9

Experience and the Arrow

L. A. Paul

A metaphysical theory of the world is designed to give us a model of the nature of the world (Paul, 2012). If the model correctly represents the world, truths in the model are truths about the world, that is, they are truths about reality.

To develop this sort of metaphysical theory, we draw on many sources, including experience. Experience of the world, or of the part of the world that the theory is targeting, is usually taken as a defeasible guide to reality. In particular, the realist metaphysician sets aside skeptical and antirealist concerns and (defeasibly) takes experience to give us evidence about the nature of the world. We start with a defeasible assumption that the world is as it seems to us phenomenologically, that is, as it seems to us in ordinary experience, and then we go from there. In other words, contemporary metaphysics starts with the manifest image—the world as it appears to us—as the basis for its account of the nature of reality, and then refines that account in response to empirical scientific pressures as well as various theoretical pressures (Maudlin, 2007: p. 127).

Now, when we draw on the manifest image to develop our metaphysics, we are drawing on how the world seems to be, given our experience. I'll describe this attention to our experiences of how the world seems to be as attention to the *manifest features* of the world. We attend to the manifest features of the world that are relevant to metaphysical questions when we attend to the way we seem to be able to experience causation, time, objects, locations, events, and so on. If, for example, we want to develop a metaphysical theory of causation, we start by attending to the manifest features of the world that involve the experience as of causation, such as the cue ball causing the eight-ball to drop into the corner pocket, Suzy's throwing a rock causing the window to break, and pressing the volume button to generate a louder sound.

But are all manifest features created equal? No. Not all of our experiences should be treated equally in terms of the evidential support they give to our metaphysical theories. In particular, evidence from attending to the manifest features of the world

I'm indebted to Ross Cameron, Matt Kotzen, Barry Loewer, Tim Maudlin, Ram Neta, Ian Phillips, Susanna Siegel, Brad Skow, Josh Tenenbaum, and Alastair Wilson for comments and discussion.

that is undercut by empirical and theoretical work in the sciences is disqualified as observational evidence. If we know, for example, that unseen, entirely unrelated physical forces caused the window to shatter just as the rock arrived, then even if it seems to us as though the direct contact of Suzy's rock with the window caused the window to break, we know that we need to discount the evidence given by this manifest feature. This experience does not give us observational evidence. Evidence for a theory can be undercut by evidence that supports an alternative possibility: our observational evidence for the theory that Suzy's rock caused the window to break is undercut by our evidence for the possibility that other physical forces caused the window to break.

This is a familiar fact: over time, as they developed, physics and other natural sciences have undermined the status of the evidence we seemed to get from many of the manifest features of the world. We no longer think that the sun revolves around the earth, that velocity is absolute, or that surfaces without visible dirt must be sanitary. However, as I will discuss in more detail, there is a less familiar way in which the evidence given by manifest features of the world can be undercut by science. In particular, the evidence given by manifest features can be undercut when we discover ways in which these features are mere effects of cognition, that is, if we discover that these experiences are merely byproducts of the way human cognitive systems respond to the world and process information to generate experience. When such experiences, for example, experiences involving visual illusions, are discovered, the observational evidence they seem to provide can be disqualified.

If we are to start with the manifest image, we need to use what we know from physics and other natural sciences about the unobserved realm to refine or correct the image. But we must also know how the features of the manifest image are cognitively generated, so that, if we need to, we can refine the image by setting aside manifest features that are illusory byproducts of cognitive processing or other psychological effects.

The suggestion is not that we should simply look to psychology, find evidence that our experience of the world is cognitively constructed, and then simply dismiss such experience as the product of some sort of psychological bias. That is too quick. And the suggestion is not that, reasoning from the possibility that the manifest image is constructed, we should adopt a pragmatic stance towards metaphysics. That is also too quick. Matters here are more subtle. The suggestion I am making is that we acknowledge some of the deep connections between our cognitive response to the world, as discovered and developed by empirical work in psychology and cognitive science, and our theorizing about the nature of the world. These connections need to be recognized and explored, and ideally used to enrich and motivate a developed metaphysical theory of the structure of reality, one that can fit an account of the nature of the fundamental metaphysical features of the world with our cognitive response to the world.

Before proceeding further, I want to highlight a feature of illusions. Consider a common optical illusion we experience when driving on a straight road on a hot

day, the illusion that there is a pool of water on the road at some distance from the car. This illusion is easily recognized as such, because it disappears when we get close enough to the spot where the water seemed to be. Many visual illusions can be eliminated or detected through this sort of closer inspection of the stimulus.

But some of the most interesting and relevant ways in which our experiences are illusory cannot be detected like this: the illusions are perceptually inflexible in a certain way. These pose a special danger to philosophers, because unless we know independently that their status as evidence has been undercut, we might not detect them as the illusions that they are. The class of cognitive illusions that I am most interested in includes such illusions; they are often described as 'modular', that is, they are fast, automatic, and immune to belief revision. Thus, our experience of them can persist even in the face of conflicting information from a different type of source, such as the viewer's beliefs or knowledge. For example, in the Müller-Lyer illusion (shown in Figure 9.1), when you experience the illusion, you either need to know how the illusion works or you need to measure the lines (or remove the arrowheads) to discover that you are having a visual illusion.

This perceptual illusion is modular in the sense that, even after you measure the lines and discover for yourself that they are of equal length, you *still* experience the illusion. There are many other modular (nonfactive) perceptions that have this sort of characteristic, illusions that we cannot eliminate (and may not even be able to discover) by just looking a bit harder, a bit closer, or a bit more carefully.

Recognizing a modular perceptual illusion can be especially difficult when illusory parts are embedded in an experience that also has non-illusory parts. In the Müller-Lyer illusion, the illusory part of the experience is easy to pick out once we know the lines are the same length: it is our experience that the lower of the two lines seems to be shorter than the upper one. But there is much about this manifest feature that is not illusory: non-illusory parts of your experience include the experience as of seeing two lines, seeing one line above the other, seeing lines with opposite-facing arrowheads, etc.

While in this case the illusory part is easy to identify, other types of modular perceptual illusions embedded in an experience can be quite subtle and easily overlooked, especially when they are just ordinary parts of the normal, everyday

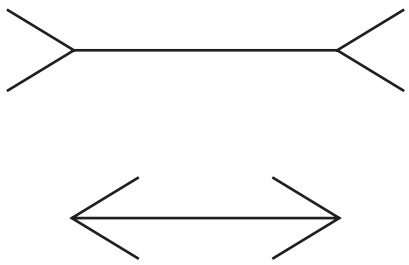

Figure 9.1 The Müller-Lyer illusion
Source: Robbins (2010).

way we experience the world. For example, as we'll see, Michotte's causal launching effect, an illusory manifest feature where an event juxtaposed with another seems to be causing it, seems like a normal, non-illusory part of ordinary causal experience, because it happens all around us, all the time, and it is often correlated in the physical world with actual causation. Until you know that the illusion is there, you won't notice it *as* an illusion (unless of course you are Albert Michotte, the brilliant Belgian experimental psychologist who discovered it).

I want to note a complication with illusions like the launching effect that will play an important role in what follows. Because the causal launching effect is sometimes correlated with causation, one juxtaposition of events that generates the launching effect can make a token launching effect a *veridical* illusion (you experience the juxtaposed events as causally related, but your experience of the sequence as causal is not due to the first event causing the second, even though it is in fact a causal sequence) while a different juxtaposition of events that generates the launching effect can make a token effect a *falsidical* illusion (you experience the sequence of events as causal, but it is not due to the first effect causing the second). Of course, there are illusions that mix veridicality with falsidicality: an illusion might be veridical in some respects but falsidical in other respects. An illusion as if *A is F* might be veridically illusory with respect to A's existence but falsidically illusory with respect to A's being F. For example, as I have noted, a launching experience might be veridically illusory in that it is an experience as of *C causing E* even though the experience is not due to C causing E, yet, in fact, C causes E. The very same experience might also be falsidically illusory in that it is an experience as of *C causing E by transmitting momentum*, but in fact, while C causes E, it does not do so by transmitting momentum to E; C causes E by other means.[1]

The metaphysician, then, must do her best to distinguish between manifest features that are illusory and manifest features that are not, and if there exist illusory manifest features, she must understand *how* they might be illusory, so as to distinguish the veridical from the falsidical. She should keep this distinction at the forefront of her mind when drawing on experiences to develop her model of the world, since she must avoid hard-to-detect illusions that are created by how we cognitively construct an experience. This means that the metaphysician is held hostage to cognitive science. If a manifest feature is illusory, but cognitive scientists (and ordinary humans like us) have not discovered this fact, the metaphysician's model of the world may not be justified. Even more importantly, if the illusion is falsidical, the metaphysician's model of the world may fail to be true with respect to the parts of the theory that draw on this feature. This hostage-taking is no more serious than what occurs when the metaphysician is held hostage to natural science, but it is much less familiar.

[1] Officially we might say that the illusory experience as of *C causing E by transmitting momentum* is not due to C causing E by transmitting momentum, but this obscures the veridical/falsidical distinction between parts of the experience.

9.1 The Temporal Arrow

These issues come into play when metaphysicians are developing theories of basic and universal entities of the world that are neither confirmed nor disconfirmed by natural science, yet seem to have observable features, such as theories of causation, persistence, and time. For example, we have many views about the nature of time that result from temporal experience. In some sense, time seems to pass, and it seems to pass at a certain rate. The present seems to have a special physical and ontological status. Temporal passage seems to be necessary for change, and passage seems to have a direction.[2] Recognizing these features of our experience can be essential to our successful functioning as agents in time, and many draw on this phenomenology to develop their metaphysics.[3] But how many of these experiences derive solely from the way we process visual and other sensory information about the world? How many of these experiences embed illusory features, or spandrels, that are the contingent result of our cognitive processing and representation?[4] We need to know whether such experience is evidence for the existence of temporally directed properties in reality, or whether the evidential status of the experience is undermined because it is merely the result of the way we respond to time as human perceivers.

I'd like to look at this issue more closely in the context of the debate over the metaphysics of the temporal direction. The debate over the temporal arrow is a debate over what fundamental ontology is needed for the temporal asymmetry of the universe, which determines the fact that time seems to be oriented or directed from earlier to later. This temporal asymmetry underlies (or, as some might argue, is the same as) the asymmetrical fact that the past is fixed while the future is open, as well as the global asymmetries of counterfactual, causal, and agential direction.[5] The issue also concerns the idea of temporal dynamism, which we can separate from the more basic idea that the universe is directed. Simply saying that time has a direction gives us a temporal asymmetry, a difference between earlier and later, and a direction of counterfactual and causal dependence where later events depend on earlier events. But part of the debate over the arrow concerns the question of whether we need more than this, that is, whether the temporal arrow isn't simply a direction, but rather is driven by something dynamic—by a primitive, dynamic, temporally forward evolution.[6]

The *antireductionist* view of time's arrow takes the arrow to be this sort of primitive, dynamic, forward evolution. On this view, there is a fundamental temporal

[2] Paul (2010a) discusses our experience as of change in detail.

[3] For the importance of temporal experience to agency, see Ismael (2011)

[4] There is a distinct and quite interesting debate over whether the content of our temporal experiences matches the structure of our temporal experience (e.g. if you experience a change, does your experience itself also change?). This is the structural matching thesis discussed by Phillips (2013). I discuss the way this issue should connect to our temporal metaphysics in Paul (MS).

[5] Albert (2000) and Ch. 8 of this volume; Price, 1996; Loewer, 2012. See also Eagle's Ch. 7 in this volume.

[6] Maudlin (2007) describes this dynamic element as 'temporal passage' but the term is used differently by different philosophers.

tendency of the universe to evolve forward by successively creating world-stage from world-stage, one stage producing the next by causally generating the [illegible] the action of productive laws of nature. It is an irreducible and primitive feature of our fundamental ontology that time passes, that is, things dynamically *happen* from earlier to later in accordance with productive laws of nature. As I understand this view, we take the notion of the world's evolving as a 'thick' or substantive notion, where successive world-stages are created or produced by prior stages, so that later stages owe their existence to earlier stages. So the forward evolution of the world involves a kind of primitively dynamic generation or production, a primitive asymmetry of reality, where successive stages come into existence, guided along by the fundamental physical laws. The Arrow plus the Laws drive reality forward, giving us the directed, productive evolution of the world.[7]

The *reductionist* about time's arrow, on the other hand, thinks that there is no need for a fundamental, primitive temporal asymmetry. Instead, time's arrow and the other asymmetries are reducible to a global entropic asymmetry. On this view, the temporal, counterfactual, and causal asymmetries reduce to a fundamental asymmetry in entropy, the one described by the second law of thermodynamics. (Loewer, 2012). The idea is that the temporal asymmetry is reducible to the worldly facts (properties and relations) that embody the fundamental physical laws, plus the fact that the initial macrostate of the universe is very low-entropy, plus, finally, a fact about a uniform, lawful, probability distribution over the possible microstates that are able to realize the initial low-entropy macrostate of the universe.[8] According to the reductionist, once we have these properties and relations, there is no further ontological need for a fundamental, primitive temporal arrow responsible for ontologically thick, productive, forward-directed evolution. The global asymmetries of entropy give us the asymmetries of causation, time, dependence, etc., and that is enough. In this way, the reductionist wields Ockham's razor to argue in favor of a simpler, more minimal ontological base.

This dispute between the reductionist and the antireductionist is not, at least in the first instance, a dispute over the physics, and both sides agree that time's arrow exists. It is primarily a dispute over the metaphysics, that is, it is a dispute over which basic or primitive metaphysical facts need to be included in the fundamental ontology.

9.2 Experience of Asymmetry

How does all of this connect to our experience of temporal asymmetry, that is, to our experience as of the direction of time?[9] I will start by describing some phenomenal features of experience and discussing their sources, and then I will explore

[7] I take Maudlin (2007) as the main advocate of this version of antireductionism, but others defend related views, such as Smolin (2013).

[8] Recent work by Wallace (this volume, Ch. 9) and (forthcoming) suggests that we may not need this probability distribution. I take no stand on the issue.

[9] Thanks to Josh Tenenbaum for discussion.

some suggestions about how these features connect to our temporal experience. By 'phenomenal features of experience' I mean to pick out the phenomenal character of experience, that is, the 'what it's like' of an experience. I will sometimes use the 'as of' locution to indicate that I am discussing phenomenology, such as 'we have an experience as of passage' to indicate that I am discussing the phenomenology of passage.

Start with the manifest features of our experience of moving in a particular spatial direction. What are some of the sources of this sort of experience? One very important source involves (nonfactive) perceptual judgments based on causal beliefs.

Imagine that you are a passenger in a moving car, looking out the window. You see a tree alongside the road, in front and somewhat to the right of you. As time passes and you continue to gaze out your window, you will experience the tree as filling up more and more of the space of your visual array. This gives you an experience that has the phenomenal character of moving towards the tree. The tree appears larger, in a sense, as a causal effect of your change in position relative to the tree. You know this, and in response, correctly judge that you are moving towards the tree.

Now, the way I just described things, it makes it seem as though your experience is largely the result of inferential perception, that is, the phenomenal character of your experience of moving in a particular direction is significantly affected by an inference you make about the causal source of some visual stimuli. But there may be a significant noninferential component here as well, because the visual system processes information in a way that takes what counts as a 'background landscape' as fixed relative to a moving observer. In other words, it isn't just your background causal beliefs about the sources of your experience that affect the character of your experience; features of the way the visual system processes information also affect it. Such features are independent of our beliefs and knowledge, and they can result in 'modular' or 'encapsulated' perceptions.

Consider a different case. You are sitting on a delayed New Jersey Transit train. There is another train next to you. As you look out the window at the train next to you, you see that your train is finally leaving: your train is moving backwards, pulling out of the station, and you feel a slight physical sensation from your gut corresponding to that.

Or so you might think. What is actually happening is that the train next to you is moving forward. Alas, your delayed train hasn't moved a millimeter. What happened? You experienced an illusion of motion in a certain direction, due to a mistaken sort of interpretation: you perceptually interpreted the sequence of images you had of the side of the train you saw out the window (the lettering on the train, the window, the door) in the wrong way. The change across the series of visual impressions was consistent with your moving backwards along a straight line causing a sequence of different qualitative experiences of the (stationary) train outside your window. But these experiences are also consistent with the other train moving forward along a straight line in the opposite direction while you remain stationary. A train moving forward that is viewed by a stationary observer can create a series of impressions in

the observer that are qualitatively indistinguishable from those created by a stationary train viewed by an observer moving backward.

Your visual illusion that your train was the one moving is an experience that was caused by your series of visual impressions, but not just by these impressions—other things mattered, such as your visual system's processing, which took the default situation to be that you were moving while your background was fixed. In this case, your causal beliefs still played a role. But you can have this experience as of moving backward even when you know your own train is not the one moving, and if you do have it even while you know you aren't moving (this is easy if there are no conflicting visual cues that disrupt your experience), your experience as of moving backward is a falsidical modular perceptual illusion.

The strength of the modular component of experience can be brought out by considering a related example of perceived direction. When a subject is placed on a chair inside a cylinder with the inner walls painted in a pin-striped pattern, and the cylinder is then rotated, initially the subject experiences herself as sitting stationary, with the wall moving. But after about thirty seconds, the phenomenal character of the experience will flip, and the subject will experience herself as spinning, with the cylinder stationary, even though she knows she is not moving. This effect is so intense that subjects will sometimes vomit. The example brings out just how strongly, when conflicting visual cues are absent, features such as the visual system's default settings can influence our experience.

The point of these examples is to show that we can have very intense experiences as of moving in particular directions, or as of other things relative to us moving in particular directions, but where these are highly constructed experiences, usually generated from a mix of inferential and modular perceptions. By focusing on the intensity of these experiences, I mean to emphasize how, when we have these experiences in the right situations, they can dominate our phenomenology, such that it can seem as though we are simply directly and transparently latching on to features of our environment in an unconstructed way. Now, in some cases, perhaps we really are latching on to features of our environment via some sort of pre-representational cognitive process. But such a process is still sub-personally constructed, even if it doesn't seem that way.

In other words, as the empirical work in psychology shows, there are many properties of our experience that, even if they are veridical representations of features of the environment (such as the case where we are moving closer to the tree), and even if, say, the phenomenal character represented by these properties of our experience intrinsically matches the features they represent, they are *constructed* features, that is, they are often the result, at least partly, of pre-representational cognitive processing. And this matters, for such cognitive processing can result in the possibility of illusions, both veridical and falsidical, when the experience we have is caused merely by normal cognitive processing as opposed to being caused by the right feature of the external world.

For example, in the case of our experience of motion on the train or inside the cylinder, we are correctly perceiving that there is motion, but we are incorrectly perceiving *ourselves* as moving, due to the fact that our visual system takes the background to be fixed as its default. Our experience embeds a falsidical illusion. In the case of our experience of moving towards the tree: we are correctly perceiving that there is motion, *and* we are *also* correctly perceiving that it is we who are moving, but our perception that it is we who are moving is not due *merely* to the fact that we are moving, it is also due to the fact that our visual system takes the background to be fixed as its default. Here, our experience embeds a veridical illusion.

Another sort of experience that is relevant to our discussion of experienced asymmetry is one I introduced earlier: our experience as of causation, the sort of experience described in Michotte's launching cases. Such cases can bring out features of the way we experience the world as causally generative. Launching cases are cases involving a special sort of experience as of causation (even though we don't directly observe the causal relation). In such cases, we have a strong visual experience as of causation. Classic versions of such cases describe things like a cue ball hitting an eight-ball and knocking it into a corner pocket, and *an image of* a cue ball seeming to 'hit' *an image of* an eight-ball on a computer screen: both create an experience as of causing motion. What Michotte discovered was that the observation of events arranged in the right spatiotemporal manner creates an illusion, whether veridical or falsidical, in the perceiver as of forwards productive causation, the sense that one event (the cue ball making contact with the eight-ball) is causally launching another event (the eight-ball's rolling into the corner pocket).

There are three features of Michotte's intensely interesting discovery that are important here. First, we have these experiences when the right spatiotemporal arrangements exist, entirely independently of whether the perception is veridical, that is, there may be causation, but there may not be causation: we will have the causal impression all the same. Thus, a description of the experience as a causal perception is not necessarily a description of a veridical perception.

The second important feature, for our purposes, is that these causal impressions are modular. We can have these experiences even when we know there is no causation involved. For example, when an experimenter shines a light on a wall, and moves it so that it seems to contact another light on the wall, and this second light moves away just as the first light stops, observers seem to see the first light causing the second to move, even when they *know* the first light did not cause the second light to move—they can see the experimenter, and they know that she moved the second light just after moving the first light. We have these sorts of modular causal impressions all the time. When we watch films or play video games, we are seeing images juxtaposed in a way that create a sense of watching one thing cause the next, and the fact that we see them as causal is part of how we play the game. Visual effects like these are the basis for the construction of special effects and video imaging, and of course for many sorts of 'virtual reality' effects.

The third important feature is that, from infancy, these felt causal impressions play an important role in the development of our understanding of causation. Psychologists have done a lot of fascinating and important work on this point, showing how infants develop the ability to ascribe causation to the world and use that to learn and manipulate causally, and these experiences are thought to play a central role in developing our sense of the causal structure of the world as adults and our ability to predict the way the world is going to unfold.[10]

The lesson to draw from this psychological work is that phenomenological experience is not just a collection of raw, unadulterated visual impressions that we experience in some unmoderated way. Phenomenology is constructed, some of the features of our experience are not a priori detectable or at least are not obviously recognizable as constructed, and much of our experience may be composed of veridical and falsidical illusions (Saxe and Carey, 2006). Thus, we may need cognitive science to bring this fact to our attention, in order to understand how such experiences are constructed and to identify veridical and falsidical illusions that might not be obvious to the ordinary observer. This is why it is important, when we are considering modular perceptual illusions, to distinguish experience of such cognitive illusions from more familiar types of illusions, such as the illusion of seeing water on the road on a hot day. We are not talking about a mere distortion of sensory information that would be detectable using ordinary methods of varying one's viewing conditions.

These reflections on our experiences as of asymmetry, direction, and causal production have important connections to the metaphysics of time. First, the metaphysics of time is deeply related to the metaphysics of causation. At a minimum, causation and time are intertwined on the assumption that causation is a relation from earlier events to later events. But a deeper connection derives from the fact that the asymmetry of the temporal direction may be the basis for the asymmetry of counterfactual dependence, which in turn underlies many of the asymmetries of causation and causally directed laws. Causation, time, and the lawful evolution of the world are bound up together as an asymmetrically directed, evolving, multifaceted process, where temporal passage is the driver for the rest of these features of the universe. While backwards or temporally symmetric causation is a conceptual and metaphysical possibility, it seems to only be *merely* possible, or at least, if it actually happens, mostly irrelevant, for in the actual world, time's arrow seems to determine most or all of the asymmetries of counterfactual dependence and the asymmetric nature of causal and related lawful processes.

The tie between our experience of causation and our experience of time may be just as deep, for we experience the direction of causation as intimately bound up with the direction of time, and we experience time's passing in tandem with events causing other events to change, so much so that many have thought that the passage of time could not be observed without observing qualitative change, although

[10] Carey, 2009; Saxe et al., 2005; Saxe and Carey, 2006; Gopnik et al., 2004.

qualitative change and pure passage are conceptually and metaphysically distinct. And a[n] experience as of [a] lawfully evolving universe involves experiences as of the causal movement of things happening over time, with causal generation from earlier to later, as events produce successive events.

With this in mind, let's look at another feature of experiences we have over ti[me]. The feature involves a phenomenon described by psychologists as 'backwards masking'. An important example derives from our experience of listening to music. When we listen to music, including when we hear a piece played for the first time, we don't hear each note in isolation, one after the other. Rather, when we hear a note of the melody, the phenomenal character of our experience of that note is affected by the properties of the notes immediately prior to that note. This is not that surprising, and may be an effect of a kind of remembered experience affecting the character of one's current experience. But what *is* surprising is that when we hear a note of the melody, the phenomenal character of our experience of that note is affected by the properties of the notes immediately *after* that note. So our experience of the note is affected by what is yet to come as well as what has already been.

In other words, we do not experience the notes of a melody in isolation: our experience of each note is affected by the notes that occur immediately 'around' that note, such that the phenomenal character of each note we experience is subtly different depending on which arrangement of notes it is embedded into. Thus, when we experience a melody composed of a temporally dense series of notes, we experience it as continuously changing and unfolding over time, and our experience of the phenomenal character of prior notes foreshadows the experience we have of the phenomenal character of subsequent notes.[11]

How can we 'see into the future' in this way? What is the basis for this experience of foreshadowing? There is debate about the mechanism involved in the cognitive processing of these temporally clustered events. Some have argued that it is a predictive effect, that somehow we are such excellent predictors that a pre-perceptual, non-inferential prediction about future experience affects the phenomenal character of our current experience. Others have argued that it is what is called a 'postdictive effect', where properties of events A and C that occur in the small window just before and after event B affect the cognitive processing of A, by affecting something in the causal path from stimulus to production of the phenomenology. I take no stand here on which approach, if either, should be endorsed.

What matters here is that these foreshadowing and filling-in psychological effects are empirically well-documented, even if their source is not fully understood. The general version of this fact is that the character of our experience of an event B is, metaphorically speaking, slightly tinged by the shades of A and C (or by our precognitive 'expectation' of the properties of C). The events A and C surrounding an

[11] Music and other auditory experiences are not the only place where this happens. It seems to occur over a wide variety of sensory modalities. Another very well-known, much-discussed case of this sort of thing occurs when we observe the 'flash lag' effect. See Le Poidevin (2007: ch. 5).

event B affect our experience of B in ways that can result in a gradual phenomenal blending from A to B to C. This contributes, I suggest, to an observer in the world having a sense of a seamless, orderly, unfolding of events whose qualitative changes are experienced as occurring in a natural, continuous fashion.

I will call this character of our experience *temporal anticipation.* Temporal anticipation is the phenomenal feel of having an experience with an anticipatory or predictive unfolding character that also seems to follow from the character of prior events. As I've described it, our experience of temporal anticipation and the unfolding of events occurs in a directed way. I will come back to this point: here, I will just note that for now we may assume that the events we experience occur in a temporal direction and assume that our phenomenology simply tracks this direction.

Here is another example that involves what I am describing as temporal anticipation. When we watch natural phenomena like a flash flood or a small avalanche of rocks tumbling down a hillside, we can have a sense of watching mechanics in motion: we sense the way the water will be affected by obstacles in its path and that it is inclined to flow downhill, and we see rocks crash into other rocks and the effects of those forces. This sense of mechanics in motion seems to support our ability to grasp and understand the physical world as unfolding in a law-governed way. And in fact, we need to have this sort of sense in order to sensibly predict what will happen in our day-to-day existence, and to survive as a species. If we couldn't learn to intuitively and quickly judge that the bus coming our way will continue on its path to intersect ours, we wouldn't last long.

You might think that this experience is just a simple grasping of the lawful structure of the world. However, Peter Battaglia, Jessica Hamrick and Joshua Tenenbaum have shown h[illegible]hese sorts of physical intuitions, while reliable, come apart from the real structure of the world in some cases. For example, imagine a situation where we see a rock balanced on the edge of a cliff, but balanced in such a way that it does not and will not fall. When we see it balanced there on the edge, unmoving, we can have the very strong phenomenological impression that the rock is about to fall—and indeed, for all we can see, it *should* fall. But the rock isn't falling, and indeed, it isn't going to fall (maybe it is glued there). Our phenomenological sense of its immanent downward motion is an illusion, created by our learned cognitive response to the world, a response where we navigate our environment by making very quick and intuitive predictions about how the physical world will evolve forwards (Battaglia et al., forthcoming).

Now recall Michotte's work on causal impressions. When events are juxtaposed spatiotemporally in the right way, we have a modular illusion as of causation, that is, we experience the series of events as a series of causes and effects, where each cause seems to generate the next event. Such causal impressions have the phenomenal character of being productive, that is, we seem to see one event producing or generating the next (we 'see' the moving cue ball 'produce' the motion of the eight-ball). As developmental psychologists have shown, from infancy we use these types of causal impressions along with other information, including perhaps innate

capacities for causal computations based on probabilistic inferences or our sense of causal efficacy, to develop a sophisticated concept of causation. These representations as of causation are an empirically documented feature of our phenomenology, and are thought to play an important role in the development of causal concepts.[12]

So we have phenomenal features of our experiences involving causal impressions and causal direction, and phenomenal features involving temporal anticipation. Putting them together, we can see the beginnings of an account of how our cognitive system could construct our experience so that it presents the world to us *as* an evolving, causally governed, productive universe. We have an experience of prior events slightly foreshadowing subsequent events, and of subsequent events slightly shading prior events, suggesting that the world is evolving naturally in a way that is consistent with the temporal anticipation built into our experience. And, when we have modular causal impressions, we experience appropriately spatiotemporally juxtaposed events in these evolutionary series as causally generating or producing later events in these series. As we develop our conceptual capacities from infancy into adulthood, we further understand and inferentially learn to experience our world as a causally governed, causally productive universe.

9.3 Conclusion

The question for the metaphysician is how to interpret these facts. Obviously, *some* of our ordinary experience, naively interpreted, gives us evidence for our metaphysical theories of the world. Setting aside skeptical views, our world *is* a causal and temporal world, involving causation, temporal asymmetry, laws, and counterfactual dependence. The question that we need to attend to, however, when considering the debate over the temporal arrow, is which features of our experience of the world provide evidence for which metaphysical theories about its temporal nature. In particular, given the psychological fact that our experience may embed unnoticed veridical and falsidical illusions, does our experience of the external world as being causally productive and evolving forward give us evidence for an antireductionist account of the temporal arrow?

In the absence of alternatives, our experience does seem to provide evidence in support of the antireductionist metaphysic. Ordinary experience provides a kind of defeasible evidence for an antireductionist view, and assigning a primitive productive capacity to the world seems to explain our experience of the world as an evolving, causally governed, productive universe. Perhaps we are simply detecting this capacity when we experience, and since this productive element is causally and temporally directed, we detect temporal asymmetry by experiencing directed productivity.

In other words, the antireductionist can argue that our experience of the evolving world is *as of* a causally governed, causally productive universe evolving forward in

[12] See Carey (2009) and Richardson et al. (2007) for excellent discussions of how our ability to represent causation might arise.

time, and use this to suggest that the world just *is* a causally governed, causally productive universe evolving forward in time.[13] She can then argue, against the reductionist, that simply finding a global entropic asymmetry in the world that aligns with the temporal arrow does not give us evidence for the view that this global entropic asymmetry *is* the arrow (or is what grounds the arrow), for it does not explain why we experience the world as dynamic in the way that we do. Our experience *as of* temporal asymmetry is not the same as an experience *as of* a world with a global entropic asymmetry, and (she will argue) the reductionist has no good explanation for why, if the temporal arrow is reducible to entropic asymmetry, our experience of the evolving world is *as of* a causally governed, causally productive universe evolving forward in time. The antireductionist thinks that merely endorsing the existence of counterfactual dependence and global entropic asymmetry without adding in a primitive temporal arrow leaves us without an explanation of our sense of time's passing, or of things *happening*. But if we endorse the antireductionist view, we do have an explanation of why we experience the world as dynamic in the way that we do—because it is dynamic in just this way. There is primitive happening and production, just as our experience suggests.

If we agree that the world is presented to us as being a causally governed, causally productive universe evolving forward in time, and that we are supposed to start with the manifest image and refine it consistent with observation and evidence, antireductionism has the advantage here. (One move the reductionist can make is to deny that experience is as of a causally governed, causally productive universe evolving forward in time. If the reductionist makes this move, then I take the argument to shift to a debate over whether the reductionist's metaphysical view is sufficient to account for scientific and ordinary claims made about causation, change, temporal direction, etc. This move just builds the need for explanations of experience in at a different point.[14]) Antireductionism says that there actually are substantive, irreducible properties of causal production and temporal passage. It can capture and explain our experience of the world as a causally governed, causally productive universe evolving forward in time by holding that, at least with respect to these features of the world, the manifest image gets it right.

But as I've shown, cognitive science gives us the resources to develop alternative hypotheses about the source of our experience, hypotheses that can undercut the antireductionists' claim that ordinary experience provides evidence for their view.

[13] Although it is not how I am framing things here, this debate can also be developed in terms of the contents of experience; see Siegel (2010) on the contents of experience and Skow (2011) for an approach to temporal experience along these lines. The argument between the antireductionist and the reductionist I've described would need to be restructured under such a recasting, at the very least because different parties to the dispute will likely agree on some of the contents of experience but disagree on others. Once a decision about the contents that are ascribed to experience are agreed upon, a debate about the evidence provided by experience can ensue. (And if the evidence is not supposed to come from experience but from elsewhere, one must explain where it comes from and why it isn't experiential evidence masquerading as 'metaphysical' evidence.)

[14] I'm indebted to conversation with Brad Skow and Ross Cameron here.

There are at least three obvious alternatives to the antireductionist view. Each of these views can draw on empirical results in cognitive science to provide an alternative explanation of our experience, an explanation that draws on evidence that undercuts the evidence from experience that an antireductionist might use to defend primitive temporal passage.

First, there is *weak primitivism*. The weak primitivist agrees with the antireductionist that at the fundamental level there is a primitive temporal asymmetry that gives us an objective temporal direction. This view grants that events are somehow primitively ordered so that events are earlier than, later than, or simultaneous with other events. But that's *all* the antireductionism this view grants. The weak primitivist denies that there is primitive dynamic temporal passage or productive lawful productivity. She explains our experience as of dynamic temporal passage and as of lawful causal productivity as falsidical illusions that are due to cognitive processing stemming from the constructed nature of our local experiences as of causation and temporal anticipation. The falsidical part of our experience of causation is the impression of productivity we find in launching cases: the 'productivity' part of the phenomenal character of an experience of causation isn't due to C's causing E, it's a spandrel generated by our sub-personal cognitive processing.

Second, we can distinguish two reductionist-friendly views. First, we have *impure minimalism*. The impure minimalist denies that there is any temporal asymmetry at the fundamental level, and, with the weak primitivist, denies that there is primitive temporal passage or productive lawful evolution. Like the weak primitivist, she argues that our sense of productivity and dynamic unfolding is a falsidical illusion that is an artifact of the sub-personal processing of the brain, perhaps combined with downstream causal learning. The impure minimalist holds that at the fundamental level, all we have is an asymmetry of entropy. But temporal asymmetry exists as a derivative ontological relation that is grounded by, but not reducible to, global entropic asymmetries.[15] So temporal asymmetry is part of the ontology, but it is derivative, not fundamental. When we experience temporal asymmetry we are detecting the derivative temporal asymmetry that is grounded by the overall entropic asymmetries.

Finally, we have *pure minimalism*, which, in addition to denying that there is any primitive temporal asymmetry at the fundamental level, also denies that there is any additional ontologically distinct, supervenient, derivative temporal asymmetry in the ontology. Since temporal asymmetry simply reduces to entropic asymmetry (plus a few more bells and whistles), the temporal asymmetry in the world just *is* entropic asymmetry. In effect, the pure minimalist holds that we should accommodate temporal asymmetry as a part of our ideology, not as a part of our ontology.

The pure minimalist agrees with the impure minimalist and the weak primitivist about the cognitive source and falsidical illusory nature of our experiences as of productivity and dynamic evolution. But pure and impure minimalists differ with

[15] On grounding, see. e.g. Schaffer (2009).

respect to the cognitive source of our experiences as of (mere) temporal asymmetry, since the impure minimalist can hold that ontologically derivative temporal asymmetries are the direct source of our experiences as of temporal asymmetry, while the pure minimalist takes the phenomenal character of our experiences as of temporal asymmetry to be the product of cognitive processing of our detection of entropic asymmetries.

For the impure minimalist, experiences as of temporal asymmetry can be veridical perceptions. For the pure minimalist, experiences as of temporal asymmetry are a neat kind of illusion. For the pure minimalist, because our experiences as of temporal asymmetry do not have the same phenomenal character as our experiences as of entropic asymmetry,[16] but our experiences as of temporal asymmetry are due to our cognitive processing of entropic asymmetries, they are illusions.

How might the pure minimalist argue for the claim that experiences as of temporal asymmetry are merely the results of our cognitive processing of entropic asymmetries? We have already seen the general approach to take. Recall that the visual system constructs our sense of the direction we are moving in by comparing the background to the foreground and deciding which one is fixed, and (in the absence of other information) takes the default view to be that the background is fixed. The pure minimalist could advance the speculative hypothesis that, in a similar way, our experience as of temporal direction is merely our visual and other cognitive systems' response to entropy gradients.

Here's one way the story could go: at some stage of cognitive processing prior to our experience, our brain could respond to a difference in entropy by constructing our experience as of temporal direction from lower to higher entropy (taking this to be the default direction in the absence of other information). Just like a contrast between spatial background and spatial foreground signals a movement in a spatial direction, a contrast between lower entropy and higher entropy cognitively generates an experience as of a temporal movement in a direction. Similar to how the cognitive system will privilege background over foreground to construct motion in certain direction, perhaps the cognitive system privileges higher entropy over lower entropy to construct the temporal direction. Take a series of three states with entropy values n, n + 1, and n + 2, where the state with entropy n + 1 is temporally 'in between' the states with entropy n and entropy n + 2: on this picture, the brain, when it processes a batch of stimuli from these states, takes a comparatively higher entropy state to define the default direction, and so represents the temporal direction by tracking the changes of entropy in a direction of entropy increase.

(An analogy: imagine three colored slides where an orange slide is in between a red slide on the left and a yellow slide on the right. Is the series of slides getting lighter or is it getting darker? It depends on what is taken as the default: is a comparatively longer wavelength of light reflectance taken as the default feature that determines the

[16] I'm assuming that an experience as of temporal asymmetry would have a phenomenal character that is different from an experience as of entropic asymmetry.

direction of color change? Or is a comparatively shorter wavelength of light reflectance the default that determines the direction of color change? If the default is that it is longer, the series is ordered in the getting-darker direction. If the default is that it is shorter, the series is ordered in the getting-lighter direction.)

Another, similar hypothesis could draw on evidence from Michotte's launching cases that suggests that our sense of the direction of causation is a veridical illusion. For not only do we have an experience as of causation in Michotte launching cases, but we have an experience as of a causal direction: in each launching effect, we determinately experience one event as the cause (the motion of the cue ball) and the other event (the motion of the eight-ball) as the effect. In the Michotte cases, we know that such experiences as of causal direction are constructed by our cognitive systems, just like the rest of the causal impression. The pure minimalist can argue that this could generalize: perhaps the cognitive system also privileges higher entropy over lower entropy to construct experiences as of a causal direction, and our widespread experiences as of causal direction determine widespread experiences as of temporal direction. (The temporal direction might be metaphysically more fundamental than the causal direction, but the experience of causal direction might still determine the experience of temporal direction.)

Defenders of each of these views, as well as of other variants that reductionists might find appealing, can draw on cognitive science to argue against the antireductionist for an alternative thesis about the source of our experiences of causal productivity, lawful evolution, directed motion, and change.[17] Such theses draw on evidence that is intended to undercut the evidence used in the antireductionist argument from ordinary experience. We can grant that our world is a causal and temporally asymmetric world, and that we observe things like billiard balls hitting other billiard balls, and streams of events such as a row of dominoes falling over, each hitting the next one in sequence, and that generally we experience one event as following another in a temporally directed, causally productive way. We can grant that our experience suggests we are discerning the lawfully productive evolution of the world as we watch the water run down the mountain. But we can use empirically documented psychological work to develop ontologically minimal explanations for why we have many of these experiences, explanations that provide empirically justifiable alternative accounts of how our experiences could be produced. These alternative explanations draw on evidence intended to undercut the observational evidence marshaled by the antireductionist who gives a naive argument from ordinary experience to support her view.

Obviously, which theories about the metaphysics of the temporal arrow we ultimately want to endorse will depend upon total theory and total evidence. We'll

[17] I haven't fully addressed our experience as of change here. In Paul (2010a), I show how the cognitive system could construct our experience as of change from a series of static frames. Roughly, the suggestion is that our cognitive system synthesizes all the data from the external world in the way it synthesizes a series of frames of a film to construct an experience as of change. For more on metaphysics and cognitive science, see Goldman (2007, 2014), and Paul (2010b).

need to weigh complete, competing versions of all of the varieties of reductionism and antireductionism in order to decide between them. But as it stands, we can draw on cognitive science in an attempt to undercut arguments from the manifest to the existence of primitive temporal passage, and the debate over the Arrow should reflect this.

References

Albert, D. Z. (2000). *Time and Chance*. Cambridge, Mass.: Harvard University Press.

Battaglia, P., Hamrick, J., and Tenenbaum, J. (forthcoming). Simulation as an Engine of Physical Scene Understanding. *Proceedings of the National Academy of Sciences*, 110(45), 18327–32.

Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.

Goldman, A. (2007). A Program for 'Naturalizing' Metaphysics, with Application to the Ontology of Events. *The Monist*, 90, 457–79.

Goldman, A. (2014). Naturalizing Metaphysics with the Help of Cognitive Science. *Oxford Studies in Metaphysics*.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., and Danks, D. (2004). A Theory of Causal Learning in Children: Causal Maps and Bayes Nets. *Psychological Review*, 111, 1–30.

Ismael, J. (2011). Decision and the Open Future. In Adrian Bardon (ed.), *The Future of the Philosophy of Time*. London: Routledge.

Le Poidevin, R. (2007). *The Images of Time: An Essay on Temporal Representation*. Oxford: Oxford University Press.

Loewer, B. (2012). Two Accounts of Laws and Time. *Philosophical Studies*, 160, 115–37.

Maudlin, T. (2007). On the Passing of Time. In *The Metaphysics within Physics*. Oxford: Oxford University Press.

Paul, L. A. (2010a). Temporal Experience. *Journal of Philosophy*, 107, 333–59.

Paul, L. A. (2010b). A New Role for Experimental Work in Metaphysics. *Review of Philosophy and Psychology*, 1(3), 461–76.

Paul, L. A. (2012). Metaphysics as Modeling: The Handmaiden's Tale. *Philosophical Studies*, 160, 1–29.

Paul, L. A. (MS). Temporal Dualism.

Phillips, I. (2013). Perceiving the Passing of Time. *Proceedings of the Aristotelian Society*, 113(3/3), 225–52.

Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford: Oxford University Press.

Richardson, T., Schulz, L., and Gopnik, A. (2007). Data-Mining Probabilists or Experimental Determinists? A Dialogue on the Principles Underlying Causal Learning in Children. In A. Gopnik and L. Schulz (eds), *Causal Learning: Psychology, Philosophy, Computation* (pp. 208–30). New York: Oxford University Press.

Robbins, P. (2010). Modularity of Mind. In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2010 edi.). http://plato.stanford.edu/archives/sum2010/entries/modularity-mind.

Saxe, R., and Carey, S. (2006). The Perception of Causality in Infancy. *Acta Psychologica*, 123, 144–65.

Saxe, R., Tenenbaum J. B., and Carey, S. (2005). Secret Agents: Inferences about Hidden Causes by 10- and 12-Month-Old Infants. *Psychological Science*, 16(12), 995–1001.

Schaffer, J. (2009). On What Grounds What. In D. Chalmers, D. Manley, and R. Wasserman (eds.), *Metametaphysics* (pp. 347–83). New York: Oxford University Press.

Siegel, S. (2010). *The Contents of Experience.* Oxford: Oxford University Press.

Skow, B. (2011). Experience and the Passage of Time. *Philosophical Perspectives*, 25(1), 359–87.

Smolin, L. (2013). *Time Reborn: From the Crisis in Physics to the Future of the Universe.* Boston: Houghton Mifflin.

Wallace, D. (forthcoming). The Logic of the Past Hypothesis, to appear in a volume edited by Barry Loewer, Eric Winsberg, and Brad Weslake, on themes from David Albert's *Time and Chance* (Harvard)