

Attention to Experience
(Talk presented at UCR)
Charles Siewert
January 11, 2004

1

Consciousness, some say, is “transparent” to first-person acts of attention; experience is “diaphanous.” You might take this to mean that when you try to focus attention on consciousness or experience itself, you find you cannot: attention “passes through” straight to the object you are conscious of—the object you experience. But the metaphor has been elaborated in various ways. Here are two recent invocations of transparency by philosophers who think it bears on the nature of experience. Gilbert Harman (1990) writes:

Look at a tree and try to turn your attention to the intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree.

And here’s one way Michael Tye (1995) puts the point:

Try to focus your attention on some intrinsic feature of the experience that distinguishes it from other experience, something other than what it is an experience of. The task seems impossible: one’s awareness seems always to slip through the experience to blueness and squareness, as instanced together in an external object. In turning one’s mind inward to attend to the experience, one seems to end up concentrating on what is outside again, on external features or properties.

Whatever the truth or import of transparency claims, they are generally offered on the basis of first-person reflection, or introspection, broadly construed. What I mean here is this. We can make first-person judgments attributing experience to ourselves. And the warrant we have for making them differs in kind from that which other people ordinarily have or would need for the corresponding

judgments about us. When Harman and Tye invite us to try to attend to our own experience, they are pretty clearly inviting us to confirm their claims about what we will find—and what we won't find—on the basis of that distinctive kind of warrant (whatever it is) enjoyed by first-person judgments about experience.

This introspective “transparency” finding has been put at the service of far-reaching philosophical views about consciousness. Tye thinks it provides an argument for what he calls “strong representationalism.” (2002, p.45) Roughly, this is the thesis that the phenomenal character of experience is to be explained by identifying it as a species of mental representation. We can explain why it looks or feels to us as it does, and why its looking or feeling to us is phenomenally conscious, if we maintain that (as Tye puts it) “phenomenal character is one and the same as representational content that meets certain further conditions.” (Tye, 2000, p.45) We might, borrowing David Chalmers’ (2004) terms, describe this as a *reductive representationalist* account of phenomenal consciousness—“R” for short. Tye’s claim is that introspection supports this sort of theory by making evident to us the transparency of consciousness.

Is this right? In just what sense, if any, does first-person reflection show us that consciousness is transparent? And what, if anything, does this reveal about its nature? These questions are interesting partly because we may want to determine whether introspection really does somehow speak for or against the kind of representationalist theory of consciousness Tye and others advocate. But I also think dealing with these questions will put us in a better position to

understand the relationship among experience, attention, and introspection. And that, I hope, will help us be clearer about what introspection itself is.

Maybe we will fasten most quickly on some basic issues, if we look at an argument that goes from transparency to Reductive Representationalism. For this I will focus on Tye's exposition, since it seems one of the more explicit and sustained discussions. Tye's core argument seems to be this. We are invited to recognize that the *representational content* of experience is not rightly regarded as a *feature of* experience. And we note that, according to R, the phenomenal character of an experience just is a certain kind of representational content it has. Therefore, R tells us that *phenomenal character of experience* is not a feature of experience. But then, it follows that phenomenal character of experience is not a feature of experience to which we can attend or have direct access. (Tye 2000, pp.48-9) Now this latter point is *also* just what we discover when we try, in introspection, to attend to the features of our experience. According to the passage cited earlier, we find we cannot: we can attend only to external objects, their features, properties or qualities. As Tye has more recently put the matter, "If you are attending to how things look to you...the only objects of which you are aware are the external ones making up the scene before your eyes." (Tye 2000, pp.46-7) And: "The qualities of which you are directly aware in focusing on the scene before your eyes and how things look are not qualities of your visual experience." (Tye 2000, p.46) Also: "Your awareness of phenomenal character is not the direct awareness of a quality of your experience. Relatedly, the phenomenal character itself is not a quality of your experience to which you have

direct access.” (Tye 2000, p.47) And finally: “The qualities to which we have direct access are the external ones, the qualities that, if they are qualities of anything, are qualities of external things.” (Tye 2000, p. 51) The upshot is this. Introspection provides us with a certain finding regarding what we *cannot* attend to or be directly aware of (qualities of experience) and what we *can* (external objects and qualities). This finding can be explained, provided that one adopts R. For the finding in question is a consequence of that theory’s account of phenomenal character. And so, this provides us with an important reason to accept that account.

Tye’s exposition of this “argument from transparency” is more involved than this brief outline indicates. But I will take this statement of it as adequate to identify the aspect of it I most want to examine: the transparency claim that figures in it. We might sum this up as follows:

T1 The phenomenal character of one’s experience is not a feature of experience to which one can attend, or of which one can be directly aware.

According to the argument from transparency, this is an introspective datum best explained by a certain representationalist theory of phenomenal character.¹ However, it seems to me that there are a number of ideas tangled up in T1, not clearly distinguished by fans of transparency. T1, I take it, entails some or all of the following:

- (i) One cannot attend to the phenomenal character of one’s experience.
- (ii) One cannot be directly aware of the phenomenal character of one’s experience.
- (iii) The phenomenal character of one’s experience is not a feature of the experience.

Now my view is that first, judging by introspection, one can, in some sense at least, attend to the phenomenal character of experience. Second, neither (ii) nor (iii) finds any *support* from introspection. I also find it pretty unclear that we have good reason from any quarter to endorse (ii) and (iii). Nor is it clear to me that any of these three would, if true, receive a satisfactory explanation from or provide reason to accept reductive representationalism. But, in any case, if I am right, an introspection-based argument from transparency to R does not even get off the ground, since first-person reflection simply does not warrant T1.

To argue for my view, I am going to propose one way of defending the idea that we can attend to our own experience, based on considering cases of visual illusion, which will yield my conclusions about (i). I will then say why, given this, introspection does not speak in favor of (ii) or (iii). Next I will offer two more examples as evidence for my claim about attention, one of which appeals to what I might call “finely discriminated sensibles,” while the other concerns what I label “judgments of pictorial resemblance.” These not only furnish additional reasons to reject the argument from transparency. They also, I will maintain, show that introspection provides no conception of the intentionality of experience that would enable one to explain phenomenal consciousness as a form of representation. So, introspection it turns out, is really no friend to reductive representationalism. I will finish up with some remarks about the sense in which I think consciousness actually *is* transparent, and also with a few thoughts about what all this might indicate about the nature of introspection.

2.

I will start with something I think is relatively uncontroversial. You can, in some sense, attend to how it appears to you—for example, to how it looks, sounds, smells, tastes, feels to you. Now, I would also say, it looks, sounds, etc., to you in these ways just in case you have certain phenomenally conscious experiences: visual experience, aural experience, and so on. Further, at least some differences in ways of looking, sounding, etc. to you, in this sense constitute *phenomenal* differences. For example, at the very least, differences in the ways colors or shapes look to you are phenomenal differences. And, where there are phenomenal differences in how it appears to you, you have experience that differs in *phenomenal character*. I think that many of those who find use for the phrases ‘phenomenal consciousness’ and ‘phenomenal character’ would grant this much at least.

Now, here is something perhaps a little more controversial. If, in the sense just invoked, it looks or otherwise appears some way to me, its so appearing to me is a feature I have: a “phenomenal feature,” in my terminology. Furthermore, when, for example, I attend to how it looks to me, part of what I attend to is what phenomenal feature I have. That is, I do not somehow attend to the way it looks to me, while leaving its *looking* that way to me out of it. When something looks blue or square to me—to take Tye’s example—and I attend to how it looks to me, I do not somehow attend just to blueness or squareness, without attending to its *looking* blue or square to me. Its *looking* to me as it does “falls within the scope of my attention,” just as much as, and together with, the figure itself and its

blueness and squareness. The figure, its properties, and its appearing to me, all come together as a package. I may only *look at* the blue square—I certainly don't look at my visual experience of it. However, I can, while looking at the blue square, attend to its looking to me as it does.

That, at any rate, is how I would describe matters. But can I say anything to persuade you it is the right way to describe them? Let's consider a case of visual illusion. Suppose you look at a pair of circles A and B. You attend to how they look to you. In doing this, you attend to the circles, of course, but that is not all. You may attend to their possession of certain qualities: their circularity for example. And presumably you can attend to *the size* they appear to have relative to one another. Can we interpret this attention in a way that excludes their *looking* to you as they do from what you attend to? We might say, for example, that you attend to *A's being larger than B*, provided that A not only *looks* to you, but actually *is* larger than B. However, this condition will not be met in a case of illusion. Suppose A looks bigger to you than B, but really isn't—as for example, in the famous Titchner (or Ebbinghaus) illusion. In that case, we should not say that what you attend to is A's actually being bigger than B. For A isn't in fact bigger than B. How then should we conceive of what you attend to, when you attend to how it looks to you in such a case? This way seems readily available: A looks bigger to you than B. *That* is what you're attending to. Or, to nominalize: you attend to *A's looking bigger to you than B*. Or, if you like, you attend to *the fact or state of affairs* that A looks bigger to you than B. To grant any of these ways of putting the matter is to grant my claim about the scope of attention.

If we think this way of describing things violates some insight into the transparency of experience, what alternative way of construing what we attend to should be endorsed? Tye often speaks in ways that suggest that certain (“external”) *qualities* themselves are to be construed as objects of attention. Someone might, in line with this, say that what you attend to here is (not A’s looking bigger to you than B), but simply: the circle A, the circle B, and the relational quality: “being larger than.” But if this is what transparency demands we say, the truth of the claim no longer seems introspectively evident. Partly this is because it is not at all clear to me what could even be meant by attending to a quality “being larger than” *all on its own*. Now I do think there is a sense in which we can attend to a quality or a relation, in a context where we *take it into consideration*, in anticipation of some claim about it, or some request to conceive of a circumstance involving it. So someone might say, “Consider the relation, being larger than. It is transitive relation.” And one might add, “Now suppose there are two circles, one larger than the other”. Here I think we can rightly speak of attending to the “larger than” relation.

However, in *that* sense surely I can also attend to its looking or feeling to me some way, when experiencing the Titchner illusion. So here I don’t find a viable way to conceive of what I can attend to, such that first-person reflection shows I cannot, in this sense, attend to its looking to me as it does.

At this point maybe someone will suggest that, just as one can, in a sense, “see” a dagger that isn’t there when hallucinating, so there is a sense in which one can attend to A’s being bigger than B, even when it isn’t. I am not sure this is

right as a comment about the ordinary sense of ‘attend.’ But perhaps we could at least introduce such a “non-factive” use of the word, and then propose that, judging by introspection, in the illusion case we do not attend to A’s *looking* to us bigger than B at all—we only “non-factively” attend to A’s *being* bigger than B. I myself don’t find this introspectively evident. Also, there is introspective reason to reject the suggestion. It would seem that, when we are asked merely to *suppose* A is bigger than B, we are also, in this (non-factive) sense, attending to A’s being bigger than B. And if we are then asked to suppose something different—to entertain the supposition that A is the same size as B—and we’re paying attention to what we’re doing, then we attend to a *change* in what we’re supposing. But now, imagine you are shown the circles illusion on a video screen, and you attend to a change in how big the circles look relative to each other, as the image of circle A is reduced. (Perhaps you do this because you are asked to report on when A no longer looks bigger to you than (but now the same size as) B.) Now notice: if it were true that, as far as introspection could tell, you could only attend to A’s being some size relative to B, then you could not introspectively tell the difference between attending to a change in what you *suppose* to be so, regarding the circles relative size, and attending to a change in how the circles *visually appear* to you, in respect of size. For, introspectively, *the contents of attention* would be the same with either task. In either case, you would discern only a switch from attending to A’s being larger than B, to attending to A’s being the same size as B.

But now ask yourself, does introspection tell you that there is no difference between attending to a change in how it looks to you, and attending to a change in corresponding suppositions you make? Would you say that, though first-person reflection allows you to tell the difference between a change in how the circles look to you, and a change in what you suppose about them, when it does this, it also tells you that what you are *attending to* in the two cases is just the same? That it seems, is what you need to say, if you hold that introspection leaves its looking to you as it does outside the scope of your attention. Since I do not find introspection tells me any such thing, I will stick with my proposed way of describing what's going on, as far as first-person reflection reveals, when I attend to how it looks to me in the Titchner circle case. And, on that proposal, introspection does not exclude its looking to me as it does from what I can attend to.

Suppose that, in light of this, you concede my “scope of attention” claim. How does this bear on the argument from transparency? If you agree that you can attend to its looking to you as it does, then, granting me my terminology, this means: you can attend to your having a certain phenomenal feature. And, since your visual experience just is your having that feature, you can attend to your visual experience. But now, can you attend to the phenomenal character of experience, as T1 suggests you cannot? I would say that to have a certain phenomenal feature—for it appears to you some way—is to have an experience with a certain phenomenal character. That, at any rate, is how I understand talk of “phenomenal character.” But then, since you can attend to what phenomenal

features you have, you can attend to the phenomenal character of your experience. If T1 entails (i) above—if it says that you cannot attend to the phenomenal character of your experience, then, to judge by first-person reflection, T1 is false.

Perhaps then we should focus on yet a different aspect of T1—thesis (ii) above the idea that we cannot be “directly” aware of—the phenomenal character of our experiences. Does first-person reflection perhaps still support *that* claim? Well, if I am *attending to* something, and I am *thinking of* it or *about* it, then, in some sense surely, I can also rightly be said to be *aware of* it. So, in that sense at least, I am indeed sometimes aware of the phenomenal character of my experience. Perhaps the suggestion will be that introspection reveals this sort of awareness is not “direct,” but “indirect.” However, this will be so only if first-person reflection shows that I attend to, and think of, A’s looking bigger to me than B, only by means of being aware of something else, entirely distinct from A’s looking bigger to me than B. It’s hard to see what that would be. At any rate, the burden is certainly on whoever wants to say introspection shows that there is this something else we are aware of that plays this role.

Maybe one will say: “OK, sure, you can attend to the phenomenal character of experience. And if that’s true, then T1 is also mistaken if it implies that introspection shows our awareness of it is at most is indirect. But maybe T1 is still right about thesis (iii) above: the phenomenal character of your experience is not a feature of the experience.” I don’t think I have any particular interest in denying this claim, but I also don’t quite see what should convince us it’s true. I

don't see what should prevent us from saying that an experience can have the feature: *being an instance (or a state, or an episode) of its appearing some way to someone*. And to be such an instance (state, what-have-you) is to have a certain phenomenal character. Now maybe there is something wrong with talking of experiences as states, episodes, or instances of features. But I don't see how we're supposed to conclude it's mistaken, based on some introspective transparency insight.

What I conclude now is that if—as visual illusion seems to show—first person reflection supports my scope of attention claim, then introspective arguments for reductive representationalism from the transparency claim T1 will fail. For if, to go by introspection, we can attend to things appearing to us as they do, then, by those lights, we can indeed attend to the phenomenal character of our experience. Once we grant that, we should grant that we can be aware of the phenomenal character of experience in a way we at least have no introspective grounds for saying is *indirect*. (Nor is it clear we have *any* good grounds for saying that.) And the same applies to this idea that the phenomenal character of an experience is not properly construed as feature that it bears. But then nothing remains of the claim that first-person reflection warrants T1.

Suppose you agree with me so far. Still, you may find this provides only a limited worry about R. Also, you might be doubtful about my description of the illusion case, and wonder whether my scope of attention point depends specifically on cases of *misperception*. I now want to argue that this point can be made on other grounds, and raises a more general concern about introspection as a source of warrant for reductive representationalism.

To appreciate this, it will help first to distinguish explicitly between *reductive* representationalist and *non-reductive* intentionalist views of phenomenal character.² Consider the following view. Its visually seeming or appearing to one a certain way is a phenomenal feature one has. Further: such visual phenomenal features are also intentional features in this sense. It is sufficient for being intentional that a feature is one in virtue of which its possessor can be assessed for accuracy or correctness—as happens when one says things such as: “The way it looks to her is accurate (or correct)” or “The way it looks to you is illusory.” And if the phenomenal feature is intentional, then the phenomenal character of the experience had, in having that feature, is inherently intentional.³

One could hold a view of this sort, and refrain from the reductive view of R. For it doesn't follow from this view that such phenomenal features can be *explained by identifying them with* a species of representation. Recall Tye's formulation of strong representationalism: “phenomenal character is one and the same as representational content that meets certain further conditions.” Now for this to work as an *explanation* of phenomenal character, the representational

content of the experience, and the further conditions mentioned in the account, must both be specified in terms other than merely as *ways of seeming* (looking, sounding, smelling). We need some way of conceiving of the special sort of content involved that does not make appeal to phenomenal modes of appearance of the kind targeted for explanation.

Let me try to make this a little clearer. Suppose I am asked to say what special form of representational content, and what further conditions explain why my visual experience has the phenomenal character it does. What makes this experience the experience of its looking a certain way to me? I might propose: “Well, it looks that way to me, because its looking that way is a state with a special sort of representational content—it represents a certain surface to have a certain shape and color.” Then the question will be: “But *what* shape, and *what* color?” At this point it would undermine R’s ambition to explain phenomenal character to say: the state represents the surface to have the *color that looks this way*, or the *shape that looks this way*. For then we would be appealing to the essentially *phenomenal look* of something to specify the *representational content*, when the representational content was supposed to explain that very phenomenal look, by telling us what it is.

Generally then, the reductive representationalist needs some way of conceiving of the representational content that is supposed to explain phenomenal character without appeal to the phenomenal modes of appearance—looking, tasting, smelling, etc.—to be explained. I will call this the “No Phenomenal Appeal” condition. Now notice, by contrast, that *non-reductive*

intentionalism does not require that this condition be satisfied. For while it says that phenomenal character is, in a way, inseparable from intentionality, it does not purport to explain phenomenal character by identifying it with a special sort of representational content.

The question I now want to pose is whether first-person reflection affords us the conception of representational content needed to fulfill the “No Phenomenal Appeal” condition to which R is committed. My contention will be that it does not. And my argument for this further supports the scope of attention thesis I used earlier to criticize the transparency argument for R.

First, consider how noises sound and how odors smell. It seems that, when I think of precisely *what* kind of odor it is that I am smelling, I am sometimes quite unable to think specifically of just *that* odor, relying on introspective resources, in any other way than simply as: *the odor of what smells this way to me*. I *may* of course classify it in some more informative way—as the smell of a wet dog, the smell of frying butter, the smell of lemon, and so on. But I can’t always or even commonly use such characterizations to distinguish all the ways of smelling to me that I can distinguish in thought. In any case, the effort to give such a characterization is guided by some sort of attention that precedes it, which focuses the question, “What is *that* odor?” on something’s smelling this way to me. And such attention permits me to recognize the aptness of the characterization offered in answer. So, to think of an odor in this manner, I attend to its smelling to me a certain way. Similarly, when I think of just what sound I hear, it seems I can sometimes think of just what sound it appears to me to be,

only as: *the type of sound that sounds this way to me*. I may of course go onto classify the sound as a hiss, a squeal, a screech, a roar, or what have you. But again, such classifications as I am prepared to offer may not capture the precise sound quality of which I am then thinking, in all its specificity. And, in any event, I have some way of thinking of what sound it is, prior to imposing such classifications, and on the basis of which I can determine their aptness.

One might wonder here: why is it not enough to express the relevant conception of an odor or sound here just to appeal to some demonstrative formulation, e.g., '*that* odor,' '*that* sound,'—and leave the matter there? And if one does this, one might claim that no phenomenal manner of appearing enters into one's conception of the quality. But notice: one can have concepts of sounds and odors expressible in that demonstrative manner, distinct from those expressible by phrases like 'what smells this way to me' or 'what sounds this way to me.' A musical composer absorbed in creative effort might think thoughts properly expressible as '*That* sound wouldn't go with *this* one,' even though he thinks of neither as: what now sounds some particular way to him. For he can compose in thought, when there is nothing he demonstratively indicates that sounds any way at all to him. (He is composing without playing, and maybe he is deaf.) Perhaps we will suppose he must be at least *imagining* the sounds—"hearing them in his head."? I am not sure that actually is essential. And anyway, 'what sounds this way to me' and 'what I *imagine* sounding this way to me' express different concepts—a difference that is unspecified by invoking demonstrative expressions like 'that sound.' And so, to express the specific

concept one is employing, it seems one cannot exclude the manner of appearance. (Similar remarks could be made in the case of odors or scents. Imagine a creator of perfumes “composing” a new fragrance.)

My point is that when we form, on a first-person basis, conceptions of the specific or fine-grained differences in sensibles, that is to say, in what it appears to us something is, we make appeal to phenomenal modes of appearance: what *smells* this way to me, what *sounds* this way to me, and so on. And evidently we often have, on the basis of first-person reflection, no available alternative. Similar observations seem to apply to other sensory modalities and “sensible qualities” (color and shape). If this is right, there are often cases in which we have, introspectively, no way to conceive of just what it appears to us something is, but by attending to its appearing to us as it does. Thus in first-person reflection, we conceive of what some would call the specific “external qualities” of things that appear to us, when we attend to how these appear to us, only relative to their appearing to us this way or that.

This provides an additional reason to accept the scope of attention thesis urged against some versions of transparency: we can, in some sense, attend to its appearing to us as it does—we can attend to our experience. But further, it gives us reason to doubt one can use introspection to justify the claim that experience has content completely characterizable without appeal to phenomenal modes of appearance. First-person reflection does not equip us to meet the “No Phenomenal Appeal” condition on R.

Here finally, is another kind of example I want to use to make the same point. It concerns phenomena of visual perspective long familiar in philosophical discussions of perception. Someone holds up before your eyes a round plate, at an oblique angle. Then she holds up a piece of paper, face-on, and on that paper is drawn an ellipse. Then she asks you, ‘Does *this*’ (referring to the figure on the paper) ‘look to you like *this*?’ (referring to the plate). To answer in the affirmative, on the basis of first-person reflection, as well you truthfully might, you must attend to how the plate and the figure look to you.

Now when you do this, and think of how it looks to you so as to answer the question, can you employ a conception of the way it looks to you that is free of appeal to phenomenal modes of appearance? Can you think of how the figure looks to you like the plate, by thinking of *what* they both look to you to be (or in Tye’s terms, what “qualities look to you to qualify something”)—but *without* appeal to their looking to you as they do? It is not easy to see how to do this. If you agree that, in some sense, the *plate* at least does not look to be elliptical, it clearly will not do to say that the quality that looks to qualify both plate and figure is: *elliptical*. Maybe one will try something along these lines. The shape the plate and the figure look to me to have is not elliptical, period, but something quite different: *elliptical-from-here*. But what does being “elliptical-from-here” mean, if it does not entail being elliptical? Tye suggests it means something like: having a shape such that “it would be occluded by an ellipse placed in a plane perpendicular to the line of sight.” (2000, p.79)

But this can't be quite right. For the plate could look to me quite *different* than the way it actually looks, and *not* like the figure, and yet still look to me to have a shape such that it would be occluded by an ellipse placed in a plane perpendicular to the line of sight. For all sorts of very different looking things could qualify as looking *that way*, provided that we alter the size of the ellipse and shape of the plate, and the placement of the occluding ellipse in the right ways.

What happens when I try to be more exact here? Maybe I should say: something looks to me to have a shape such that an ellipse of *a certain size* placed in a plane perpendicular to my line of sight, at *a certain distance* from me, would *exactly occlude it*. But now: what size and distance am I thinking of here, and what do I understand by 'exactly occlude it'? All I can understand by 'exactly occlude' here is: occlude it in such a way that everything right around it still looks some way to me, while it does not itself look any way to me at all. And I don't know how to identify the relevant specific size and distance here, but by reference to my experience: '*This size*—I mean, the size that this now looks to me to have.' But notice then that I am specifying what shape, size and distance the plate looks to me to be, by appeal to a phenomenal mode of appearing of the sort to be reductively explained (its looking some way to me). What I think this shows is that when I attend to how it looks to me so as to make first-person judgments of pictorial resemblance ('*This looks to me like that*') I have no conception of what things look to me to be, adequate to allow me to make such judgments, *without appeal to phenomenal modes of appearing* (its looking to me

a certain way). But this goes to show first, again, that when I attend to how it looks to me, its looking to me as it does falls within the scope of my attention. This further justifies what was said earlier about transparency. Second, first-person reflection affords me no conception of the representational content of my visual experience that meets the “No Phenomenal Appeal” condition on R.

Now reductive representationalists will, of course, want to appeal to non-introspective sources of warrant for their theories. But the question is whether—supposing I am correct here—the absence of introspective support to which I have pointed is significant. I think it is. Typically, when it has been urged that R will have difficulty in accounting for some introspectible difference in phenomenal character, philosophers have assumed that the challenge will be met, as long as the problem cases provide no clear examples of a phenomenal difference without a representational difference. But part of what I am arguing here is that this misconstrues the nature of the challenge to reductive representationalists. Non-reductive intentionalists can accept that there are no phenomenal differences without intentional ones. But R, by its very nature, takes on a major additional commitment regarding the relevant intentional differences. For it holds that these are differences in representational content, of a sort specifiable without appeal to phenomenal modes of appearing. But this conception of the intentionality of experience gets no support from introspection. Thus, to the extent that first-person reflection tells against non-intentional qualia, pure sensational features and the like, it can give us reason only to adopt non-reductive intentionalism, not R.

But now, to return to the transparency issue: isn't there *some* sense in which it is correct to say experience is transparent? I think there is. Note the role played by the notion of "turning your attention" in the quotations I gave at the outset from Harman and Tye. The idea in Harman seems to be that you can't turn your attention *to* your experience, and *away from* the tree. And in Tye: you cannot turn your attention *inward* to experience, away from the *external* things experienced. I agree that there is something to this. It seems to me that I cannot attend to what distinguishes one experience from others, while *taking attention away* from what the experience is *of*. Thus I would endorse this general formulation of transparency:

T2: You cannot attend to how it appears to you, by turning your attention *away from* something that appears to you (or away from what—and where—it appears to you something is), and *towards* its appearing to you as it does.

Now this seems to be part of what moves people to speak of the transparency of experience. Since one cannot turn attention to experience by turning it away from the objects experienced, it sounds right to say that the experience does not (and cannot) "block" attention to the object. And since the experience does not block the object, but in fact reveals it, experience, we might say "transmits" attention, somewhat as transparent things transmit light.

T2 invites us to recognize this important disanalogy between attending to visible objects and attending to visual experience. You can turn your attention

away from one visually apparent thing, and *to* another, so as to *ignore* the first, in *favor* of the second. Directing attention to the second thing excludes attending to the first. But if you attend to how some object looks to you on some occasion, you don't (and can't) do so, by turning your attention away from it or diminishing how much attention you devote to it, while increasing your attention to its looking to you as it does. Thus an injunction to turn your attention "inward" on experience seems especially misleading. For if there is some thing in public space your experience is of—as there will be if you're not hallucinating—attending to your experience equally will be attending to this thing. Also, even if you are hallucinating, in attending to how it appears to you, you will inevitably be attending to where it (incorrectly) appears to you something is. And where is it that it appears to you something is? In no sense I can think of is it "*inward*."

Notice that all this is entirely consistent with my "scope of attention" thesis. Clearly T2 does not contradict my claim that we can attend to its looking, feeling, sounding, etc., to us certain ways. There is simply no need to take this to mean that one can attend to one's experience to the *exclusion* of attending to things in one's surroundings and one's own body. We can recognize our ability to attend to its appearing to us in various ways, without conceiving of this as some kind of mental withdrawal from an "outer" to an "inner" realm. Once we reject the notion of private mental items as the immediate objects of appearance or perception, we would do better, I think, simply to dispense with this inner/outer talk altogether, which perhaps is partly to blame for errors, both among those who might counsel us to turn our attention inward on experience, and those who insist that we can

attend only to “external” objects and qualities. For this habitual metaphorical opposition of inner and outer may obscure from us the possibility that we can attend to the sorts of things that attract the “outer” label—things at least with spatial location and unperceived aspects—while also (and indivisibly) attending to their *appearing* to us—that is, to what is traditionally taken to fall on the “inner” side of the contrast. The mistake is not to suppose that one can attend to or be aware of one’s experience. The mistake is to suppose that to do this, one must withdraw attention from objects on the “outside” to objects on the “inside.”

But now the question resurfaces: does transparency, in *this* guise at least, if not the others, support R? An argument would have to be made that it did, and there is reason to think it wouldn’t be successful. For it may be that this sort of transparency gets some explanation from the inseparability of the phenomenal character of experience from its intentionality. But it seems non-reductive intentionalism would explain this as well as reductive representationalism, and with fewer commitments.

5

What do I conclude about the much vaunted transparency of consciousness? If the thesis of transparency states that we cannot attend to, or are at most indirectly aware *of* our own experience or its phenomenal character, or if it says experience has no feature that *is* its phenomenal character, then—going by what first-person reflection has to tell us—the thesis is either false, or

groundless. This emerges from consideration of examples that illustrate my scope of attention thesis.

Now there are other problems in arguments from transparency to R. But even so, I think it is worthwhile to focus on problems in the transparency claims with which they start. Partly this is because I believe this can encourage a certain reconfiguring of current debates about consciousness. Reductive representationalists are likely to see their main challenge coming from those who argue for non-intentional qualia or non-representational phenomenal differences. But, if I'm right, attention to experience shows that even if first-person reflection allows us to isolate no pure sensational qualities, it also yields no conception of pure representational contents. This suggests reductive representationalism faces a rather different sort of challenge: that while phenomenal features may indeed be inherently intentional, phenomenal experience does not have non-phenomenal representational content.

I also suggested at the start that part of why I wanted to investigate the idea that consciousness is transparent was that this would bear on efforts to understand what introspection is. Now even though I have just been casting aspersions on the notion that in introspection we turn our attention inward, and so suggesting it is misleading to use the term 'introspection' at all—some of what I've said might seem to speak in favor of an "inner sense" view of how we know our own minds. For some people (Bill Lycan, for one) would like us to think of attending to experience as involving a special "higher order" sense-like form of representation, distinct from the "first order" sensory representation of our

environment. But we shouldn't assume that recognizing the form of attention my examples illustrate should lead us to some such inner sense view. However, if we are to examine the notion of inner sense adequately, we need to come to terms with this sort of attention, and figure out just what role it plays in how we know our own minds.⁴

¹ For my interpretation of 'phenomenal character,' see Siewert (1998, Chapter 3).

² I prefer the term 'intentionalist' to describe the non-reductive position, since I subscribe to that position, and I have doubts about whether the intentionality of experience is always best thought of as its having *representational* content, properly speaking. For example, I think there may be differences in how it looks to us that are not specifiable by attributing either linguistically or imagistically expressible content to visual experience, but which are not, for all that, rightly regarded as non-intentional, merely qualitative differences. So I want to leave open the possibility that the notion of intentionality extends more widely than that of representation.

³ I work out this version of non-reductive intentionalism in more detail in Siewert 1998, Chs.6-8. For other views I would classify as non-reductive intentionalism, see Crane (2002) and Horgan and Tienson (2002). For a detailed discussion of different "intentionalisms" about consciousness, reductive and non-reductive, see Chalmers (2004).

⁴ I would like to thank Amy Kind, whose July 2002 talk on transparency at the NEH Summer Institute on Consciousness and Intentionality (forthcoming: Kind 2003) did much to stimulate my thoughts on this. I am also much indebted to questions and comments from (among others): Dave Chalmers, Terry Horgan, Kirk Ludwig, Alan Thomas, Amie Thomasson, Michael Tye, and Dan Zahavi.

REFERENCES

- Block, N. (1998) Is Experience Just Representing? *Philosophy and Phenomenological Research* 58: 663-70.

-
- Chalmers, D. (2004) The Representational Character of Experience. Forthcoming in *The Future for Philosophy*, ed. Leiter. Oxford: Oxford University Press.
- Crane, T. (2002) The Intentional Structure of Consciousness. In *Consciousness: New Philosophical Perspectives*, ed., Smith and Jolic. Oxford: Oxford University Press.
- Dretske, F. (1995) *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Harman, G. (1990) The Intrinsic Quality of Experience. In *Philosophical Perspectives*, ed. J. Tomberlin, vol. 4. Atascadero: Ridgeview Press.
- Horgan, T. and Tienson, J. (2002) The Intentionality of Phenomenology and the Phenomenology of Intentionality. In *Philosophy of Mind: Classical and Contemporary Readings*, ed. Chalmers. Oxford: Oxford University Press.
- Kind, Amy. (2003) What's So Transparent About Transparency? Forthcoming in *Philosophical Studies*.
- Peacocke, C. (1983) *Sense and Content*. Oxford: Clarendon Press.
- Shoemaker, S. (1996) *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.
- Siewert, C. (1998) *The Significance of Consciousness*. Princeton University Press.
- . (2001) Consciousness Neglect and Inner Sense: Reply to Lycan. *Psyche* 7: 07
- Tye, M. (1995) *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Tye, M. (2002) *Consciousness, Color and Content*. Cambridge, MA: MIT Press.