

A Theory of Phenomenal Concepts

1) There is widespread agreement that consciousness must be a physical phenomenon, even if it is one that we do not yet understand and perhaps may never do so fully. There is also widespread agreement that the way to defend physicalism about consciousness against a variety of well known objections is by appeal to phenomenal concepts (Loar 1990, Lycan 1996, Papineau 1993, Sturgeon 1994, Tye 1995, 2000, Perry 2001) . There is, alas, no agreement on the nature of phenomenal concepts.

2) Concepts are mental representations of worldly entities—things, events, states, properties etc. They are exercised whenever we undergo cognitive mental states. One cannot notice something, recognize it, make a judgement about it without conceptualizing it in some way, without bringing it under a concept. A child who is unable to count may see four pieces of candy but he/she cannot notice that four pieces are present. A dog may hear a Beethoven symphony, but it cannot recognize the sounds as being a Beethoven symphony.

3) Phenomenal concepts are the concepts we exercise when (but not only when) we notice or become aware of the phenomenal character of our experiences and feelings via introspection. Our experiences have phenomenal character whether or not we attend to them, but when we notice how an experience feels, what it is like, in doing so we are bringing it beneath a phenomenal concept. Without phenomenal concepts, we would be 'blind' to our phenomenal feels (Dretske 1995, Tye 1995), just as the child who cannot count is 'blind' to the fact that there are four pieces of candy in front of her.

4) Physicalists about consciousness typically agree with the following claims:

a) Absent qualia are conceivable. We can conceive of physical duplicates, one of whom has experiences while the other has no experiences at all. Such duplicates may be metaphysically impossible, but they are conceivable (just as it is conceivable that I am not Michael Tye even though, given the actual facts, it is metaphysically impossible).

b) Frank Jackson's Mary (1982)—the color scientist imprisoned since birth in a black and white room and possessed of all the physical information about color and color vision—doesn't know what it is like to experience red, green, etc while she remains in the

room. When she is freed and she starts to undergo color experiences, she makes some important discoveries.

c) Presented with the full physical story about pain (or any other experience), we can still intelligibly ask, "Why do those physical states feel like that? Why do they feel any way at all?"

5) The natural way for the physicalist to explain (a) is to say that phenomenal concepts are not physical concepts¹. Since phenomenal concepts are different from physical concepts, we can conceive of absent qualia. There is no contradiction or incoherence in the thought that a given organism meets whatever are the relevant physical conditions for consciousness and yet feels nothing — any more than there is a contradiction or incoherence in the thought that I am not Michael Tye, or that water is not H₂O, or that now is not 2:15pm.

6) The natural way for the physicalist to explain (b) is to say that Mary in her room does not possess the phenomenal color experience concepts the rest of us possess. She acquires these concepts as she notices the colors of flowers, trees, houses, etc and as she attends to her color experiences in doing so. Once the new concepts are acquired, Mary can come to think new thoughts, and thereby she is able to make new discoveries. I shall return to this point later.

This line of reply to the Mary example requires again that phenomenal concepts not be physical concepts. For if they were, Mary would possess them in her room, given her complete knowledge of all the physical facts.

There is a further conclusion to be drawn. Phenomenal concepts are not demonstrative concepts utilizing physical sortals. To appreciate this, suppose that Mortimer is undergoing an experience of red and that Mary is viewing the physical state in Mortimer with which this phenomenal experience is identical through a cerebroscope suitably attached to her black and white room. She conceives of the state she sees as that

1. This is not to say, of course, that phenomenal concepts do not refer to physical entities. The concept THIS is not a physical concept, nor is the concept I, but it does not follow that these concepts pick out nonphysical items.

F state, where '*F*' is a physical predicate expressing the appropriate physical property. Patently, when she leaves her room and attentively experiences red, she still makes a significant discovery.

7) The intelligibility of the questions in (c) requires again that phenomenal concepts not be physical concepts. It also requires that there be no physical concepts that are a priori co-referential with any phenomenal concepts. To see this, suppose that 'pain*' below is used purely phenomenally for a state whose essence is the specific, unpleasant phenomenal character of pain and that '*F*' is a physical predicate. Now consider the following argument form:

- i) Pain* is the *F*
- ii) Physical state so-and-so is present
- iii) Physical state so-and-so is the *F*

Therefore,

- iv) Pain* is present.

ii) and iii) are straightforwardly empirical, physical claims. Thus, if i) is an a priori truth knowable by anyone who possesses the phenomenal concept PAIN*, then since iv) is a priori deducible from i) - iii), there will be an explanation for why physical state so-and-so feels the way pain* does that is available without further empirical investigation to anyone who has that concept and who also has the requisite physical information. And this will be the case, note, even if i) is not a *necessary* a priori truth.

8) For the physicalist, then, any satisfactory account of phenomenal concepts must allow that, although phenomenal concepts refer to physical properties, a) they are not physical concepts, b) they are not demonstrative concepts utilizing physical sortals, and c) they have no a priori associated co-referential physical concepts. c) entails that phenomenal concepts are not concepts that designate their referents rigidly but whose reference is fixed by an a priori associated physical description.

9) Corresponding points can be made with respect to phenomenal concepts and the view that consciousness is a functional phenomenon that is realized physically. For ease of exposition, I shall not consider this view separately.

10) Having said what phenomenal concepts are not, what positive alternatives remain open? One possibility is that phenomenal concepts are concepts having explicitly

non-physical definitions. A second possibility is that phenomenal concepts are primitive rigid concepts whose reference is fixed by an explicitly nonphysical description. A third alternative is that phenomenal concepts are indexical concepts utilizing explicitly nonphysical sortals. All three of these alternatives entails that physicalism about consciousness is false.

Another alternative is that phenomenal concepts are concepts having phenomenal definitions. This sets off a vicious regress and so gives us no satisfactory account of how phenomenal concepts operate. The same is true if we say that phenomenal concepts are primitive rigid concepts whose reference is fixed by a phenomenal description. For how do the concepts expressed in the phenomenal description refer? Given that phenomenal concepts have their reference fixed by a phenomenal description, the answer must be by further associated phenomenal descriptions; and so on without end.

A further alternative is to hold that phenomenal concepts are demonstrative concepts utilizing phenomenal sortals. Prima facie, this proposal sets off a similar regress. But the threat of such a regress is staved off by a recent proposal by Ned Block (APA presentation, 2001²) that phenomenal concepts paradigmatically have the form THAT PHENOMENAL PROPERTY, where the indexical or demonstrative THAT refers to the phenomenal property exemplified in an associated mental sample (presumably an image or quasi-image³). For example, suppose that I think in a phenomenal way of something's looking red. On this proposal, the image of red accompanying my thought exemplifies the phenomenal property RED* and my thought refers to the same phenomenal property by conceiving of it as that property—the one exemplified in my image.⁴ If this account is

2. I should add that this proposal may not reflect Block's current view.

3. The notion of an image or a quasi-image is to be understood broadly here so that it covers a phenomenal memory of pain, for example. The latter is a phenomenal state that faintly echoes real pain, a state that may elicit a mental shudder or grimace.

4. Here and throughout the paper, I write as if I accept the dogma that phenomenal properties are intrinsic properties of images and experiences. That, of course, is not my real view. See here Tye 1995, 2000. For present purposes, whether phenomenal qualities

applied to the concept PHENOMENAL PROPERTY as itself having the form THAT PHENOMENAL PROPERTY, the regress is stopped.

Leaving aside the point that it is far from obvious that there is always an associated mental sample when a phenomenal concept is exercised⁵, there are two insuperable difficulties for this proposal. One is that a mental sample that exemplifies one phenomenal property will exemplify many. My image, when I think of something's looking red, will not only exemplify RED* but also (let us suppose) SCARLET*, DARK RED*, HAVING A COLOR*, and so on. Which of the exemplified properties is the one to which the demonstrative concept THAT PHENOMENAL PROPERTY refers? It seems that appealing to a mental sample does not help to fix the reference of the phenomenal concept at all.

A second related difficulty concerns the phenomenal concept PHENOMENAL PROPERTY. What is the relevant sample for this concept? It appears that *any* phenomenal image or quasi-image will do, in which case the problem of too many eligible candidates for reference rears its head again.

11) Perhaps it will be replied that the problem of too many eligible candidates goes away on the supposition that the property to which the demonstrative THAT refers is the phenomenal property (exemplified in the sample) to which the imager is attending. This does not help, however. Attention to a property is not like training the eyes on a point in space. Given the multiplicity of exemplified properties, attention to one of those properties rather than another requires noticing the relevant property. And that involves bringing it under a concept. The appropriate concept here will surely be a phenomenal one. So, the proposal is now circular.

12) The conclusion to which we seem driven as physicalists is that phenomenal concepts refer directly. They have no associated reference-fixers, no descriptive content

are qualities of experiences or qualities represented by experiences does not matter. The story I have to tell about phenomenal concepts will apply to either view with minor (and obvious) modifications.

5. More on this later.

at all. For concepts of this sort, the referent is presented without the assistance of associated features distinct from the referent which the thinker a priori associates with it. There is no separate guise that the referent takes in the thinker's thought. Intuitively, independent of the truth of physicalism, this seems to me the right approach. If I focus introspectively on the feeling of pain, as I experience it, I form a conception of how it feels, and the concept that enables me to do that is not one that I apply to the feeling by discerning non-phenomenal features (or for that matter *other* phenomenal features) that aid in the identification of its phenomenal character. Intuitively, I know that I am in pain just by attending to how my state feels, not by knowing something *else* connected to it.

13) The natural picture, it seems to me, is as follows. Our phylogenetic nature determines which experiences we undergo. We are hard-wired to experience various bodily sensations and to undergo various perceptual experiences. We cannot experience what a bat experiences when it uses echo-location, since we lack the appropriate sensory system. We are also equipped by evolution and nature to respond cognitively to our experiences in a certain range of ways once we undergo them. In responding cognitively, we bring the experiences under phenomenal concepts. Which concept is applied may depend on a number of factors: how our attention is directed, previous experiences, learning, attention span. But there are limits set upon the phenomenal concepts available to us by our nature.

I am happy to allow that there could be other creatures capable of undergoing the same experiences as us but who conceive of their experiences differently on a first person basis. For example, they might be capable of much finer grained classifications with respect to their color experiences than we are. Such creatures are equipped with a different battery of phenomenal concepts, one that no doubt partly overlaps with ours.

In my view, introspection of phenomenal character is a *reliable* process that takes phenomenal character as input and yields awareness *that* a state is present with a certain phenomenal character as output. It is the reliability of this process that underwrites knowledge of phenomenal character. In this respect, introspection of phenomenal character is like introspection of thought contents. Let me explain.

If I think that water is wet, and I introspect, I become aware *that* I am thinking that water is wet. This awareness is not based upon an inference from other propositional

states. Nor is it the result of attention to an internal auditory image of myself saying that water is wet, though such an image may accompany my thought. Intuitively, my introspective access to what I am thinking is direct. It seems plausible to suppose that introspection of thought contents is a reliable process that takes as input the content of the thought and delivers as output a belief or judgement that one is undergoing a state with that content.

On this view of introspective knowledge of thought contents, the concept of a thought that *P* is, in its first-person present-tense application, a *recognitional* concept.⁶ Those who have mastered the concept can introspectively recognize that an occurrent thought that *P* is present without going through any process of reasoning. In cases involving what Tyler Burge has called “Cogito thoughts” (that is, cases in which one consciously thinks to oneself that one is thinking that *P*), there is a conscious act of recognition. But in the typical case, one's recognition of what one is occurrently thinking does not involve a conscious act. One can recognize that one is thinking that water is a liquid, when the only conscious thought one is having is that water is a liquid.⁷

In much the same way, we do not have introspective knowledge of phenomenal character by inferring that character from something else. We acquire introspective knowledge of what it is like to have such-and-such an experience or feeling via a reliable process that triggers the application of a suitable phenomenal concept or concepts. Phenomenal concepts—the concepts that enable us to form a conception of phenomenal character via introspection—are, in my view, recognitional concepts of a special sort.

14) My proposal, then, is that phenomenal concepts refer via the causal connection they have with their referents. In first approximation, a phenomenal concept *C* refers to a phenomenal quality *Q* via *C*'s being the concept that is exercised in an introspective act of awareness by person *P* if, and only if, under normal conditions of introspection, *Q* is tokened in *P*'s current experience and because *Q* is tokened. I say ‘in first approximation’

6. For more on recognitional concepts, see Brian Loar 1990.

7. Introspection of thought content is discussed in more detail in McLaughlin and Tye 1998.

here since a further condition is needed to handle the possibility that *C* not only causally covaries with *Q* but also with a further non-phenomenal, indeed non-introspectively accessible, quality of the experience under normal conditions.

This difficulty is, of course, part of a more general one for causal covariation accounts of representation, whether that representation is conceptual or not. The hair shedding of cats (under normal conditions) is causally correlated with the lengthening of days; and lengthening days correlate (roughly) with increasing temperature. Thus, shedding in cats causally covaries with both day length and temperature. Even so, given what we know of the relevant biological mechanisms, it seems wrong to say that the shedding of hair represents temperature as well as (or instead of) day length.

In the cat case, the causal covariation between the shedding of hair and increasing temperature arises because the hair shedding causally covaries with day length and day length covaries with temperature. Were the covariation link between temperature and day length broken (by, for example, keeping cats indoors at a constant temperature or moving them to higher altitudes at the same latitude), the hair shedding would continue to covary with day length (albeit artificial day length for the indoor case generated by varying the hours of artificial light), but not with temperature. For this reason, hair shedding is best taken to represent day length—provided that we are prepared to talk of representation in this context at all.

Intuitively, then, what is needed to supplement the basic causal covariation approach is a further asymmetric dependence condition. For state *S* to represent feature *F* not only must *S* causally covary with *F* under optimal or normal conditions but it must also be the case that if there is some other feature *G* such that *F* covaries with *G* under optimal conditions then were *F* to fail to covary with *G*, the causal covariation link between *S* and *F* under optimal (normal) conditions would still hold but that between *S* and *G* would be broken.

This qualification handles the case of phenomenal concept *C* covarying with both phenomenal quality *Q* and non-phenomenal quality *N*. *P* can be held to represent *Q* and not *N* so long as it is held that were the covariation link between *Q* and *N* broken, *C* would continue to causally covary with *Q* but not with *N*.

15) Perhaps it will be objected that it is surely possible for a concept to refer directly to a phenomenal quality without being a phenomenal concept. Suppose, for example, that the distinctive phenomenal character of pain is a brain state and that Fred is a 21st century neuroscientist who is incapable himself of feeling pain in virtue of a neurological defect he has had since birth. Fred has a device partly wired into his brain that causes him to think that another person is feeling pain when and only when the external part of the device is directed at the other person's brain and the relevant brain state is present there. Fred's thought exercises a concept of pain, but that concept isn't a phenomenal concept. For Fred does not know what it is like to experience pain, and intuitively one cannot grasp the phenomenal character of pain, one cannot have a *phenomenal* concept of pain, without knowing what it is like.

16) This example shows that we need to distinguish the question, "What is it that makes a phenomenal concept of quality *Q* be about or of *Q*?" from the question, "What makes a phenomenal concept phenomenal?". Not all concepts that refer directly are phenomenal. Concerning the latter question, the thesis in Tye 1998 was that a concept that directly refers to a phenomenal quality is phenomenal if and only if it functions in the right sort of way. I denied, however, that this functioning could be specified *a priori* in a way that eschews any phenomenal language. My view was (and is) that the concept of a phenomenal concept is conceptually irreducible: no *a priori* definition or analysis is possible in non-phenomenal terms. This should come as no surprise. If such an analysis were possible then a suitably cognitively informed automaton, totally without any experiences, would be able to acquire the concept of a phenomenal concept simply by reflecting upon the analysis. Intuitively, however, that isn't possible. Such an automation could glean no phenomenal notion of an experience; and without such a notion the concept of a phenomenal concept would be beyond its grasp.

This does not have the consequence that we cannot say *a priori* anything illuminating about the relevant functioning of phenomenal concepts. Quite the contrary. A concept is phenomenal, I maintain, if and only if 1) it is laid down in memory as a result of undergoing the appropriate experiences (barring miracles, etc), 2) it tends to trigger appropriate conscious images (or quasi-images) in response to certain cognitive tasks, and 3) it enables its possessors to discriminate the phenomenal quality to which it refers

directly and immediately via introspection.⁰ This proposal was originally made in Tye 1999 and it was motivated by what I take to be a priori links between phenomenal concept possession and knowing what it is like (one cannot possess a phenomenal conception of a given experience type unless one knows what that experience type is like) and further a priori links between the latter and certain phenomenal abilities underpinning the stated conditions (abilities to imaginatively recreate the experience, to remember it, to recognize it directly when it comes again). Moreover, the proposal, though non-reductive, is not vacuous or trivial. It imposes real requirements on a concept's being phenomenal, requirements that are not met by most of the concepts we have.

The proposal also entails that phenomenal concepts are, in a certain sense, perspectival. Intuitively, possessing the phenomenal concept PAIN* requires having a certain perspective on pain, the one conferred by experiencing pain oneself (barring miracles, etc). Why should this be? Answer: because the phenomenal concept PAIN* would not be a phenomenal concept at all if it didn't function in the right sort of way, and that functioning brings with it a distinctive first person perspective on pain. This is why the 21st century pain detector wired into the head of our neuroscientist Fred does not provide Fred with a *phenomenal* concept of pain, and why Mary in her black-and-white room does not have the phenomenal concept RED*. These individuals, given their special conditions, don't have any internal mental representations that function in the appropriate ways.

17) The imagistic dimension of phenomenal concepts deserves further comment
Consider first the following example of a phenomenal-physical identity claim:
The visual experience of red = brain state *B*.⁰

0.7. Note incidentally that this account itself entails that phenomenal concepts refer directly.

0.9. I myself do not accept identities of this sort. In my view, the objective states with which phenomenal states should be identified are complex representational states (Tye 1995, 2000). In the present context, however, this does not matter.

One reaction some philosophers have to claims of this sort is that they must be mistaken, since the the phenomenology isn't captured by the right-hand side. From the present perspective, this reaction involves a sense/reference confusion. When we think of the referent of the designator on the left-hand side in a phenomenal way, we bring it under a concept that has a distinctive functional role. In reflecting on the identity claim and what is puzzling about it, the phenomenal concept we deploy is apt to trigger in us a visual image of red. In this event, if the identity is true, our brain actually goes into brain state *B*. But, of course, when we think of the referent of the designator on the right hand side as brain state *B*, nothing like that happens. Exercising the neurophysiological concept is not apt to trigger a visual image of red. It may then be tempting to infer that the right-hand side has left out the phenomenology of the left, that there is a huge gap that the physicalist has failed to close. This conclusion clearly does not follow, however. There is indeed a striking difference in the roles that the concepts play, in their functioning, but not (so far as is shown here) in their referents.¹⁰

It should be emphasized that the view I am proposing of phenomenal concepts does not require that a conscious image or quasi-image always be present when a phenomenal concept is exercised. Some philosophers take a stronger position. We saw earlier that, according to Ned Block, phenomenal concepts have the structure THAT PHENOMENAL QUALITY, where the indexical refers to the phenomenal quality exemplified in an associated mental sample. David Papineau (1993) has a similar view. His claim is that our brains are wired to form copies or replicas of the experiences we undergo, and these replicas play a role in fixing the reference of phenomenal concepts. Specifically, Papineau's proposal is that phenomenal concepts have the structure THAT EXPERIENCE, where the demonstrative refers to the experience type exemplified in an associated image or copy of the experience. On Papineau's account, exercising a concept of a phenomenal state involves recreating it or simulating it, and thinking of it as that state, the one tokened in the simulation.

10. Cp. Papineau 1994.

This cannot be correct. For one thing, when I deploy a phenomenal concept in an introspective act—when I introspectively recognize the feel of a tickle, say—the only experiential state present is surely the one I am recognizing, the tickle feeling. There isn't a *further* image or copy of a tickle within my introspective act of recognition or associated with it. If there were, it would be accessible to me via my further awareness that I am engaging in an act of introspectively recognizing a tickle. But no such tickle copy or replica reveals itself to me.

A second problem is that the earlier objection to Block applies to Papineau too. What Papineau calls (following Hugh Mellor (1992)) “exemplificatory reference by secondary experience” (10993, p. 112) is not reference at all. Consider the idea when I think of pain in a phenomenal way, I exercise the concept THAT EXPERIENCE (TYPE), where the demonstrative picks out the type of experience tokened in an associated pain image or replica. There is, alas, no such thing as *the* type of experience so tokened. There are many types. My pain replica exemplifies the phenomenal quality PAIN* but it also exemplifies such phenomenal qualities as THROBBING PAIN* or DULL PAIN* or PRICKING PAIN* as well as such phenomenal properties as HAVING A PHENOMENAL QUALITY, HAVING AN UNPLEASANT PHENOMENAL QUALITY, and so on.¹¹

Block and Papineau's primary mistake, then, is to suppose that images play a *reference-fixing* role for phenomenal concepts. As far as reference goes, associated images play no role at all. Still, it does seem right to say that sometimes exercising a phenomenal concept triggers an image or faint replica of the relevant experience, and that this connection is one that is essential to phenomenal concepts generally. Why should there be such a connection? What is it about phenomenal concepts that explains their essential imagistic associations?

18) Consider the following possible model. Suppose that an explorer sees an animal belonging to a hitherto unknown species. He takes a picture of the animal, and assigns a name to the species. He then keeps the photograph in a file with the name written

11. Appealing to attention with respect to one of the properties does not save the proposal. See 11) earlier.

beneath it. On later occasions, in talking about the animal to others, he opens the file, takes out the photo and holds it up as he uses the name.

The name refers to the species, not to the particular animal the explorer saw. The picture is a picture of a member of that species, a token of the type.

Alternatively, the explorer, instead of taking a photograph might carve a replica of the animal in wood. The wooden replica is a replica of a particular animal, though the explorer can certainly use it to represent the species too, as he uses the name to discourse about the species.

We might be a bit like the explorer. Consider the following proposal. Suppose that we each have a phenomenal character detector wired into our heads. The detector is set up to register phenomenal qualities in our experience that are the focus of our attentional mechanisms. The detector does this by outputting a name or simple symbol, a different symbol for each different phenomenal quality. If the phenomenal quality is an unfamiliar one, the device places the name in memory and it makes a copy or a faint replica of the phenomenal token of that quality which is present in our experience. The copy is then placed in memory along with the name. These processes, we may suppose, are automatic and unconscious.

On some occasions, as when we introspectively recognize a phenomenal type, we make no use of the stored copy or image. On other occasions, when the phenomenal type is absent from our experience but we are thinking phenomenally about it, we use the name and we retrieve a copy of a token of the phenomenal type too, an image or quasi-image, which we may then put to cognitive use.

For example, when asked the question, "Which is darker green, grass or a Scottish fir?," people typically think about the color of grass and that of Scottish firs in a phenomenal way. This generates phenomenal images of the two greens, images that are then inspected and compared by the subjects before they reply to the question. Those without the capacity to form such images, visual agnosics, for example, are unable to answer.

19) Two worries remain. First, it may be objected that my account faces a regress problem just as some of the earlier proposal do. For in saying what makes a concept

phenomenal I have used phenomenal concepts. These concepts must refer to physical states and properties, if physicalism about phenomenal consciousness is to be true.

My reply is that there is no vicious regress set off by this requirement, since the reference of phenomenal concepts is direct. It is *not* the case, on my view, that in order for a given phenomenal concept to *refer* successfully, other phenomenal concepts must do the same, where these concepts refer successfully only if other phenomenal concepts do, and so on without end.

Nor is there any regress in my account of what makes a concept phenomenal. The account is not proposed as a reductive one. So, the fact that it uses phenomenal concepts does not create a regress. What the use of these concepts reflects is simply the conceptual irreducibility of the concept of a phenomenal concept.

20) The final worry I want to consider pertains to the case of Mary and whether the proposal I have made about phenomenal concepts allows Mary to make any new discoveries about the phenomenal character of color experience after she leaves her black and white room.

The answer, in brief, is yes. In order for Mary to think a phenomenal thought, she must exercise a phenomenal concept. She does not have phenomenal color experience concepts in her room. So, when she leaves her room, she starts to have *new* phenomenal thoughts. Content-wise, it must be granted, in one standard sense of the term 'content', these thoughts will not be new, given Mary's complete physical knowledge in her room. For, on my proposal, phenomenal concepts refer directly and thus, the contribution they make to thought content is given by their referents alone—referents that are physical, if physicalism about consciousness is true. But thought-types need not be individuated by their contents alone.

Intuitively, phenomenal thought types play a different role in rationalizing explanations than non-phenomenal thought types. If their contents are identical, then a second factor must account for this difference. And intuitively that factor is simply that phenomenal thoughts exercise different concepts—*phenomenal* concepts (whose difference from non-phenomenal concepts, on my account, is given by their functional role). Accordingly, I claim that the identity of a phenomenal thought type may be traced both to its content and to the fact that it employs concepts that function in a certain

characteristic way. This two-factor view of phenomenal thought types permits the physicalist to maintain that there is a perfectly good sense in which Mary discovers that so-and-so is the case after she is released. For she comes to think new thoughts and thereby instantiate cognitive thought-types (knowing-that types) she did not instantiate before, even though, given her exhaustive knowledge of the physical facts, the contents of her thought-types before and after remain unchanged. And if Mary or anyone else knows that p at time t without knowing that p before t , then surely it is correct to say, in ordinary parlance, that the person has made a discovery at t .

21) The theory of phenomenal concepts sketched above seems to me both to respect anti-physicalist intuitions and to give the physicalist everything she needs. It allows for the conceivability of those states of affairs pertaining to experience that are dear to the hearts of anti-physicalists. It preserves the intuition that we know phenomenal character in a direct, non-inferential way. It acknowledges that undergoing a new experience and attending to it yields a discovery, and that this is the case regardless of how much physical knowledge we have. It finds a place for images or faint copies of experiences in the exercise of phenomenal concepts. Moreover, it does all this while holding that the referents of phenomenal concepts are physical. The theory thus discharges a heavy burden on the physicalist; and it leaves those who insist that we still don't have a good account of phenomenal concepts (Levine 2001) with the burden of explaining why.

Michael Tye

The University of Texas at Austin

REFERENCES

- Jackson, F. 1982 "Epiphenomenal Qualia," *Philosophical Quarterly*, 32, 127-136.
- Levine J. 2001 *Purple Haze: The Puzzle of Consciousness*, Oxford: Oxford University Press.
- Loar, B. 1990 "Phenomenal States," in *Philosophical Perspectives*, 4, J. Tomberlin, ed., (Northridge : Ridgeview Publishing Company).

- Lycan, W. 1996 *Consciousness and Experience*, Cambridge, Mass: The MIT Press, Bradford Books.
- McLaughlin, B. and Tye, M. 1998 "Externalism, Twin-earth, and Self-Knowledge" in *Knowing Our Own Minds: Essays on Self-Knowledge*, ed. C. Macdonald, B. Smith, and C. Wright (Oxford: Oxford University Press).
- Mellor, H. 1992 "Nothing Like Experience," *Proceedings of the Aristotelian Society*, 93.
- Papineau, D. 1993 *Philosophical Naturalism*, Oxford: Blackwells.
- Perry, J. 2001 *Possibility, Consciousness and Conceivability*, Cambridge, Mass: The MIT Press, Bradford Books.
- Sturgeon, S. 1994 "The Epistemic View of Subjectivity," *Journal of Philosophy*, 91, pp. 221-35.
- Tye, M. 1995 *Ten Problems of Consciousness*, Cambridge, Mass: The MIT Press, Bradford Books.
- Tye, M., 1999 "Phenomenal Consciousness: The Explanatory Gap as a Cognitive Illusion," *Mind*, 108, 705-25.
- Tye M., 2000 *Consciousness, Color, and Content*, Cambridge, Mass: The MIT Press, Bradford Books.