal of a criterium focused on detectability seems to suggest that a unidirectional connection (of object to subject) would indeed suffice. If this is the case, either focusing on the causal nature of the connection is redundant (since a

unidirectional connection would imply detectability), or it's superfluous (since its being causal is trivial if all that matters is that the entity in question detectable). In either case, it raises the question of why it, and not simply detectability itself is taken as the criterium for existence, since detectability would seem to suffice. Even if it is practically the case that all detectable entities must be causally connected to us, that seems to be missing the point, at least from a metaphysical point of view. If the actual criterium is detectability, then that is the essence of the matter, and so that should form the basis of the Principle, and existence (as detectability) itself being causal in nature must be, metaphysically speaking, accidental. That is, it would mean that detectability is the (metaphysically) principal matter. It needing to be causal would in this respect just be a product of it needing to be detectable.

Obviously, what we're working towards here is that there is reason to take causality over mere detectability. While detectability certainly plays a role in the appeal of causality, the true value of causality lies, and must lie, in the essence of causality itself. Detectability is an important "consequence" of possessing causal properties, but it's not the only one. The first indication of this comes from an intuition one might have in types of situations in which we find ourselves affirmed in a pre-existing belief. To give an example of this, consider a situation in which we see a ball lying on the ground in front of us. We may (justly) believe it to be filled with air, and would suppose that, if we were to give it a swift kick, it would fly away in a nice arc. Thus, if we test our hypothesis by kicking the ball, and then observe that it does indeed fly away in such an arcing trajectory, we would not only feel ourselves strengthened in our belief that our hypothesis was true; we would be more convinced of the fact that the ball, as we thought it was, existed.

Nevertheless, this doesn't yet really give us any satisfying answers. After all, wasn't it already sufficient to see the ball in order for us to believe it existed? Sure, kicking the ball tested and thereby strengthened the belief, but so would additional observation. Besides, not

everything we believe to exist is something we have the liberty of (meaningfully) causally affecting in this way.

Perhaps it's more revealing to consider the situation in which we aren't affirmed in our initial belief. What if, when we kick the ball, instead of it flying away, our foot passes right through it. It seems obvious that we would now conclude that the ball did not exist, a clear case in which causally affecting an object is relevant to determining its existence. Yet, at the same time, something fundamental is revealed here. After all, our initial observation isn't fully disregarded. Sure, the ball wasn't actually a ball, but certainly it was *some*thing. We might consider that the ball was really just a hologram, a spherical cloud of gas, or perhaps even a hallucination, but whatever the case, something most certainly did exist.

What this reveals is that there are two different questions being answered here. The first question is about whether something exists that is (responsible for) our observation, which is (immediately) answered by the fact that we have the observation. The second question is about *what it is* that exists. It's this latter question in which causality plays the central role, and so to connect the concepts of existence and causality is to argue that these two questions are fundamentally connected to one another. So to say that something must be causal, then, is to argue that it must have certain identifying properties (identifiable through its causal connection to us). Therefore the argument made by the Principle is that, to say that something exists, is to imply that it has an *essence*, something that identifies it for *what* it is, which manifests in an ability to affect or be affected. The Principle essentially implies, among other things, that nothing just "is"; it is always something (more) specific.

Objecthood

One might at this point be tempted to raise a counterpoint: we don't need to causally affect a thing in order to help determine its nature. Consider, for example, the situation in which you don't attempt to kick the ball, but simply wait for something to happen that may

contradict or support the idea that the ball is what we think it is. This seems to provide us with the same kind of information about the ball. In fact, we might witness someone else attempting to kick the ball, and it's obvious that in such a case we'd have the same expectations, and the same considerations should those expectations be failed.

To a degree, this is true, but only insofar as we can make certain assumptions about these observation, in particular: that someone else kicking the ball is (enough) like us ourselves kicking the ball. This may seem trivial, but it's actually an important distinction. Consider that, fundamentally, our distinction of the ball as an object separate from the rest of the universe is largely trivial. There's little beyond practicality driving us to consider the ball to be a separate entity, and not just a pattern in an ocean of "pixels." Perhaps that is simply what objects fundamentally are, but specifically when it comes to distinguishing objects there's something important about our own interaction with them to take note of. The one thing that occurs when we interact with objects that doesn't—at least not within the scope of our experience—when we witness someone else doing so, is that we fundamentally distinguish ourselves from the ball when we interact with it.

This should be understood not so much in the sense that we distinguish our body from the ball, but rather in the sense that there is (i.e. we experience) an inner will imposing itself on the outer world whenever we ourselves interact with that world. Obviously, the objecthood of the ball is still in large part defined by our observations of it. We still consider balls separate entities even when we never interact with them after all. But how this takes place is defined fundamentally by how we see the world, and the objects in it, as something to be acted upon. In every event we see the shadow of our capacity to interact with the involved objects.

Of course, it's not that we don't experience this distinction between ourselves and outside objects when we're just observing them—observing is also a form of interacting—but

rather that we have this fundamental distinction between these two otherwise identical events. As far as objecthood is concerned, we don't learn the same things from observing someone kicking a ball as we do from kicking it ourselves, not *unless* we can assume these things are interchangeable to some degree. We need to see the possibility of our own interactions with the ball reflected in these observations. To distinguish these types of events then, is to make a distinction similar to that between the experience of a movie versus that of a video game. The sense of reality a movie invokes is premised on the implication of the possibility of our interacting with it (that is, with the objects and people portrayed in it) *had we been* part of it. Video games simply give us (part of) that reality. The movie's sense of reality is premised on the illusion of the possibility of something which is simply granted by the video game. Without that possibility, a movie is nothing more than an arbitrary two-dimensional pattern of pixels, and that possibility can only be implied because it can remind us of our own interaction with the world. Thus this possibility of interactivity is central to understanding things as objects, as individual entities.

Time

We may now be compelled to conclude that causal connections do really need to be bidirectional, but as we'll find, it's not that easy. All this talk of causal interaction makes a lot of intuitive sense when we're dealing with objects in the here and now, but things become a little more difficult to explain when we consider objects that are fundamentally beyond that, that is: past and future objects. The problem with past and future objects is that it seems that either kind has only a unidirectional causal connection to us, and, perhaps more problematically, each in the way opposite to the other. Past objects can't be affected by us; we can only be affected by them. Future objects, on the other hand, can only be affected by us; we can't be affected by them.

This might seem unproblematic. We might be tempted to say that, like before, if we simply have the possibility of causal interaction with these objects, we may have enough. But the problem is that that really only helps give us a *sense* of reality, not reality itself. It's a means to objecthood, to essence, which may be necessary for existence, but it is not itself existence. A ball in a video game might, in that video game, have all the causal properties of an actual ball, but in our world it still is just a bunch of code and pixels on a screen. Pixels we can affect to the point that we can—and in fact feel almost forced to—apply the concept of a ball to, but pixels nonetheless. Virtual reality is still virtual. Existence requires an actual causal connection, and if this causal connection must be two-way, then that seems to imply that neither past nor future objects exist.

Perhaps the simple solution is to say that they don't *actually* exist. That is to say, with perhaps a bit of a double meaning, that their existence is not actual, they don't exist *right now*. Past objects used to exist, and future objects will at some point exist, but right now, they don't. However, that seems to be taking a rather definitive stance in the debate about the philosophy of time, that is, the discussion about whether the things that exist always exist or whether they exist only at specific points in time. This is something we should be inclined to avoid, partly because it's simply rhetorically more powerful to base ourselves on fewer assumptions, but mostly because the Principle is supposed to actually underlie these sorts of discussions. In a way, it should be explaining something we already know. If there are schools of thought within philosophy that fundamentally disagree, then that's a problem.

Luckily, we may have an alternative solution that covers the other side of the debate. In the *Critique of Pure Reason*, Kant, in his efforts to explain the relationship between time and causation, actually finds himself relying on the concept of substance, specifically, the persistence of substance (Kant 1998, 299–316).[3] The basic idea is that time is nothing more

[3] A 182-211, B 224-256

than change, and causation is nothing more than the laws by which this change occurs. Thus causation *is* succession; the fact that events succeed one another *is* (Kant's answer to Hume). And this process of change requires a backdrop of things that fundamentally don't change in their being, i.e. things that always are. These things are substances.

So here we seem to have what we need: one account for each side of the discussion that fits into the Principle. But we can go a step further. Consider Newton's Third Law of Motion, which is that when a body exerts a force on another body, that second body always also exerts an equal, but opposite force on the first. What this tells us is something that we've been itching to say for a while now: causal interactions are, truly, *inter*active, to causally affect something is to be causally affected by it; it's a bidirectional connection. The reason we're bringing this up now, is that this implies that past objects can't really cause present or future objects, at least not directly. To put this more clearly: we can't simply say that the previous state of our self "caused" our present self, for example. By definition, these two selves at no point exist at the same time: one is exactly when the other isn't. Kant is right, at least to some degree: being requires persistence, perhaps not eternal persistence (we'll leave this open to discussion), but at least to the degree that insofar as we can determine objects (causal entities), there must be temporal overlap with other entities in order to interact with them, at least to the degree that the past can causally affect the future.

The real insight here is that Kant is also right on another front: causal connections and temporal succession are *the same thing*, or rather, the latter is constructed out of the former. It exposes the mistake of dividing up an entity, such as our selves, into momentary states, because we don't—can't exist in any individual moment, causally isolated from anything else. That's not to say that we don't exist *at* those individual moments, but temporally speaking, our existence is one-dimensional rather than zero-dimensional. Insofar as we are (or insofar as anything is for that matter), we persist. In a way, one might say that we really

exist “in between” moments, that is to say, in the causal relations that link our various states together, rather than in those moments themselves. Regardless, the point is that isolating something in an individual moment is isolating it from causation itself, because causal relations manifest themselves in temporal ones. To be a past or future object already implies a causal connection. And so, there is no real problem.

Indirection

We’ve discussed cases about moving beyond temporal limitations, but what about spatial ones? In a classical view of time and causality, spatial limitations aren’t necessarily relevant, because there are no limitations, beyond practical ones at least, to how fast information can travel. But modern physics tells us that information—and by extension causality—can only travel at speeds up to the speed of light. It isn’t immediate, and in fact it’s very absolutely limited. This poses a problem, because it seems to sever our causal connection to objects we would definitely be inclined to claim to exist.

This problem is exactly what Mark Colyvan brings up when criticizing Armstrong’s Eleatic Principle, arguing that it implies that objects outside our light cone don’t exist, and they definitely do (Colyvan 1998, 317). “Our light cone” is referring to the area of spacetime within which events take place that we could physically affect or be affected by. If one were to make a graphical representation of all the events in the universe, with us (that is, the here and now) at the center, then the area representing the events from which information could travel to us and those to which information could travel from us would have a conical shape, with one half representing the former extending into the past, and the other half representing the latter extending into the future. The borders of this cone represent the events that require light speed to affect or have affected us, and so anything beyond them fundamentally can’t affect or be affected by us.

We might be tempted to argue that this doesn't change our previous argument. Each of these events still takes place in time, and so is by definition causally connected to us. This isn't false, but it simply invites the counterargument that we're relying on an outdated understanding of (space)time. We would essentially be Begging the Question. Alternatively, we might argue that Kant is right and the universe essentially consists of eternal substances. This might seem problematic if one considers that the universe is expanding at a faster rate than the speed of light can keep up with, but since all of the universe is expanding from a single point in spacetime, in which all these substances would've been contained, it would still make every object in the universe causally connected to us.[4] We might even simply consider that this isn't really as much of an objection at all. Perhaps things outside our light cone simply really don't exist. It's not that strange of a thought. After all, we frequently refer to the observable universe as simply "the universe," and it is subject to essentially the same limitations. For all intents and purposes, this might as well *be* all that exists.

But none of these solutions are really satisfying, each coming with its own set of things we should be reluctant to admit, and each really trying to circumvent the problem rather than tackling it head-on. The real solution is perhaps much more obvious: the connection needn't be a direct one. Indeed, this seems to be exactly what our solution to past and future objects was implying: that past objects are causally connected to us because they're directly causally connected to objects that are causally connected to us, directly or indirectly. The same then applies to objects outside our light cone. We're connected to objects at the edge of our light cone. These objects have their own light cones with objects

[4] To expand on this: consider that looking at a distant object is also looking back in time. You're not so much seeing a distant star, as you are seeing that star as it was when the light was emitted; that's the fundamental idea the concept of the light cone is based on. Thus, considering that the universe is about 13.8 billion years old, when we look that far away—that is, 13.8 billion lightyears away, plus the distance added by expansion in the meantime (which adds everything up to about 46.3 billion lightyears in total)—we should see the birth of the universe, which took place at a single point. (We can't really see the exact birth for, for lack of a better term, practical reasons, but that's beside the point.) Essentially, in terms of spacetime, in whichever direction you look, you're looking in the same direction, in a similar sense as when you're standing on the north pole, every direction you can walk is south.

inside it that are outside our light cone. We're connected to these objects through the object here. If time manifests itself in a chain of causal connections, spacetime manifests itself in a network of causal connections.

This solution addresses a number of other problems we've been dealing with as well. Recall that earlier in this chapter we had considerations about the epistemic value of referring to causal connections to an observer, but that the value of this seemed to undermine the need for using the concept of causality instead of simply detectability. We also concluded that, depending on whether the causal connection must be uni- or bidirectional, the collection of things that exist is respectively either relative to the observer or absolute. We've since established the value of a bidirectional connection—and by extension causality over detectability—and this has some important implications. The reason a bidirectional connection leads to an absolute set of existing things is because it makes the reference point (the observer) interchangeable. Specifically, it makes it interchangeable with any other existing thing, any other causally connected thing. But what's important is that we established (the necessity of) this bidirectional nature specifically to *direct* causal connections. This gives us the best of both worlds: it doesn't require us to have a bidirectional connection to every existing thing, but it does give us all the benefits of being able to conclude that the set of things that exist is absolute. After all, if we are interchangeable with an object at the edge of our light cone, and it in turn is interchangeable with something at the edge of its light cone, but outside our light cone, then we are interchangeable with that object as well. This line of reasoning leads to interchangeability between every object in the entire universe, be it past, present, or future.

What this also does, is address some of the vagueness in our terminology. We established earlier that situations like these could be dealt with by arguing that a causal connection not so much implies actual causal interaction, but rather the possibility of such.

But “possibility” is a vague term. What is possible varies depending on certain conditions, and these conditions weren’t very clear until now, amounting merely to the idea that we “could have” causally interacted with a e.g. past or future object had we been there to do so. But this was really just moving the goalposts; in what sense “could we” have had that ability? In what sense was it a possibility in any meaningful way? Now we have the tools to answer this question: the Principle implies the possibility of interchangeability between reference points. This is the “possibility” we’ve been referring to; we “could have” been there to interact with these indirectly connected objects in the sense that the reference point could have been one of the objects with a direct causal connection to the object in question. To call it a possibility is simply to emphasize that on the one hand, as reference points, we are inside observers to the greater whole of the network, but on the other hand, that our specific role as a reference point is arbitrary. It’s a description of the reality of existence that does justice to our role as internal observers, while at the same time not reducing existence to mere detectability.

What Is Connected

Thus far we’ve primarily been talking about the existence of objects (and people), but we should consider that these aren’t the only types of things that could be said to exist, nor the only ones that can be said to have causal connections to one another. David Hume characterizes causation as a (necessary) connection between events, for example (Hume 1999, 134–47).[5] Another common conception of causal relations is one that describes them as connections between universe-states. That is, causal laws are the laws that describe the specific sequence of states of the universe that make up the totality of existence. So which is it? Which things are those things that are connected to each other by way of causality?

[5] In his two definitions (Hume 1999, 146), Hume still uses the term “object,” but these, to him, are the “objects” of experience, which are impressions, or events as they are perceived.

The answer is that, once again, we don't want to pick a side. Luckily though, we don't even need to. The truth is that none of these answers is necessarily incompatible with what we're defending. In fact, they're not even necessarily incompatible with each other. Certainly Hume, when describing billiard balls colliding into each other as an example of (supposed) causal interaction, isn't denying that these billiard balls are things supposedly having causal interactions with each other. True, he may be arguing that *fundamentally* we're describing patterns between events, or impressions really, but that's really the result of another metaphysical theory about what objects really are. In fact, he pretty much says this quite explicitly himself, arguing that "[t]he existence, therefore, of any being can only be proved by arguments from its cause or its effect; and these arguments are founded entirely on experience" (Hume 1999, 210). Not only does he, like us, link the concepts existence and causation together, he clearly argues that knowledge thereof must be founded on experience, i.e. impressions. He's (the pinnacle of) an empiricist, and so everything must be fundamentally founded on experience, including any idea of physical objects.

So the distinction between these theories isn't so much about which things are connected by causality as it is about which *fundamentally* are. The reason we've been talking specifically about objects thus far, is because it emphasizes the perspective which the Principle—or really the concept of causality in general—appeals to: that of an agent. An agent will have a tendency to see the world in terms of objects to be manipulated, for reasons that should be obvious, but that doesn't mean that that is what the world "truly" is (as Kant, for example, would likely have pointed out[6]). Besides this (though perhaps because of it), objects are the most obvious examples of things that are generally claimed to exist. An account of existence that can't at the very least make it seem plausible that such things might exist would've been a fairly bad one. There's a good reason Armstrong says they "certainly

[6] What the world "truly" is would be referring to the world as a thing-in-itself.

exist" (Armstrong 1978, 1:131); it's an obvious standard. The fact is that what the "true" entities, (substances perhaps) are, is a matter outside our purview, not simply in the sense that we have neither the reason nor the means to be the deciding factor, but in fact because it undermines the very purpose of the Principle. For it to make a claim about which entities are fundamentally those between which the causal connections occur would amount to arguing which causal connections are "real," which is just another way of making a claim about which causal connections exist. Obviously, it can't do this. Causal connections define existence; they can't themselves be said to exist or not.

Yet at the same time we do need to say at least something meaningful about those things that exist. There are still things that need to be distinguished from one another, else the Principle is meaningless. The goal is to walk the subtle line between implying that everything that exists must have (causal) properties, but without giving any meaningful hint beyond that as to *what* those properties are, in doing so avoiding saying anything meaningful about what those *things* are. But how can a Principle have meaning when its goal is to describe everything that is, well, a thing? The term "existence" must cover everything, or perhaps better put: every possible thing, i.e. *any*thing, and so describing it must amount to essentially saying nothing, or rather: "not nothing." How is something like that not by definition completely meaningless?

This is where the example of kicking a ball comes back into play. Recall that we made a distinction between two questions. On the one hand there was the question of whether something exists, a question of existence, whether something has causal properties. On the other hand we had the question of what it is that exists, a question about essence, i.e. what its (causal) properties are. These questions are in fact intimately related. The nature of the connections might still be up for debate (by what law does kicking the ball make it fly away?), as might the nature of the things that would be connected (which entities are

interacting with each other to make the ball fly away?), but whatever story you tell about what the universe fundamentally consists of and how its parts are connected to one another, it must in the end still amount to the same reality. As should be plain, the answer to the question about existence implies that there is an answer to the question about essence; that's what the Principle puts forward as the relevance of causality (as a marker of essence). The point is then to provide the criterium by which any hypothetically existent thing must be identified: to exist means to have causal properties, specifically causal properties that are ultimately linked to us, the point of reference. This is no different than stating that to exist means to have an essence, because to have (real) essence is to have causal properties. The Principle simply gives us the means to identify this from the reality of our perspective as inside observers, as agents, as parts of the greater whole.

A Causal Precondition

We started this chapter off with the observation that it's one thing to know what you're saying, and quite another to understand what you've said. If we were going to defend a principle, we would have to know precisely what it was we were going to be defending, because this process of coming to know it is in fact part of defending it. What we've since found out is that while this principle does have a certain appeal because of a clear desire to exclude the fundamentally undetectable, it moves beyond that; causality is more than merely a prerequisite for detectability, for experience. Conceptually it might simply be a precondition for the ability to detect, but practically, it's a condition for our ability to interfere. We don't see the world in terms of causal laws simply to make sense of it; we see the world in terms of such laws to make sense of it in a way that allows us to manipulate it. And so, as much as it acknowledges the reality of our fundamental role as observers, it acknowledges that of our role as *agents*.

That it can do both these things is the result of the fact that there's a distinction between how direct and indirect causal interactions work. A direct causal connection is always bidirectional, but an indirect one can be unidirectional, at least in the sense that something can affect us without us affecting it, or vice versa. But this can only occur via something else. This results in a universe that is much bigger than simply what we can detect, or even that which could have detected us. It sets up a universe that is essentially a network of causal connections, within which the limits of detectability are defined specifically by the limits of direct causal connections. That is not to say that we can't detect indirectly causally connected entities, but rather that we only can insofar as our direct causal interactions are traceable back to the indirectly connected entity in question. The only way to detect an entity is by detecting its (causal) properties. Thus their causal properties (i.e. essence) must be reflected in our direct causal interactions with entities through which we are indirectly connected to them. This actually helps us explain why there would be more than simply that which we can detect, while still placing a specific ontological value in detectability.

Causality, then, is the means by which we can come to know the universe for what it is. This may seem counterintuitive, given that we specifically wish to focus on the question of whether something is, rather than what that something is, but it is in fact a necessity given the fundamental connection between the two. To know that something is, is to know that it is *some*thing. Specific causal properties, or laws, like metaphysical theories about substance or empirical observation, might actually identify that something, but that's not what the Principle gives us. All it does is give us that simple identifier: that it must be causal, specifically, that it must be causally connected to us.

This is just another way of saying that for something to exist, it must have an essence, it must have a form of being. By characterizing this as having a causal connection to us, it characterizes the preconditions for having an essence, and that precondition is existence.

Existence can only ever manifest itself in things that have an essence, and essence is defined by causal connections. We can emphasize that this characterization demonstrates our nature as observers, or as agents, but really what it does is emphasize our being part of the world, our being *participants in* the world. It describes the whole of existence as a network of causal connections, from the perspective of a node in that network ("me"). But it also implies that this perspective, as a reference point, is interchangeable with any other node. It recognizes that being can only be understood as relational, because being is only ever understood by *a* being. Everything that exists, exists in some form *to us*, and so it is as much identified by what it is as it is by what we are. And what defines this relation is causality.

Chapter 2

Essence

Disagreement

In the previous chapter, we made the claim that if there are schools of thought that fundamentally disagree with what the Principle has to say, then that's a problem. If that's the case, then it seems we actually have a problem. After all, doesn't the very fact that there's apparently a need to write about this imply that such discussion is possible? Of course, that need can in part be explained by our distinction between knowing something and really understanding it—the premise of the first chapter—but this doesn't really explain the degree of disagreement principles like ours have been subjected to. Previously, we used the example of mathematics as an analogy to explain how principles might be absolutely true while not necessarily known. But most such discussion has been fairly conclusive. Discovering a principle of mathematics isn't generally a matter subjected to the kind of controversy as is often found in matters of metaphysics, yet this matter seems as controversial as ever (Armstrong 1978; Field 1989; Colyvan 1998; Oddie 1982).

We've mentioned and addressed some criticisms of Armstrong's Eleatic Principle already, but what's perhaps most problematic is the suggestion of alternatives. Some of the more reactionary criticisms we might've been able to dismiss as matters exposing the faults of these specific formulations or defenses, a result of not fully understanding the Principle after all, but the suggestion of an alternative implies *principal* disagreement. This goes against the very premise of our defense. Generally speaking, the Principle should be something we all actually agree with. If there is disagreement, it can only be the result of a lack of clarity, either about what the Principle actually says, or about our own thoughts on the matter, akin to making an arithmetical error.

But it's difficult to see how a principal disagreement such as introducing an entirely alternative theory could ever be akin to an arithmetical error. It would be equivalent to introducing an entire alternative to mathematics, something that obviously doesn't really happen. We might of course consider arguing that the problem is that these discussions are philosophical, and philosophers do have a tendency to remove themselves so far from the reality of, well, reality, that what they come up with no longer has any meaningful connection to real life. Simply put, the concept of "existence" they're discussing and finding principles for, might simply no longer have any bearing on what people in practice are actually talking about when they're saying that something exists. In such a case the Principle might well still apply, it would just be obscured by conceptual confusion.

While there may well be some truth to this, and this does technically provide an answer to the question how disagreement might be possible, it would be a rather cheap dismissal of these criticisms. It's one thing to accuse a philosopher of being too far removed from reality to properly conceptualize specific aspects of it, but to accuse one of being so far removed from reality itself is really rather absurd. After all, they do still themselves exist. This isn't like an art critic versus an artist analyzing art; the philosopher is as much a participant in existence as anyone else, at least fundamentally speaking, which is precisely the point here.

Yet there is something to be gained from this consideration, because it does still suggest that it may be possible for there to be conceptual confusion about existence. The question is whether we can find a more satisfying answer than that philosophers have locked themselves away from the real world, from existence itself, in their ivory tower. Perhaps there's something specific about the concept of existence that makes it easy for us to slip into a state of conceptual confusion. Our goal then should be to identify specifically what this is,

why it would be so easy to become conceptually confused about this matter, and how this can occur in spite of the fact that it is about something so universal.

Esse est Percipi

Our next order of business would appear to be to take a look at some of these alternative accounts of existence and see whether we can identify a source of conceptual confusion and explain not only why it is a confusion, but also why it occurred. To start off, let's consider one we've come across already: to exist means to be detectable. In Chapter 1, we put this type of account aside simply by differentiating it from an account based on causality, the implication of course being that there was enough value to what causality brought to the table to choose it over detectability.

But perhaps there isn't, one might consider. Many an empiricist has made such a suggestion, most notably Berkeley, who directly argues that *esse est percipi*: to exist is to be perceived (Berkeley 1996, 25). There is a spectrum of possibilities an account like this can lead to, depending primarily on whether entities perceived by others count. Berkeley himself appears to defer to the status of God, as an observer of everything, to account for the fact that objects continue to exist while no-one is perceiving them (Berkeley 1996, 55). But, of course, appealing to something being (potentially) perceived by yourself as a criterium of existence is one thing; appealing to that of others is quite another. After all, while we can perceive others, we can't really perceive their nature as perceivers. After all, perceiving them doesn't mean perceiving what they perceive. This means we run into trouble when we try to appeal to our own ability to perceive perceivers to justify the existence of the things these other perceivers (can) perceive.

The reason this is problematic is not that it would imply that certain things wouldn't be considered to exist which we generally do consider to; that would be reasoning ad hoc. Rather, the problem is that it raises the question of where our reliance on the idea of other

perceivers even comes from, if all our (empirical) knowledge is accessed by means of our own ability to perceive. If we can't justify our knowledge of (any) other perceivers, then how can we justify the concept of "existence" to be referring to anything beyond our own (individual) ability to perceive? Of course, one might consider that we simply shouldn't pretend it does, but that would seem to call the concept of "perception" itself into question. After all, the concept itself would imply that there is, or at least might be, something to perceive, that there is the possibility of something beyond what we perceive, beyond our impressions. It implies that there is something else that "marks"[7] existence, and that, if one is married to the idea that what exists is limited by what we perceive, this is merely a metaphysical coincidence. And what is this mark if not causality? It would seem that rather than the concept of existence relying on other perceivers, the notion of other perceivers would have to rely on the concept of (external) existence, and that this in turn is founded on causality. In other words, the only way to argue that the only things that exist are the things that are perceived (or perceivable) is by arguing that this just "happens" to be the case, that it's metaphysically accidental, a reversal of the situation we had in Chapter 1.

This seems to be exactly the direction Hume takes empiricism into: ideas are derived from impressions (Hume 1999, 96–100), and so any idea of other perceivers would have to be founded on either a direct impression of them, or be constructed from ideas derived from other impressions. Seeing as the former is impossible, the latter must be the case, resulting in a similar situation to that of causation as he finds himself bound to define it (Hume 1999, 146), which means that it faces the exact same problems. In fact, the idea of external perceivers would be subject to the same restrictions as the idea of external objects, which we've already seen Hume argue to be *reliant on* that of causation (Hume 1999, 210).

[7] Armstrong in fact refers to causality as the "mark of being" (Armstrong 1978, 2:46), taking this characterization from Plato's *Sophist* (247d-e).

So it seems, insofar as detectability could form a basis for an account of existence, it would effectively end up saying the exact same thing as an account based on causality, to the point where it would have to appeal to the concept to make its point clear. Thus it's not really an alternative account at all; we merely confuse it for another, because at face value such an account can have much different implications. The question then remains how this confusion could occur. Well, as we pointed out in Chapter 1, detectability relates specifically to direct causal connections, and confusion about an account that bases itself on causation quickly occurs when the distinction between direct and indirect causal connections isn't properly made. Essentially, an account based on detectability falls into the same kind of trap we carefully evaded at the beginning of Chapter 1: failing to recognize that causation brings more to the table than mere detectability. This demonstrates that it's easy to confuse detectability and causality because it's easy to confuse direct causal connections with indirect causal connections.

Spatiotemporal Existence

Another, and rather obvious example that we've discussed before is one that defines existence as being spatiotemporal. Luckily for us, we already identified the source of conceptual confusion for this case in Chapter 1 by pointing out that, similar to the above, these accounts are in fact identical. To state that something exists when it is spatiotemporal is to state that it exists when it is causal. Nevertheless, an objection now arises. In a discussion on the existence of mathematical entities, Hartry Field addresses the notion of spacetime potentially being causally inert, which he argues to be an outdated view (Field 1989, 69–70). This argument would've worked in our favor before, since it would simply imply that spacetime serves as another entity through which indirect causal connections are manifested, which would include past and future objects and objects outside our light cone. But now it

seems to work against us. After all, if spacetime is itself a causally active entity, then it can't itself be a manifestation of causal connections.

Interpreted as such, existence defined as being spatiotemporal does seem to imply more than a mere conceptual disagreement. Obviously, it doesn't really change what counts as an existing entity (except spacetime itself of course), hence why it wasn't a problem before, but it does change what defines existence. We essentially face the same kind of issue we saw with detectability at the beginning of Chapter 1: if merely being spatiotemporal is enough to characterize existence, then something being causal would be a metaphysically accidental property, even if it happens to be that being spatiotemporal implies being causal.

However, this needn't necessarily be as much of a problem as it seems. Recall that in Chapter 1 we also discussed that the precise nature of the entities is not part of the purview of the Principle, or really any principle of existence. In other words, it can only tell us whether an entity exists, given that it has certain properties, not which specific properties the entities that exist have. Hartry's argument is premised on an account of spacetime as a causal entity, and so perhaps this premise doesn't so much offer a different account of spacetime than Kant's as much as it distinguishes an entity which can justifiably be named spacetime, which is then distinguished from causal interconnectedness. According to such a distinction, causally inert spacetime would be spacetime considered as the manifestation of causal connections (thereby explaining why it is itself causally inert to begin with). Causally active spacetime, on the other hand, is then simply an entity identified by modern physics, still subject to the conditions imposed by causally inert spacetime (now merely identified as causal interconnectedness), but meaningfully distinguished from it.

Considering we've largely followed a Kantian line of thinking with this conceptualization, this actually sheds light onto the fact that Kant tries to argue that space and

time, as forms of intuition (*Anschauung*), aren't technically concepts (Kant 1998, 155–71).[8] Essentially then, the point is a semantical one (hence the conceptual confusion), and the fact that we may (nowadays) be inclined to argue that spacetime is itself an entity actually reinforces the notion that causal interconnectedness should be taken as the defining quality. After all, what all this implies is that, insofar as being spatiotemporal could be interpreted to *not* mean being causally connected, it precisely isn't something that can help define what it means to exist, because in that case it defines itself as an entity, something that can exist, instead of a condition for the possibility of such a thing.

So why then might we be inclined to consider an account based on spatiotemporality rather than causality? Well, part of the answer might lie in the specific way these kinds of principles are usually formulated. As opposed to a "causal connection to me," most accounts rather define these principles in terms such as being "capable of participating in the causal process" (Colyvan 1998, 314), "being causally active" (Oddie 1982, 286), or as having "causal powers" (Kim 1993, 348). The problem is that such formulations don't necessarily exclude entities we would certainly wish to. A fictional ball has causal properties; it seems to adhere to each of these formulations in that sense. What we would wish to say about fictional entities is that they are part of a different universe, and the easiest way to make that clear is obviously to refer to a separate time and space in which these entities are situated. They have no spatiotemporal relation to us, just to each other.

But of course, the same problem applies as much to referring to spatiotemporal properties as it does to referring to causal ones. After all, a fictional ball is as much spatiotemporal as it is causal.[9] The problem is the same: they need to stand in spatiotemporal relation *to us* in order to actually exist. The difference is that this is much more obvious when

[8] A 19-49, B 33-66

[9] In fact, as we've established, it should be *exactly* as much.

referring to spatiotemporal properties than when referring to causal "powers" or "activity," because we automatically think of connections, spatiotemporal relations to us, when we think of spatiotemporal properties. The major drawback of such an account, however, besides the semantics mentioned before, is that it fails to identify the empirical reality of dealing with "real" entities, that is, that we deal with them *as* causal ones. An account based on spatiotemporality can't explain why fictional balls can sometimes *seem* real, not without further involving the concept of causality. The same isn't necessary for using an account based on causal connections to us to explain why fictional balls aren't real: its causal properties (at least insofar as they define its nature as a ball) aren't traceable back to us. They are confined to the physics of the ball's fictional world. Divorced from our (causal) world and thereby divorced from existence. Thus, even as far as semantics go, our principle is the superior account.

Inference to the Best Explanation

Thus far, we've been able to explain, with some level of consistency, that conceptual confusion appears to occur as a result of effectively trying to say the same thing as the Principle. In the case of detectability this manifested in an account that effectively amounts to relying on direct causal connections, but falls short because it can't account for indirect connections; either detectability is a meaningless criterion, because there can be only one verifiable perceiver, or the idea of other observers is meaningful, which requires a more fundamental basis in causal reasoning, meaning that it, and not detectability, serves as the actual criterion. In the case of spatiotemporality, it manifested in an account that can only meaningfully serve as a criterion for existence precisely when it means exactly that a causal connection to us is required. Even if it may effectively be true, it only can be when it happens to be subjected to causal reasoning. The only way spatiotemporality can serve as a criterion

for existence, is precisely if it's conceptually reduced to mean the manifestation of causal interconnectedness.

What we see in both these cases, is that in their attempts to serve as criteria for existence (i.e. "marks of being"), they inevitably need to fall back on causality to actually make it happen. As a result, causality demonstrates itself to be something that plays a more fundamental role in the story. And so what seemed to be a potential failing of the criterion of causality at the beginning of Chapter 1, when it was considered to potentially be metaphysically "accidental" compared to detectability, actually turned out to be a failing of detectability instead, it by itself actually falling short of meeting the required standard and in the end being forced to rely on causality to truly make its case. It just happens to be capable of doing so implicitly.

But what if the same kind of accusation can actually be leveled at causality? What if there's something that is even more fundamental than causality that's really the criterion we're looking for? That is exactly the suggestion Mark Colyvan makes when he proposes his alternative: inference to the best explanation (Colyvan 1998, 326). Essentially, the argument for this case can be understood as a kind of indispensability argument,[10] the idea being that something exists when it is indispensable for explaining reality. In this particular case, the argument is somewhat more restrictive in that this explanation must fit certain specific criteria; it should be the "best" explanation.

What exactly makes something the "best" explanation is left somewhat open, and this may give us the inclination to argue that to be able to explain anything, something must have a causal connection to our world. And this is, in fact, exactly what Colyvan accuses his opponents of, arguing that "the motivation for the Eleatic Principle […] rested ultimately on

[10] As is evidenced by him repurposing the essay as part of an argument in support of the Quine/Putnam indispensability argument in *The Indispensability of Mathematics* (Colyvan 2001, 39–65).

the claim that only causally active entities can have explanatory power" (Colyvan 1998, 326). And this poses a huge problem, not so much because, as Colyvan tries to argue, causally active entities aren't the only entities that can have explanatory power, but rather that merely making the claim seems to admit to the fact that Colyvan's suggested alternative is more fundamental.

What's worse, it seems as though it might be the case that inference to the best explanation actually lies at the core of justifying the use of the Eleatic Principle to begin with. While it may be a point of discussion what exactly would constitute the "best" explanation, a principle often suggested to play a cardinal role in determining this is Ockham's Razor. And it just so happens that Armstrong himself invokes Ockham's Razor precisely as the justification for the need to invoke his principle (Armstrong 1978, 1:131). There we seem to have it then: irrefutable evidence that using causality as the criterion is in fact only possible when it is interpreted to be identical to an appeal to inference to the best explanation, and that therefore the latter is the actual criterion, whereas causality is, at best, metaphysically accidental.

If defenders of the Eleatic Principle haven't been much concerned with this, it's merely because they're not actively trying to fundamentally describe the nature of existence. For any other argumentative purpose, it's perfectly fine if it's metaphysically accidental, so long as it is *effectively* necessary. However, that should make us curious as to what exactly that argumentative purpose is. Armstrong seems to be very well aware of the fact that he's justifying the use of his principle with something else. It seems something odd is going on here: why introduce something as a fundamental principle, if you're immediately going to undercut its fundamental nature by grounding it on something else? Perhaps this justification is purely argumentative. In other words, it could be that Ockham's Razor is only more fundamental insofar as it is more central to the purpose of his line of reasoning. Perhaps he

needs to justify the use of his principle with Ockham's Razor not because that latter is more fundamental per se, but rather because he's concerned with something else than describing the nature of existence. Perhaps Armstrong is really concerned with a matter that simply requires him to involve the relevant theory in his line of reasoning. Much then as the use of mathematical principles must be justified when the use of mathematics itself has to be justified, he needs to justify his reliance on the theory and its principles. It seems then that we may have to take a closer look at exactly what Armstrong—and other supporters of his principle—are trying to achieve. Specifically, we need to look at the greater context of the debate his arguments are a part of: what problem is he trying to solve? Why (and how) does he think his principle can help solve it? And why do his opponents think it can't? These are the questions we need answers to.

Argumentative Purpose

If we're going to try and find answers to these newfound questions, the most obvious place to start would be the source. Armstrong introduces his (at that point yet-to-be-named) Eleatic Principle in *Universals and Scientific Realism* (Armstrong 1978). As its name suggests, the work is largely devoted to the matter of whether—and if so, to what degree—universals exist. At the end of the first volume, we find out that Armstrong's specific position is actually somewhat nuanced, with him describing it as a "minimal Realism." What exactly this entails becomes clear when he notes that he specifically wishes to argue that "we have no reason to postulate anything over and above particulars, their properties and their relations" (Armstrong 1978, 1:126–27). We can see here that, while he isn't a materialist, he does limit the existence of universals insofar as it must be based on that of particulars. His goal still appears to be to exclude certain types of things from being allowed the status of "existent," to find some sort of standard by which one might adjudicate which things exist and which don't.

It should be obvious what kind of role his principle might serve in this. Clearly, something that doesn't correspond to a principle of existence fundamentally cannot exist. Someone like Armstrong who's taking a stance on the existence of certain types of entities will have a need for such a criterion, especially considering that, as a minimal realist, he can't dismiss universals altogether, meaning he can't dismiss them simply on the basis of them being universals (and then arguing what the inherent problem with universals is). And considering the focus on the basis of particulars, a criterion based on causality makes sense, since causal interaction can only manifest in interaction between particulars. The Eleatic Principle, then, would essentially be a criterium that states: if something is to exist, it must play some causal role in our world. This criterium is then used to exclude a number of (types of) entities from the list of existent things. As we've already seen, this is exactly what Armstrong does, arguing that "[i]f any entities outside this realm are postulated, but it is stipulated further that they have no manner of causal action upon the particulars in this realm, then there is no compelling reason to postulate them" (Armstrong 1978, 1:130).

And indeed, we can see how this has affected the discussion as a whole. As an argument intended to imply the non-existence of certain types of entities, we would expect its opponents demonstrate themselves to be advocates of the existence of these, which is exactly what we see. They end up taking one of the two argumentatively possible courses of action: either the Eleatic Principle is incorrect (Colyvan 1998), or it is (potentially) correct, but fails to achieve its intended purpose (Oddie 1982). Additionally, we would expect his principle to be invoked in other discussions about the existence of certain types of entities, and again we see this to be the case: e.g. for mathematical entities (Colyvan 2001), negative facts (Barker and Jago 2012), and for mereological properties (Cowling 2014). In fact, what we can see in the discussions on mathematical entities, is that the idea of such a causal criterion has been

brought up seemingly independently (Azzouni 1997; Cheyne 1998; Field 1989), suggesting that this discriminatory function is a common use of such a principle.

Perhaps more notable, however, is the occurrence of a similar principle in an (initially, at least) entirely different discussion, under the name of Alexander's Dictum. First coined by Jaegwon Kim (Kim 1993, 348), the name is a reference to Samuel Alexander's stance on epiphenomenalism, specifically with regard to mental states. Alexander himself invokes the argument in order to criticize the self-defeating concept of an epiphenomenon. Essentially, his argument boils down to the fact that something the presence of which makes no difference in the world by definition cannot exist (Alexander 1966, 2:8). Kim, in turn, distills this reasoning into a principle: "[*t*]*o be real is to have causal powers*" (Kim 1993, 348), and then uses this as an argument in favor of (materialist) reductionism. In other words, this discussion takes place within the realm of the discussion on Cartesian dualism, i.e. whether—and if so, in what way—there is a fundamental division between mental and physical entities. Though arguably related, this is clearly a separate discussion, as is evidenced by the fact that a seemingly identical principle is discussed under an entirely different name, even though it's serving a similar purpose.

The reason all this is relevant, is that since we're looking for an explanation for the possibility of these discussions in general, we should take an interest in any similarities between them. And from these different examples we can now see an emerging pattern: invariably, these causal principles are invoked for the purpose of depriving certain types of entities of a certain ontological status. However, what's interesting is that this doesn't necessarily amount to the full exclusion of the entities in question. Kim argues in favor of reductionism, but explicitly not what he calls eliminativism, and perhaps most interestingly, Armstrong's minimal realism amounts to a similar stance. This is evidenced by the fact that he doesn't exclude all properties and relations from ontological status, but that he argues that

they can only exist insofar as they are properties and relations *of* particulars. In other words—his own, in fact (Armstrong 1989, 6)—what he's saying is that these properties and relations, at least insofar as they're not themselves causal, exist only insofar as they *supervene* on particulars. Thus, just as Kim does very explicitly (Kim 1993, 337–38), Armstrong establishes a kind of hierarchical ontology, with particulars having the highest ontological status.

And this is where things get interesting for us, as this highest ontological status is essentially no different from the traditional concept of substance, that is, something that exists independently of other things. Armstrong openly acknowledges this in fact, defining a substance as "what is logically capable of independent existence," and going on to state that "all particulars are substances, [and] all substances are particulars" (Armstrong 1978, 1:132). But as we saw at the end of Chapter 1, when pointing out the problems with trying to determine which things the world "truly" consists of, this type of discussion takes place in a realm beyond the purview of the Principle, and thus, insofar as they would be the same, beyond that of the Eleatic Principle and Alexander's Dictum, or really any other principle of existence. These arguments we've been discussing aren't about existence, identifying the "mark of being," but rather about determining the nature, or essence, of the things that exist. These principles, insofar as they are meant to define existence, should have nothing to say about the latter, beyond to effectively state that to exist means to have an essence. So now, we're faced with a new question: how has this apparent fallacy been allowed to occur?

Conceptual Confusion

Perhaps it's best to first review why exactly it was problematic to have a principle of existence be a determining factor in identifying these properties. Now, strictly speaking, this isn't entirely true. There is one thing that these principles do get to say about it, and that is that to have such properties, i.e. to have an essence, something *must* be causal. But of course,

this must effectively be meaningless, at least in terms of determining the nature of things, because it really only expresses what essences (entity-defining properties) are: causal relations. Essence is about the manner in which things exist, matters like ontological hierarchy, and so a causal principle of existence must imply that essence is about the manner in which things are causal. It doesn't really say anything about the essences themselves; it simply determines how they must be understood in relation to the concept of existence. It's simply because these concepts are so closely related that how one deals with one can have such a major effect on how one deals with the other. Thus, if the principle in question ends up not actually applying, the same essences should still hold, they simply need to be understood differently, in accordance with another principle of existence.

What this implies is that the only types of entities that a principle of existence could ever exclude are those of the kind of Alexander's epiphenomena, things which exclude themselves by definition, oxymorons. That this is indeed really what Armstrong is also trying to use it for, is perhaps most clearly visible when we look at a type of entity he seems to particularly enjoy singling out: possible worlds. Specifically, he criticizes the existence of the entities within them. Using his principle, referring to it as the Causal Argument, he rejects the notion of such entities altogether. One can essentially summarize his argument as follows: causality defines what is part of our (spatiotemporal) world and thereby existence. Entities in possible worlds appear to be causally disconnected from the entities in our world and as a result from our world altogether, therefore such entities by definition cannot exist (Armstrong 1989, 7). This also shows how Ockham's Razor really comes into play. The reasoning is not that we have no reason to postulate causally inert entities and therefore should assume they don't exist, but rather that we have no reason to postulate that there are causal connections

between entities in possible worlds and ourselves,[11] and therefore we should assume they don't exist.

So why does Armstrong choose to make his case in specifically this way? Well, if not for the same purpose as we have for such a principle, then it seems at least for the resulting intuitive appeal. However, perhaps the primary reason is because of his application of the principle when it comes to universals and particulars. We should consider that in such a context it makes sense to rely on a principle based on causality, because causal interaction can only take place between particulars. And since we know that particulars are a lot less controversially considered to exist when compared to universals, this seems an obvious, relevant identifier.

One can imagine the line of reasoning being as follows: causal interaction defines existence; causal interaction takes place between particulars; particulars, therefore, undeniably exist, but universals can only exist insofar as they play a role in the causal interactions between these particulars. It makes sense because not only does it exclude entities with properties so defined that they make no causal difference, it also excludes supposed properties of existent entities which make no causal difference. But this has a hugely important implication, which is that universals wouldn't be considered such if they could themselves be causally interacted with (because that would make them particulars). Armstrong's (implied) argument is really that universals cannot be meaningfully distinguished from particulars without appealing to this distinction; what makes particulars particular is inseparable from their propensity for physical (i.e. causal) interaction. The reason this is relevant for us, is that this demonstrates that the hierarchy arises not from the causal principle, but from the distinction between universals and particulars. It's simply

[11] In fact we have reason not to postulate this, as David Lewis points out (Lewis 1986, 78–81).

because particulars have a primal causal status, that, given Armstrong's Eleatic Principle, they can be argued to *by definition* have a primal ontological status.

And the same applies to the hierarchical nature of reductionism that Kim discusses. The reason it's hierarchical, or "layered," is because "the distinctive properties of entities at a given level are *reducible to*, or *reductively explainable in terms of*, the properties and relations characterizing entities at lower levels" (Kim 1993, 338), meaning that the entities at the lowest level are the ones between which the causal relations fundamentally occur. It's what defines these entities to begin with; electrons, quarks, and neutrinos are all fundamental (causal) entities posited to fundamentally explain the causal interactions between (directly observable) higher-level entities. Essentially, the point is analogous: higher-level entities only exist insofar as they are constituted of lower-level entities; they *supervene* on them. Thus we can see that the basic reasoning is the same across the board: these causal principles don't serve as arguments in these discussions themselves, they merely serve as a precondition for the ability to have these discussions by helping to define the relevant terms.

To demonstrate that this truly is the case, perhaps we should look at yet another discussion that invokes a similar principle, one that actually served as the inspiration for Armstrong's Eleatic Principle (Armstrong 1978, 2:45–46). In Plato's *Sophist*, we witness a dialogue between Theaetetus and an unnamed man from Elea, generally referred to as the Eleatic Stranger. In their discussion, they eventually come to face a problem similar to that of Armstrong: how to come up with a definition of being which would satisfy their opponents. The definition they come up with, which has served as Armstrong's source of inspiration, is the following:

> I suggest that everything which possesses any power of any kind, either to produce a change in anything of any nature or to be affected even in the least degree by the slightest cause, though it be only on one occasion, has real existence. For I set up as a

> definition which defines being, that it is nothing else than power. (Plato 2015, 7:379)[12]

As we can see, though it doesn't invoke any (familiar) concept of causality, but rather one of power, this principle is very similar, and considering the similar context, one that may well be an attempt to make the same point, but in the context of a vastly different conceptual landscape.[13] What's also interesting is that their opponents are in fact the Atomists, who they claim "define existence and body, or matter, as identical, and if anyone says that anything else, which has no body, exists, they despise him utterly" (Plato 2015, 7:373),[14] a position clearly very similar to that of a material reductionist. The other position, rather, they describe as the idea that "real existence consists of certain ideas which are only conceived by the mind and have no body" (Plato 2015, 7:373).[15] Nevertheless, they argue that while their opponents might see themselves affirmed in their position by such a definition, in fact it serves the other just as well, that the act of knowing is actually also an exercise of power upon things (Plato 2015, 7:378–87).[16] In fact, Armstrong even acknowledges Plato himself suggests this; he simply argues that this brings about problems due to the unchanging nature of Plato's Forms (Armstrong 1978, 1:128).

What this demonstrates, is that it is entirely possible for the outcome of this kind of debate to be dependent on a discussion on the nature of things like Plato's Forms rather than that of a principle of existence. On top of that, it contains an acknowledgement of the fact that one might be inclined to think that it does depend on such a principle, that that is what makes the difference. It's the answer we've been looking for, and that it truly applies to our situation can in fact be seen in Graham Oddie's response to Armstrong, who similarly argues

[12] 247d-e

[13] The term 'landscape' is chosen deliberately in favor of 'framework' in order to avoid using a more philosophically loaded term.

[14] 246a-b

[15] 246b

[16] 248a-249d

that Armstrong's conclusions don't hold even if the Eleatic Principle is true. And he does this by attacking precisely the thing that we've argued must then be the actual point of contention: Armstrong's conceptions of particulars (Oddie 1982). Whether this criticism truly holds may still be up for debate, but what should now be clear is that the outcome of this issue, and any other like it, will hinge not on the truth of any causal principle of existence per se, but rather on how the contesting entities are distinguished from one another. Principles of existence only play a role insofar as they serve as (part of) a conceptual framework within which this discussion is allowed to take form. That is why it may seem like they make an effective difference, since they fundamentally affect the semantics, but also why they ultimately never could.

The Best Explanation

So how exactly does this affect Colyvan's suggested alternative of inference to the best explanation? Well, if it were truly a proper alternative, it should similarly be meant to perform the function of conceptual framing. Yet this doesn't appear to be the case. In fact, as we already observed before, its function would have to be involved in the same points of discussion as Ockham's Razor, to the point that Ockham's Razor is in fact often part of the story of what constitutes the best explanation. We've seen that this isn't the case for the Eleatic Principle or other principles like it; it's not part of the story of what constitutes a causal connection to us.

What we can now see is that Colyvan's accusation, that defenders of the Eleatic Principle are really arguing that only causal relations have explanatory value, and thereby appealing to inference to the best explanation as a more fundamental principle, is only partly true. It demonstrates not that inference to the best explanation is a more fundamental principle of existence, but rather that it is one of essence. The reason this confusion can happen so easily is because both types of principles seem to function as the foundation for

determining the answer to the question: what exists? However, the difference is that one answers the question by defining existence, while the other answers it by defining the qualities of the things that exist. Moreover, even in the latter sense, inference to the best explanation isn't actually the most fundamental. That would rather be what actually turns out to be the best explanation: the laws of physics. Of course, that's a little unfair. What it actually does is give us the principle that determines the process of finding the best explanation, just as a principle of existence is (also) a principle that determines the process of how a discussion about essence should be framed. It provides empirical science with the conceptual framework necessary to describe the actual (causal) relationships between entities. But all this just demonstrates how easily one can get confused between these matters. The relevant answer depends on the intended purpose of asking the question, which isn't clear prima facie.

Regardless, we can now see that Colyvan's accusation is really mistaking one question for another, at the very least insofar it could ever address us, as well as how easily such a mistake can be made. However, that still leaves the other horn of Colyvan's attack: the accusation that a principle like ours can't account for all things we intuitively would consider to exist. But with the groundwork we laid in Chapter 1, this is easily addressed: a principle as such can't be refuted in the classical sense. It can't be disproved with intuitive counterexamples because it's not itself something to be determined by such intuitions, but rather itself something which determines them. These intuitions have more to do with the conceptualization of the problem than with the problem itself. They simply depend on which principle is assumed. In this sense it's simply arbitrary; any other principle would've sufficed as well. But as we've since determined, the reason it is nevertheless not so arbitrary after all, is because it serves the same kind of function as a Kantian precondition. To understand this

more thoroughly, we might consider something which Kant says in the first version of the *Critique of Pure Reason*:

> Thus we ourselves bring into the appearances that order and regularity in them that we call **nature**, and moreover we would not be able to find it there if we, or the nature of our mind, had not originally put it there. For this unity of nature should be a necessary, i.e., *a priori* certain unity of the connection of appearances. But how should we be able to establish a synthetic unity *a priori* if subjective grounds of such a unity were not contained *a priori* among the original sources of cognition in our mind, and if these subjective conditions were not at the same time objectively valid, being the grounds of the possibility of cognizing any object in experience at all? (Kant 1998, 241)[17]

What this says is that any order we may find in the world ("nature") we only find because we ourselves put it there. And this order is placed there in the form of (universal) principles. Regardless of Kant's own specific aspirations with this, our takeaway from it should be that a principle derives its meaning from the way it is applied to reality (ordering it). Their necessary truth is not a coincidental product of the nature of reality, but rather an assumption, which simply still lacks meaning before it is applied. Once it is applied, it gains meaning by how it structures reality. In this sense it's still contingent on experience, but nevertheless it remains an a priori truth. Empirical and logical statements have contingent truths and set meanings, whereas principles have set truths and contingent meanings.

Essentially, what we may now conclude, based on how alternatives seem to "normalize" to the Principle when properly explained, is that this normalization occurs because any principle still has to conform to the same reality, just not in way we're generally used to. Eventually, once the principle is normalized to fit reality, the discussion turns into one about semantics, which leaves the remainder of the issue a consideration about the best terminology. This brings us to the final argument for the Principle in favor of any others: that, unlike its alternatives, it doesn't pretend to be anything other than it is, than what existence is.

[17] A 125-126

This is essentially what we now hope to have demonstrated in Chapter 1: that, unlike its alternatives, the Principle is true to form.

Chapter 3

Existence

From the Logos to the Phainomena

At the end of the previous chapter, we concluded that the final argument for what makes the Principle superior to its alternatives, is that it is true to form. What this means is premised on the outcome of our investigation into how a principle can both express an absolute truth, yet be arbitrary. We now know that, because principles have set truths and contingent meanings, their formulation is arbitrary, while the truth they express is absolute. This means that there are considerations to be had about which formulation to favor. The appeal to the Principle being true to form is then that one should favor the formulation the form of which most neatly fits the form of what the principle is about. Besides the appeal of the kind of elegance that brings, the argument in favor of such a formulation is that it makes clear the crux of the discussions it frames, as we saw in Chapter 2. This leads to a more familiar type of defense one might initially have expected for the principle: an appeal to the merits of adopting it.

But our newfound knowledge also has another important consequence, which is that adopting a principle also means adopting the resulting meaning of its formulation and the terms it uses. After all, this meaning was contingent, and the formulation arbitrary. One might argue that this meaning may be limited to apply only within the context of the principle's application, but as we've seen, such application can—and in our case certainly does—have far-reaching consequences for other matters of discussion, precisely with regard to conceptual framing.

What this means is that the job we set out to do in Chapter 1 isn't finished, because we have only covered these implications in abstract terms. In order to understand the meaning that arises from applying the Principle to reality, we need to explicate these

consequences in concrete, tangible cases. Only these can demonstrate the *effective* impact of adopting the Principle. Thus our final assignment actually has two purposes. We need to demonstrate the merits of adopting the Principle, which comes down to seeing how it makes clear the cruxes of the discussions it frames, and we need to work out the specific consequences of adopting it, which comes down to examining how discussions are framed by the Principle. Since the former is embodied in the latter, we can do this in one fell swoop simply by examining the specific types of situations within which the concrete consequences of a principle of existence manifest. These situations are those within which there is a need to explain why something does or doesn't exist. Of these, there are three types: positive cases, negative cases, and controversial cases.

Positive Cases

Thankfully, when it comes to explaining the existence of the things that we're already inclined to believe to exist, a lot of this work is done for us by Armstrong. His claims about the certain existence of particulars in relation to the Eleatic Principle (Armstrong 1978, 1:126–32) exemplify a large portion of the story we would wish to tell here. There is an obvious link between the nature of causal interaction and particulars, and thus, there is an obviousness to the existence *of* particulars. The reason the existence of the edge of my bed (or at least the fundamental entities it is composed of) is uncontroversial is because it seems obvious that it is the thing I'm (causally) interacting with when I bump my toe into it. The same applies to the example of the ball we saw in Chapter 1: its existence, that is, the existence of a spherical object filled with air, hinges on my ability to kick it and make it fly away. That *something* there exists is still true if it turns out the ball was an illusion, but only because other causal interactions—specifically, my seeing the (alleged) ball—have confirmed the existence of at least something that could effect them.

What one may start to be seeing here as the core manifestation of the Principle's merits becomes even more apparent in how, by appealing to indirect causal connections, we addressed Colyvan's suggestion that the Eleatic Principle implies that objects outside our light cone don't exist (Colyvan 1998, 317). Specifically, it gives us the means to determine under which conditions an entity *still* exists even when it is not only fundamentally unobservable, but in fact fundamentally unaffectable by us (and vice versa).

What we can see emerging from this is that the ultimate role—and with it merit—of the Principle becomes apparent in how it provides the means by which one might "design" a hypothetical existing entity. It provides us with the basic instructions: all we need to do is causally link it back to us. Alternatively, it allows us to know which possibilities to leave open if we wish to leave open the possibility of something existing. One can see this reflected in the concept of Suspension of Disbelief as used in works of fiction, when they deliberately fall back on resemblances of our own world. The familiarity creates (the illusion of) a possibility of a causal link back to us, in doing so invoking a sense of realism, or rather, reality. Here we can see then that this merit is in fact a very directly practical thing: a blueprint for (metaphysically) designing real entities.

Negative Cases

Inasmuch as the Principle offers us a blueprint for designing real entities, it also offers us a blueprint for designing fake ones. Consider the following examples:

1. Santa Claus's nihilistic cousin, who flies around the globe on Christmas Eve on his magical sleigh without anyone noticing, but gives nobody presents because God is dead and thus there can be no good children[18]

[18] If it's any consolation, he also doesn't put coal in anyone's stocking, because there can be no bad children either.

2. An elementary particle which contains no charge, mass or energy, nor has any interaction with the nuclear forces
3. A large number of invisible balls which only ever physically interact with each other, flying around and bumping into each other in space

What these three have in common is that the existence of each of them seems highly implausible (at minimum) and that this was achieved by deliberately trying to minimalize or completely sever their causal connection to us. The first is a fictional character, somewhat of the kind we discussed earlier (just less convincingly). We've seen that Suspension of Disbelief can make the existence of a character *seem* plausible, at least in the moment, by making the causal connections to us seem plausible. However, what ultimately makes them not, or highly implausible to exist, is the ultimate impossibility or implausibility of these causal connections when we reflect on them. That it's so apparent in Santa's nihilistic cousin, is because the effort to causally "neuter" him is clearly so deliberate. Precisely the things one would want to test in order to confirm his existence have been complicated. He leaves no trace, and there is no way to morally appeal to him through one's actions. He might as well not exist.

The second example is more absolute. It's very explicitly (and, of course, deliberately) defined in such a way that all causal pathways are blocked. If there is any semblance of a possibility of its existence, it's only because it isn't explicitly defined as a particle that can't causally be interacted with. It simply blocks all possible pathways defined by physics.

What makes the third example interesting is that it is by definition causally severed from us. Yet, the balls aren't causally inert per se. After all, they can interact with each other. But they are explicitly causally severed from *us*. This makes them very precisely not existing entities according to the Principle, and the idea is that even the slightest possibility of

interacting with one of them would make them real, as one might consider with the vast amount of neutrinos passing through us right now. The fact that this possibility is excluded, invokes the idea that their world is really separate from ours. True, one can allocate spatiotemporal coordinates to them, but because they are causally isolated from us, there is nothing tethering their internal spatiotemporal relations to ours. What this really does is raise the (now familiar) issue of what spatiotemporal relations are if not fundamentally causal connections, and that, if this is the case, these spatiotemporal coordinates in reference to us are really lies, merely an illusion of something that would have been *if* certain criteria that aren't met, had been.

It's important to understand that this really relates back to the connection between epistemology and ontology we saw in Chapter 1. When we consider that spatiotemporal relations can only be given form in causal interactions, what this really means is that we can only know spatiotemporal relations fundamentally as causal relations. This is most obvious in scientific experiments where we can't (directly) see where the objects in question are, and have to use tricks like bouncing particles off of them to determine their location, similar to a bat's echolocation. But the reality is that fundamentally this applies just as much to our "seeing" spatiotemporal relations as well. We don't so much see spatiotemporal relations as we rely on basic causal laws to represent reality in a spatiotemporal framework, just like one might represent data in a graph or table.

And we do this to a fault, because the laws we base these on aren't our determinations of the fundamental laws of physics, but just the patterns in nature that serve a practical purpose to us. When we see a fish in the water at an angle, we see it in a different place than it is, to the point that we have to aim at a different spot if we wish to catch it with a spear. Similarly, when during an eclipse we see a star next to the sun, it might well actually be behind it, because through General Relativity the sun's gravity causes the path of the light

coming from the star to change. Thus we can see that where and when something "actually" is, fundamentally depends on the conditions under which they can be causally affected. Importantly, this doesn't mean that they don't exist if they can't be causally affected by us, but rather that we can't *know* they exist if they can't. The reason the relation to us (and thereby epistemology) is ontologically relevant is not because it restricts what exists to what is affectable (or detectable) by us, but rather that it acknowledges the fact that causal relations—and thereby spatiotemporal relations, and existence itself—can inherently only ever be expressed from a point of reference.

What's also important is that the involvement of causality carries with it the implication that one can't measure the properties of a thing without affecting it. In other words, to be able to find out whether something exists one has to change the manner of its existence, because the criterium for existence is that one *can* change the manner of its existence, or at least that it can be changed by something the manner of existence of which one can change. The point of reference isn't just any point of reference, it is a part of the whole of existence. This is the explanatory power, if you will, that the Principle brings along with it: the ability to conceptualize entities in a way that acknowledges all these inescapable realities, and therefore avoids situations in which impossibilities are discussed as if they were possible, e.g. viewing entities from a God's Eye point of view.

Controversial Cases

Armed with the tools we've now acquired, we can now both properly assess why there is so much discussion on what Armstrong considers to be unnecessary additions to a naturalist conception of the world: "transcendent universals, a realm of numbers, transcendent standards of value, timeless propositions, non-existent objects such as the golden mountain, *possibilia* and/or possible worlds, 'abstract' classes which are something more than the aggregate of their members, including unit-classes and the null-class"

(Armstrong 1978, 1:128). For each of these, the case must be that their causal connection to us is unclear, that this is largely the result of a lack of clarity on the proper (causal) conditions of what makes something existent, and that thus the proper means to resolve the lack of clarity would be to redefine them in terms of their (potential) causal connection to us (if they have any). Since there is a lack of clarity, oftentimes one would be presented with multiple possibilities, some of which may define a causal connection to us while others may not. This provides us with clear conditions for their existence: if they are meant to be understood as one of the former, then they do exist, otherwise they do not.

Of course, we have seen in the *Sophist* (Plato 2015, 7:378–87)[19] that this might not result in as clear-cut a case as Armstrong would like to have it. If it does, it can only be for one of two reasons: either Plato's arguments face more difficulties than were apparent to him, or the conceptual landscape has changed so much since then that Plato's Forms themselves would need to be redefined. Coincidentally, these correspond exactly to Armstrong's own response (Armstrong 1978, 1:128–29). It seems then that, at the very least, a principle based on causality leaves less wiggle-room for those that advocate the existence of these entities, which is probably what prompts them to attack such principles to begin with. However, as we've seen in Chapter 2, it really appears to be the only option. Alternatives seem to either be forced to fundamentally make the same claim, or, like inference to the best explanation, can't actually serve as an alternative, in fact themselves requiring a principle of existence to function, bringing one back the problem that there is no viable alternative.

To see how this would work in practice, perhaps we should examine a more familiar example of a controversial case when it comes to debates about whether certain things exist: God. Conceptions of God can very wildly, ranging from a physical being that walks around in Paradise, to a disembodied voice or all-seeing eye in the clouds, to an abstract idea. It

[19] 248a-249d

should be noted that this corresponds to a range of increasingly philosophical conceptualizations, which is to be expected, considering that the debate about God's existence becomes increasingly metaphysical as he becomes less physical, so to speak. The problem, of course, with more physical conceptions of God is that it implies that determining his existence would be a matter of physics, but history has driven physics to be unable to offer a definite answer to this question. Armstrong appears to agree, stating:

> [R]eligious thinkers often used to think of God as outside but intervening freely in Nature. He might give victory to the righteous or answer prayers for rain in defiance of the way that matters would have shaped if nature had been left to its own devices. But those who still believe in a transcendent God are increasingly reluctant to believe that he acts upon nature in this way. They may hold that God created Nature, and created it for a purpose which is still working itself out. But they are likely to believe that he does not intervene. (Armstrong 1978, 1:129)

What we see is that as one moves along this trend, eventually the relationship between us and God becomes akin to that between a fictional character and their writer, which would imply that the existence of such a God must imply that the nature of the actual causal relations between things are beyond what we can know, in the same sense that the causal relations between characters in a story are, in reality at least, fundamentally rooted in the causal laws of the writer's world, not the causal laws of the fictional world they live in (even if it must still abide by them). This means that one seems to become tied to a Kantian conception of God, in which belief in God is a matter of faith on a fundamental level. The only alternative is to somehow reconnect God with nature. Perhaps this is what prompted philosophers such as Spinoza and Hegel to adopt a conceptualization of God that equates him to Creation itself, as, to put it a bit crudely, a last-ditch attempt to salvage the philosophical understanding of God. Regardless, the result is that the existence of God is reduced to essentially two possibilities: either he exists outside our ability to detect, or at least know, similar to objects outside our light cone, or he exists within our ability to detect and we simply don't usually recognize his manifestation(s) as him. Of course the Principle isn't

necessary to get to this. Rather, it's that when we end up here we find that it's the most clear expression of the crux of the issue. What we can see then is not only that the Principle helps us in these discussions by reducing them to be dependent on very concrete conditions, but also that these conditions are exactly the ones that have manifested in the history of the debate. Philosophical discussion will always be bound to the crux of the debate, which appears to be precisely where the Principle predicts it to be.

One unique exception to this might be causality itself. It, perhaps along with causal connections, requires a separate explanation, since the Principle would imply that such a discussion should be impossible. After all, if causality is itself a concept necessary for defining what existence is, then the existence of causality is a paradoxical consideration, like thinking of an event that occurred before the beginning of time, or outside the confines of space. However, that this debate can genuinely occur we've already explained in Chapter 2, since it simply depends on the possibility of considering that another principle of existence applies. Perhaps the most telling manifestation of this is the fact that the primary argument used against the existence of causality is Hume's skeptical argument. As we've seen, Hume himself actually endorses the idea that "[t]he existence … of any being can only be proved by arguments from its cause or its effect" (Hume 1999, 210). He doesn't himself believe his argument to even be able to refute the existence of causality, likening such a suggestion to what he calls excessive skepticism (Hume 1999, 207).

The problem with skeptical arguments, he points out, is that they can make us doubt pretty much anything. But such doubt is only temporary; it only manifests as a lack of our ability to justify our belief rather than an actual lack of belief. He points this out by arguing that while a skeptic "may throw himself or others into a momentary amazement and confusion by his profound reasonings; the first and most trivial event in life will put to flight all his doubts and scruples" (Hume 1999, 207). The fact remains that there are simply some

things, like causality, which we can't help but believe in. No matter how much of a skeptic you are about causality, you'll still try to dodge someone trying to hit you in the face. Hume, advocating an empirical science akin to psychology to serve as a new metaphysics, a first philosophy, seems to think this is a psychological manifestation; we can't help but think in terms of causal laws for some reason. We would argue, essentially as Kant did, and as we've seen at the end of Chapter 2, that in fact it is the result of how principles work.

Self-reference

There is one final merit to the Principle which actually relates back to what it means to be true to form. To understand this, we need to delve back into how exactly principles are applied to reality. Recall that in Chapter 1 we determined that, despite the fact that there is an (apparent) arbitrariness to the truth of a principle, principles can be absolutely true in the form of a kind of hidden assumption, something one becomes aware of by reflecting on the meaning of what one says when the principle is expressed. In Chapter 2, we determined that this meaning is actually conditional, contingent upon the principle's application to reality. Reflecting on what one says when one expresses a principle *is* the act of applying it to reality, the act of making it fit, whereby it *gains* meaning. Thus the arbitrariness is actually a product of the fact that principles lack meaning until applied to reality, and the absolute truth they express is not a product of their presumed meaning (as with logical truths), but rather their meaning is a product of their presumed absolute truth.[20] That it is able to express this truth, is only possible for the same reason a logical or empirical question is able to express its truth: it can be applied to reality. Simply put, principles are specific expressions of the more general statement that reality is real, and in their specificity they shape a realm of possibilities,

[20] In German this can be beautifully expressed as a *Voraussetztung* (premise), i.e. something placed ahead.

creating a kind of conceptual space within which reality can (and must)[21] take form as one of these possibilities. This is how they're able to serve as transcendental preconditions.

The way this takes form should be more clear when we return to the example of mathematics if we try to understand it in a way somewhat reminiscent of Wittgenstein's Picture Theory. Consider that, while mathematical principles would imply that it is necessarily the case that 7 + 5 = 12, the principles of mathematical symbology certainly do not. That is, in mathematical symbology it's perfectly "possible" that 5 = 12, and it's in fact *because* of this possibility that 7 + 5 = 12 is a meaningful truth, since now it can be wrong. The same applies to the Principle; it is "possibly" not true that something exists if and only if it is causally connected to me, in the same sense that 7 + 5 possibly doesn't equal 12, or perhaps more convincingly, that 35782 × 67821 possibly doesn't equal 2426771022. That this truly is a genuine kind of possibility should be evident from the fact that we simply do not know at first glance whether it is. It's "possibly" false; we have to check. The Principle, then, is only not necessarily true insofar as 7 + 5 = 12 isn't necessarily true, that is, it's entirely possible for another system of representation to be adopted.[22]

Now let us consider a similar situation, concerning the principles governing electric circuits, specifically logic gates. Logic gates can essentially be described using the same principles as propositional logic. There is, however, one glaring exception to this, which becomes apparent, when one, for example, decides to connect an inverter (equivalent to the negation symbol) to itself. In an electric circuit, this results in what is called a clock, an oscillating signal, turning itself on and off at a consistent frequency, dependent on the time it

[21] There is no escaping this. Reality *must* be one of the possibilities. This is an inevitable product of the interplay between the principle and reality. The act of its application to reality is the act of implying that reality is one of the "coordinates" in its conceptual space.

[22] In another mathematical "language," for example, it might be gibberish. But more relevantly, in what would be another jargon it would be incorrect. This could be as simple as switching the functions of the "×" and "+" symbols, but also much more complex. The limiting factor is that it should still be possible to recognize the statement as a proposition, something that expresses a (possible) truth.

takes for the signal to travel from the output back to its input. This is impossible in propositional logic; it would result in a paradox, equivalent to a proposition that is false if and only if it is true, inexpressible in its logical symbology,[23] and even if it were expressible, unresolvable, akin to a division by zero in mathematics. Thus one might say that even though propositional logic can accurately and almost perfectly describe the nature of electric circuits, it falls short, because it fails to express a fundamental aspect of electric circuits. It fails to be true to form, because its form doesn't neatly represent the nature of the truth.

The interesting thing is that this also applies to language, which allows us to express the statement: "this statement is false." Barring any context that would imply that "this statement" is referring to any other than that which it is a part of, such a statement is perfectly analogous to the clock we described earlier. As such, propositional logic runs into the same issue; it can't adequately reflect the reality of linguistic expression. This issue isn't restricted to propositional logic either; *any* logic that tries to include the possibility of a statement being about itself runs into the same issue. Invariably, it leads to the possibility of paradoxes (or the similarly problematic possibility of circular statements, e.g. "this statement is true") for the simple reason that such self-reference allows for the possibility of something to be (part of) its own truth condition.

Now there is a rebuttal to this, which is that one might say that the fact that statements like these are paradoxical *means* they aren't (meaningful) statements at all, that either "this statement" actually does, and must, refer to something else, or that the statement is nonsensical. This is nothing else than to argue that language is defined by logical principles, that language is, principally, logical. And as we know, that's perfectly legitimate; a principle can be applied like that without issue per se, and by definition state truth. The problem is that it is in denial of the reality of language. We all know that "this statement" is referring to the

[23] Though this might easily be solved by introducing a symbol for an infinite sequence of negations.

statement it is part of; such is evidenced by the very possibility of the discussion we're having right now, if not the fact that we have a meaningful concept of a paradox. A logical principle of language simply redefines language to exclude the illogical, and the problem with that is not that it is incorrect (which would be impossible), but rather that it obfuscates what's actually going on when we use language, when we employ paradoxes. It hides the practical reality of (linguistic) meaning behind a veil formed by a logical interpretation of language.

But what's so significant about this is the reason why self-reference is possible in the first place, which is the simple fact that linguistic expressions are themselves part of the reality they are meant to refer to. This is really unavoidable; the alternative would be for language to only be able to refer to reality *minus* linguistic expressions, but that is not reality. Reality is the totality of existence, which must include even statements about it, lest they not exist, and lest we have to explain why including such statements is anathema, even though all other statements would be fine. It's making the fundamental error of Cartesian dualism, and Newtonian physics, the idea that we can pretend to adopt a God's Eye point of view, that the world is a system isolated from us when we try to fundamentally describe it. Is it useless? Of course it isn't! It has proven to be extremely practical. But eventually, we run into not just practical issues, but fundamental ones, where we need to account for the fact that we are inherently part of the system we try to describe, and that the means by which we acquire knowledge of (i.e. measure) the system are inherently its own manifestations.

And this, finally, is where the true value of the Principle shows itself, because it is exactly this reality that "a causal connection to me" expresses. The reason we can't adopt a God's Eye point of view, why we fundamentally can't pretend to be outside reality, is our inability to be affected by reality without affecting it. It is because these causal connections to us are both a prerequisite for our ability to know things as well as the reason that we can't

know them as they are in themselves, that self-reference is possible. Thus the Principle is itself the condition for the possibility of self-reference, and thereby an expression of this inescapable reality of reality. Causality expresses the fact that to detect something means to affect that thing; the connection to me expresses the fact that there is a greater whole understood from the reference point of a part of that whole. Any principle that fails to acknowledge these facts will inevitably enable the illusion of a reality we can both be and not be a part of, and thereby a reality that cannot be.

Conclusion

At the beginning of this paper we asked the question of what it means to exist. We can now see why addressing such a question as a principled one was an effective approach. Since principles have an absolute truth and a contingent meaning, instead of vice versa, the meaning of existence could be revealed by its application to reality. In large part this was already achieved in Chapter 1, where we were able to establish that existence is to be understood as being part of a large network of causal relations. It is a being-part-of-the-world, understood in the specific sense that what makes one part of the world is one's ability to affect, and be affected by, the things in it.

The trick then was to determine how to distinguish this from alternative principles, especially with relevance to what degree they could even be alternatives to each other, which was achieved in Chapter 2. There we could see that in effect each alternative is forced to conform to the same fundamental reality, and thus must end up expressing the same fundamental truth that we established in Chapter 1. This is because this reality isn't captured in the formal truth of the principle, which as with any principle is merely assumed, but rather in its application to the reality it is meant to be a principle of: in this case existence itself.

The final task was then to finish the story in Chapter 1 by laying bare the specific consequences of applying the Principle, and as a result demonstrating the merits of adopting it. This we achieved in Chapter 3, in the first sense simply by examining how this would manifest in practical ontological considerations, where it could demonstrate itself to take the form of concrete criteria to which the resulting explanations had to conform. The Principle's greatest merit, however, was its ability to avoid fundamental problems that would be the result of a disingenuous representation of the nature of reality, by avoiding the problems of an account which fails to acknowledge two fundamental facts: that detecting something is inherently tied to affecting it, implying causality, any being is only ever understood from a

reference point that is part of a greater unity, that existence is only ever manifested in relation to a thing that exists. And it's precisely this which naturally takes form in the principle that something exists if and only if it is causally connected to me. Thus the Principle makes existence itself more understandable.

Perhaps this seems like a fairly dry and analytical understanding of existence, but a keen observer may have noticed that our central question, formulated as it is, bears a striking resemblance to Heidegger's *Frage nach dem Sinn vom Sein*. This formulation, like many others in this paper, was no accident. In fact, when we look at the very first words on the very first page of *Sein und Zeit*, we find ourselves confronted with a quote from an eerily familiar source: Plato's *Sophist* (Heidegger 2006, 1). What we mean to imply with this is that the account of existence we've put forward here in fact in many ways bears a striking resemblance to his account of being, and can perhaps even be understood as an analytical reconstruction of his ontology.

This carries with it the implication that, just like Heidegger's ontological phenomenology, the Principle can form a metaphysical basis for dealing not only with analytical issues, but existential ones. Therefore, as much as it can help us deal with fundamental metaphysical questions about things such as universals or possibilia, it can serve as a means by which to frame and understand existential questions and struggles. This means that the Principle has very real and concrete implications for everyday life. An understanding of what it means to exist implies an understanding of what makes for a meaningful existence.

And we would suggest that this sense of meaning is captured in a particular experience, an experience that embodies the very core of the Principle, because a causal connection to me is as much a prerequisite for the possibility of a thing's existence of as it is a prerequisite for the possibility of my knowing the thing. To be able to know something is to be able to affect, and be affected by, it. But this means that in order to know something, we

need to interfere with it, we need to involve ourselves in its being. We can never know it for what it is in itself, because we can only ever know what it is *in relation to us*. It's because we are part of the same causal system that we have epistemic access to things, but it's for this same reason that we can't have epistemic access to what they truly are. Thus our experience of existence is characterized by the fact that we can only know things by virtue of the fact that our ability to truly know things is fundamentally undermined. In other words: what we here suggest is that the experience that most directly captures that of existence, is irony.

Reference List

Alexander, Samuel. 1966. *Space, Time, and Deity: The Gifford Lectures at Glasgow 1916-1918*. Vol. 2. Palgrave Macmillan.

Armstrong, David. 1978. *Universals and Scientific Realism*. 2 vols. London: Cambridge University Press.

———. 1989. "The Causal Argument." In *A Combinatorial Theory of Possibility*, 3–13. Cambridge: Cambridge University Press.

Azzouni, Jody. 1997. "Applied Mathematics, Existential Commitment and the Quine-Putnam Indispensability Thesis." *Philosophia Mathematica* 5 (3): 193–209.

Barker, Stephen, and Mark Jago. 2012. "Being Positive About Negative Facts." *Philosophy and Phenomenological Research* 85 (1): 117–38.

Berkeley, George. 1996. *Principles of Human Knowledge and Three Dialogues*. Edited by Howard Robinson. New York: Oxford University Press.

Cheyne, Colin. 1998. "Existence Claims and Causality." *Australasian Journal of Philosophy* 76 (1): 34–47.

Colyvan, Mark. 1998. "Can the Eleatic Principle Be Justified?" *Canadian Journal of Philosophy* 28 (3): 313–35.

———. 2001. *The Indispensability of Mathematics*. Oxford University Press.

Cowling, Sam. 2014. "No Simples, No Gunk, No Nothing." *Pacific Philosophical Quarterly* 95 (2): 246–60.

Field, Hartry. 1989. *Realism, Mathematics, and Modality*. Oxford: Blackwell.

Heidegger, Martin. 2006. *Sein Und Zeit*. Tübingen: Max Niemeyer Verlag.

Hume, David. 1999. *An Enquiry Concerning Human Understanding*. Edited by Tom L. Beauchamp. New York: Oxford University Press.

Kant, Immanuel. 1998. *Critique of Pure Reason.* Edited by Paul Guyer and Allen W. Wood.

Cambridge: Cambridge University Press.

Kim, Jaegwon. 1993. “The Nonreductivist’s Troubles with Mental Causation.” In *Supervenience and Mind: Selected Philosophical Essays*, 336–57. Cambridge: Cambridge University Press.

Lewis, David. 1973. “Causation.” *The Journal of Philosophy* 70 (17): 556–67.

———. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.

Oddie, Graham. 1982. “Armstrong on the Eleatic Principle and Abstract Entities.” *Philosophical Studies* 41 (2): 285–95.

Plato. 2015. *Plato in Twelve Volumes*. Edited by Jeffrey Henderson. Vol. 7. Cambridge, Massachussets: Harvard University Press.

Putnam, Hilary. 1981. “Brains in a Vat.” In *Reason, Truth and History*, 1–21. Cambridge: Cambridge University Press.