

Experience as action *

BENJ HELLIE

December 25, 2009

Joan is conscious. She has experiences. There is something it is like to be her.

But *what* is it like to be Joan? What kinds of properties characterize her way of experiencing? Is there something general we can say about them, that would tell us the nature of consciousness? I say: to find out what it is like to be Joan, ask what Joan is doing. What it is like to be her is what she does; and what she does is what it is like to be her. The nature of consciousness is: agency. Experience and action are one and the same.

More precisely:

E=A

The kinds of experience are the kinds of action.

I intend both ‘experience’ and ‘action’ in a perfectly ordinary way: no tricks, I promise! I also intend **E=A** to mean the same when read backward: ‘the kinds of action are the kinds of experience’.¹

*Thanks first and foremost to Jessica Wilson. Thanks also to: David Chalmers, Alex Byrne, Susanna Siegel, Jim John, Brad Thompson, Nico Silins, Geoffrey Lee, Imogen Dickie, Charles Siewert, Sergio Tenenbaum, Kieran Setiya, Sebastian Watzl, Chris Mole, Declan Smithies, Wayne Wu, Ian Philips, Fiona MacPherson, Mike Brady, John Campbell, Mohan Matthen, Dustin Stokes, Stephen Biggs, Matt Fulkerson, Tim Bayne, Matt Nudds, Charles Spence, Casey O’Callaghan, Farid Masrour, Jason Stanley, Jim Brown, Anthony Bruno, Ulrich Schloesser, Luke Gelinias, Devlin Russell, Matt Habermehl, Diana Raffman, Louis DeRosset, Jessica Moss, and Cian Dorr. Thanks also to the members of my Fall 2009 graduate seminar on Conscious Life, including Luke Roelofs, Adam Murray, Andrew Max, Zack Irving, Steve Pearce, Perry Ge, Elana Koren, Michael Sudnick, Dominic Alford-Duguid, and Alden Jerrett-Enns. And thanks finally to various members of the blogosphere.

¹This view has no explicit proponents in the analytic literature, so far as I can tell. Antecedents of the view include Ryle (1949) (of whom there will be extensive further discussion below) and White (1964) (who writes: “‘consciously’ and ‘unconsciously’ can qualify only our intentional or unintentional actions, not our results” (52)). Mole (2009) also comes very close to advocating **E=A**: he explicitly advocates Premiss 1, and seems to advocate Premiss 2. (Discussion with Chris contributed greatly to my initial formulation of the core ideas in this paper.)

My view is very different from the “enactivist” perspective in the philosophy of perception (Noë *et al.* 2000). As far as I can tell, this view has certain core similarities to “representationalism” in the philosophy of perception. Both are views about the nature of the phenomenal character of a perceptual state; but whereas the representationalist thinks the phenomenal character is constituted by a rich inner representation, the enactivist thinks it is constituted by certain facts about how movement of the perceptual apparatus would influence an impoverished inner representation (or perhaps by the subject’s *knowledge* of such facts). Two crucial differences to my view stand out. First, my view is not a view about phe-

So: one aim of this paper is to advocate $E=A$. (The other is to discuss a certain argument for $E=A$: more on that in a bit.)

If you have taken a course in philosophy of mind, you probably believe something incompatible with $E=A$. So I should say up front why you should take $E=A$ seriously enough to read on. I say you should read on because $E=A$ is obviously true. After all, as Thau (2002, 203) reminds us, intending to tautologize, “to have a bank teller’s experiences is to have done the things that a bank teller does”. That’s obviously true, right?

“Well yes”, you say, “but that is the *job-market* sense of ‘experience’; what I am concerned with is experience in the sense of *consciousness*”. But I am too! I think the job-market and *consciousness* senses of ‘experience’ are one and the same. You continue: “that can’t be true: I meant *phenomenal* consciousness—such phenomenal experiences as states of pain and seeing as of a red thing”. But I say that you have a false opinion about the nature of the experiences (in the one and only reasonable sense): phenomenal properties like *pain* are not kinds of experience. You have this opinion, moreover, in part because of a wrong turn in relatively recent philosophy of mind: your philosophical tradition has bamboozled you into rejecting an obvious truth for a falsehood. “Scandal!”, you say, “but it is evident that our disagreement is substantive, and perhaps you are correct that I ignore as a philosopher what I know in ordinary life. If so this would force me to reconsider much of what I think I know about consciousness—I must read on!”. Splendid—let me not delay you any further then.

I mentioned above that a second aim of this paper is to discuss an argument for $E=A$ —namely, the following *argument from attention*:

1. The K is a kind of action just if a K can occur attentively;

 \Rightarrow The grammar of Premiss 1 patterns with that of ‘the tiger is a kind of mammal just if a tiger can give live birth’.
2. A K can occur attentively just if a K must occur consciously;
3. A K must occur consciously just if the K is a kind of experience.

The logic of the argument is the following: the premisses trivially entail that the K is a kind of action just if the K is a kind of experience; universal quantification over

nominal character or perception, but of consciousness: as we will see in (K), I sharply distinguish these phenomena. And second, my view is not restricted to a certain small class of actions, but is about *all* actions.

K is assumed; $\mathbf{E}=\mathbf{A}$ is to be read as equivalent to this universally quantified biconditional. Premiss 3 will be seen to be relatively trivial. Premisses 1 and 2, by contrast, are somewhat more substantive, but still, I think, initially highly plausible: the idea behind Premiss 2 is that *attention is “high-grade” consciousness*, a doctrine endorsed by most of the interested philosophers over the last century; the idea behind Premiss 1 is that *“attending” is always doing something attentively*, a doctrine endorsed by roughly half of the interested philosophers over the last century.

I don't advance the argument from attention in the spirit of talking my audience into accepting something controversial: after all, in my view, we accept its conclusion tacitly already. The aim is instead twofold: first, throwing the notion of *attention* into our conceptual mix sets off an explosion of interesting predictions and theoretical tools; second, when $\mathbf{E}=\mathbf{A}$ has been “split”, by the argument from attention, into four substantive conditionals, these conditionals can serve as convenient loci onto which to focus anxieties about $\mathbf{E}=\mathbf{A}$: alleged counterexamples to $\mathbf{E}=\mathbf{A}$ can be sorted into counterexamples to the forward and backward halves of Premisses 1 and 2.

Roadmap: Part I of the paper performs groundwork; Part II describes strategies for rebutting counterexamples to the argument's substantive Premisses 1 and 2.

Caveat: the paper is oriented more toward breadth than depth; this will lead to frequent use of expressions like ‘sketch’ and ‘space prohibits more detailed consideration of this point’.

Part I

Groundwork

This first part of the paper explains terms, states background assumptions, argues for central supporting theses, and the like.

(A) expands on certain extremely general hypotheses about experience that I share with most philosophers of consciousness. (B) and (C) lay out metaphysical assumptions about kinds and their instances; this discussion concludes with a quick argument for Premiss 3. (D) illustrates the use of these metaphysical assumptions in the context of a more-or-less orthodox view of what the kinds of experiences are. (E) explains what I mean by ‘action’, and briefly situates $\mathbf{E}=\mathbf{A}$ in the contemporary literature in action theory.

(F) returns to the literature in philosophy of consciousness, comparing my view to the popular doctrine of “representationalism”; (G) then explains the relation between

E=A and the concept of “phenomenal consciousness”.

I then turn to the argument from attention. (H) discusses the adverbial use in Premises 1 and 2 of ‘attend’. (I) bundles all of the preceding discussion together, sketching a metaphysical theory of attention and experience/action. This will enable me to address, in (K), a potential source of anxiety about **E=A**, the case of perceptual “experience”. (J) runs through the argument from attention at a low level of detail, briefly illustrating how it allows us to make decisions about particular cases.

Finally, in (L), I sketch six additional arguments for **E=A**: (i) an argument from ordinary language discourse; (ii) an argument from the unification of mental phenomena; (iii) a phenomenological argument; (iv) an argument from the unification of the epistemology of the self; (v) a historical argument; and (vi) a “meta”-argument bundling the other arguments together.

(A) My doctrines concern *experience*. I am thinking of experience in the most inflationary sense reasonable. The terminology of “consciousness” is often used in the contemporary literature to characterize my subject-matter. When we think of zombies, we think of creatures like us in all material respects but *not conscious* or *lacking in consciousness* (Chalmers 1996). My uses of ‘experience’ and cognates aim to express concepts in the same family as those we express with this sort of talk. I could just as well characterize zombies as creatures just like us in all material respects, but “lacking in experience” or “having no experiences”.

My purposes do not require that zombies be possible, merely that they be conceivable. Indeed—though I believe, and will assume, that they are—they do not even require that: merely that there be a common concept of consciousness which each of David Chalmers and David Lewis thinks the other misuses to some degree (Lewis 1995, Chalmers 1996). It is this sort of concept which generates an “explanatory gap” (Levine 1983) and is therefore “conceptually fundamental” (Chalmers 2003) that my uses of ‘experience’ and cognates aim to express. So let there be no mistaken sense that what I aim to convey can be expressed using some explicitly etiolated merely functional concept such as “access consciousness” (Block 1995), which only a deeply unreasonable person could withhold from a zombie.

Understood as I intend it, then, **E=A** entails that “actions”, in a (perfectly ordinary) sense to be elucidated further below at (E), are the source of the explanatory gap, are those features in respect of which we differ from zombies; that concepts of action are conceptually fundamental; and perhaps also that action-kinds are themselves fundamental ingredients of the world. (For further discussion of conceptually fundamental concepts of action, see below: 2L(c).)

Philosophers of consciousness think of experiences as occurrences which are “like something”, and of the kinds of experience as what, in particular, those occurrences are like (Nagel 1974). I do too. But I think that what, in particular, experiences are like are kinds of action. This claim is initially plausible: when queried what something was like, we often state what we did: What was your trip to Barcelona like? —I spent the whole time working, didn’t do any sightseeing.²

‘Experience’ is frequently used in the philosophical literature on consciousness to abbreviate ‘perceptual experience’: for but one source of this use, see my Hellie 2010. So don’t be misled! I no longer intend this restriction: rather, my usage in this paper is intended to cover all conscious occurrences—not just perceptual experiences. In fact, it isn’t even required to cover the sorts of states I once called “perceptual experiences”. And, in fact, it doesn’t: I no longer think that these states are in fact *experiences*. After all, seeing a red thing, or being in some state which is “as of” seeing a red thing (or having a pain, or an itch, or being sad), is not *doing* anything.³ So whatever kind of thing it is, it is no action. So by **E=A**, it is not an experience, either. I expand on these remarks at length in (K).

Am I changing the subject? Am I talking past my past self? Not if I was trying to characterize the nature of consciousness in general: and I think I was. If so, I have meant the same thing by ‘experience’ all along: I have merely changed my mind about its extension.⁴

²We also sometimes state what we perceived or underwent: —I saw the Sagrada Familia; —I was jostled by hordes on the Rambla. We will see below at (L) that this pattern of ordinary use is compatible with **E=A**.

³Of course seeing is underlain by various psychological processes, such as opponency processing of data from the retina. But these are not actions. Actions require agents. I am an agent; my opponency channels are not agents. They process this data; I do not. I merely benefit from their activity.

Of course what I see depends on what I am doing: if I am not looking in a certain direction, I will not see what is over there. But the dependence is causal: my looking over there caused me to see what is over there. It is not constitutive: my looking over there is no *part* of my seeing what is over there, any more than my hitting my thumb with a hammer is *part* of my feeling pain in it.

Of course what I see influences what I am in a position to do: if I do not see something, I cannot look at it. I am inclined to think that this may be a matter of what I am seeing *constituting* (in part) what I am doing. But there are uncontroversial cases in which a nonaction can constitute an action: if I am not in contact with a cup, I am not picking up a cup.

⁴The same goes for other philosophers of consciousness, of course. Philosophers in other subdisciplines may have different aims: for instance, epistemologists might speak of perceptual experiences without requiring anything in particular to be true about the nature of consciousness in order for their theories to be significant.

(B) My doctrines concern *kinds* of experience.⁵ I will assume certain fairly robust metaphysical doctrines about kinds—or at least the pragmatic virtue of speaking in line with these doctrines. In laying out the argument from attention I write with great fastidiousness about ordinary language usage of “the *K*” as a kind of experience, patterning with our ordinary talk of “the tiger” as a kind of mammal: referring, perhaps, to a certain abstract individual. However, in general, I will not be at all fastidious in distinguishing kinds and properties, and will elide the distinction between predicates and general terms with lackadaisical blitheness.

My robust metaphysics of kinds will include the following assumptions. First, if *F* is a kind of *G*, *F* is something like a variety, sort, species, or determinable of, or a way of being, *G*. All of these notions express at least metaphysically necessary entailment: $\Box\forall x(Fx \supset Gx)$. Very plausibly, they express more, as well. We could appeal to the apparatus of “naturalness” or “carving at the joints” to provide further elucidation: if *F* is a kind of *G*, then if *G* is carved at the joints, we find *F*; or *F* is a natural way of making *G* more specific. We appeal to familiar cases to provide paradigms and foils of this sort of relationship: oxygen is a kind of element and tigers are a kind of mammal, but oxygen in Chicago is not a kind of element and tigers within five miles of a burning barn are not a kind of mammal. These notions express an entailment of sorts in the opposite direction as well. If *F* is a kind of *G*, then anything which is *G* is also some kind of *G*: $\mathbf{K}(F, G) \supset \Box\forall x(Gx \supset \exists H(\mathbf{K}(H, G) \wedge Hx))$. I assume that the “kind of” relation between properties, **K**, is a strict partial ordering: irreflexive, antisymmetric, and transitive. I will assume that if anything is a kind of *G*, *G* is also a kind: so, in particular, *experience* is a kind. It follows that anything that is of any kind is of at least

⁵Does the class of kinds of consciousness have a clear line between its extension and its antiextension? Some candidate fence-straddlers are the familiar theoretical and practical attitudes: belief, wonderment, desire, intention. We speak of “conscious belief” and “unconscious belief”, and so forth (well, not very many of us, and not very frequently: ‘consciously believe’ and ‘conscious belief’, excluding ‘philosophy’, collectively turn up about 1/2000 as many Google hits as ‘believe’ and ‘belief’). What is meant by this discourse?

The answer varies from case to case. Some theoretical attitudes, such as belief, are not actions, and thus cannot be experiences. Talk of “conscious belief” is either a weird way of talking about thought or judgement, or a somewhat affected way of talking about belief that one is in a position to recognize that one has.

Other theoretical attitude talk, such as talk of “wonderment”, varies in its application among a range of phenomena: an act of wondering is an action (a sort of thinking, I would imagine) and hence an experience and hence conscious; a state of credential uncertainty, which is a bit like belief, is not an experience but might from time to time manifest in acts of wondering. Perhaps “conscious wonderment” is the act of wondering and “unconscious wonderment” is a state of credential uncertainty which one is not in a position to recognize oneself to be in.

Practical attitude talk is another beast entirely. Arguably (Thompson 2008), intending to *A* (properly so-called) is an action, hence never unconscious; if so, talk of “unconscious intention” is either talk of a sort of dispositional state or talk of *inattentive* intention (for more on inattentiveness, see (H) and (I)).

two kinds—perhaps more as well.

Second, if F is a kind of anything, then anything which is F is “essentially” F . Once again, essentiality has familiar modal consequences: if o is essentially F , then $\Box Fo$. Once again, there is very likely more to essentiality than this: familiarly, it plausibly both essential to $\{\text{Socrates}\}$ that it contain Socrates and not essential to Socrates that he be contained in $\{\text{Socrates}\}$ (Fine 1994).

(C) My doctrines concern *experiences*. I am thinking of experiences as *occurrences*: temporal particulars, along the lines of the familiar states and events. (For a somewhat congenial perspective on temporalia, see Parsons 1990.) An experience begins, lasts for a while, and then ends: never to return. It has no gaps: if it begins at one time and ends at another, then for any intermediate time, it is occurring at that time.

Experiences are not properties or kinds or universals or in any way general: the relevant general things in the neighborhood are the *kinds* of experience. If an experience of kind K is going on in Chicago and another of kind K is going on in New York, that makes for two experiences and one kind of experience.

If some occurrence is an experience, then, as we just saw, it is thereby an experience of some kind; and therefore essentially of that kind and essentially an experience.

We can now pause to make a case for Premiss 3: recall, that a K must occur consciously just if the K is a kind of experience. For the right-to-left direction, let me stipulate that for o to “occur consciously” is just for o to be conscious. If K is a kind of experience, then it is also a kind of consciousness (since, as we have seen, these notions are one and the same), so that anything which is K must be conscious. For the left-to-right direction, let K be restricted to the kind-properties. Then suppose that, necessarily, anything of kind K is conscious: one needn’t be all that Humean to think that this is best explained if K is a kind of consciousness.

(D) Let us now consider some examples that illustrate these metaphysical assumptions.

According to the orthodoxy in the philosophy of consciousness, *pain*, *itch*, and “*seeing as of*” a red thing—perceptual (or sensory) “experiences”, more generally—are kinds of experience. As we saw in (A), this orthodoxy is incompatible with the left-to-right direction of $\mathbf{E}=\mathbf{A}$ (if K is a kind of experience, K is a kind of action), but let us play along temporarily for pedagogical purposes. Suppose that Bill’s head is in pain: then there is an experience, in Bill’s head perhaps, of pain. This experience is an occurrence in Bill’s head: very plausibly, it is a *state* in or of Bill’s head. This state is essentially a pain, and essentially an experience. It is a different experience than that which Joe has when his head is pain—though it might be an experience of the same

kind. Bill's toe also itches: so there is an experience, perhaps in Bill's toe, of itch; this occurrence is also, plausibly, a state (in or of Bill's toe): essentially an itch and essentially an experience, and distinct both numerically and in kind from Bill's pain. I will discuss the relation between perception and experience at length in (K).

(E) My doctrines concern *actions* and *kinds of action*. I port the metaphysics of kinds and their instances I have sketched over to this case: "kinds of action" are something like properties, which carve the highly general property *being an action* at its joints, and which are essential to their instances; their instances are the "actions", which are particular occurrences. (For a discussion of the metaphysics of action with which I am in many ways sympathetic, see Thompson 2008.)

By 'action', I intend to track ordinary use. Ordinary use encompasses such occurrences as Babe Ruth's pointing to center field in the 1932 World Series, Obama's campaigning for President of the USA in the 2008 election cycle, and my cycling to work this afternoon. All three of these are actions: particular occurrences. Each of these three falls under a kind of action. I won't provide, and don't have, a precise theory of how to extract from a natural language characterization of an action, what kind that action belongs to. Instead, I will take a "seat-of-the-pants" approach, aiming only to maintain a sort of consistency in my practice. So, for example, the kinds of the three examples above are (or perhaps include): *pointing to center field*); *campaigning for President of the USA*; and *cycling to work*.⁶

These three actions (and their kinds) all involve a significant "bodily" component: moving arms and legs, emitting noises in stadiums, and the like. I wish to include in the discussion also such "mental" actions as my current act of thinking about how to phrase the sentence I am writing (and the kind, *thinking about how to phrase a sentence*).

Waxing a bit more metaphysically, I am inclined to suspect that an occurrence is an action just if it is performed by an agent; and that a property is a kind of action just if it is possible for an agent to acquire moral credit or demerit for performing or failing to perform an occurrence of that kind. So for example, plausibly, one can acquire moral credit or demerit for pointing (or failing to point) at something—for example, in a police line-up. By contrast, one cannot acquire moral credit or demerit for one's finger's affecting a pointing motion: although one may acquire credit or demerit in such circumstances, if one performed the pointing motion with one's finger.

⁶Each of these actions is performed in a specific way: I cycle to work in part by cycling up Palmerston. Considerations of the mereology of actions are complex and would divert us from the main line of the paper, so for the most part I will ignore them: for an approach with which I have a great deal of sympathy, see Thompson 2008.

The reader may have a certain mainstream “Davidsonian” position in the philosophy of action in mind (a locus classicus of this position is Davidson 1978); if so, it will be of pedagogical value to place it in contrast with my approach.

- The kind of every (actual) occurrence is, fundamentally, physical (or at least nonpsychological); accordingly, action-properties are not kinds, but are rather extrinsic and inessential features of their instances.
- For an occurrence to have the action property *A-ing* is in part for it to be of a certain physical kind; in part for it to be caused by a certain sort of psychological state; and in part for it to be caused in a certain way. More specifically:
 - To coin a phrase, the physical kind to which acts of *A-ing* belong is *display of A-ing behavior*;
 - Every act of *A-ing* is caused by an *intention to A*—an internal, psychological, representational state;
 - There is a certain causal relation, “right-causation” (the nature of which may be a priori opaque to us), such that for an occurrence to be an act of *A-ing* is for it to be a display of *A-ing* behavior that is right-caused by an intention to *A*. “Right-caused”? Here I allude to the familiar difficulty of “causal deviance” (Davidson 1973): causal chains come cheap, so mere causation of display of *A-ing* behavior by an intention to *A* is insufficient for the behavior to count as an act of *A-ing*. Davidson’s solution is to posit an irreducible concept of “right-causation”, which has as its a posteriori referent some physicalistically acceptable causal relation (see also Mele 1992 and Setiya 2007).

To see this position at work, consider the famous example: Davidson’s finger moves in manner *M*; Davidson moves his finger in way *M*; Davidson flips a switch; Davidson illuminates the room; Davidson alerts a burglar. According to the Davidsonian position, the story involves only one occurrence: its kind is the physical property, *displaying finger M-motion behavior*. Because Davidson does not intend to alert the burglar, this event does not have the action-property *alerting a burglar*; but because he does intend to move his finger in manner *M*, to flip the switch, and to illuminate the room, the event can (if caused in the right way by these intentions) have the action-properties *moving a finger in manner M*, *flipping a switch*, and *illuminating a room*.

I see things differently. First, in addition to the nonpsychological occurrence of displaying finger *M*-motion behavior, there are at least three additional occurrences:

an act of *moving one's finger in manner M*; an act of *flipping a switch*; and an act of *illuminating a room* (the actions are distinct because the kinds are not mutually entailing, and the actions therefore differ in modal profile).⁷ Davidson's system of characterization is motivated by an up-front desire for physicalism and ontological parsimony. I am inclined to take a "let the chips fall where they may" approach on the former; and I am inclined to regard the latter as a matter of bookkeeping, which we might wish to sort out once we have a workable theory up and running. Since, in my view, my ontologically more profligate system of description is easier to use, I see no immediate reason to lash my discussion to the reductive language of mainstream action theory.

My second difference is more substantive, and remains even against a weaker Davidsonian approach according to which action properties are genuine kinds. I do not accept the mainstream view of the ground of action properties: recall, that for a display of *A-ing* behavior to be an action of *A-ing* is for it to be right-caused by an intention to *A*. Two quick points. For one thing, it is contentious whether *A-ing* reduces in part to intention to *A* rather than the other way around (Thompson 2008): see footnote 10 for further discussion. For another, my view is that $\mathbf{E}=\mathbf{A}$ is a priori: entailed a priori by canonical concepts of types of actions and experiences, in the sense of being knowable by "phenomenological reflection" alone. It is plausible that the mainstream view in action theory is also intended to be a priori. If it is a priori, display of *A-ing* behavior right-caused by an intention to *A* is conceptually sufficient for consciousness. Behavior and intention together are conceptually insufficient for consciousness, but there is still some wiggle room in the concept of right-causation. Many action theorists have followed Davidson in reacting to the difficulty of causal deviance by claiming the true nature of right-causation to be knowable only a posteriori (for example: Mele 1992, Setiya 2007). So perhaps they regard the concept of right-causation as compatible with its referent being a primitively psychological sort of causation. It is not easy to reconcile this suggestion with the vigorously physicalist spirit of much of the literature in this field.⁸ (I discuss a potential source of the conceptual irreducibility

⁷There may also be an occurrence of *alerting a burglar*, but it is not an action. This property is a mere property, rather than a kind, of the act of illuminating the room. Since it is a mere property, this property cannot be an action kind, and must be a different property than the action kind *alerting a burglar*. The former property would perhaps be more perspicuously labeled *causing a burglar to become alert*: after all, Davidson's illuminating the room did cause the burglar to become alert; by contrast, I think we can recognize a strict sense of "alert a burglar", in which it can only be sensibly predicated of agents and not of actions. In this case, ordinary language blurs over metaphysics somewhat. I discuss the relation between an action's kind and its effects below at 2L(c).

⁸An alternative heterodox position would endorse a nonreductive view of the sort of intention involved in action.

of action below, at 2L(c).)

(F) $E=A$ purports to tell us which properties the kinds of experiences are. Another theory with the same purport is *representationalism*, according to which the kinds of experience are at least *among* the representational or intentional kinds (Siewert 1998, Byrne 2001, Chalmers 2004, Thompson 2009). As we have seen, there is representation without experience; so in order to provide a class of properties which are *exhausted* by the kinds of experience, the representationalist needs to posit a certain distinctive sort of representation, *E-representation*—occurrences of which are “E-representings”—such that the following is then true:

Representationalism

The kinds of experience are the kinds of E-representing.

Let us not worry about the nonreductive character of this proposal, and focus instead on the status of *compatibilism*, the doctrine that neither $E=A$ nor Representationalism entails the denial of the other.

Compatibilism, of course, entails that the kinds of action are the kinds of E-representing. According to the (weakened) Davidsonian view, for an occurrence to be an act of *A-ing* is for it to be a display of *A-ing* behavior that is right-caused by an intention to *A*. If an intention to *A* is an internal representational state, the compatibilist could endorse the view that the kinds of experience are properties like *being a display of A-ing behavior that is right-caused by an intention to A*: this would preserve both $E=A$ and Representationalism. This version of representationalism would be rather different from approaches canvassed in the literature, which have to my knowledge exhaustively regarded the relevant representational properties as akin to *theoretical* attitudes like belief rather than *practical* attitudes like intention.⁹

Of course, this compatibilist position is only available if: first, action reduces to intention rather than the other way around; second, intention is representational. Both doctrines can be contested.¹⁰

(G) The literature on representationalism displays a somewhat striking terminological habit: the view is often expressed not quite as just discussed, but rather as the view that the *phenomenal* kinds are representational kinds (Shoemaker 1994, Tye 2000,

⁹Alternatively, a theorist of a more internalist bent could preserve representationalism and this elevation of practical attitudes by identifying kinds of experience with kinds of intention.

¹⁰Concerning the first: according to Thompson (2008), intending to *A* is roughly being at least in an early stage of *A-ing*; in my view, intending to *A* rather doing something indiscriminable from (*A-ing* in the good case). Concerning the second: perhaps the main stream of philosophy and cognitive science has gone wrong in regarding representation as the key to the mind.

Chalmers 2004). Representationalism as just discussed is a consequence of that view together with the following:

The phenomenal state conception

The kinds of experience are the kinds of phenomenal state.

Evidently the reason for the striking terminological habit is that the phenomenal state conception is widely accepted.¹¹ I think it should be rejected.

Many authors speak as if it could not be intelligibly rejected.¹² If it cannot be intelligibly rejected, then ‘phenomenal state’ is being used as equivalent to my use of ‘experience’. That is fine: we may use our jargon however we like.

However, ‘phenomenal state’ is often used in a way that seems to presuppose substantive analytic entailments. For example, the following often seem to be presupposed:

- a’. Perceptual “experiences” are paradigm phenomenal states;
- b’. Phenomenal states are *narrow*: the kind of phenomenal state a subject is in supervenes (at least nomologically) on the intrinsic character of that subject’s brain;
- c’. Phenomenal states are *passive*: no phenomenal kind is a kind of action;
- d’. Phenomenal states are *states*: no phenomenal state has any essential “temporal flow”.¹³

¹¹Turn to nearly any work on consciousness from the last decade and one will find the familiar budget of (c) passive (a) perceptual (d) states—pain, seeing (“as of”) a red thing as exemplary of phenomenal experience. The extensive literature on (b) narrowness reflects a very strong consensus on at least the prima facie narrowness of phenomenal kinds (Shoemaker 1994, Chalmers 2004, Horgan and Tienson 2002), even among the opponents of narrowness (Lycan 2001)—or at the very least a consensus *that there is* such a (near) consensus (Martin 2004).

¹²I estimate that about 90% of work in the philosophy of consciousness simply elides the distinction between kind of experience and phenomenal kind (prominent examples: Lewis 1988, Chalmers 1996). Dissenters are largely clustered in the “higher-order” tradition: for example, in the 1991 introduction to part IV of his anthology Rosenthal 1991, Rosenthal takes pains to distinguish two problems of consciousness, those of “qualia” and of “point of view”; later, Levine (2006) and Kriegel (forthcoming) distinguish from qualia a “for-me-ness” in consciousness. I am highly sympathetic to the need for a distinction along these lines—though I feel that this approach overstates the connection between qualia and consciousness and understates the many ways in which subjectivity is shaped by agency. Martin (2006) also distinguishes between point of view and phenomenality. (See also footnote 14.)

¹³Let us say that a kind *K* is “dynamic” just if either: (i) necessarily, if *o* is *K* and occurs at a certain time, *o* does not occur at any other time; or: (ii) necessarily, if *o* is *K* and occurs at a certain time, *o* also occurs at some later time—and that *K* is *static*, otherwise. Let us say that an occurrence has temporal flow just if its kind is dynamic, and is a state otherwise.

Testing the definitions: —*seeing (“as of”) a red₁₇ thing* is a static kind (assume that red₁₇ is a point

If the phenomenal state conception were true, then the following would be facts about *experience*:

- a. Perceptual “experiences” are paradigm experiences;
- b. Experiences are *narrow*;
- c. Experiences are *passive*;
- d. Experiences are *states*.

Obviously not every theorist who approvingly uses the terminology of ‘phenomenal state’ is explicitly committed to each of these doctrines; and yet they certainly seem to be strong implicit commitments much of the literature.¹⁴

in color space): against (i), one could see (“as of”) a red₁₇ thing over some interval; against (ii), one could see (“as of”) something changing color at a constant rate, the color of which passes through red₁₇ for just an instant. —*noticing Jane* is dynamic by condition (i): it is “over as soon as it has started” (Rothstein 2004). —*skating* is dynamic by condition (ii): if one is *skating* at *t*, then at *t* one’s activity of skating is continuing; it can’t stop cold at *t*.

Accordingly, any occurrence of seeing (“as of”) a red₁₇ thing is a state; any occurrence of noticing Jane has temporal flow; any occurrence of skating has temporal flow.

¹⁴For just two especially visible examples, consider first the following passage from Lewis 1988, 266:

In so far as I don’t know what it would be like to drive a steam locomotive fast on a cold, stormy night, part of my problem is just that I don’t know what experiences I would have. The firebox puts out a lot of heat, especially when the fireman opens the door to throw on more coal; on the other hand, the cab is drafty and gives poor protection from the weather. Would I be too hot or too cold? Or both by turns? Or would it be chilled face and scorched legs? If I knew the answers to such questions, I’d know much better what it would be like to drive the locomotive. So maybe “know what it’s like” just means “know what experiences one has”.

Another part of Lewis’s problem (that goes unmentioned) is that he doesn’t know which actions he would perform in driving the locomotive—and, for that matter, in taking in these perceptual states.

And consider also Chalmers’s “catalog of conscious experiences” at the outset of Chalmers 1996, 6–11. This opens with a lengthy discussion of perceptual (in the strict sense) states: visual, auditory, tactile, taste, and olfactory states. It continues with somewhat more brief discussions of emotions, and of sensory (in the strict sense) states: “experiences of hot and cold”, pain, and “other bodily sensations”; and of (the somewhat analogous case of) mental imagery. Potentially a bit outside the mold is a discussion of “conscious thought”, but this is assimilated to the previous cases: thinking of a lion has a “whiff of leonine quality”.

Notably, Chalmers also acknowledges “the sense of self”: “something to conscious experience that transcends all these specific elements: a kind of background hum, for instance, that is somehow fundamental to consciousness and that is there even when the other components are not”. Hear hear: although I’m inclined to think that the hum Chalmers hears is actually the multi-voiced chorus of the very many actions he is performing, which shape and contour, highlight and ignore the other components. (See also footnote 12.)

$E=A$ is incompatible with each of these doctrines. To use the terminology of the “square of opposition”,¹⁵ $E=A$ is the contrary of (c). For reasons I relegate to a footnote,¹⁶ (d) entails (c), so that $E=A$ entails the contrary of (d). As we have seen, pain, itch, and “seeing as of” a red thing are not actions: so (a) would witness the contradictory of $E=A$. Indeed, such perceptual “experiences” are intuitively *states*, so that (a) would witness the contradictory of an entailment of $E=A$. And finally, some kinds of action are not narrow; so that $E=A$ entails the contradictory of (b). In my view, then, each of (a)–(d) is at least a false theoretical claim (perhaps even, for reasons discussed above, false a priori).

If you think that each of these claims is true and a theoretical claim, I hope my argument will convince you otherwise. But if you are among those who treat both the phenomenal state conception and any of (a′)–(d′) as analytic, then you treat at least one of (a)–(d) as analytic. In this case, each of us regards the other as conceptually confused. If so, I hope my argument will both re-educate you about the correct use of the concept of experience and convince you that $E=A$ is true.

The discussion of phenomenal properties will continue in our discussion of perception, in (K).

(H) Premisses 1 and 2 use ‘attend’ as an adverb. This grammatical structure reflects a certain view the metaphysical structure of attentional phenomena—a view I endorse.

What metaphysical form could attentional phenomena take? One possibility is that “attending” is a kind of action: for instance, attending to the teacher, the garlic one is chopping, the problem one has been set in an exam. In this case, if *a* is an *attending* of one sort or another, it is essentially so, and so forth. The other is that *attentiveness* is a “mere property”—and not a kind—of actions: so that if *a* is *attentive*, it has some further kind *K*, and is essentially a *King* but not essentially an attentive *King*. Examples of attentive action would be attentively gardening or attentively chopping garlic: *gardening* and *chopping garlic* are kinds of action, some possible instances of which are performed attentively, others inattentively.

We have been assuming a sharp distinction between the kinds and the mere properties: accordingly, *attending* and *attentiveness* are distinct properties. Therefore, there are three metaphysical forms for attentional phenomena to take: we could posit either exactly the kind *attending*, or exactly the mere property *attentiveness*, or both: call

¹⁵‘All *F* are *G*’ and ‘no *F* are *G*’, recall, are contraries, or could not both be true but could both be false: contrariety is a weaker relationship than contradictoriness, or equivalence to denial, which requires exactly one to be true.

¹⁶Kinds of action seem to be one and all “dynamic”, in the sense of footnote 13. It is a very nice question why this should be.

these the *kind* view, the *adverbial* view, and the *mixed* view. I sketch a case for the adverbial view based loosely on remarks in that overlooked little classic of late-stage ordinary language philosophy of mind, White 1964.¹⁷

Against the mixed view, either nonmixed view would simplify our metaphysics of attentional phenomena. To paraphrase Anscombe on intention (Anscombe 1963, 1) in the “material mode” —it is implausible to say that attentional phenomena are utterly disunified in the way specified by the mixed theory: that we pre-theoretically regard them as a single phenomenon is manifest in our unified discourse. To take the mixed theory as a starting point for metaphysics would be to consign ourselves to remaining in the dark about the character of the underlying phenomenon.

Against the kind view, first: the adverbial view provides a plausible account of such “kind talk” as ‘attending to’ such “objects of attention” as the teacher or the garlic one is chopping. But second: the kind view falls down on accounting for such “adverbial talk” as ‘chopping garlic/thinking attentively’.

On the first point. We can think of kind talk as concerned to discuss doing something salient to the objects of attention. For example, *attentive observation* (Mole 2009): I listen attentively to the teacher; I watch the garlic attentively as I chop it (to insure that I shave its cut side minutely thinly). Or *attentive thinking*: I think attentively about the problem on the exam. I see no obvious cases that cannot be accommodated using this strategy, and no in principle problem with the underlying metaphysics.

Objection: there is no such thing as *inattentively* listening to the teacher or inattentively watching the garlic; accordingly, the attentiveness of these actions cannot be a mere non-kind property. Reply: even if talk of “listening inattentively” made no sense, this would not have obvious metaphysical significance for what is possible. And—supposing it did—even if it were impossible for one to listen inattentively, this need not obviously be explained by asserting that attentiveness is part of the *kind* of listening.¹⁸ Without a theory in hand of what attentiveness *is*, there is not much point in speculating about what an alternative explanation would look like. Finally, it seems that it *is* possible to listen inattentively: one can imagine a dull telephone conversation in which one listens inattentively while checking up on something on the internet. In (I), I will provide a characterization of attentiveness that suggests that such a “multi-tasking” withdrawal of attention from actions in progress is always a possibility.

On the second point: recall, that the kind view falls down on accounting for such

¹⁷For another case for adverbialism, see Mole 2009.

¹⁸Perhaps the inference from modality to kind is defeasibly a good one; in the present instance we could combine the case (yet to come) against the kind view with the case against the mixed view into a case that the impossibility is due to something else.

“adverbial talk” as ‘chopping garlic/thinking attentively’. The friend of the kind view would naturally interpret such talk as concerning actions of *attending to* one’s chopping of garlic, or one’s thinking. However, there is a pitfall: talk in ordinary language of “attending to one’s chopping of garlic” is not univocal, so it can be easy to be misled into thinking that this proposal provides a plausible genuine alternative to the adverbial view. Discussing a related proposal, White (1964, II.(c)) distinguishes between attention to one’s action “as subject” and “as spectator”: his discussion provides a template for our interpretation.

Attending to one’s chopping of garlic “as subject” is just our chopping garlic attentively: a plausible interpretation, but no genuine alternative.

The genuine alternative involves attention to one’s action “as spectator”, and, as I shall argue, is not especially plausible. The paradigm of attention as spectator is attention to the teacher, understood as a sort of essentially attentive observation. But, as Ryle (1949, 138) noted, it is not plausible to regard attentive chopping of garlic as performing two actions at once: chopping garlic, and directing this sort of observational attention on one’s action. Now, there is no doubt that one can at times take a sort of alienated or spectatorial perspective on one’s own action: in giving lectures, one can sometimes “watch” the words tumbling out of one’s mouth while reflecting on one’s uncanny situation. This would seem to be a case of doing two things at once: lecturing somewhat inattentively, while attentively observing one’s lecturing with a sense of the uncanny. But this is absolutely not the ordinary case: an ordinary case of chopping garlic is not accompanied by any uncanny sense of self-spectatorship. After all, if self-spectatorship were the ordinary case, why would one regard it as uncanny? Note that my concern is not that attentive observation of one’s own chopping of garlic as spectator is impossible. Rather, I make two points: first, that the spectatorial stance is not *necessary* for attentively chopping garlic; and second, that it seems to be *incompatible* with chopping garlic attentively (what one does attentively is instead *observe oneself* chopping garlic).

It does not seem to me, therefore, as though the proponent of the kind view has an adequate account of the phenomena going under adverbial talk, so we should reject it. Because the mixed view is a view of last resort, and because the adverbial view can accommodate the phenomena going under kind talk, we should reject the mixed view as well. So we should accept the adverbial view.

(I) How does it all fit together? What do the core “non-framework” phenomena in the argument—action, attention, consciousness—have to do with one another?

I see the interaction as being roughly the following. Actions, understood as dis-

cussed in (E), can be long-lived occurrences: Obama's campaign for the presidency lasted nearly two years (perhaps even longer); my working toward my PhD lasted a bit over six years; Rockefeller's assembling the Standard Oil company lasted decades. They can also be short-lived occurrences: my typing an individual character lasts a fraction of a second, my taking a sip of tea lasts not much longer.

Actions overlap one another. We multitask: not just now in our hectic age of information-saturation, but always. Right this instant, I am writing a paper; but I am also working to reform the curriculum of my undergraduate program, teaching a graduate seminar, sitting on a search committee, planning dinner, thinking about genealogy, reading several books, maintaining my health, planning for my upcoming undergraduate courses, and so forth. All this is going on right now: accompanied at various points, of course, by the occasional stretch, sip of tea, or scratch.

Why think this is all going on simultaneously? Two principles: (i) the doctrine, from (C), that a single occurrence does not have gaps, but occupies an entire stretch of time from its commencement to its end; (ii) straightforward, "seat-of-the-pants" judgements about the beginnings and ends of actions. I started writing this paper two weeks ago and have not stopped: why would I have? I haven't given up, and haven't finished the paper. I started teaching my graduate seminar three months ago, and have not stopped: why would I have? The semester is not over yet, and I wish to keep my job. If so, I am doing at least two things right now, writing this paper and teaching my seminar. And we can see how this readily expands into my doing an indefinite range of things right now.¹⁹

It would distort the matter, however, to assert that the writing of this paper and the teaching of the seminar are always precisely on a par as regards consciousness from start to finish. We might put the phenomenological point as follows: at this moment, writing is *foregrounded* and teaching is *backgrounded*; last Wednesday evening, the order was reversed: teaching was foregrounded and writing was backgrounded.

What is the difference between foreground and background? There is a clear *functional* impact: when an action is foregrounded, one performs it *more intensely*, so to speak. I am not, at the moment, making a great deal of progress on preparing or delivering my lecture, marking papers, or communicating with students; last Wednesday evening, I was making a lot of progress in doing these things, but was not making a great deal of progress in writing this paper.

But the functional impact is not all there is to it, at least not notionally. There is also a phenomenological impact. Teaching and writing are not now (and were not last

¹⁹This overlap is central to the argument of Thompson 2008.

Wednesday night) on a par phenomenologically: right now, I am ignoring teaching; last Wednesday evening, I was ignoring writing. Indeed, the point that not all moments in the course of an action are on a par phenomenologically can be made still more dramatically. Last Wednesday night about six hours after the end of my seminar, I was ignoring *everything*: I was asleep, dreamlessly. The phenomenological contrast between my present condition and my condition then could not be more stark. Now, some action is foregrounded; then, nothing was.

I want to suggest that the difference between an action's status as foregrounded or backgrounded at a time is just the difference in its *attentiveness* status. For an action to be foregrounded at a time is for it to be then performed *attentively*; for it to be backgrounded at a time is for it to be then performed *inattentively*. Accordingly, right now I am writing my paper attentively and teaching inattentively; last Wednesday evening I was teaching attentively and writing my paper inattentively; in the thick of night afterward, I was doing nothing attentively (though both actions were still ongoing). In a case in which I am deeply multitasking (carrying on a conversation while reasoning about something unrelated), I am perhaps performing more than one action attentively. The foreground–background distinction is *just a distinction in ways of being related to other actions*, however: it is not some further ingredient piled atop the structure of ongoing actions.

I am aware that my view of attention—as whatever it is about the pattern of independently existing actions that brings some small number of them to the foreground—does not square perfectly with ordinary usage, as I will discuss further when defending my premisses. For now, I will make two further points about the view I have described, before assembling the metaphysical package. First, on my view, attentiveness is *metaphysically light*. Attention is not an action, like directing a spotlight or a zoom lens; it does not have any intrinsic power to restructure the stream of consciousness (Watzl in preparation); it does not come in degrees. It just comes in “on” and “off” variants, as one independently ongoing action rises to the foreground, followed by another, then another. Second, it is a nice question why the phenomenological foreground is associated with its functional aspect, of making them be performed “more intensely”. I think these questions can perhaps be answered, but explaining my position would necessitate a long and detailed discussion of the mereology of action (see my Hellie in preparation for the view).

Now to the metaphysical package (simplified somewhat to ignore the very important issue of the mereology of actions). We can think of a person's life as involving a flow of many simultaneous ongoing actions, typically progressing in parallel. Since, in

the spirit of $E=A$, the highly determinable kind *action* is just the same property as the highly determinable kind *experience*, the kinds of the totality of actions that a person is performing at a time are also the kinds of experiences that the person is undergoing at that time.

This totality of kinds, together with certain relationships they bear to one another, determines all the facts about one's consciousness at a time, by determining both the kinds of one's experiences and whether they are going on attentively.

We can see that my premisses fall out of this picture. Premiss 1 is true because, first, as we have seen, if some kind can stand in the relationships required for it to have an attentiveness status, it is an action/experience; and second, it is simplest not to have to explain why there should be any asymmetries among the actions/experiences in regard to whether they can be performed attentively. Premiss 3 is, as we saw in (C), relatively trivial; and, with its right hand side understood in accord with Premiss 3 (and assuming $E=A$), Premiss 2 is just Premiss 1 all over again.

That is not a defense of either the premisses or the picture, of course! The aim of that little exercise was to show that the premisses and the picture are proportional: just as the premisses entail $E=A$, $E=A$ in turn—in combination with the theory that attentiveness is just a relationship an action stands in to other ongoing actions—entails the premisses.

(J) Now that we know what I mean by my theses, let us run through the argument for $E=A$, in order to see how it applies to particular cases.

First let's run through the argument from left to right. According to Premiss 1, the K is a kind of action just if a K can occur attentively. Suppose that Jane is crossing the street attentively, watching out for cars: Jane's action is plausibly a *crossing the street*; so it is a crossing the street that occurs attentively. Joan is crossing the street inattentively: that's OK, all that is required for Premiss 1 is that some *crossing the street* occurs attentively. According to Premiss 2, a K can occur attentively just if a K must occur consciously. Jane's *crossing the street*, because attentive, occurs consciously. Joan's action doesn't occur attentively but, according to Premiss 2, occurs consciously nonetheless: after all, it is of the same kind as Jane's and Jane's action occurs attentively. And according to Premiss 3, a K must occur consciously just if the K is a kind of experience. Jane's and Joan's acts of *crossing the street*—like every act of this kind—are both conscious; so it follows that they are also experiences, and that the kind of experience to which they belong is *crossing the street*. No zombies, Jane and Joan.

Now let's run through the argument from right to left. Something is going on in

Bill's life: quite an experience indeed, that *K* he is undergoing. What can we learn about *K* from this? Because the *K* is a kind of experience, according to Premiss 3, Bill's *K* (like every *K*) is occurring consciously. So according to Premiss 2, Bill's *K* (like every *K*) is the kind of case that someone could undergo attentively. And so according to Premiss 1, we can conclude that Bill's *K* is something she did, one of his actions. Not merely passive, that Bill.

That is how you will reason if you accept my premisses. But why should you? Your mind was doubtless flooded with counterexamples as I ran through my chains of reasoning. But if I succeed in Part II, you will see two things: first, that the substantive Premisses 1 and 2 are at least *prima facie* plausible; and second, that it is not especially difficult to corral those counterexamples which come down the pike.

(*K*) Standing outside of experience, in my view (as we have seen several times), are *perceptual states*. At last, I give my opponent a chance to speak on this issue:

OK: this has gone far enough. You've been saying that kinds of *perceptual*, *sensory*, or *emotional* state—"perceptual" state, for short—are not kinds of experience? Let's give some examples here: seeing something red, having pain, being said. Maybe *those* aren't kinds of experience—perhaps because they aren't "narrow" enough. But certainly instances of those kinds have narrow parts, which are of *phenomenal* kinds. We might have to coin some terminology here: *seeing "as of" something red, being "as" having pain, being "as" sad*. And these phenomenal kinds are absolutely kinds of experience! Anyway, let's not be too picky about whether we're concerned with perceptual or phenomenal kinds: you think that neither are kinds of experience, and it's time for you to stop hinting at why you think this and *come out with it*.

My opponent is right: the time is ripe for me to explain why I think perceptual (or phenomenal) kinds are not kinds of experience; and to explain what they *do* have to do with experience, such that my opponent's position is so alluring. I proceed as follows. First, I advance a metaphysical theory of perceptual consciousness which distinguishes perceptual *states*, which are not experiences, from perceptual *acts*, which are; but which also recognizes certain tight links between the states and the acts. Second, I argue that the perceptual states are not experiences. Third, I diagnose the allure of the view that they are. (Questions about "narrowness" are complex and not easily addressed in the course of this paper: this is why I elide the distinction between the perceptual and the phenomenal. See my Hellie 2010 for some treatment of narrowness.)

Let us begin the **first part** of the strategy. I think about types of total perceptual states as complex properties of organisms. These properties are at least partly intrin-

sic: for example, pain is an intrinsic property of an organism. They are also partly relational: for example, *seeing a red tomato* is a property which cannot be instantiated by an organism in the absence of a correctly situated tomato. These properties are probably partly functional; and I see no reason to deny that they are also partly physiological (optic, acoustic, textural, chemical, reflective, etc.). They may well be partly representational: certain strands in vision science strongly suggest that perceiving involves representing the perceived scene. Finally, they may well also be partly *primitively qualitative*. This is to say that they may have parts which generate an explanatory gap with respect to the physical and are moreover fundamental: for example, the surface property of tomatoes, *redness*, may well be a fundamental non-physical non-mental property of the sort envisaged in Chalmers's "Eden" (Chalmers 2006); the same may be true of *pain*, a fundamental non-physiological non-mental property of organs; and the operation of the exteroceptive senses may well involve fundamental non-physiological non-mental properties of the eyes, ears, brain, and so forth.

For one to be in a perceptual state is for one to instantiate one of these complex properties (or for the organism one "rides around in" to instantiate it). Such properties are not, to repeat, kinds of experience; in virtue of instantiating such a property, one is not thereby conscious. I defend this claim shortly. However, instantiating such a property both bestows considerable powers of action on one and subjects one's courses of action to considerable coercion; so by **E=A**, such properties both enable and constrain the flow of consciousness.

Let me elaborate. In addition to the perceptual states, we should recognize the perceptual *acts*. There are five highly determinable kinds of perceptual act: looking, listening, tasting, smelling, and feeling.²⁰ These determinables fall away into an indefinite range of determinates: looking falls away into looking *for* something, looking *at* something, looking *around*, and so on; *looking at something* in turn falls away into such indefinitely specific acts as looking at Tweedledee, looking at a small red-winged blackbird of a certain shade and shape sitting on a wire, and so forth.

I don't think that there need be anything distinctively *behavioral* about any of these highly determinable acts: changing the way in which one is looking doesn't always involve turning one's head or eyes. A shift in which of these sorts of action one is performing may result merely from a shift in one's perceptual state: if the red circle at which I am staring is replaced by a green circle, I cease performing the act *looking at a red circle* and begin performing the act *looking at a green circle*. Even in the

²⁰These highly determinable kinds of act obviously correspond to the traditional five senses; I am inclined to think that they are the metaphysical basis of the traditional categorization.

absence of any change in one's perceptual state, one may shift the perceptual action one is performing merely by undergoing what the psychologists call a mere "attentional shift": when one "shifts attention" from the color of the wine in one's glass to its flavor in one's mouth, one has stopped looking and begun tasting; when one "shifts attention" without moving one's gaze from the dot on the left to the dot on the right, one has stopped looking at the dot on the left and begun looking at the dot on the right. Or at least one has stopped looking at the one on the left *attentively* and begun looking at the dot on the right *attentively*.

The determinable perceptual act of *feeling* is especially noteworthy in the present context. We speak in ordinary language not just of feeling the temperature, texture, or shape of an object in one's hand, but also of feeling the alignment of one's hips (for instance, in a yoga class). I maintain that this discourse has a common subject-matter: in each case, determinates of the act of feeling are under discussion. One doesn't just feel objects in our environment by touching and handling them: one also "mentally probes" or feels *around in* the condition of one's body. The psychologists might put this (or something similar) by saying that one is "shifting attention" around through various parts of the body; if the adverbial theory of attention is correct, then the subject-matter of this way of speaking is shifting of the part of one's body which one is attentively feeling around in.

So in particular, feeling pain is a kind of action, a determinate of the act of *feeling*. My claim is not that the type of perceptual state, *being in pain*, is a kind of action. In my view, it is possible to be in pain without feeling the pain. Being in this condition does not require any distinctive intervention, whether scientific, supernatural, or science fictional. Rather, it is the normal case: most of the time, when one is in pain, one is not feeling the pain. Thus I disagree with Lewis's (1980, 130) claim that "Pain is a feeling. Surely that is uncontroversial. To have pain and to feel pain are one and the same. For a state [property] to be pain and for it to feel painful are likewise one and the same" (on one reading, anyway). One is (or at least I am) typically in some degree of mild pain; most of the time, I do not dwell on the pain by feeling it attentively, but am rather doing something else attentively. There is nothing bizarre about my views. Everyone agrees (or should agree) that there is a perfectly good semi-ordinary sense in which one can be in pain, while "attending to" something else. My doctrine is a metaphysical claim about the subject matter of this semi-ordinary claim: under such circumstances, one is in a perceptual state which has the property of pain as an aspect or part; but when chalking up an inventory of what one is doing attentively, we should not chalk in the action of feeling (around in the part of the body afflicted by) the pain.

I have distinguished perceptual states from perceptual acts. Although these are distinct parts of reality, they are tightly intertwined, in two respects. Here is the first respect. One cannot look at Tweedledee if one does not see Tweedledee. If a red tomato is in a certain location, and one looks at the color of whatever is in that location, one will be looking at the red color of the tomato. If one does not feel a pain in one's toe, then if one feels around in one's toe for pain, one will not feel pain.²¹ These claims would seem to be true by metaphysical necessity. What could explain this metaphysical necessity? The evident answer is that one's perceptual state partly constitutes one's perceptual actions. When one performs a perceptual action of a certain determinable type, for example looking at the color of the object in a certain location, one's perceptual state as it were "leaps in" to make it determinate: for example, by determining it as an action of looking at something *red* (if that is the color of the object there).²² So perceptual states are entwined with perceptual acts by *partly constituting* them, in particular cases.

Here is the second respect in which perceptual states and perceptual acts are tightly intertwined. Often, one's perceptual state has a significant *causal* influence over which kind of perceptual act one performs. Psychologists characterize this influence by saying that some stimuli "capture attention". Here is how I understand this phenomenon. Suppose that, in a quiet library, one is engrossed in a book, when a loud alarm goes off. One hears the alarm: its sound, therefore, is partly constitutive of one's perceptual state. One's entry into a perceptual state that is so constituted causes one to perform a certain perceptual action: an act of *attentively listening* to the alarm. In just this way, the sudden onset of pain can cause one to change what one is doing attentively: hammering away, one strikes one's finger, and a pain in one's finger comes to partly constitute one's perceptual state. This change causes one to start feeling (around in the site of) the pain, attentively.

Aspects of the perceptual state compete for influence over what one does attentively with one another, and with one's ongoing projects. Some of us have amazing powers of concentration, and can refrain from attentively listening to the conversation at the next table, continuing one's project of writing a paper by attentively typing away. Perhaps we can think of the "intensity" of a stimulus of a given sort as amounting to its ability to control the character of one's attentive action in this way: the onset of a

²¹These claims are true simpliciter; and yet considerations of illusion, hallucination, twins, and spectral inversion complicate our efforts to provide a complete description of these situations which rules out all possible misunderstanding. See my Hellie 2010 for various nuances.

²²In this sense, perceptual actions are partially "spontaneous" but also partially "receptive". It is because the world constitutes a perceptual state which in turn constitutes a perceptual action that perceptual actions can be regarded as an epistemic guide to the character of extramental reality.

mild ache has much less power to distract one from hammering than does the onset of an intense pain. So perceptual states are entwined with perceptual acts, not just (as we have seen) by partly constituting them, but also by *exerting a significant causal power* over them.

In sum, the core doctrines of my metaphysical theory of perception are that there are both perceptual acts and perceptual states (of the sorts discussed); and that the latter both partly constitute and exert a significant causal power over the former (in the manner discussed). These core doctrines seem to me to be (upon reflection) relatively obvious. Moreover, they can be accepted completely independently of one's views on whether either the states or the acts (or both) are experiences.

So let the metaphysical theory be granted, and suppose that the kinds of action are *among* the kinds of experience. Now to the **second part** of the strategy: should we grant that the sorts of perceptual state are also among the kinds of experience? No. I provide three arguments.

First, according kinds of perceptual state this status would disunify the kinds of experience. Perceptual states and actions are highly dissimilar. Actions are dynamic, at least partly “spontaneous”, potentially attentive, and morally significant; perceptual states are static, purely “receptive”, not potentially attentive, and morally neutral. If consciousness and its determinates are in at the fundamental level, it would be oddly displeasing if the determinates fell away into two such different categories.

The *second argument* that perceptual states are not experiences is based on a “pessimistic meta-induction in the philosophy of perception”. As Martin (2004) has remarked, it is a surprising fact about the philosophy of perception that there is disagreement about the nature of perceptual “experiences” (he doesn't use the scare-quotes). Martin has in mind the debate among direct realists, sense-datum theorists, and intentionalists concerning the natures of the object of perception and our relation to that object; we could add to this the endless early twentieth-century debates about whether the surface or the whole object is the object of perception (Swartz 1965); and the endless contemporary debates among Russellian, Fregean, rich, poor, singular, general, and gappy intentionalists, and over the status of the qualia theory (Siegel 2005). A pessimistic meta-induction on the philosophy of perception is suggested: we cannot ascertain the kinds of perceptual “experiences” merely by first-person reflection. But how can this be, if these “experiences” are genuine *phenomena*, or forms of consciousness, the kinds of which should be simply revealed? Although Martin raises the question, it is not clear to me how he answers it. But isn't the answer obvious? Take the *modus tollens*! Perceptual “experiences” are not in fact forms of consciousness;

they are not phenomena; their nature is not simply revealed.²³ Martin cannot accept this resolution (he seems to endorse aspects (a), (c), and (d) of the phenomenal state conception). But if $E=A$ is true, it becomes available.²⁴

The *third argument* that perceptual states are not experiences is that, if perceptual states are not experiences but perceptual acts are, we have a pleasing resolution to a somewhat vexing puzzle. Consider the “refrigerator light trilemma”: a range of cases *seem* to show that experience is “blind” to certain phenomena which perceptual states *seem* to be required to register (Block 2001). For example: while one is counting the passes of the basketball, the gorilla strides to center stage, dances a jig, and departs unnoticed (Simons and Chabris 1999). What just happened? If perceptual states are not experiences, no problem: one saw the gorilla but did not experience it (did not perform any action in regard to the gorilla).

By contrast, if perceptual states are experiences, we face an unfortunate choice. Either one *does experience* the gorilla, or one *does not see* the gorilla. If the former, then why does one not notice the gorilla, why does one not endorse the presence of the gorilla, why does one not recall the gorilla? The option renders experience very strangely opaque to judgement in a way that upends even relatively modest doctrines of privileged access. I leave this point here and push on to the second option. I begin with an example. Suppose that at the table next to ours, a hushed but prurient conversation proceeds. It is readily discernable when listened to, but its causal pull on my attentive perceptual acts is rather weak. If engaged in a captivating conversation, I do not attentively listen to it, and the conversation is as nothing to me; as soon as my conversation partner gets up to order more coffee, the conversation recaptures my attention. How does it do this? Presumably because it remains part of my perceptual state. If I ceased to perceive it when I did not attentively listen to it, why would it have any more significant a pull on my attentive perceptual acts than a conversation in Timbuktu? Clearly I do perceive it. But then parity of reason suggests that I see the gorilla. If seeing the gorilla fails to exert a causal pull on my perceptual acts—rather surprisingly, in light of the ordinary strangeness of such a stimulus—this is because my ongoing project of counting the passes of the basketball overwhelms the causal

²³Indeed, the epistemology of consciousness developed in Chalmers 2003 would seem to make the sort of dispute Martin describes impossible in the long term among reasonable people.

²⁴A related consideration: the perceptual acts fall into the categories of the traditional five senses (they are the determinates of looking, listening, feeling, smelling, and tasting), but the perceptual states do not: the familiar Gricean ways of individuating the senses (Grice 1962), all tied to perceptual states, as is well-known, do not work. The tenacity of the theory of five senses strongly suggests that they are available to reflection in the manner of experience; if so, then perceptual acts are experiences and perceptual states are not.

pull that seeing a gorilla would ordinarily exert.

Either complaint can of course be resisted, resulting in patterns of dialectic that are relatively easy to envisage. I do not claim to present a knock-down argument for my case. But the argument is still strong: we can simply *avoid* either dialectic if we reject the claim that perceptual states are experiences. Why not save ourselves the effort?

Third and finally, I turn to explaining the allure of the *perceptualist* doctrine linking kinds of perceptual state to kinds of experience—an allure to which, as we have seen, I have not always been immune. A strong version of the perceptualist view claims that the aspects of one's perceptual state *exhaust* the kinds of one's experience (perhaps with a bit of fiddling around the edges); a weak version claims merely that every aspect of one's perceptual state is *among* the kinds of one's experience. I detect three sources of the allure of perceptualism to contemporary philosophers: (a) certain phenomenological aspects of perception as filtered through the Moore-Russell tradition; (b) a transformation worked on the relation between experience, action, and sensation in a chain of influence stretching from Ryle to Lewis; and (c) a desire for perceptual experience to serve as an epistemological "neutral court of appeal".

I discuss *first* (a) the phenomenological source of the allure of perceptualism. There is no doubt that one's perceptual state is closely tied to experience. (i) For one thing, one never performs a perceptual action without its incorporating some aspect of one's perceptual state, where the aspect renders the action determinate in the way we have discussed. In this way, aspects of one's perceptual state become part of one's perceptual act—which is itself part of one's experience. So there is at least an abstract notion of 'part' on which such aspects are part of one's experience. (Of course, such aspects are only a part of one's experience "as conditioned" by one's perceptual act and thus not as kinds of experience; and the remainder of one's perceptual state is no part of one's experience.) (ii) Another respect in which one's perceptual state is closely tied to experience is that *only* one's perceptual state can determine a perceptual action in this way. In this sense, one's *total* perceptual state is special: every aspect of one's perceptual state is related to experience in a way that is very different from the way in which any other aspect of the world is related to experience. Conditions in Timbuktu (assuming one does not perceive Timbuktu!) have no capacity whatever to determine the character of one's perceptual actions. (iii) A final respect in which one's perceptual state is closely related to experience is that it is *hard to ignore*: one is almost always performing some perceptual action. We almost always have a reason to incorporate perceptual action into our experience: most of one's actions involve one's body in interaction with perceptible objects in one's near environment, and are thus performed

only poorly when perceptual action is not performed. And, as noted, one's perceptual state exerts a strong causal influence over one's actions. It is only in a state of deepest concentration on matters of pure thought that one can mask this influence entirely.

In sum, one's perceptual state (i) *can in part be*, (ii) *can alone be*, and (iii) *almost always in part is* incorporated into a perceptual act one is performing, and thereby into one's experience. These facts are, I think, tacitly known by all of us. When constructing theories, it is easy to conflate these facts with the strong version of perceptualism, if we ignore the following phenomena: non-perceptual acts; diversity of perceptual acts that does not stem from diversity in underlying states (and, in turn, the possibility of perceptual multitasking); and the 'almost' in (iii). The result of all this ignoring is a picture on which one's perceptual state can alone be and always in part is incorporated into the only (relevant—shhh!) sort of act that is ever performed, and thereby into one's experience. This looks an awful lot like the familiar picture in which experience consists of a perceptual state with a somehow superadded attentional spotlight—the operation of which is the only (relevant—shhh!) sort of act that is ever performed.

But why would anyone engage in this pattern of ignoring? Straightforward answer: our philosophical tradition of thinking about consciousness is heavily influenced by the Moore-Russell approach to perception; on this tradition, we care about perception thanks solely to its serving as a referential link to the external world. To discuss this referential link, the “attentional spotlight” is perceptual action enough.

As we will see in (L), once we stop engaging in this pattern of ignoring, the phenomenological allure of strong perceptualism evaporates; and with it, any reason to endorse even weak perceptualism.

A *second source* of the allure of strong perceptualism stems from a transformation wrought in the main stream of discussion of consciousness on the way from Ryle to Lewis. In brief outline the story is this (see (L) for more detail): in the mid-1950s, UT Place proposed an alternative to Ryle's view on the nature of the actions that are experiences. According to Ryle, they are ordinary actions; according to Place, they are a special sort of internal part of perceptual action (“sensational action”, perhaps). Somewhat later, JJC Smart took up certain aspects of Place's view, but largely overlooked the role of agency in the view, mistaking it for the view that experiences are a lot like sensational states. Somewhat later, David Lewis overlooked a certain nuance in Smart's view, and started using ‘experience’ synonymously with ‘sensational state’, thus rendering strong perceptualism in effect analytic. The great influence of Lewis's approach, in turn, burned this usage into the collective linguistic practice of

the philosophy of mind.²⁵

Third and finally, some rather broad-brush reflections on the significance of $E=A$ for epistemology will point to a source of the allure of weak perceptualism stemming from (c) the desire for perceptual experience to serve as an epistemological neutral court of appeal.

Suppose that we have in hand a notion of “propositional justification” (Pryor 2000), according to which a proposition has one of (at least) three justificatory statuses (these statuses may be relative, perhaps, to a choice of “linguistic framework” or something along those lines). First, the *a posteriori* propositions are justified only relative to a course of experiences, in at least the following sense: there is a nonconstant function mapping a type of course of experience into a proposition, such that a subject who has just undergone the course of experience has justification for the proposition. Second, the *a priori* propositions are justified “absolutely”, in the sense of not relative to any given course of experiences; the members of a third class of *a priori* falsehoods are “anti-justified” absolutely.

This picture is highly abstract. In particular, the assumption that there is a *function* mapping a type of course of experience into a proposition it justifies does not by itself entail anything about what that function is like. Still, if $E=A$ is correct, it may be that this function is highly sensitive to which perceptual action one is performing. So for example, a pair of subjects in perceptual states of the same type, looking at the same thing, might nevertheless diverge in their propositional justification concerning that thing because of *how* they are looking at it. Suppose that perceptual justification is “action-dependent” in this way: then what?

The influence of ways of looking on propositional justification records an obvious article of common sense: the expert learns more from observation than the novice; the cardiologist and the nephrologist learn different things from looking at the same films. But if I read the contemporary literature in epistemology correctly, the literature shrugs its shoulders at this phenomenon, and it shouldn't. My sense is that the literature is on the side of a “rich” understanding of propositional justification, according to which one has justification for a very large range of propositions concerning one's perceptual state: the novice and the expert, like the cardiologist and the nephrologist, do not differ in the propositions for which they have justification. Where they differ is in their ability

²⁵Other sources: Carnap's usage; a lack of a well-worked-out and widely-known mathematical framework for discussing temporally dynamic phenomena and a consequent privileging of the momentary and stative; a general fondness stemming from Frege and Kant for representation, culminating in the computer model of the mind and its consequent minimization of the role consciousness to providing “color” for representational states (Harman 1990).

(or tendency) to *access* this justification: they differ merely in the judgements they are able (or liable) to form on the basis of this justification. The action-dependence of justification suggests a different characterization. Propositional justification is “poor”: the cardiologist and the nephrologist differ in the propositions for which they have justification, not merely in their ability or tendency to access this justification. The justificatory power inherent in one’s course of perceptual experiences is constitutively dependent on what one has been doing.

If this characterization is correct, significant implications for our understanding of epistemic rationality may follow. After all, what one has been doing is constitutively dependent on one’s aims, goals, projects, other ongoing actions, and the like—by one’s “values”. So if the epistemic justification afforded by one’s perceptual experiences is constitutively dependent on what one has been doing, this justification is constitutively dependent on one’s values.

This “value-ladenness” may undermine the “objectivity” of perceptual justification. If I understand the purpose of the concept of propositional justification correctly, its role in epistemology is in part to undergird the doctrine that there is a value-neutral characterization of one’s epistemic position by reference to which we may assess whether one has engaged in optimal epistemic practices. This propositional justification puts teeth in advancing certain perceptual practices over others as a sort of “regulative ideal” which subjects should attempt to live up to, more-or-less independently of what they value. By way of comparison, the doctrine that some propositions “are a priori”—that everyone always has justification for such propositions—provides a sort of regulative ideal by reference to which certain practices of reasoning may be evaluated: one reasons well just to the extent that one’s reasoning reaches out and grabs the justification that one has already.

Putting the point a bit more vividly if crudely, if there is value-neutral justificatory power in perceptual experiences, we may say to a subject “you’re looking at things wrong”—and legitimately have in mind a respect of wrongness which does not vary from subject to subject on the basis of their values, but is rather an impartial wrong of bare epistemic rationality. By contrast, if perceptual justification is value-laden, this value-neutral basis may be a mirage. We may find that sometimes, when we think someone is looking at things wrong, this boils down to nothing more than our disapproval of their goals.²⁶

²⁶Here I am thinking less of people reaching inconsistent conclusions by looking and our being unable to say who got to the truth and who got to falsity, and more of people reaching conclusions about different aspects of a situation and our being unable to say which is more important.

Summing up: if $\mathbf{E}=\mathbf{A}$ is correct, the perceptual state is distinct from the perceptual experience. The perceptual state is value-neutral but not an aspect of consciousness; the perceptual experience is an aspect of consciousness but is value-laden. So unless we are willing to regard the normative force of perception as grounded in something “blankly external” to experience, we may need to acknowledge a significant limitation on the power of appeals to value-neutral “reason” to settle disputes. It is plausible, then, that the allure of the perceptualist doctrine that perceptual states are part of experience is in part a result of our fondness for such appeals.

(L) I finally will briefly sketch further motivations for $\mathbf{E}=\mathbf{A}$, less with the aim of saying anything convincing than with the aim of displaying that there is more to my position than a neat structure that can be built by using ordinary language in certain ways if we are willing to be sufficiently procrustean. At a number of points, doctrines hastily sketched in this section will receive more extensive discussion later.

The motivations I will discuss are (i) that $\mathbf{E}=\mathbf{A}$ unifies mental phenomena; (ii) that $\mathbf{E}=\mathbf{A}$ unifies the epistemology of the self; (iii) that $\mathbf{E}=\mathbf{A}$ provides a good explanation of the meaning of ordinary discourse about experience; (iv) that $\mathbf{E}=\mathbf{A}$ is superior phenomenologically to its competitors; (v) that in light of the post-WWII history of the philosophy of consciousness, $\mathbf{E}=\mathbf{A}$ should be regarded as a default position; and (vi) that collectively, the previous considerations (plus the argument from attention) suggest that $\mathbf{E}=\mathbf{A}$ has broad and deep explanatory power.

First, $\mathbf{E}=\mathbf{A}$ unifies mental phenomena. I develop this point by showing how a certain case for representationalism functions better as a case for $\mathbf{E}=\mathbf{A}$.

Representationalism is sometimes provided with prima facie motivation by its unificatory power: minds characteristically both represent and are conscious; it would be unpleasant if these had nothing to do with one another; so a view on which there is some essential connection between the two tidily binds together a range of phenomena (Chalmers 2004).

This argument works better on behalf of $\mathbf{E}=\mathbf{A}$. After all, it is prima facie highly plausible that both consciousness and agency are both necessary and sufficient for mentality. By contrast, it is a commonplace that representation is not sufficient for mentality. And conversely, representation is not obviously even *especially pervasive* in mentality. Consider shopping: that is something most people do a lot of. Is shopping a kind of representing? No doubt shopping involves representing in the form of the occasional appeal to opinions about what one needs and where one can get it. But hiking involves eating in the form of occasional handfuls of gorp: this doesn't make hiking a kind of eating. So a reasonable attitude would be that, while representation

is sometimes pressed into service by mentality, the essence of mentality is conscious agency.

Of course, in (E), we adumbrated a representationalist version of $\mathbf{E}=\mathbf{A}$ based on a Davidsonian approach to agency. Still, as we have seen, that approach is contentious. Instead of betting on the success of a certain project in the philosophy of action, unifiers should rather push for $\mathbf{E}=\mathbf{A}$ up front and let the chips fall where they may with representationalism.²⁷

Second, $\mathbf{E}=\mathbf{A}$ unifies the epistemology of the self. Privileged access doctrines of some degree of strength about both consciousness and action are plausible. Arguably (Chalmers 2003), one is generally in a position to know roughly which kind of experience one is having, using a special approach which others cannot use to know which kind of experience one is having. And arguably (Anscombe 1963), one is generally in a position to know roughly which kind of action one is performing, using a special approach which others cannot use to know which kind of experience one is performing.

If $\mathbf{E}=\mathbf{A}$ is correct, the content of the former such knowledge is the same as that of the latter such knowledge; and the special approaches by which one attains this knowledge are the same. Awareness of consciousness just is action-awareness. There is no need for concern about a proliferation of special methods; insights from the literature on each sort of knowledge can be bundled together into a more thoroughly articulated unified theory.

Third, $\mathbf{E}=\mathbf{A}$ provides a good explanation of *ordinary language discourse about experience*.²⁸ Recall Thau's claim that "to have a bank teller's experiences is to have done the things that a bank teller does". This is a deep point, not merely a superficial feature of the technical language of the job market.

When we discuss our experiences, a great deal of what we say has to do with what we did: "I strolled around Manhattan all afternoon". We also make claims about our perceptual states: "I saw the Empire State Building and my feet hurt"; and about purely objective matters: "the streets were thronged". This raises two puzzles: first, why do we make claims about objective matters of fact when we are trying to characterize events in our mental life? And second, *prima facie*, experience is a unified phenomenon; but the disunity of such discourse might suggest otherwise.

If $\mathbf{E}=\mathbf{A}$ is true, we can resolve both puzzles by providing a pragmatic story on which audiences can extract information about one's actions from one's claims about (i) perceptual states and (ii) the objective: the primary aim of all such discourse is then

²⁷Thanks for discussion here to Sebastian Watzl.

²⁸A lot of this can be found at experienceproject.com.

to communicate information about one's actions, which, by $E=A$, is information about experiences.

Here's a sketch of the pragmatic story. On claims about *perceptual states*: if knowledge is the norm of assertion (Williamson 2000), then every assertion that p implicates that one knows that p ; so if one asserts that one saw the Empire State Building, one implicates that one knows that one saw the Empire State Building. One's audience recognizes this implicature, and wonders how one came to know that one saw the Empire State Building. One's audience, let us suppose, knows that typically, a person can only know that they see o by looking at o . So one's audience can conclude that one looked at the Empire State Building: this is an action. On claims about *the objective*: one's assertion that the streets were thronged implicates that one knew this. Since the context has established that one was in Manhattan, the most natural explanation of how one came to know this is that one observed thronged streets. So one's audience can conclude from one's assertion that one observed thronged streets: this is an action.

Working this into a full-fledged argument would require canvassing and undermining solutions to our puzzles that do not presuppose $E=A$; I leave this task for another day.²⁹

Fourth, $E=A$ can accommodate certain *phenomenological* points that its competitors cannot.

What is it like for me, as I type this sentence? Can this be characterized solely by citing my phenomenal properties, or solely by citing my perceptual "experiences"? Not plausibly. For one thing, it is evident and widely recognized that *attentional* shifts can alter the character of one's consciousness in ways that seem independent of changes in one's perceptual "experience": if one sees a pair of red dots against a black background, an attentional shift from one dot to the other would influence what it is like for one without any change in one's perceptual "experience" (Chalmers 2004).

But on the adverbial theory of attention, an attentional shift is a shift in what one is *doing attentively*: in the two-dots case, one shifts from attentively *looking at* the left dot to attentively *looking at* the right dot. So the case of a "perceptual action" (in the sense of (K)) like looking at something, change in action can influence what it is like for one.

But now the floodgates are open: *thinking* is an action, can influence what it is

²⁹Such approaches would need to accommodate the mounting evidence that the "folk" lack the concepts used to articulate the phenomenal state conception (Sytsma 2009). (Sytsma draws the conclusion from this evidence that there is no hard problem of consciousness. But if $E=A$ is correct, there might still be a hard problem of consciousness, if it is a problem about action. And there might be a hard problem of qualities, even if it is not a problem about consciousness.)

like for one, and is not a perceptual action. So perceptual and cognitive actions can influence what it is like for one. But *moving one's body* is an action, can influence what it is like for one, and is not a perceptual or cognitive action. And such actions can be hard to dissociate: playing Defender involves looking at aliens and people, thinking about how to save the people from the aliens, and moving my hands so as to shoot the aliens where I see them and save the people in the way I planned without shooting them. It would not be easy to isolate a class of actions that are experiences from a class that are not.³⁰

The remarks so far are compatible with a mixed view, on which perceptual states are experiences and so are actions. Against this, I suggest: a unified view is superior on theoretical grounds and should be strived for; since actions have proved indispensable as experiences, we should not push for a perceptual states-only view; and, as we have seen in (K), perceptual states should not be included in experience, so that there is no reason to push for a mixed view. Accordingly, we should endorse the actions-only view; which is to say that we should endorse $E=A$.³¹

Fifth, $E=A$ is *smiled upon by Clio, the Muse of History*. If we roll the tape back sixty years to the dawn of the contemporary philosophy of mind, we find the premisses in the argument from attention in place: there was enough in the zeitgeist to support $E=A$. But within twenty years, the phenomenal state conception was firmly entrenched. What explains this shift? Not good epistemic reasons! In an alternate history, where we had made fewer mistakes, we would now endorse $E=A$. Surely we should use the judgements of our better selves as a guide to belief: so we should endorse $E=A$. That's the form of the argument: now to fill in the relevant historical details.

Once upon a time there was a philosopher who wrote “a person's thinking, feeling, and purposive doing cannot be described solely in the idioms of physics, chemistry, and physiology”. This philosopher was Gilbert Ryle, writing in *The Concept of Mind* (Ryle 1949, 18); his name for the view of the mind that he developed there was “logical behaviorism”. This doctrine presumably had something to do with “behavior”—but what?

Much attention has focused on Rylean analyses of mental states such as belief in terms of dispositions to behavior. But what has been overlooked is Ryle's understanding of “behavior”. In the popular mind, Ryle thinks of behavioral properties as

³⁰Internalists will no doubt suggest that the “tryings” are internal and the only actions which count as experiences. Against this, it is arguable that trying to *A*, like intending to *A*, is just a sort of incomplete *A*-ing: see footnote 10.

³¹Footnotes 12 and 14 also discuss phenomenological matters.

conceptually reducible to the nonmental. But the passage above shows (to my mind conclusively) that he did not. Rather, Ryle's behavioral properties seem to be quite a bit like our kinds of action: as including thought, perceptual action, and bodily action; and as conceptually irreducible.

This is in line with Ryle's purposes. Logical behaviorism was supposed to supplant the then state of the art "official doctrine" of the "ghost in the machine". Ryle's problem with this doctrine stemmed from its endorsement to the narrowness of mental properties: this rendered all of them imperceptible by others (see for example Ryle 1949, 60). Which kind of action someone is performing is often perceptible: I see Bill mowing the lawn. Of course when I am in a condition that seems like seeing Bill mowing the lawn, it is coherent to suppose that I am subject to an illusion here: that what I am really seeing is Bill merely displaying lawn-mowing behavior, in the sense of (E). But this doesn't undermine the perceptibility of Bill's *genuinely* mowing the lawn unless the perceptibility of redness is undermined by the coherence of seeing something white under red light (it isn't).

Ryle endorsed the adverbial theory of attention (pushing a reductive view on which A-ing attentively is A-ing in a way that maximizes the chance of success: 137–9). So I am inclined to think that Ryle endorsed our Premiss 1. Sadly, for some reason—bloody-mindedness?—Ryle equated consciousness and attention (157–8). Since he acknowledged that sometimes one As inattentively, he did not endorse our Premiss 2.

We find support for Premiss 2, on the other hand, in the writing of UT Place, the father of central state materialism. The epoch-making clash that set in motion the modern history of the philosophy of mind, eventually spawning the phenomenal state conception, was Place's assault on Ryle in 'The concept of heed' (Place 1954)—the readership of which is utterly dwarfed by both its historical significance and the readership of its mopping-up successor paper, 'Is consciousness a brain process?' (Place 1956). In Place's view, attention is high-grade consciousness (244).

But Place rejects our Premiss 1. He rejects the the adverbial theory of attention, endorsing instead the *kind view* of attentional phenomena: so, recall, he recognizes the action kind *attending* (as he puts it, "heeding") and rejects the non-kind property *attentiveness* (or "heedfulness"). Why? Not for good reasons. The crux of the argument (in sec. VIII) is this: (i) Ryle thinks that having a pain is necessarily coextensive with consciously observing the pain; (ii) Ryle equates consciousness and attention; so observing pain is necessarily attentive; (iii) so observing pain is essentially attentive; so there is such an action-kind as *attending* (found at least in cases of observing pain). Of course this only gets us to the mixed view; perhaps Place thought that having gone this

far, we should just eliminate *attentiveness*. That’s a lot of bad argument: (iii) slides from necessity to essentiality in a way we warned against in (H); steps (i) and (ii) are ad hominem (*we* don’t accept either of those claims)—and Place himself rejects the equation in (ii)!

Since Ryle didn’t have a good case against Premiss 2, and Place didn’t have a good case against Premiss 1, the sensible reaction to this clash would have been to endorse Premiss 1 and Premiss 2 and then push on to $E=A$. Sadly, that is not what happened. Instead, we wound up with the phenomenal state conception: the trail of confusion that led to this is the second part of our story.

Recall the phenomenal state conception:

- a. Perceptual “experiences” are paradigm experiences;
- b. Experiences are *narrow*;
- c. Experiences are *passive*;
- d. Experiences are *states*.

What did Ryle and Place think about all this? As we have seen, Ryle did endorse a tie between action and consciousness, though a weaker tie than in our $E=A$. In Ryle’s view, a *K* can occur consciously just if the *K* is a kind of action: this follows from Premiss 1 and the equation of attention and consciousness. Place follows Ryle in seeing a tie between action and experience, if a rather different one: in his view, the kind *attending pain* is a paradigm kind of experience; it is a high grade of a more general sort of action, which we could call *consciousnizing pain*. So Place thinks that every kind of experience is a kind of action, and not vice versa. It is in the ballpark of their views, then, to say that Ryle rejects all of (a)–(d), while Place endorses (a) and (b), rejecting (c) and therefore (d).

Why does Place endorse (b)? On grounds that resemble the argument from hallucination for the narrowness of experiential kinds (Place 1954, 254). This leads to our next point, Place’s most famous contribution, the “identity theory”. Place was confused about what Ryle didn’t like about the official doctrine. In Place’s view, the worry was *metaphysical*, about dualism, rather than *epistemological*, about perceptibility: in effect, he put too much weight on the “ghost” part and not enough on the “in” part of the “ghost in the machine” rhetoric (254–5). So Place proposed to resolve this by letting the act of consciousnizing pain be a “brain process” rather than a ghostly process. Voila: the birth of the identity theory.

Place wrote the follow-up paper to block an objection to his view that experiences are neural acts. The worry stemmed from the “phenomenological fallacy”, according to which

when the subject describes his experience, when he describes how things look, sound, smell, taste, or feel to him, he is describing the literal properties of objects and events on a peculiar sort of internal cinema or television screen, [...] the “phenomenal field”. (Place 1956, 49)

Instead,

we describe our conscious experience [...] by reference to [...] physical properties of [externalia] which normally [...] give rise to the sort of conscious experience which we are trying to describe. In other words when we describe the afterimage as green, we are not saying that there is something, the afterimage, which is green; we are saying that **we are having the sort of experience which we normally have when [...] looking at a green patch of light.**³² (49; my boldface emphasis)

I am not completely clear on why Place thought the “phenomenological fallacy” was supposed to threaten his neural action theory of experience, but note the reply: it is going to blow up.

In the next stage in our story, Place’s friend JJC Smart decides that Place is on to a really excellent theory but has not done such a great job of explaining it. So Smart follows up by in effect rewriting Place’s paper: the result is ‘Sensations and brain processes’ (Smart 1959). Smart was right: the writing in Place’s paper is not very clear. So much so, in fact, that Smart—who seems not to have read ‘The concept of heed’—did not understand Place’s theory: which makes for a momentous expository mistake.

To see the source of the mistake, reflect on how weird Place’s theory is. He is correct to say that experiences are actions: this squares perfectly well with ordinary understanding, to which, like Ryle, he is trying to be faithful. But his case against adverbialism has led him to say that pains and other sensations are experiences; so he ends up saying that sensations are actions. Of course they aren’t! They are states.

Smart ends up saying both things about them. They are states and sensations at the beginning of the paper; but they are processes and experiences in the parts of the paper most clearly under the influence of Place: in the title, and in the reply to Max Black’s famous “Objection 3”, when Smart adapts Place’s stuff on what we are saying when we talk about the color of an afterimage (149).

³²Here Place reiterates a point made earlier in ‘The concept of heed’, page 253.

Smart is aware of the tension. “Objection 5” runs:

It would make sense to say of a molecular movement in the brain that it is swift or slow, straight or circular, but it makes no sense to say this of the experience of seeing yellow. (151)

That is to say, of what is evidently a state, Smart is saying it is a process.³³ Summing up, Smart brings us closer still to the phenomenal state conception: following Place, he endorses (a) and (b); in his confusion about Place’s view, he sometimes accepts and sometimes rejects (c) and (d).

In the next stage in our story, David Lewis writes ‘An argument for the identity theory’ (Lewis 1966), adapting Smart’s stuff (ultimately Place’s) on what we mean when we talk about the color of an afterimage into the view that all concepts of kinds of sensation fix reference through reductive functional characterizations (20). Following Smart, he uses ‘experience’ to refer to sensation. Lewis seems not to have read Place’s paper, and briskly cleanses his discussion of all traces of process: experiences are now unequivocally states. The phenomenal state conception is complete: Lewis endorses (a)–(d) in this paper and for the rest of his career.³⁴

And in the final stage of the story, due to the immense influence of Lewis’s paper, we all developed the bad habit of calling sensations ‘experiences’. Thanks to several decades of obsession with dualism, philosophers of mind never looked back to ask why we were constantly saying such strange things. The phenomenal state conception became entrenched.³⁵

Summing up: Ryle and Place both denied that anything but an action could be an experience. But Place endorsed some bad arguments that caused him to think—contra

³³His reply is lame:

All I am saying is that ‘experience’ and ‘brain process’ may in fact refer to the same thing, and if so we may easily adopt a convention [. . .] whereby it would make sense to talk of an experience in terms appropriate to physical processes. (151–2)

Whatever “conventions” we may adopt, it remains metaphysically impossible that some occurrence is both a state and a process.

³⁴The examples of kinds of experience discussed in the Lewis corpus are: pain (1966, 1972, 1980, 1994); sensation of red/seeing red/green/other colors (1972, 1988, 1997), tasting vegemite (1983*a*, 1988, 1995), smelling a skunk (1988), being hot/cold (in various body parts) (1988), having a sweet taste (1995), perception in general (1996).

³⁵Epilogue: in Lewis’s final word on these issues (Lewis 1995), he winds up accepting in effect both that almost everyone finds zombies conceivable, and that they are not. The source of this widespread irrationality is not explained; nor does he address the obvious questions about the bearing of the existence of such widespread irrationality on projects in metaphysics involving an extensive role for conceptual analysis. If we had only paid attention to Ryle, this tragedy could have been averted!

Ryle—that the experiences were the sensations: he concluded that sensations were actions. This weird idea confused Smart, who ended up keeping Place’s view that every experience is a sensation, while vacillating about whether sensations are actions (as Place thought) or states (as they obviously are). Lewis brushed aside this agonized aspect of Smart’s discussion, and ended up keeping the Placean view that every experience is a sensation while also endorsing the obvious claim that every sensation is a state.

So the doctrine that experiences are sensations entered the story thanks to bad arguments; and thanks to the presence in the discussion as examples of sensations properly so-called, these frankenstein occurrences gradually became less experience-like and more sensation-like. This is how we all wound up stuck with the phenomenal state conception of experience.

Let’s fix this mess. Rewind the tape: with both Ryle and Place, say that nothing but an action can be an experience; grab Premiss 1 from Ryle and Premiss 2 from Place. Ditch Place’s bad arguments: that way, sensations never contaminate the theory of the kinds of experience. Wave goodbye to the phenomenal state conception; endorse **E=A**.

Finally, there is the obvious holistic meta-argument: the existence of such a wide range of diverse considerations on behalf of **E=A** suggests two things. First, to avoid **E=A**, one has one’s work cut out for one: one must rebut all of the arguments sketched. Second, by accepting **E=A**, one can explain a great deal. **E=A** is, of course, incompatible with the phenomenal state conception, a powerful article of dogma in the contemporary philosophy of consciousness; and there is much to be said for remaining on familiar ground. But I am inclined to think that the evident risks of wasted effort involved in any attempt to head out for the territory are, in the present case, outweighed by the likelihood of digging up rich unmined veins.

Part II

In defense of the argument from attention

Recall the argument from attention:

1. The K is a kind of action just if a K can occur attentively;

2. A *K* can occur attentively just if a *K* must occur consciously;
3. A *K* must occur consciously just if the *K* is a kind of experience.

I turn now to defenses of Premises 1 and 2: in each case, first left-to-right, then right-to-left—Premiss 3, recall, was defended above, at (C). I am inclined to regard each of the premisses as *prima facie* quite plausible, for reasons we will see. (Recall also that the premisses follow from **E=A** together with the theory of attention gestured at in (I), which is a natural fit with the adverbial theory defended in (H): this lends the premisses plausibility from other directions.) And yet there are a range of purported counterexamples to the premisses. Most of the discussion, therefore, will be judo: redirecting the dialectical force of these objections.

Premiss 1 Action and attention

Premiss 1, recall, was this:

1. The *K* is a kind of action just if a *K* can occur attentively.

Let me be explicit about how this is to be understood. I presuppose the adverbial theory of attention, as discussed in (H). Then the right hand side presupposes that—and is therefore true only if—possibly, there exists an occurrence which has *K* as its kind and *attentiveness* as an inessential property. If there is no such inessential property, then the left-to-right direction of Premiss 1 is not true.

In (I), we suggested that for an action to occur attentively is for it to be so related to other actions as to be in the “foreground”. If this is correct, the right-to-left direction of Premiss 1 is clearly also correct, and the left-to-right direction is *prima facie* plausible: for any particular kind of action, what is to prevent its instances from occupying the foreground, if only for a while?

The left-to-right and right-to-left directions of Premiss 1 are as follows:

- 1L. If the *K* is a kind of action, a *K* can occur attentively;
- 1R. If a *K* can occur attentively, the *K* is a kind of action.

I discuss strategies for responding to counterexamples first to 1L, then to 1R.

1L From action to attention

A counterexample to 1L would be a kind of action which could only be performed inattentively. In the following sense: it is not possible that an action of that kind has the non-kind property *attentiveness*. (Note that it would not suffice for there to be a kind of action that *could* be performed inattentively: as we saw in (I), plausibly any kind of action could be performed inattentively.)

1L(a) “Attending”?

Recall the discussion in (H), in which we distinguished three theories of attentional phenomena: the *kind* theory, on which all attentional phenomena are actions of kinds like *attending to a tree*; the *adverbial* theory, on which all attentional phenomena are actions with the non-kind property *attentiveness*; and the *mixed* theory, on which both are observed.

Let’s assume that the kind property excludes the non-kind property. Then according to the mixed theory, *some* kinds of action are counterexamples to 1L (namely, the attentional kinds); while according to the kind theory, *every* kind of action is a counterexample to 1L (because the non-kind property does not exist).

But, as we saw in (H), neither of these theories is especially plausible. The adverbial theory seems to be correct; and the adverbial theory does not require the existence of *any* counterexamples to 1L.

1L(b) Essentially inattentive actions?

A second sort of counterexample to 1L would be a kind of action which is not essentially an attentional kind, and the instances of which also cannot possibly have the inessential property *attentiveness*. Such a kind of action could not possibly be an attentional occurrence.

Such a kind would certainly be unusual. Nearly everything I can do, I can do attentively.³⁶ I can talk attentively, walk attentively, eat attentively, think attentively, rub my

³⁶Everything *sustained*, anyway. Perhaps some actions are “achievements”: events which, by their type, can only last for an instant. Examples: reaching the summit, spotting Paula, flashing on a memory. There is perhaps something a little odd thinking of these as things that can be *performed*, let alone *attentively*. Doesn’t attention require *sustained focus*? I am inclined to think that there are several ways that achievements can be treated before abandoning 1L becomes necessary. First, perhaps they are actions: one can be held responsible for reaching or failing to reach a certain point, spotting or failing to spot someone or something, or failing to recall some piece of information (can one be held responsible for recalling something? “you better forget you saw me!”). Second, perhaps they are actions which are always performed attentively. Third, perhaps they are merely very short-lived sustained actions (Zucchi

foot attentively, drive attentively, work on a project attentively, watch a movie attentively, calculate attentively, count attentively, type attentively, play guitar attentively, garden attentively, and chop garlic attentively. How simple it would be if *exactly* everything that can be done, can be done attentively.

And how elusive would such a counterexample be. If attention foregrounds, then anything that is performed inattentively is in the background: from a phenomenological point of view, as far as that action is concerned, one might as well be asleep! This lends considerable plausibility to the thought that in order to recognize what one is doing from the first-person perspective, it does not suffice to merely be doing it. Chalmers (2003), for example, builds the distinction into his theory: no self-knowledge of consciousness without attention.³⁷ (This doesn't require me to fully engage with the act, of course: attention is a butterfly, shifts can be evanescent and very cheap.) But if not, a kind of action that could only be performed inattentively would be a kind, no instance of which was amenable to first-person inspection. But in what sense then could it be an *action*? How could one intend or desire or try to perform it if one were unable to recognize its instances from the inside? How could one voluntarily regulate its performance? Such a type of occurrence would seem to belong more in the realm of the *biological process*—respiration, circulation of the blood, metabolizing of nutrients—than of the genuine action.

We should, therefore, take a most jaundiced eye toward any purported counterexample. Let us consider some.

First, consider some property (not all properties are “kinds”, in my sense) such as “inattentively (carelessly, cluelessly) spilling milk”. Granting that one cannot attentively inattentively spill milk, the adverbial theory of attention implies that this property is not a kind; but then it is not a candidate counterexample to 1L. (Of course, this is consistent with supposing that *spilling milk* is a kind. But acts of this kind can be performed attentively: if the aim is to distract one's dinner companion, for example.)

Second, consider a property such as *daydreaming*. Is this something one can do attentively? I am inclined to think so. The daydreamer's position is not that of the dazed truck driver, who (like the sleeping person) is doing nothing attentively. Rather, the daydreamer is thinking desultorily about a range of subject-matters that are perhaps inappropriate to the task they are supposed to be performing—still, they are doing so attentively.

If this claim seems odd, this may be due to our tendency to consider a school-

1993). Finally, the cases of concern constitute a distinctive metaphysical category, so we could restrict 1L without cross-cutting any joints.

³⁷Chalmers seems to regard the relevant form of attention as a kind of action, however.

teacher or foreman reprimanding the daydreamer: “pay attention!” What the authority figure wants the daydreamer to do is not to do something or other attentively, however: it is rather to perform the *assigned task* attentively. Or the claim may sound odd because of a tendency to conflate attentiveness with attentive highly focused action: if one’s ship-in-a-bottle building activity is utterly immune to distraction by the telephone, one’s children, and so forth we are inclined to describe it as highly attentive. But even here we can think that one might be subject to short-lived flashes of distraction without abandoning any of the actions constituting building the ship-in-a-bottle: one briefly entertains a nautical memory before shaking it away and returning to one’s tweezers. Attentiveness does not seem to require utter focus: it only requires foregrounding. Since daydreaming can occur in the foreground, it can occur attentively.

Third and finally, one might object that the end of training in meditation is to be able to perform a certain sort of act in which attention is absent: call it “final stage meditation”. Not being an adept, I can merely speculate about the nature of final stage meditation. I discuss three hypotheses, on none of which is final stage meditation a necessarily inattentive kind of action.

One hypothesis is that the condition of the final stage meditator is that of the dazed truck-driver. In that case I would be inclined to think that final stage meditation is not a sort of action, but a sort of state of utter inattention; the point of training in meditation is to be able to cause oneself to enter this state voluntarily (compare: being asleep is not an action, but those who have struggled with insomnia know that one can perform actions that cause one to fall asleep).

A second hypothesis about the nature of final stage meditation is suggested by my yoga teacher’s instruction: “thoughts will come and thoughts will go: let them come and let them go”. Perhaps what one aims to do in final stage meditation is to allow attentiveness to flow freely over a range of undisciplined mental acts: instead of attentively engaging in stringently controlled strings of dialogical reasoning, one attentively thinks about this, then attentively thinks about something unrelated, and so forth. Final stage meditation would thus be not an action, but a dispositional state in which one does little to regulate which acts of thinking are foregrounded.

A third hypothesis about the nature of final stage meditation is that it is a certain distinctive sort of thought or observation: perhaps thinking about a certain subject-matter with penetrating intensity, or observing with penetrating intensity the condition of one’s body (or some further phenomenon observable only by the adept). Here final stage meditation would indeed be a kind of action, to be performed by the adept attentively.

1R From attention to action

A counterexample to 1R would be a kind of occurrence which can have the non-kind property *attentiveness*, but which is not a kind of action.

1R(a) A foregrounded state?

Could a perceptual state, such as pain or seeing (as of) an apple, or a propositional attitude, or a nonmental condition, be attentive? Ordinary language militates against this: it is bad English to speak of “seeing an apple attentively”, “being in pain attentively”, “attentively believing that the solstice is approaching”, or “attentively being on a train”. If these locutions have any sense, what they suggest are properties like *attentively looking at a seen apple*, *attentively mentally probing one’s pain*, *attentively pondering one’s view that the solstice is approaching*, and *attentively observing the events of one’s train ride*.

Perhaps, though, such states could be a certain way which motivates regarding them as attentive for theoretical reasons, even if doing so does not sit well with ordinary language. If such states should be thought of as in some sense *foregrounded*, this could provide the necessary motivation. Should they? A case which might seem to provide this motivation is intense pain. Being in pain is not an action, but there is perhaps a reasonable sense in which an intense pain is “foregrounded”: it crowds out other phenomena from the scope of consciousness.

However, I have been urging, at (K), a picture on which theory sharply separates in its ideology one’s perceptual state from one’s experience, while recording the tight causal and constitutive links between these phenomena. Intense pain foregrounds itself “indirectly”, on my picture, by exerting intense causal force over one’s actions—one is compelled to feel around in one’s intense pains—and by partly constituting those actions: the action is a relational kind with the pain as a relatum. But in the absence of such an action, there is nothing it is like to have the pain; but if consciousness is a contingent feature of pain, pain is not a kind of experience. So if we distinguish the “direct” foregrounding of actions performed attentively from the “indirect” foregrounding of properties which compel and part-constitute those actions, we can maintain that what is characteristic of attentiveness is “direct” foregrounding. According pain the status of attentiveness would conflate phenomena that should be kept separate; accordingly, we have theoretical reasons not to regard intense pains as attentive.

Premiss 2 Attention and consciousness

Premiss 2, recall, was this:

2. A *K* can occur attentively just if a *K* must occur consciously.

A few words on how to understand this are in order. The left-hand side, of course, says that some possible occurrence of kind *K* has the non-kind property *attentiveness*. The right-hand side reports a somewhat different metaphysical structure: it says that any possible occurrence of kind *K* is also of the kind *consciousness*; and so, by Premiss 3, *K* is a determinate of that kind.

Premiss 2 is *prima facie* plausible. Recall William James's famous claim that "focalization, concentration, of consciousness are of [the] essence [of attention]" (James 1890/1981, 403–4).³⁸ UT Place agrees: "in paying attention to something the individual is exercising a measure of control over the vividness or acuteness of his consciousness" (Place 1954, 244). The core idea here is that attentiveness is a "high-grade" form of consciousness. A high-grade form of *X* is a *form* of *X*: something that cannot exist without *X*. So attentiveness entails consciousness. And high-grade contrasts with low-grade: it is easier to be conscious than to be attentive; the conscious phenomena could outnumber the attentive phenomena: so at the level of occurrences, consciousness does not entail attentiveness. Still, simplicity of theory is on the side of consciousness entailing potential attentiveness at the level of kinds: otherwise we would need to justify the existence of a cut in the kinds of experience separating those which can be attentive from those which cannot.

Premiss 2 divides into two conditionals:

- 2L. If a *K* can occur attentively, a *K* must occur consciously;
- 2R. If the *K* must occur consciously, a *K* can occur attentively.

Same drill: I will discuss strategies for absorbing counterexamples first to the left-to-right direction (2L), then to the right-to-left direction (2R).

2L From attention to consciousness

A counterexample to 2L would be a kind of occurrence which can take place attentively but might not take place consciously.

³⁸I have borrowed this citation from Watzl in preparation, who provides further citations of this perspective from James's day and more recently. (Discussion with Sebastian contributed greatly to my initial formulation of the core ideas in this paper.)

2L(a) Unconsciously asleep?

A first sort of concern stems from the discussion in (I) about attentive foregrounding. Actions are long-lasting occurrences; seat of the pants judgements about beginnings and endings of actions support the view that some of them straddle periods of dreamless sleep; and we settled early on (C) on the view that occurrences occupy every time between their beginning and end. So for instance, if one's *studying for a PhD* continues through a period of dreamless sleep, then while one is dreamlessly asleep, one is still performing an act of the kind *studying for a PhD*. By 1L, this act can be performed attentively; so if 2L is correct, it must occur consciously. So it follows that, while one is dreamlessly asleep, one continues to consciously study for a PhD. And it would seem to follow from this that during that period of dreamless sleep, one remains conscious. "Scandal!" says my opponent: "during dreamless sleep, one is *unconscious*".

I bite the bullet: I endorse the prediction of my theory that that during dreamless sleep, one is still conscious—in the sense of 'conscious' that generates the explanatory gap. To soften the bullet, I adopt a twofold strategy. First, I distinguish several senses of 'conscious', explaining the verdict of my view on each of these on the difference between dreamless sleep and waking life. And second, I compare my view with the position that results from abandoning 2L on the present grounds while endorsing Premiss 1 and 2R. I argue that the resulting position is either indistinguishable from or significantly inferior as a theory of the nature of experience to my view; therefore, anyone seeking a theory of the nature of experience must either abandon other aspects of the argument from attention or join me in biting the bullet.

To the **first part** of the strategy. Many discussions of consciousness register that the term 'consciousness' is used to express a number of disparate concepts (Block 1995, Rosenthal 2002). This seems right. Those most relevant to our discussion are the following:

- i. The concept of *subjectivity*—roughly, the existence of a "personal perspective" or "point of view";
- ii. The philosopher's concept of *phenomenality*;³⁹
- iii. A sort of quasi-medical concept, tracking such familiar judgements as that if one is in a dreamless sleep, in a coma, under general anaesthetic, or concussed, one is not conscious; while under normal conditions, during waking life, one is conscious.

³⁹For more on the distinction between subjectivity and phenomenality, see (G), (K), 2L(b), and 2L(c), and footnotes 12 and 14.

Accordingly, the claim that one is conscious during dreamless sleep could mean either: that one has subjectivity during dreamless sleep; or, that one has phenomenality during dreamless sleep; or, that one is conscious in the medical sense during dreamless sleep. I tentatively endorse the first of these, remain somewhat neutral on the second, and reject the third.

On the claim that one has *subjectivity* during dreamless sleep: the notion of subjectivity is not completely easy to get a handle on, and allying it with such notions as “perspective” and “point of view” runs the risk of conflation with notions from projective geometry. This risk can perhaps be diminished through appeal to a notion of *personal* perspective, where persons are understood as something like loci of “I-as-subject” (Shoemaker 1968). These are deep waters. But if this is the right way to think about subjectivity, then the claim that one retains subjectivity during dreamless sleep can be understood as the claim that I do not go out of existence when I fall asleep. I most certainly do not: so I am happy to endorse the claim that I am conscious during dreamless sleep, if this is understood as the claim that I retain “subjectivity” in the sense of existence as a person during dreamless sleep.

On the claim that one has *phenomenality* during dreamless sleep: this is perhaps not obviously false. Perhaps intense pain can cause one to enter a sleep-like state, and then (once this effect runs its course) wake one back up again. A coherent hypothesis about such a situation is that the intense pain persists throughout the story: after all, how else would it have woken one up? One is of course not performing any perceptual actions in regard to the pain, but as we have seen repeatedly, this does not mean that the pain is absent. So I regard it as a coherent possibility that I am conscious during dreamless sleep, if this is understood as the claim that I continue to instantiate phenomenal properties during dreamless sleep.

On the claim that one is conscious in the *medical sense* during dreamless sleep: this is of course false, but I do not think that my theory predicts otherwise. Recall that we have discussed dreamless sleep in (I), above. The verdict there was that in dreamless sleep, while one continues on in performing a range of actions, one performs none of them *attentively*. As we have seen, facts about what one is doing attentively are phenomenologically relevant: what one is doing attentively is foregrounded, everything else is backgrounded. So when one is doing nothing attentively, nothing is foregrounded. Now, the phenomenon of “inattentional blindness”—our loss of any current or retrospective reflective access to aspects of consciousness outside of the scope of attention—is familiar (Mack and Rock 1998). So in my view, in dreamless sleep, one falls victim in effect to a sort of *global inattentional blindness*. I hypothesize that this

sort of state global inattentive blindness⁴⁰ is what the medical notion of consciousness amounts to. A full unpacking of this medical notion would require a thorough treatment of its prudential and moral role, for which there is no room here. Still, a swift preliminary assessment of this role suggests that global inattentive blindness may have many of the desired prudential and moral properties: if one can't attentively probe one's body until the anesthesiologist removes the mask, one needn't worry for the future concerning, care at the time about, or look back regretfully on any pains that may be present in one's body; and a surgeon needn't concern herself with her role in provoking them.⁴¹ So, summing up, I deny that one is conscious in the medical sense—in effect, *attentive* (doing something or other attentively)—during dreamless sleep.

As discussed in (A), our subject-matter is consciousness in the sense which generates an explanatory gap, so we should assess whether the prediction that one is “conscious” during dreamless sleep is true in an explanatory gap-generating sense. As we have just seen, this prediction is true of consciousness as subjectivity; perhaps true of consciousness as phenomenality; and false of consciousness as attentiveness. I am inclined to think that subjectivity and phenomenality both generate explanatory gaps, and that the former is of greater significance for the philosophy of consciousness: I discuss these points further below in 2L(c). I am uncertain whether attentiveness generates an explanatory gap over and above the subjectivity gap and the phenomenality gap. It may be that my account of attentiveness (alluded to in (I)) removes any such further gap; either way, for purposes of this paper, we are not in a position to explore this question any further.

In sum, then, the prediction that one is conscious during dreamless sleep should not trouble us, if ‘conscious’ is understood in the sense most relevant to our concerns.

The **second part** of the strategy for allowing us to live with bullet-biting on consciousness during dreamless sleep is to argue that a *rejectionist* theory of the nature of consciousness—the theory reached by ditching 2L alone to avoid this prediction—is either clearly worse than or no different than my view.

What would the rejectionist position look like? By 2R and 1R, if a *K* must occur consciously, *K* is a kind of action. By 1L, every kind of action can be performed attentively; but if 2L fails, some kinds of action need not occur consciously.

⁴⁰Accompanied perhaps by a disposition to remain in this state of a somewhat greater degree of stability than that of the dazed truck driver.

⁴¹—Why not think that a rock is inattentively conscious? —A rock is not doing anything. —Sure it is, it's holding my papers down. —That is not an action. We don't hold a rock morally responsible for doing this or not doing it.

To make this rejectionist position more vivid, consider an example of a kind of action that could persist through sleep: working on a paper. On the position we are envisaging, it is possible to work on a paper attentively (by 1L). Moreover, any feature which may qualify the way in which one is conscious is a type of action (by 2R and 1R), so perhaps working on a paper may qualify the way in which one is conscious: let us suppose that it may. And finally, it does not follow merely from the fact that working on a paper is a kind of action that at every moment of every course of working on a paper, working on a paper qualifies the way in which one is conscious (by the failure of 2L): let us suppose that it does not.

To make the rejectionist position more vivid still, consider its predictions in the following case:

Arnold is writing a paper (on *On the Plurality of Worlds*) due on Wednesday at 6PM. Arnold begins writing on Tuesday at 11PM, and from time to time gets distracted by his Facebook account, text messages coming through, various friends dropping by his room, trips to the fridge for cokes, and, at around 3AM, a quick round of beer bongs. Finally, as the sun rises at 6AM on Wednesday, Arnold crashes, his paper still incomplete. He is unconscious until 2PM, when he awakens, refreshed and inspired: after a quick tooth-brushing, an advil, and a coke, Arnold is back at his desk for another three hours of uninterrupted hard work. After a half hour of polishing, Arnold emails his excellent paper to his professor, who will assign it an A+ and suggest Arnold develop it as a writing sample for the upcoming year's round of Philosophy PhD program admissions.

Let w be Arnold's particular act of writing, and let c and c' be Arnold's particular stretches of consciousness running from 11PM Wednesday to sleep, and from waking to 6PM Thursday.

Since w outlasts each of c and c' , these three occurrences are jointly distinct. *Consciousness* is clearly not among w 's kinds: we assume that w persists while Arnold is unconscious. Is *writing a paper* among either of c or c' 's kinds? No: if it were, these stretches of consciousness would be acts of writing a paper, counter to our supposition that *writing a paper* is not a kind of consciousness.

So *writing a paper* qualifies the way in which Arnold is conscious (when he is), and *consciousness* qualifies the way in which Arnold writes his paper for its initial and final stretches. However, they do so merely extrinsically: w is conscious for its initial and final stretches, it would seem, because it is accompanied then by c and c' , respectively; and c and c' are qualified by *writing a paper*, it would seem, because each is accompanied by w .⁴²

⁴²If we continue to regard attentiveness as high grade consciousness, we may think of attentiveness

These rejectionist predictions raise a range of questions:

- i. Kinds of action do not determine *consciousness*, but rather merely qualify stretches of consciousness. (Nor does attentiveness status, for reasons discussed in footnote 42). And there is nothing else on the table to do the job. It would seem that there is only one kind of consciousness. That is counterintuitive. We recognize not just that sometimes it is like *something or other* to be one, but also that it is like *something in particular*—something *variable*—to be one.

What is more, the question that motivated the writing of this paper presupposed that there are kinds of consciousness: the game was to state what they were. The rejectionist therefore bows out on the question that initially motivated us.

(Note also that the way we have been understanding ‘*o* occurs consciously’ and thus our defense of Premiss 3 no longer works. For a *K* to occur consciously is no longer for a *K* to be a kind of experience: it is now for *K* to qualify a stretch of consciousness. Accordingly, at the very least the rejectionist must also reject Premiss 3.)

- ii. If consciousness comes in only one kind, what could it be about this kind that makes it be the case that only kinds of action can qualify it? For any occurrence of an arbitrary kind *K* that one is undergoing, why wouldn’t co-occurrence of that occurrence and a stretch of consciousness make *K* qualify the stretch of consciousness? I do not see how to answer this question.

as qualifying *c* and *c'* in the first instance, and *w* only secondarily. But we can’t think of it this way. The picture requires attentiveness to be associated with some further reducible property *T*, such that to *A* attentively is to *A* both *T*ly and consciously. When Arnold takes time out of writing for his round of beer bongs *b*, he is doing beer bongs attentively, and therefore not writing his paper attentively: *b* is attentive, *w* isn’t. But why is this? Our initial thought was that attentiveness statuses qualify, in the first instance, *c*: not either of *b* or *w*. So if *b* is attentive and *w* isn’t, this must be explicable in terms of *c*’s relation to *b* and *w*.

But that won’t work. Throughout *c*, Arnold is doing *something or other* attentively: a bit of writing, a bit of facebooking, a bit more writing, a bit of sipping coke and chatting, and so forth. It starts to seem likely that if attentiveness qualifies *c*, it does so at all moments at which it persists. So it might be simpler just to ditch attentiveness as a property of stretches of consciousness. But this is at odds with the doctrine that attentiveness is high grade consciousness, and is thus perhaps mildly counter-intuitive.

We don’t want to posit attentiveness as an intrinsic property of actions, because it would then be hard to explain why any attentive action is also a conscious action. Perhaps we should then think that whether an action is attentive for a period is a matter of whether during that period it is “perfused” by some stretch of consciousness. But if consciousness is simple, what could make it the case that a stretch of it for a time it perfuses *w*, then for a time perfuses *b*?

The only place left to go would seem to be to think of Aing attentively as a conjunction of being conscious and Aing in a certain special reducible way: the *T* I discuss in the main text.

What is more, why does consciousness need to be qualified in any way in particular? If *c* is a self-standing occurrence, why can't it just be unqualified? Why does it seem incoherent that there might be merely something (but nothing in particular) it's like for one—simpliciter?

More generally, what is it for a stretch of consciousness to be qualified by some other occurrence? Suppose we were to posit a relation of “perfusal”, such that this happens just when the stretch “perfuses” the occurrence (or has the occurrence “as its object”). This would be to explain the dim with the dark. What is the nature of the relation of perfusal? I don't know.

These questions may be setting off metaontological alarm bells: is there something substantive to distinguish between my view and rejectionism? Or are they rather merely terminological variants of one another? In the former case, rejectionism is weighted down with heavy metaphysical and theoretical baggage: it can't answer the real question, asserts that consciousness only comes in one variety, posits a mysterious relation of perfusal. In the latter case, rejectionism is not a genuine competitor to mine: rather, it is merely a way of putting my view in a way that raises confusions in light of the systematic metaphysical choices we made in (A)–(C).

So rejectionism does not offer an adequate alternative to my theory. The bullet of endorsing the claim that one is conscious during dreamless sleep must, somehow or other, be bit.

2L(b) Attention suffices for consciousness?

The second sort of concern would stem from a view on which consciousness is *identical* to attentiveness.⁴³ Since there is considerable initial plausibility to the claim that there are strictly *fewer* attentive phenomena than there are conscious phenomena, the friend of this doctrine has the options of *expanding* the range of attentive phenomena to fit the initially plausible range of conscious phenomena, and *contracting* the range of conscious phenomena to fit the initially plausible range of attentive phenomena. The expansionist approach poses no threat to either direction of the premiss: must implies can. But the contractionist approach threatens the left-to-right direction. If Premiss 1 is correct, then *chopping garlic* can occur attentively (and moreover this is initially plausible). But it is initially plausible that it need not occur attentively. On the contractionist approach, when it does not, nor does it occur consciously. Accordingly,

⁴³As we have seen in (L), Ryle endorsed this view; in contemporary literature, Prinz (2005) has taken up the banner.

chopping garlic is a kind of occurrence which can take place attentively but might not take place consciously, contra 2L.

Why endorse the doctrine that consciousness is identical to attention? The best-developed case is given by Prinz (2005): roughly, the search for the neural correlate of consciousness will run most smoothly if this *identity doctrine* is assumed. My strategy in reply has two components: first, I will lay out in detail Prinz's argument as I understand it and then criticize the reconstructed argument; second, I will explain how the search for the neural correlate might run if my view is correct.

I begin the **first component of the strategy** by expanding on my understanding of Prinz's argument.

- i. Suppose that our "direct" epistemic access to phenomenal consciousness is equal to the range of attention.
- ii. Then, if we accept the identity doctrine, we would have a clear epistemic criterion for the presence or absence of phenomenal consciousness: if one is in a position to directly access phenomenal consciousness, it is present; otherwise, it is absent.

By contrast, if we reject the identity doctrine and accept that phenomenal consciousness is present outside of the range of attention, the extent of phenomenal consciousness is open to debate, and we will lack a clear epistemic criterion for the presence or absence of phenomenal consciousness.

- iii. But, unless we have a clear epistemic criterion for the presence or absence of phenomenal consciousness, we will be in a position in which we cannot run through the neural properties to assess which of them lines up in presence or absence with phenomenal consciousness: we will be unable to find the neural correlate of phenomenal consciousness.
- iv. And finally, finding this correlate is important, because phenomenal consciousness is consciousness, and it is intrinsically important to know the neural correlate of consciousness.

The argument is questionable at each step.

Against (i), it is arguable that the "transparency of experience" shows that we almost never have epistemic access to phenomenal consciousness: the object of attention is almost always in the external world.⁴⁴

⁴⁴For a proposal concerning how to extract such access from an outer-directed gaze, see my Hellie 2010.

Against (ii), the identity doctrine is in no better a position to ground a clear epistemic criterion of the presence or absence of phenomenal consciousness than is any other clearly stated claim about the extent of consciousness. What is more, as we have seen, it is plausible that attention is a high grade of consciousness; if this is moreover *analytic*, the thesis changes the subject and therefore *cannot* be a criterion for the presence or absence of consciousness.

Against (iii), even if we have not decided on a clear criterion, we might have a range of plausible criteria. We might order these criteria along a pair of dimensions: superiority on philosophical grounds; predicting the presence or absence of phenomenal consciousness in a way that lines up better with some or other neural property. If there is a criterion which is dominant along both dimensions, or which is in the aggregate superior to all others, we select the neural property assigned to it as the neural correlate of consciousness (perhaps revising our initial opinion about the extent of consciousness as a result of this). The search will only fail if two or more criteria are in the aggregate “tied” with one another. But why suppose at this very early stage that this pessimistic outcome looms so large that we must reconfigure our theory of consciousness around the implausible identity doctrine in order to ward it off? Indeed, doing so seems *counterproductive*, in that doing so might result in our missing out on a criterion which is in the aggregate superior to the identity doctrine.

Finally, against (iv), as we have seen, consciousness is not phenomenality. But if not, the intrinsic importance of knowing the neural correlate of consciousness is compatible with the relative unimportance of knowing the neural correlate of phenomenal consciousness.

The **second component of the strategy** briefly explains how the pursuit of the neural correlates of consciousness and phenomenality would run if $\mathbf{E}=\mathbf{A}$ is correct. The neural correlate of *consciousness* would turn out to be the neural correlate of *action*. This would put us in the methodologically advantageous position of being able to appeal to (initial and theoretically revised) judgements about the presence or absence of action in establishing a criterion for the presence or absence of consciousness. These judgements have a rich range of connections with legal and moral theory, as well as in ordinary narrative practice, to buttress our assessment of any such criterion.

$\mathbf{E}=\mathbf{A}$ makes no independent predictions about the presence or absence of phenomenality. However, the question of the neural correlate of phenomenality is diminished in significance. Let me draw an analogy. The question of the nature of consciousness is widely regarded as a “hard problem”, perhaps as one of the last frontiers of scientific understanding. Moreover, pushing back the frontier has the potential to be immensely

rewarding, in light of the ties of consciousness to the self and to value. There is also a strong case to be made that the nature of *color* and other “qualities” should also be regarded as a hard problem (Byrne 2006). Still, the ethical significance of any degree of scientific progress on the nature of qualities is rather limited. I am inclined to class the phenomenal together with other qualities, as being of limited ethical significance.

So **E=A** can be employed as a framework principle in the role Prinz imagines for the identity doctrine: as nailing down a key piece of conceptual structure that will enable us to make headway on the important hard problem.

2L(c) Attentive zombies?

Let us turn to a final source of concern about 2L, stemming from the assumed possibility of zombies. With that assumption in hand, we argue as follows: gardening is a kind of action; so by 1L, one can garden attentively; but a zombie does not do anything consciously; and some zombies garden. Accordingly, gardening is a kind of action instances of which can be performed attentively, and which might yet not be performed consciously—against 2L.

My response is to deny the last premiss in the argument, namely that some zombies garden. Instead, I think that while some zombies “exhibit gardening behavior” in the sense of (E), this behavior does not suffice for *gardening* in the strict sense. Substantiating this response is a major task, but I sketch the following twofold strategy of defense: first, I argue that zombies do not act; second, I develop a theory of fundamental first-person concepts of kinds of action.

First, the claim that zombies do not act is *prima facie* plausible. As Siewert (1998) notes, zombification plays havoc with our evaluative judgements: in particular, a zombie’s situation is as bad as death. I would extend this observation to the following: we are reluctant to provide a zombie with moral credit or demerit for any effects brought about by its behavior. But, in line with the link between action and moral responsibility discussed in (E), this suggests that we do not regard those behaviors as stemming from or embodying actions.

Second, where there are zombies, there are concepts of the properties which we have and they lack which are *conceptually fundamental*—in the sense that these concepts are not entailed a priori by concepts of the properties we share (Chalmers 2002). So if I am right that zombies do not garden, our concept of gardening is not entailed a priori by any merely causal or physiological concept of gardening behavior.

As we saw in (L), Ryle endorsed this doctrine that action concepts are fundamental. But how could it be true? My answer to this question recapitulates Chalmers’s

explanation of the source of irreducible concepts of phenomenal properties (Chalmers 2003). The core of this explanation is the posit of a primitive relation of *acquaintance* in which one stands to any variety of consciousness instanced in one's current experience (as well as any such variety one can recall or imagine). Armed with such a relation, a subject in the right frame of mind⁴⁵ is in a position to deploy a fundamental concept of that variety of consciousness. This explanation does not prejudice what the varieties of consciousness are;⁴⁶ accordingly, I am free to help myself to it in explaining how action concepts could be fundamental. Accordingly, I claim that our ordinary concepts of kinds of action are in the main based in this way in acquaintance with our own actions.

I should highlight a respect in which my view of the relation between kinds of experience and other concepts diverges from Chalmers's. According to Chalmers, acquaintance-based fundamental concepts of phenomenal properties are not entailed by functional concepts: in regard to acquaintance-based concepts of kinds of action, I say the same. According to Chalmers also, no functional concepts are entailed by acquaintance-based fundamental concepts of phenomenal properties.⁴⁷ Chalmers can conceive of an *invert*: a subject physically and functionally just like me when I see a red thing but whose experience is "phenomenal-green". By contrast, I deny that "action-inverts" are conceivable: someone exhibiting gardening behavior can be determined a priori not to be chopping garlic (when *chopping garlic* is conceived of under the ordinary acquaintance-based concept: note that my claim is not about *seeming* to chop garlic: it concerns *really* chopping garlic; this is ruled out a priori by my exhibiting gardening behavior). So in my view, canonical functional concepts of exhibition of behavior are entailed a priori by acquaintance-based concepts of action.

Because of this entailment, we can often allow our concept of exhibiting gardening behavior to stand in for our acquaintance-based concept of gardening (just as, for many purposes, the nomological necessitation of a certain functional property by a given phenomenal property permits us to often allow our concept of the former to stand in for our concept of the latter).

Recall that we wished for an explanation, compatible with its falsity, of the allure of

⁴⁵What's that? Something too complicated to go into here: something like Chalmers's notion of the act of attention to consciousness against a certain cognitive background (Chalmers 2003, 3.2.2).

⁴⁶Modulo the point that Chalmers's discussion, without further embellishments, seems strongly to militate in favor of a "narrow" understanding of the kinds of experience. My Hellie 2010 takes steps in the direction of extending Chalmers's approach to externalist theories of the kinds of experience.

⁴⁷Or at least there is considerable play: perhaps our acquaintance-based concept of a phenomenal sound entails a priori the absence (from the relevant region of reality) of any functional property that is actually correlated with a phenomenal visual shape.

the doctrine that zombies garden. I suggest now that the source of this allure is the ease with which we in ordinary cases can swap between functional and acquaintance-based concepts. The ordinary practice of being sloppy about which concept applies *strictly speaking* is fine for ordinary “low-resolution” tasks; after all, in ordinary tasks, they both apply. But when we carry this practice of sloppiness forward to the current “high-resolution” task of assessing whether zombies are agents, it makes false predictions.

At the same time, if we are careful to be strict with these concepts, we can see that the functional concept applies to the zombie, but the concept of action does not: the association of action with moral evaluation as discussed in (E) can be used to prise the concepts apart.

Summing up: the argument against 2L from the possibility of zombies was based on the assumption that it is strictly correct to apply the concept of gardening, considered *as an action* and not as a mere behavior, to a zombie. I have sketched an argument *that* this is not strictly correct, and a theory of *why* it is not.

2R From consciousness to attention

We turn finally to 2R. A counterexample to this principle would be a kind of occurrence which must take place consciously but can’t take place attentively.

2R(a) “Perceptual experience”?

This is the point in the argument threatened by “perceptual experience”: when one has pain, runs the objection, one has an “experience” of the kind *pain*; hence, by Premiss 3, all instances of pain are conscious; but having pain is not an action, so by Premiss 1, one cannot have pain attentively; so pain is a kind of occurrence which must take place consciously but can’t take place attentively, against 2R.

The reply here appeals to the discussion in (K): I argue there, recall, that instances of pain are not experiences, and that *pain* is not a kind of experience.⁴⁸

⁴⁸A dialectical question: prefatory to the second part of the strategy with which I argued that pain is not a kind of experience, I assumed that the kinds of action are among the kinds of experience. Why was this OK? In the present context, the assumption is safe because we are assuming that 1L and 2L are in place. This assumption, in turn, is OK because if 2R can be defeated by the present strategy only if 1L or 2L has already been defeated, it is dialectically pointless for my opponent to take this stand against 2R.

References

- Anscombe, Gertrude Elizabeth Margaret, 1963. *Intention*. Cambridge, MA: Harvard University Press.
- Block, Ned, 1995. 'On a Confusion About the Function of Consciousness'. *Behavioral and Brain Sciences*, 18:227–47.
- Block, Ned, 2001. 'Paradox and Cross Purposes in Recent Work on Consciousness'. *Cognition*, 79:197–219.
- Byrne, Alex, 2001. 'Intentionalism Defended'. *The Philosophical Review*, 110:199–240.
- Byrne, Alex, 2006. 'Color and the Mind–Body Problem'. *Dialectica*, 60:223–44.
- Chalmers, David J., 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Chalmers, David J., 2002. 'Consciousness and Its Place in Nature'. In David J. Chalmers, editor, *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- Chalmers, David J., 2003. 'The Content and Epistemology of Phenomenal Belief'. In Quentin Smith and Alexander Jokic, editors, *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press.
- Chalmers, David J., 2004. 'The Representational Character of Experience'. In Brian Leiter, editor, *The Future for Philosophy*. Oxford: Oxford University Press.
- Chalmers, David J., 2006. 'Perception and the Fall from Eden'. In Tamar Szabó Gendler and John Hawthorne, editors, *Perceptual Experience*. Oxford: Oxford University Press.
- Davidson, Donald, 1973. 'Freedom to Act'. In Ted Honderich, editor, *Essays on Freedom of Action*. London: Routledge. Reprinted in Davidson 1980.
- Davidson, Donald, 1978. 'Intending'. In Yirmiyahu Yovel, editor, *Philosophy of History and Action*. Dordrecht: D. Reidel. Reprinted in Davidson 1980.
- Davidson, Donald, 1980. *Essays on Actions and Events*. Oxford: Oxford University Press.
- Fine, Kit, 1994. 'Essence and modality'. *Philosophical Perspectives*, 8:1–16.
- Grice, Herbert Paul, 1962. 'Some Remarks About the Senses'. In R. J. Butler, editor, *Analytical Philosophy*, first series. Oxford: Blackwell.

- Harman, Gilbert, 1990. 'The Intrinsic Quality of Experience'. In James Tomberlin, editor, *Action Theory and the Philosophy of Mind*, volume 4 of *Philosophical Perspectives*, 31–52. Atascadero: Ridgeview.
- Hellie, Benj, 2010. 'An Externalist's Guide to Inner Experience'. In Bence Nanay, editor, *Perceiving the World*. Oxford: Oxford University Press.
- Hellie, Benj, in preparation. 'Attention and Gunk'. MS, University of Toronto.
- Horgan, Terence and John Tienson, 2002. 'The Intentionality of Phenomenology and the Phenomenology of Intentionality'. In David J. Chalmers, editor, *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- James, William, 1890/1981. *The Principles of Psychology*. Cambridge, MA: Harvard University Press.
- Kriegel, Uriah, forthcoming. 'Naturalizing Subjective Character'. To appear in *Philosophy and Phenomenological Research*.
- Levine, Joseph, 1983. 'Materialism and Qualia: The Explanatory Gap'. *Pacific Philosophical Quarterly*, 64:354–61.
- Levine, Joseph, 2006. 'Conscious Awareness and Self-Representation'. In Uriah Kriegel and Kenneth Williford, editors, *Self-Representational Approaches to Consciousness*. Cambridge, MA: The MIT Press.
- Lewis, David, 1966. 'An Argument for the Identity Theory'. *Journal of Philosophy*, 63:17–25. Reprinted in Lewis 1983b.
- Lewis, David, 1972. 'Psychophysical and Theoretical Identifications'. *Australasian Journal of Philosophy*, 50:249–58. Reprinted in Lewis 1999.
- Lewis, David, 1980. 'Veridical Hallucination and Prosthetic Vision'. *Australasian Journal of Philosophy*, 58:239–49. Reprinted, with postscripts, in Lewis 1986.
- Lewis, David, 1983a. 'New Work for a Theory of Universals'. *Australasian Journal of Philosophy*, 61:343–377. Reprinted in Lewis 1999.
- Lewis, David, 1983b. *Philosophical Papers*, volume I. Oxford: Oxford University Press.
- Lewis, David, 1986. *Philosophical Papers*, volume II. Oxford: Oxford University Press.
- Lewis, David, 1988. 'What Experience Teaches'. *Proceedings of the Russellian Society*, 29–57. Reprinted in Lewis 1999.
- Lewis, David, 1994. 'Reduction of Mind'. In Samuel Guttenplan, editor, *Companion to the Philosophy of Mind*. Oxford: Blackwell. Reprinted in Lewis 1999.

- Lewis, David, 1995. 'Should a Materialist Believe in Qualia?' *Australasian Journal of Philosophy*, 73:140–4.
- Lewis, David, 1996. 'Elusive Knowledge'. *Australasian Journal of Philosophy*, 74:549–67. Reprinted in Lewis 1999.
- Lewis, David, 1997. 'Naming the Colours'. *Australasian Journal of Philosophy*, 75:325–42. Reprinted in Lewis 1999.
- Lewis, David, 1999. *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Lycan, William G., 2001. 'The Case for Phenomenal Externalism'. In James Tomberlin, editor, *Metaphysics*, volume 15 of *Philosophical Perspectives*, 17–36. Oxford: Blackwell.
- Mack, John and Irwin Rock, 1998. *Inattentional Blindness*. Cambridge, MA: The MIT Press.
- Martin, Michael G. F., 2004. 'The Limits of Self-Awareness: Disjunctivism and Indiscriminability'. *Philosophical Studies*, 120:37–89.
- Martin, Michael G. F., 2006. 'On Being Alienated'. In Tamar Szabó Gendler and John Hawthorne, editors, *Perceptual Experience*. Oxford: Oxford University Press.
- Mele, Alfred R., 1992. *The Springs of Action*. Oxford: Oxford University Press.
- Mole, Chris, 2009. *Attention Is Cognitive Unison*. Oxford: Oxford University Press.
- Nagel, Thomas, 1974. 'What Is It Like to Be a Bat?' *The Philosophical Review*, 83:435–50.
- Noë, Alva, Luis Pessoa, and Evan Thompson, 2000. 'Beyond the Grand Illusion: What Change Blindness Really Teaches Us About Vision'. *Visual Cognition*, 7:93–106.
- Parsons, Terence, 1990. *Events in the Semantics of English: A Study in Subatomic Semantics*. Cambridge, MA: The MIT Press.
- Place, Ullin T., 1954. 'The Concept of Heed'. *British Journal of Psychology*, 45:243–55.
- Place, Ullin T., 1956. 'Is Consciousness a Brain Process?' *British Journal of Psychology*, 47:44–50.
- Prinz, Jesse, 2005. 'A Neurofunctional Theory of Consciousness'. In Andrew Brook and Kathleen Akins, editors, *Philosophy and Neuroscience*. Cambridge: Cambridge University Press.
- Pryor, James, 2000. 'The Skeptic and the Dogmatist'. *Nous*, 34:517–49.

- Rosenthal, David M., editor, 1991. *The Nature of Mind*. Oxford: Oxford University Press.
- Rosenthal, David M., 2002. 'Explaining Consciousness'. In David J. Chalmers, editor, *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- Rothstein, Susan, 2004. *Structuring Events*. Oxford: Blackwell.
- Ryle, Gilbert, 1949. *The Concept of Mind*. London: Hutchinson and Co.
- Setiya, Kieran, 2007. *Reasons Without Rationalism*. Princeton: Princeton University Press.
- Shoemaker, Sydney, 1968. 'Self-reference and self-awareness'. *Journal of Philosophy*, 65:555–67.
- Shoemaker, Sydney, 1994. 'Phenomenal Character'. *Nous*, 28:21–38.
- Siegel, Susanna, 2005. 'The Contents of Perception'. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*.
- Siewert, Charles, 1998. *The Significance of Consciousness*. Princeton, NJ: Princeton University Press.
- Simons, Daniel J. and Christopher F. Chabris, 1999. 'Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events'. *Perception*, 28:1059–74.
- Smart, John Jamieson Carswell, 1959. 'Sensations and Brain Processes'. *The Philosophical Review*, 68:141–56.
- Swartz, Robert J., editor, 1965. *Perceiving, Sensing, and Knowing*. Berkeley: University of California Press.
- Sytsma, Justin, 2009. 'Phenomenological Obviousness and the New Science of Consciousness'. *Philosophy of Science*, 76:***–***.
- Thau, Michael, 2002. *Consciousness and Cognition*. Oxford: Oxford University Press.
- Thompson, Brad, 2009. 'Senses for Senses'. *Australasian Journal of Philosophy*, 87:99–117.
- Thompson, Michael, 2008. 'Naive Action Theory'. In *Life and Action*. Cambridge, MA: Harvard University Press.
- Tye, Michael, 2000. *Consciousness, Color, and Content*. Cambridge, MA: The MIT Press.
- Watzl, Sebastian, in preparation. 'Awareness of Attention and the Structure of Consciousness'. MS, Columbia University.

White, Alan R., 1964. *Attention*. Oxford: Basil Blackwell.

Williamson, Timothy, 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Zucchi, Alessandro, 1993. *The Language of Propositions and Events*. Studies in Linguistics and Philosophy. Dordrecht: Kluwer.