Philosophy and Phenomenological Research
Vol. LXIX, No. 1, July 2004

The Phenomenology of Cognition Or *What Is It Like to Think That P?*

DAVID PITT
California State University, Los Angeles

A number of philosophers endorse, without argument, the view that there's something it's like consciously to think that p, which is distinct from what it's like consciously to think that q. This thesis, if true, would have important consequences for philosophy of mind and cognitive science. In this paper I offer two arguments for it.

The first argument claims it would be impossible introspectively to distinguish conscious thoughts with respect to their content if there weren't something it's like to think them. This argument is defended against several objections.

The second argument uses what I call "minimal pair" experiences—sentences read without and with understanding—to induce in the reader an experience of the kind I claim exists. Further objects are considered and rebutted.

It is a traditional assumption in analytic philosophy of mind that intentional states, such as believing, doubting or wondering that *p*, have no intrinsic phenomenal properties, and that phenomenal states, such as feeling pain, seeing red or hearing middle C, have no intrinsic intentional properties. We are, according to this view, of two metaphysically distinct minds, the intentional and the phenomenal. Both of these assumptions have been challenged in the recent literature. Block (1996), Loar (2001), Peacocke (1992) and Tye (1995), for example, have argued that purely phenomenal, nonconceptual states have intentional (or proto-intentional) properties. And a fair number of philosophers and psychologists, e.g., Baars (1988), Chalmers (1996), Flanagan (1992), Goldman (1993), Horgan and Tienson (2002), Jackendoff (1987), Kobes (1995), Langsam (2000), Levine (1993; 1995), Loar (1987; 1998), McGinn (1992), McCulloch (1999), Moore (1962), Peacocke (1998), Schweizer (1994), Searle (1990), Siewert (1998) and Strawson (1994), have expressed the view that conscious intentional states have qualitative character.

This paper concerns the latter thesis. It is notable that, though apparently widely endorsed, it has not been widely argued for.[1] Perhaps those who think it is true think it is simply too obvious to require argument. Yet, those who reject it tend to think it is just as obvious that it is false. Clearly, arguments on both sides are called for. Moreover, those who accept the thesis would do well to provide a way to focus attention on a few instances of the qualitative character of conscious thought – in order to forestall the Humean objection that no such argument could be sound, because no such phenomenology exists, because the objector cannot discover it within him- or herself (cf. Nelkin 1996: 142-43).

I shall be defending a rather strong version of the thesis.[2] In addition to arguing that there is something it is like to think a conscious thought, I shall also argue that what it is like to think a conscious thought is distinct from what it is like to be in any other kind of conscious mental state, that what it is like to think the conscious thought that *p* is distinct from what it is like to think any other conscious thought, and that the phenomenology of a conscious thought is constitutive of its content. I shall also attempt to acquaint the reader with some instances of the phenomenology of cognition.

1. Consciousness and Phenomenology

Though I do not think the thesis that there is a phenomenology of conscious thought should be assumed to be obviously true (*or* obviously false), I think there is a rather obvious argument for it – to wit:

(P3) If a mental state is conscious, then it has phenomenal properties

(P2) Conscious thoughts are conscious mental states; therefore,

(P1) Conscious thoughts have phenomenal properties

I take it there is no difference between the conscious occurrence of a thought and consciously thinking a thought. Thinking a thought is like having a pain, in the sense that the thinking and the having are not something in addition to the mere occurrence of the states. Hence, thinking (in the sense of *entertaining*) is not a propositional attitude, but merely a having-in-mind. (Compare thinking a thought (entertaining a content) and having a pain with, respectively, *believing* the thought (content) and *disliking* the pain.) (P1) should therefore be distinguished from the claim, which I do not defend here, that there is a phenomenology

---

1 Siewert 1998 (chapter 8) is an important exception. Flanagan 1992 (69) and McCulloch 1999 (20) also offer arguments; but they are very brief, and neither is developed or defended.

2 One which I would hesitate to attribute to every author in the foregoing list – with the exceptions of Siewert (1998) and Horgan and Tienson (2002).

distinctive of consciously bearing a particular *attitude* to a particular proposition (believing that or wondering whether *p*, for example: cf. Flanagan 1992: 67; Goldman 1993: 23-25; Horgan and Tienson 2002).

It has been objected that (P3) is true by definition ('conscious' *means*, or analytically entails, 'phenomenal'), and that, consequently, the argument is trivial. To say that a state is conscious *just is* to say there is something it is like to be in it; and it follows immediately from the description of a thought as conscious that there is something it is like to have it.[3]

But this is really no objection at all. For even if (P3) is true by definition, or trivially true, it is still *true*. Hence, given that (P2) is true and that the argument is valid, it follows that (P1) is true as well. What this objection could show, at best, is that those inclined to deny the conclusion of the argument have missed something that ought to have been obvious.

(P3), however – though it does seem to me to be obvious (and perhaps even *necessarily* true) – is *not* true by definition. For, unconscious phenomenal states and non-phenomenal conscious states are conceivable.[4] And if such states are conceivable, then it cannot be that *phenomenal* is (or is part of) what 'conscious' *means*. Indeed, Eric Lormand (Lormand 1996) has argued that (P3) is *false* – i.e., that consciousness does not presuppose phenomenality.[5]

Lormand maintains that there are conscious states – in particular, thoughts and propositional attitudes – that are *never* phenomenal. (It is clear that Lormand is not arguing that there are states that are access-consciousness without being phenomenally conscious.) He claims that any phenomenology that might be associated with such states is the phenomenology of *accompanying* states of familiar kinds, such as perceptual representations, bodily sensations, images or inner utterances.

---

[3] According to Block (1997), for example, on one sense 'conscious' *just means* 'phenomenal'. The objection was offered by Sara Worley in her comments on an ancestor of the present paper presented at the Pacific Division APA meetings in March 1998.

[4] That is, it is not conceptually necessary that a state is phenomenal if and only if it is conscious: the concept of consciousness and the concept of phenomenality are distinct concepts. (Cf. Rosenthal 1991; Burge 1997; Lormand 1996.) I do not mean to be allowing here for states that are (in Block's (1997) sense) phenomenally conscious without being access conscious. (In my own idiolect, 'conscious' *never* means 'access conscious'. I suspect the term is, for me, semantically primitive.) What I am claiming are conceivable are states with phenomenal properties that are not conscious in *any* sense (cf. the distinction Burge (1997: 432) makes between "phenomenality" and "phenomenal consciousness"). Such states would be states with phenomenal properties that are not like anything *for* the individual whose states they are.

(Note that if unconscious phenomenology is possible, there are in fact *two* "hard problems of consciousness" – the problem of phenomenality, and the problem of consciousness itself.)

[5] Burge (1997: 431) also suggests that there are non-phenomenal conscious states, though he provides no arguments or examples.

Lormand's arguments, however, do not establish that (P3) is false. The claim that there are conscious states that are never phenomenal must be distinguished from the claim that there are conscious states that have no *proprietary* phenomenology. Though Lormand's stated target is the first claim, his arguments are clearly aimed at the second. (I consider these arguments in section 2.) In order to discredit the first claim, one would have to show that there are conscious states that can occur without any accompanying conscious states of the familiar types Lormand mentions. For it might be that such states cannot be conscious *unless* they occur with accompanying conscious states; and there might be distinctive accompaniments for each type of conscious thought, in which case there would be something it is like to think it consciously. But this Lormand does not do. Hence, his arguments, even if they were successful, would not show that (P3) is false.

Simple denial of the phenomenality of conscious thoughts does not constitute an argument against (P3). What is required is a *reason* for thinking that a mental state could be conscious without having any phenomenality at all. Given the strength of the intuition that it is impossible for mental states of so many other kinds (sensations, perceptions, proprioceptions, emotions) to be conscious but not phenomenal, to offer conscious thoughts as examples begs the question. An explanation of *why* conscious thoughts should be different (how they *could be* different) is required. Since I know of no such explanation, I shall take it that the argument of this section, obvious though it might be, does establish that there is a phenomenology of conscious thought. (Nonetheless, the argument of the next section, the conclusion of which *presupposes* (P1), will also serve to establish the claim that there is something it is like to think a conscious thought – though it goes far beyond this.)

2. Immediate Knowledge of Content

If there is a phenomenology of conscious thought, it remains to be determined whether it is just a phenomenology of familiar sorts (visual, auditory, tactile, olfactory, etc.), or a unique and distinctive sort of phenomenology, as different from the familiar sorts as they are from each other. In this section, I shall argue that what it is like consciously to think a particular thought is (1) different from what it is like to be in any other sort of conscious mental state (i.e., *proprietary*) and (2) different from what it is like consciously to think any other thought (i.e., *distinctive*). That is, any conscious token of a thought-type *T* has a unique phenomenology different from that of any other sort of conscious mental state, and different from that of any other conscious thought.[6]

---

[6] Throughout this paper I use 'think a thought' and 'entertain a (representational) content' interchangeably. (I shall clarify what I mean by "representational" content momentarily.) I assume that it is sufficient for two

Moreover, if conscious thoughts have distinctive and proprietary phenomenologies, it becomes natural to ask whether their phenomenology bears any relation to their content – just as it is natural to ask the corresponding question with respect to the phenomenology of perceptual states. (Given that, for example, there is something it is like to hear a particular note issuing from a particular piano, what place does that phenomenology have in an account of the content of that state?[7]) I shall also argue in this section that (3) the phenomenology of a thought *constitutes* its *representational* content (i.e., is *individuative*).

All three of these claims can, I believe, be established by a single argument. Hence, I collect them in the thesis (P):

(P) Each type of conscious thought – each state of consciously thinking that *p*, for all thinkable contents *p* – has a proprietary, distinctive, individuative phenomenology.

Apart from its intrinsic interest, (P) has important consequences for the philosophy of mind and cognitive science. If it is true that conscious intentional states have a distinctive phenomenology *qua* intentional, then any story about them that leaves out what it is like to be in them – a story exclusively in terms of, say, neurophysiological events as presently understood – will be incomplete. It will be as unsatisfactory as an account of conscious visual perception that leaves out what *it* is like (cf. McCulloch 1999: 21).[8] Thus, if (P) is correct, it extends the problem of qualia to a realm whose presumed immunity from it has fed hopes for a complete naturalistic theory of the cognitive mind using only the resources of current philosophy and neuroscience. (As I emphasize in section 4, it also makes trouble for a certain brand of cognitivist

---

thoughts to be type-distinct that they have different (representational) contents. Thus, my claim that *thoughts* have proprietary phenomenology is the claim that *content-entertainments* have proprietary phenomenology.

7 Though I cannot argue for this in detail here, it seems to me (and many other philosophers) that the qualitative character of one's perceptual experience is at least partly determinative of its content. One reason for thinking this is that the more the qualitative character of a perceiver's experience differs from the qualitative properties of its external cause(s) (or, at least, from the qualitative experience most perceivers would have under the circumstances), the more implausible it is to attribute perception of the cause(s). (If, for example, when faced with a ripe persimmon on a grey table I have a visual experience as of a small nervous dog, it would be incorrect to say that I see the persimmon on the table (or that there is a persimmon on the table).) (See also Loar 2001 and Peacocke 2001; see Brandom 2002 for a dissenting view.)

8 There are of course those who would argue that the latter sort of account is *not* unsatisfactory – that nothing has been left out of an account of conscious visual experience in purely neurophysiological terms as presently understood, since there are no such things as qualia (see, e.g., Tye 1995, 2000 and Lycan 1996). In the more familiar sorts of cases, however, I think it *is* just obvious that there is something it is like to have a conscious experience, and that this is a property *of the experience*. Hence, I am presupposing realism about qualia in this paper.

eliminativism about *perceptual* phenomenology.)

Before the argument for (P) is presented, some clarificatory remarks are in order.

(1) By the "representational content" of a thought, I mean those of its properties in virtue of which it represents (expresses) the proposition it does. The proposition it represents, in contrast, I shall call its "propositional content." (P) is a thesis about the representational contents of thoughts: it is the claim that conscious thoughts with distinct representational contents have distinct phenomenologies of a cognitively proprietary sort, that these phenomenologies constitute their representational contents, and, hence, that conscious thoughts have their propositional contents in virtue of their cognitive phenomenology (equivalently: a thought's having a particular representational content *is* its having a particular phenomenology).[9] Note that (P) allows that thoughts with different representational contents (phenomenologies) might (for the same or different thinkers) express the same proposition[10] – i.e., it types thoughts by their *representational* content rather than by their *propositional* content.

References herein to "contents" should (unless it is otherwise specified) be taken to be to representational contents, and references to "the thought that *p*" should be taken to be to the thought with the representational content that *p* (i.e., with representational content expressing the proposition that *p*).

(2) I take it that the unique, proprietary phenomenology of an occurrent conscious thought, *qua* representational content, plays the role Husserl and Searle specify for their respective notions of "matter" (Husserl 1900/1970) (or "noema" (Husserl 1913/1962)) and "aspectual shape" (Searle 1990) – neither of which is, in my view, sufficiently clarified or explained by its originator. (See Pitt In Preparation (a).)

(3) (P) does not claim, entail or presuppose that the phenomenology of a particular type of conscious thought is the *same* for everyone (any more than an argument for perceptual phenomenology need claim that what it is like to see an object of a particular color is the same for everyone).

(4) To say that conscious thoughts have cognitively proprietary phenomenology is not in and of itself to say that their phenomenology is *exclusively* of a cognitively proprietary sort. The phenomenology of a

---

[9] If you think externalism is true (I do not: see Pitt In Preparation (b)), take the thesis to be that thoughts have their *narrow* propositional contents (e.g., "primary intensions" in Chalmers's (1996; In Preparation) sense) in virtue of their phenomenology, and to allow that thoughts with different phenomenologies may have the same narrow propositional content.

[10] If, again, you think externalism is true, read (P) as allowing that distinct thought tokens with the same representational (narrow) contents (phenomenologies) might (for the same or different thinkers) express different *wide* propositions (e.g., "secondary intensions" in Chalmers's sense).

conscious thought might, for example, be *partly* constituted by some sort of linguistic phenomenology – e.g., auditory or visual syntactic imagery (this possibility is discussed in section 2.2.3, below). (P) does claim, however, that the phenomenology of conscious thought cannot be *completely identified with* any other sort of phenomenology.

(5) For convenience, I shall refer to distinct representational contents expressing the same proposition using distinct sentences that express that proposition. I do not intend by this to prejudge the issue of whether or not representational contents are internal utterances of sentences.

(6) It will no doubt occur to some readers that the claim that the representational content of a thought is its phenomenology immediately raises the issue of the representational contents of *un*conscious thoughts. I address this issue elsewhere (Pitt In Preparation (a)). Here I shall only mention that I do not think the view I defend in this paper commits me to the existence of unconscious phenomenology (though I think there are compelling reasons – e.g., certain sorts of blindsight cases – to think that phenomenology is not always conscious).[11]

The argument for (P) is as follows.[12]

Normally – that is, barring confusion, inattention, impaired functioning, and the like – one is able, consciously, introspectively and non-inferentially (henceforth, "Immediately") to do three distinct (but closely related) things: (a) to distinguish one's occurrent conscious thoughts from one's other occurrent conscious mental states; (b) to distinguish one's occurrent conscious thoughts each from the others; and © to identify each of one's occurrent conscious thoughts as the thought it is (i.e., as having the *content* it does[13]). But (the argument continues), one would not be able to do these things unless each (type of) occurrent conscious thought had a phenomenology that is (1) different from that of any other type of conscious mental state (proprietary), (2) different from that of any other type of conscious thought (distinct), and (3) constitutive

---

[11] Some other pressing questions are: how does a view of the sort I am arguing for account for the *compositionality* of thought contents; and: *how* does phenomenology determine content? I intend to address these questions in future work.

[12] Cf. Flanagan 1992: 69. An argument of this sort is also considered, and dismissed, in Nelkin 1996 (142-43). I think Nelkin's dismissal is too quick. I address the reason he gives for rejecting this argument (viz., that he detects no cognitive phenomenology in his own case) in section 3, below.

[13] Boghossian (1994), following Dummett 1978, refers to this as the "transparency" of mental content.

I take it that the claim of Immediacy commits me to neither infallibility nor self-omniscience; and neither is required by the present argument. What I am committed to is the claim that we are, typically, able to identify and distinguish our conscious thoughts (at least with respect to their representational content) Immediately, and that any entertainable thought content is at least potentially Immediately available to its possessor.

of its (representational) content (individuative). That is, it is only because conscious thoughts have a kind of phenomenology that is different from that of any other kind of conscious mental state that one can Immediately discriminate them from other kinds of conscious mental states; it is only because type-distinct conscious thoughts have type-distinct phenomenologies (of the cognitive sort) that one can Immediately distinguish them from each other; and it is only because a conscious thought that *p* has a phenomenology that constitutes its (representational) content that one can Immediately identify it as the thought it is. Hence (the argument concludes), each type of conscious thought has a proprietary, unique phenomenology, which constitutes its representational content.

In brief:

(K1) It is possible Immediately to identify one's occurrent conscious thoughts (equivalently (see below): one can know by acquaintance *which* thought a particular occurrent conscious thought is); but

(K2) It would not be possible Immediately to identify one's conscious thoughts unless each type of conscious thought had a proprietary, distinctive, individuative phenomenology; so

(P) Each type of conscious thought – each state of consciously thinking that *p*, for all thinkable contents *p* – has a proprietary, distinctive, individuative phenomenology.

The two sorts of abilities (a) and (b) are analogous to what Dretske (1969) calls "non-epistemic seeing," or (1979) "simple seeing," and the ability © (which presupposes the other two) is analogous to what Dretske (1969) calls "epistemic seeing" (with an important qualification, to be discussed below). For Dretske, an object *O* is simply seen by a subject *S* if *S differentiates O* from its immediate environment *purely* on the basis of how *O looks to S* (how it is *visually experienced* by *S*), where an object's looking some way to *S* neither presupposes nor implies that *S* has any beliefs about it.[14]

Dretske argues that we must suppose there is such a thing as simple seeing given that objects that cannot be visually identified may nonetheless be seen – i.e., visually discriminated from their immediate environment. One need not be able to identify what one is seeing (know what it is) in order to be able to distinguish it from its environment.

---

[14] There is an important difference between (1) *S*'s differentiating *O* from its immediate environment and (2) *O*'s being differentiated from its immediate environment in *S*'s experience. The former requires that *S* be *attending to O*, while the latter does not. (Consider Armstrong's inattentive driver (Armstrong 1968). Presumably, objects are differentiated from their immediate environment in his experience – he is not, after all, *unconscious*, and it is reasonable to suppose that his articulated experience of his environment explains his ability to keep his car on course. *He* does not differentiate the objects of his experience from their immediate environment, however, since he is not paying attention.)

What Dretske (1969) calls "primary epistemic seeing," in contrast, necessarily involves belief, and, Dretske argues, amounts to *knowledge*: one may know that an object is *F* by *seeing that* it is *F*. According to Dretske (*id*.: 79-88), *S sees that* an object *O* is *F* only if (I) *O* is *F*, (ii) *S* simply sees *O*, (iii) the conditions under which *S* simply sees *O* are such that it would not look to *S* as it does unless it were *F*, (iv) *S* believes that the conditions in (iii) obtain, and (v) *S* believes that *O* is *F*. When conditions (I)-(iv) are satisfied, *S* has a *conclusive reason* for believing *O* to be *F*. Hence, if condition (v) is also satisfied, *S knows that O* is *F* by *seeing that O* is *F*. (There are, of course, *other* ways to know that *O* is *F*.) *O*'s looking the way it does to *S* provides *S* (in the relevant circumstances) a conclusive reason for believing that *O* is *F*. Thus, to *see that O* is *F* is to believe that it is *F because* of the way it looks.[15]

Dretske's account of simple and primary epistemic seeing may be generalized to other modes of perceptual experience. Just as one may *see that* an object is brown (or rotten), one may (though this phrasing might be uncommon) *smell that* it is musky (or dead), *taste that* it is sour (or unripe), *hear that* it is loud (or hollow), or *feel that* it is rough (or broken).[16] I take it such modes of primary epistemic perception are forms of what Russell (1910-1911) calls *knowledge by acquaintance*, and simple perception a form of what we can call *simple acquaintance* (or, simply, *acquaintance*): knowledge by acquaintance is knowledge of the properties of an object *O* grounded, in the way described above, in simple acquaintance with *O*; simple acquaintance with *O* is attentive discriminating experience of *O*. Hence, in what follows I shall use the terms *acquaintance* and *knowledge by acquaintance* as general terms for, respectively, attentive discriminating experience and the knowledge based on it in the way Dretske describes for simple seeing and primary epistemic seeing (seeing that).

The argument from Immediate knowledge of content claims that one may also have acquaintance with and knowledge by acquaintance of one's occurrent conscious thoughts. *Introspective acquaintance*, which I take it is what is operative in the abilities described in (a) and (b), above, is a form of simple acquaintance, and may be understood in a way analogous to the way Dretske explains simple seeing. A subject *S* is

---

[15] The belief that *O* is *F* is not *inferred from* the belief that *O* looks some way to *S* and the belief that the conditions in (iii) are satisfied. Seeing that *O* is *F* is *immediate*: "[t]he *immediacy* associated with seeing arises precisely because no intermediate discursive process mediates the seeing of [*O*] (which is [*F*]) and the consequent conviction that (in cases of seeing that [*O*] is [*F*]) that [*O*] is [*F*]." (*Id*.: 127.) (The belief that the conditions in (iii) are satisfied is a "background belief," where background beliefs are *"manifested* in perceptual achievements [like seeing that; but] are not *used* as premises or principles of inference." (*Id*.: 119.))

[16] What one usually says is that one can *tell that* an object is musky, sour, loud or rough, *by* smelling, tasting, hearing or feeling (touching) it.

introspectively acquainted with a conscious mental particular *M* (a state, a thought, an image, a feeling, a sensation, etc.) if *S* differentiates *M* from its mental environment purely on the basis of *how it is experienced* by *S*, where a mental particular's being experienced in some way by *S* neither presupposes nor implies that *S* has any beliefs about it. For *M* to be experienced in some way by *S* is a matter of its qualitative properties – its phenomenology. Thus, one cannot simply introspect a conscious mental particular unless it has some definite phenomenal character – unless, that is, there is something one's introspective experience of it is *like*.

I do not mean to suggest here that simple introspection is simple *perception* of mental particulars, nor that the experience of an occurrent conscious mental particular *M* is a state distinct from *M*. I take it that simple introspection of a conscious mental particular is simply attentive experience of it, and that for a mental particular to be experienced is simply for it to be conscious. Simple perception is attentive experience of external objects; simple introspection is attentive experience of internal objects. But to say this is not to say that conscious mental particulars are the *objects of* introspection in the way that physical particulars are the objects of perception. A perceived external particular (one may suppose) is distinct from an experience of it. An introspected conscious mental particular, in contrast, is *part of* the introspective experience of it: to say that one simply introspects a conscious mental particular is to say that one has a conscious experience of which the mental particular is itself a differentiated constituent. (This is, of course, to be distinguished from introspective *belief* and *knowledge*, which *are* states distinct from the introspected particulars.)

We must suppose that there is such a thing as introspective acquaintance given that mental states that cannot be identified (e.g., unclassifiable, fleeting or vague moods, thoughts, feelings, sensations, etc.) may nonetheless be experienced, and experienced as different from other mental states. And this is not a matter of one having any beliefs about the states, but of the states having distinctive qualitative properties of which one may be introspectively aware.

Knowledge by acquaintance of the identity of a thought – knowing that it is the thought that *p* – is a form of primary epistemic introspective acquaintance. Let us call this form of (introspective) knowledge by acquaintance *grasping*. Thus, just as s*eeing that* is the form of (perceptual) knowledge by acquaintance appropriate to visible objects (*mutatis mutandis* for the other sensory modalities), and *feeling that* is the form of (introspective) knowledge by acquaintance appropriate to conscious sensations, *grasping that* is the form of (introspective) knowledge by acquaintance appropriate to conscious thoughts. To Immediately identify

a thought as the thought that *p* is to grasp that it is the thought that *p* (that it has the content that *p*).[17]

Grasping a thought is different from merely *thinking* it, in the same way that seeing that *O* is *F* is different from simply seeing *O* (which is *F*). It is also different from simple introspection. To know in this way that a thought is the thought that *p* involves having a (higher-order) thought *about* it. What is distinctive about Immediate knowledge of content, on this view, is that the belief that the thought *t* has the content that *p* is conclusively justified by the experience of *t*. The phenomenology of the experience of *t* grounds the ability to grasp that it has the content it does.

Thus, the argument is that Immediate identification of a thought is introspective knowledge by acquaintance (primary epistemic introspection) that it is the thought it is, and that this is not possible without simple acquaintance, which itself depends upon the introspected state having phenomenal character. Immediate identification of a particular thought requires Immediate discriminative awareness of its distinctive phenomenal properties. Since each conscious thought-type has a distinctive phenomenology, there is something it is like to entertain it; and since there is something it is like to entertain it, it is possible Immediately to identify it.

2.1 Individuative phenomenology

These parallel accounts of simple perception and introspection appeal to the ways perceived and introspected objects appear to the perceiver or introspector. But there are crucial differences between the appearance of perceived objects and the appearance of introspected objects. The simple perception of an object that is a necessary condition on perceiving that it is *F* – which involves its appearing in some way to the perceiver – does not necessarily involve its appearing *F* to the perceiver. Perception that an object is *F* might be due to its appearing *G* ($\neq F$) in circumstances under which it would not appear *G* unless it were *F* (cf. Dretske 1969). For example, an apple may look brown in circumstances under which it would not look brown unless it were rotten, and one may thereby be able to see that the apple is rotten on the basis of its looking brown. Indeed, one may be able to see that an apple is *green* on the basis of its looking brown in circumstances under which it would not look brown unless it were green. In general, then, perceived objects can appear to have properties that they do not have, and perceptual knowledge can be based on such false appearances.

---

[17] See Conee 1994 for an argument that all phenomenal knowledge is knowledge by acquaintance. Such knowledge should (*pace* Kant (1781/1929: 35)) be distinguished from knowledge *that one is thinking* the thought in question. What I am claiming I can know by conscious introspection is not *which thought I am thinking*, but, rather, something like *which thought this thought is*.

Moreover, the relation between the ways perceived objects can *appear* and the ways they can *be* is philosophically problematic. An objectivist would maintain that the ways perceived objects can appear are *identical to* ways they can be: the apparent properties of objects are not properties of experiences, but of objects. (Objects may appear to be ways that they are not in abnormal circumstances (as when a green apple appears brown); but the ways they appear in abnormal circumstances are ways they *might be* (apples can be brown), and the ways they appear in normal circumstances are the ways that they *are* (the perceived greenness of an apple in normal circumstances is a property *of the apple*).) Subjectivists, on the other hand, would maintain that the ways perceived objects can appear are properties of experiences, and, at best, functionally dependent upon the ways they can be. (For dispositionalists (e.g., those who identify colors with unperceivable dispositional properties of objects), objects *always* appear to be ways that they are not, whereas for categorical physicalists (e.g., those who identify colors with possibly independently perceivable surface properties), objects *typically* appear to be ways that they are not.)

But there can be no false appearances in the case of conscious mental particulars, and, hence, no introspective knowledge based on false appearances. There are no conditions under which an orange after-image does not look orange; and there are no conditions under which a painful sensation does not feel painful.[18] Thus, Immediate knowledge by acquaintance that an after-image is orange can only be based on its looking orange, and Immediate knowledge by acquaintance that a sensation is painful can only be based on its feeling painful. Though one may be able to see that an apple is green on the basis of its looking brown, one cannot see that an after-image is green on the basis of its looking brown, or feel that a sensation is a pain on the basis of its feeling pleasurable. In general: necessarily, if a conscious mental particular *is F* then it *appears F*; hence, knowledge by introspective acquaintance that a conscious mental particular is *F* can only be based on its appearing *F*.

Furthermore, since there is no distinction between a conscious mental particular and the experience of it, the question of the relation between the way a conscious mental particular appears and the way it is has only one possible answer: the way it appears *is* the way it is. Naive realism is the only possible view of the apparent properties of conscious mental particulars: they are properties of the particulars themselves. If an after-image *looks* orange, then it *is* orange, because its looking orange and its being orange are the same property; and if a sensation *feels* painful, then it *is* painful, because its feeling painful and its being painful are the same property.

---

[18] This is a familiar point. (See Kripke 1980.)

These facts about conscious mental particulars do *not*, as is sometimes supposed, imply either introspective infallibility or omniscience. It simply does not follow from the fact that conscious mental particulars cannot appear other than as they are that one's *beliefs* about the way they are/appear cannot be mistaken – any more than it would follow from an external object's necessarily having the properties it appears to have that one's knowledge of *its* properties is infallible. Nor do these facts about conscious mental particulars imply that one is omniscient about the contents of one's conscious mind. It is perfectly consistent to suppose that one has simple introspective acquaintance with conscious mental particulars about which one has no beliefs (or knowledge) at all.

It *is* implied, however, that to know by introspective acquaintance that a conscious particular has a property is to believe that it has that property while one is attending to it, and, moreover, that the properties one is acquainted with in introspection on the basis of which one knows *what* a conscious mental particular is (an orange after-image, a painful sensation, the thought that snow is cold) are *individuative* – i.e., they *constitute* the particular's being the sort of particular it is.

Since conditions are *always* such that a conscious mental particular would not appear *F* unless it were *F*, Dretske's condition (iii) is superfluous, and his condition (iv) drops out.[19] Hence, knowledge by conscious introspective acquaintance that a mental particular is *F* consists in satisfaction of Dretske's conditions (I), (ii) and (v) – viz., believing that it is *F* while it is the object of conscious attention.

Furthermore, since only phenomenal properties can be conscious, and conscious introspective discrimination of a mental particular is discrimination of its conscious properties, it follows that only the phenomenal properties of conscious mental particulars are simply introspectable. If, therefore, conscious introspective acquaintance with a mental particular provides conclusive justification for believing that it is the particular it is, its having the phenomenal properties it does must constitute its being the particular it is: that is, its phenomenal properties must be individuative. One can know Immediately that a mental particular is a sensation of pain or an orange after-image only because pains and orange after-images have phenomenal properties that are introspectably discriminable from those of other type-distinct mental particulars and which *make them* the particulars they are. The argument from Immediate knowledge of content claims that the same is true of conscious thoughts: one can know Immediately that a mental particular is the thought that

---

[19] If (iii) is necessarily satisfied, one need not believe it is in order to be conclusively justified in believing that a conscious mental particular is *F*. In the case of perception one is justified in believing that *O* is *F* only if one believes that the conditions in (iii) obtain because it is possible that conditions are not such as described in (iii).

snow is cold only because thoughts that snow is cold have phenomenal properties that are introspectably discriminable from those of other type-distinct mental particulars and which *make them* the particulars they are.

2.2 Objections to (K1)

The argument (K1)-(P) is obviously valid; so the way to resist its conclusion is to resist one or both of its premises.

(K1) seems very hard to deny. One does not normally have to *infer* what it is one is occurrently consciously thinking: one can know what one is thinking simply by attending to the contents of one's mind. And it seems entirely inappropriate to call on someone to provide inferential justification of his attributions of self-knowledge of content.[20] Hence, a very strong motivation would be needed for rejecting (K1). I shall argue in this section that there is none.

According to (K1), the identification of any occurrent conscious thought as the thought that *p* requires neither observation nor inference: simply by attending to it one can know which thought it is. Hence, the denial of (K1) would be that it is not possible to identify a conscious occurrent thought in this way. There are three ways to support this claim. The first is to establish that the identification of an occurrent conscious thought necessarily requires observation; the second is to establish that it necessarily requires inference; and the third is to establish that it necessarily requires both observation and inference.[21]

Externalist considerations are perhaps the most obvious way to try to motivate the claim that the identification of an occurrent conscious thought as the thought that *p* cannot be Immediate, since the identification of any thought always requires both observation and inference.[22] In order to know that a thought is the thought that, e.g., water is a liquid (that arthritis is a disease), and not that twater is a liquid

---

20 Cf. Boghossian 1989 (152). Boghossian argues that "the epistemic norms governing ascriptions of self-knowledge do not require supplementary [i.e., observational or inferential] evidence ...."

21 There are also strong and weak versions of each of these claims, depending on whether it is held that thought identification is *entirely* observational, inferential or observational/inferential, or only *partially* observational, inferential or observational/inferential. I shall not address all of these possible objections individually here. Instead, I shall consider in detail what I take to be the best-motivated and most familiar of them, viz., that self-knowledge of content is always partially observational and partially inferential, and then discuss some considerations that seem to me to militate, individually and severally, against all of the others.

22 Conceptual/inferential-role theories on which self-knowledge of content is essentially inferential (i.e., most of them) are addressed in note 28, below. Those that are not "solipsistic" (e.g., Harman 1987) are of a piece with the externalist views here under discussion. Moreover, I agree with Boghossian (1989: 152) that the Rylean view that one comes to know one's own thoughts in the way one comes to know those of others – viz., via observation of one's behavior – "carries no conviction whatever."

(that tharthritis is a disease), I must make an inference from known facts about how content is determined and about my natural (social) environment. That is, I can identify a thought *only* by consulting background knowledge and (together with externalist principles of content determination) drawing a conclusion from it. Moreover, since knowledge of the background environmental (social) facts that determine thought contents is not obtainable introspectively, neither is knowledge of those contents themselves. One must consult external sources for information about the properties that determine the contents of one's thoughts. Further, this is consistent with allowing that one's knowledge that a conscious thought is occurring is both introspective and non-inferential: externalism only shows that knowledge that a thought is the thought *that p* requires observation and inference. Hence, contrary to (K1), the identification of a thought is never Immediate, but is always a complex process involving observation and inference as well as introspection and direct apprehension.

However, to base the denial of (K1) on the claim that externalism is true and implies that knowledge of content is observational and inferential is to put it on shaky ground.

There is considerable controversy over whether or not externalism has this consequence.[23] Most externalists – e.g., Burge (1988), Davidson (1984) and Heil (1988) (to name just a few) – think that it does not. They hold that a thinker need not know the causal history (or any other relational properties) of a thought in order to know its content, since that history also determines the content of the thought that he is having that thought. Thus, the contents of the first- and second-order thoughts will always necessarily correspond, and the second-order thought will automatically be true.[24] Since one need not know the causal history of either thought in order for their contents to correspond, neither inference nor observation is required by externalism for self-knowledge of content.

Others, however, think this response is inadequate. Boghossian (1989; 1994), for example, has argued that it can account for neither self-knowledge of thoughts not simultaneous with such second-order

---

[23] See, e.g., the essays in Ludlow and Martin 1998 and Wright, Smith and Macdonald 1998. There may be further controversy over whether or not the contents of *all* types of thoughts (including, e.g., those with logical or mathematical or moral content) are determined externally. If they are not, then some other way of motivating the denial of (K1) would have to be found for such thoughts. (It seems to me, however, that externalist thought experiments (of the Putnamian or Burgean type) can be formulated for any sort of content whatsoever. See Pitt 2000: 234.)

[24] Burge (1988) stresses that the second-order thought *I am thinking that p* contains the first order thought *p*, and so is self-verifying.

judgments (i.e., the object of the judgment that *I just now thought that p*; cf. note 26, below) nor for the ability to *distinguish* thoughts in virtue of their contents.[25]

Whether or not this controversy is eventually decided – and however it may turn out if it is – it is not at all advantageous to the opponent of (K1) to pin hopes on it now. At least at present, the claim that externalism implies that (K1) is false has, *prima facie*, far less going for it than (K1) itself. (Boghossian (1989) takes the incompatibility of externalism and Immediate knowledge of content to be something like a *dilemma*.)

In any case, the account of self-knowledge of content one would be left with if (K1) were to be abandoned is simply unbelievable. For, surely, one's knowledge that one is occurrently consciously thinking a thought is Immediate if anything is. However, assuming the externalism-inspired story just sketched, one could have such knowledge without having any idea *which* thought one is thinking, and, furthermore, be completely incapable of determining which thought it is no matter how much one introspects.

All of the other ways of objecting to (K1) seem to me to founder upon one or more of the following considerations.

It could not be the case that thought-identification is *entirely* observational. One's thoughts are not (or, in any case, not entirely: some externalists believe that thoughts *include* external objects) in the external world to be observed. In order to know that one's (occurrent, conscious) thought is the thought that *p*, one must know that it is occurring. But such knowledge could not be observational; it is *essentially* introspective.[26] Thus, self-knowledge of content always requires at least some introspection.

Moreover, given that it is, as argued above, highly implausible to suppose that the non-inferential component of self-knowledge of content is exclusively observational, it follows that at least *some* self-knowledge of content must be introspective. But, I shall now argue, it cannot be that *all* introspective knowledge of content is inferential, on pain of vicious infinite regress.

On an inferentialist account, to know that a thought *t* is the thought that *p* I must infer the thought *t is the thought that p* (call this thought *t'*) from some other thoughts, *q, r*, ... However, if I am to know by introspection that a thought is the thought *that p*, then I must also know by conscious introspection that I

---

[25] Of course, if Burge, Davidson and Heil are right, and externalism does not imply that self-knowledge of content requires observation or inference, then it does not support the rejection of (K1).

[26] I might reason that if I *were* to think a thought (one I am *not* at the moment thinking) that I *would* express by uttering 'water is wet', then I *would* be thinking that H_2O is wet – purely on the basis of observation. (Someone else could obtain this knowledge about me in this way.) But this is, by stipulation, *not* knowledge of the content of a *conscious occurrent thought*, which is what the argument of this paper is concerned with. It does not claim that we can have Immediate knowledge of the contents of thoughts we are not thinking.

have performed the inference *q, r, ..., therefore, t'*. And I cannot know introspectively that I have performed this inference unless I can introspectively identify it. But I cannot identify the inference unless I can identify its constituent thoughts, *q, r*, ... and *t'*. And since introspective identification of thoughts requires inference, an infinite number of inferences will have to be performed in order to achieve introspective knowledge of the content of any thought. Such knowledge would therefore be impossible. But, *ex hypothesi*, it is not; so it must be possible to identify by introspection at least some thoughts non-inferentially.[27]

It may be objected that knowledge that I have drawn the primary inference (i.e., an inference of the form *q, r, ...; therefore, t'*) is required to *know that I know* what the content of my thought is, but is not required for first-order knowledge. The thought that a particular thought has a particular content has to *have* the property of being inferred from some other thought or thoughts, but I need not *know* that it has this property in order for it to count as knowledge. So I need not know that I have drawn the inference, and the regress does not get started.[28]

The objection fails to distinguish between *knowing that* one knows something and *knowing what* one knows. Whereas the former implies the latter, the latter does not imply the former. One can know what it is one knows – i.e., what the content of the state that counts as knowledge is – without knowing that one knows it. (Indeed, one might even *doubt* that one knows it.) But the latter alone is sufficient to generate the inferentialist regress: you have to know that the inference has been drawn in order to know *what* you know.

Moreover, the knowledge generated on the alternative account would be *introspective* knowledge in name only. If introspection is to be a *source* of knowledge – that is, if one is to know *by virtue of* introspection the identity of a particular thought – then there must be introspectively accessible properties of the thought detection of which is sufficient to ground knowledge of its content. According to the inferentialist, however, there are no intrinsic properties (introspectable or otherwise) of a thought knowledge of which is sufficient for knowledge of its content. What is sufficient to ground such knowledge is that the

---

[27] In fact, I think this argument shows that *any* thought must be identifiable non-inferentially, for reasons I give below.

This argument is similar to one given in Boghossian 1989 (10). It differs from Boghossian's in its emphasis on conscious introspection and its requirement of introspective knowledge that the relevant inference has been made. Boghossian requires that the premises of the inference *themselves* be known; the present argument requires only that it be known what the *content* of the premises is.

[28] Boghossian (1989) does not address the parallel objection to his argument, which assumes justificatory internalism. (An assumption I do not make: introspection does the work in my argument that internalism does in Boghossian's.) Of course, if one does assume justificatory internalism, the present objection has no force.

thought *t is the thought that p* be inferred from some other thought or thoughts. But, then, if introspection is to deliver knowledge, and if it is not by virtue of introspective knowledge of my having made the inference, then it is not by virtue of *anything*. In which case I have no knowledge at all. So it follows that if conscious introspective knowledge of content is possible, then, given the inferential theory, it must be conscious introspective knowledge that an inference of the form *q, r, ...; therefore, this thought is the thought that p* has been made; and this leads to the regress. Hence, it cannot be the case that introspective self-knowledge of content is entirely inferential.[29]

Finally, there can be no *a priori* restrictions on which thoughts could play an inferential role with respect to a thought to be identified. That is, it would seem that *any* thought might, in principle, figure inferentially in the identification of some other thought. But if this is the case, then, since at some point in a process of introspective/inferential thought identification inference-constitutive thoughts must be known non-inferentially, a non-regressive inferentialism would allow that any thought could be identified introspectively and non-inferentially – which is, of course, just what (K1) claims.

I conclude that there is no plausible way to motivate the denial of (K1).

2.3 Objections to (K2)

2.3.1. Perhaps a better objection to the argument would attack the second premise, by offering an alternative explanation of Immediate knowledge of content. (K2) claims that we could not Immediately identify and distinguish our conscious thoughts unless they had distinctive phenomenologies: the phenomenology of a conscious thought can provide immediate, conclusive justification for believing that it is the thought it is. Hence, if there is an account of how we could do it that does not appeal to phenomenology, (K2) is false. The following familiar sort of view is perhaps the best candidate.

My introspective knowledge of the content of a particular thought *t* is explained by appeal to a reliable process of belief formation – for example, the activity of a mechanism whose job it is to deposit a token of a thought *t'* with the content *the content of t is that p* in my "belief-box" whenever, or *because*, *t* consciously

---

[29] The considerations in this passage militate equally against theories on which the relevant inferences are from premises about, *inter alia*, one's environment *and* those on which inferential relations to other thoughts are constitutive of content (i.e., conceptual- or inferential-role theories). The general point is that an object's relational properties are not detectable by mere inspection of it (cf. Boghossian 1989: 162).

Note that a regress problem would also beset any attempt to characterize knowledge of content in behavioral or dispositional terms. Knowing that one is thinking that *p* in virtue of knowing that one is doing or saying, or is disposed to do or say, certain things is knowing that one is in certain intentional states, whose content one is, *ex hypothesi*, capable of distinguishing introspectively. (It is not enough simply to *be* disposed to do or say certain things. In order to have Immediate knowledge of content (on this view), one would have to *know what* one is disposed to do or say.)

occurs. Such a state is functional/computational: to know that a thought is the thought that *p* (and not that *q*) is just to have the appropriate higher-order thought reliably tokened in one's belief-box. Moreover, this process does not require that the thought that *t* have phenomenology of any kind: even supposing *t* did have phenomenal properties, they would not be what justified the belief that *t* has the content that *p*.

This counts as introspective knowledge of the content of *t* because *t* is a mental particular, the content of *t′* is that the content of *t* is that *p*, and *t′* plays the role characteristic of a belief due to the activity of a reliable mechanism. What justifies the belief that *t* is the thought that *p* is not the experience of *t*, but the etiology of the belief about *t*. And since this belief is *caused*, not inferred, such knowledge is non-inferential. Finally, when *t′*, so caused, is conscious, knowledge of the content of *t* is Immediate (conscious, introspective and non-inferential).

But to think that *t* is the thought that *p* while *t* is occurring – even *because t* is occurring – is not to *identify it* as the thought that *p* in the sense at issue in this paper. Introspective identification of occurrent conscious thoughts is analogous to perceptual identification of objects and introspective identification of sensations: it is a form of knowledge *by acquaintance*. Such identifications have the canonical form *this [that] is (an/the) F*; they require simple acquaintance, in the relevant mode, with the object identified. That is, the object identified – "this" – must be *experientially discriminated by the perceiver from its environment*. And, as pointed out above, this requires that the object appear to one in some determinate way, and that one be attending to it. One cannot, say, visually identify a thing as a dog (see that it is a dog) unless one has an attentive discriminative visual experience of it – that is, unless one *simply sees* it (in the sense of section 2). Merely to think *this animal is a dog* when a dog is within visual range and is causing one to have a visual experience of it and to think that it is a dog, is not to *see that* it is (visually identify it as) a dog. Likewise, merely to think that *s* is a pain when *s* is a pain, is occurring, and is causing one to think that it is a pain, is not to *feel* that it is (introspectively identify it as) a pain. And merely to think that *t* is the thought that *p* when *t* is the thought that *p*, is occurring, and causing one to think that it is the thought that *p*, is not to *grasp that t* is (introspectively identify it as) the thought that *p*.

In addition, if the conscious occurrence of *t′* is not sufficient to ground conscious knowledge of *its* content (i.e., does not provide conclusive reason to believe that its content is that the content of *t* is that *p*), then one could consciously think it without thereby being able to know which thought it is. But if one cannot know thereby which thought it is – what its content is – then its conscious occurrence cannot ground

conscious knowledge of the content of *t*. This is because it is supposed to be *by virtue of* the fact that *t′* is conscious (and reliably caused) that its occurrence can ground knowledge *of which thought* t *is*. That is, *t″*'s conscious occurrence is supposed to ground conscious knowledge of the content of *t because* (it is reliably caused and) its content is that the content of *t* is *p*. But if *t″*'s being conscious is not sufficient to ground knowledge that it has this content – i.e., if it is not sufficient to ground conscious knowledge of *which* thought it is – then its being conscious does not explain the conscious introspective identification of *t*.

To put the point slightly differently, *t″*'s consciousness is supposed to make the content of *t* Immediately knowable. But this is because the content of *t′* is that the content of *t* is *p*. Hence, if the consciousness of *t′* is not sufficient to make its content Immediately knowable – i.e., to make it Immediately knowable that it says that the content of *t* is *p* – then it is not sufficient to ground Immediate knowledge of *t*'s content.[30]

Hence, conscious occurrence of *t′* must, if it is to be sufficient to ground Immediate knowledge of the content of *t*, be sufficient to ground conscious knowledge of its own content as well. Since the theory under consideration denies this, it is false.[31]

---

[30] This argument must not be confused with the following, subtly different and fallacious one:

> *t″*'s conscious occurrence could not account for Immediate knowledge of the content of *t* if its conscious occurrence were not sufficient to ground Immediate knowledge that the content of *t* is that *p*. But *that the content of t is that p* just is *t″*'s content. Hence, the conscious occurrence of *t′* must, if it is to be sufficient to ground Immediate knowledge of the content of *t*, be sufficient to ground Immediate knowledge of its own content.
>
> *This* argument equivocates on the phrase 'knowledge of the content of *t*', which may be read as meaning either (where *p* is the content of *t*) *knowledge that p*, or as *knowledge that the content of t is that p*. The equivocation occurs between the first premise, in which 'knowledge of the content of *t*' has the former reading, and the conclusion, in which it has the latter reading. The same equivocation is more obvious in the following: *A* knows what the content of the thought *t* is. The content of the thought *t* is (the proposition that) *p*. Therefore, *A* knows that *p*. The equivocation is fallacious since it is false that knowing in the first sense is, or is entailed by, knowing in the second sense. (Arguments that equivocate are not *ipso facto* fallacious, since it is possible for a term τ to have two senses, σ_1 and σ_2, such that σ_1 *entails, or is entailed by,* σ_2, and for this entailment to be presupposed by, or a tacit premise of, the argument.)
>
> The argument in the text, in contrast, neither equivocates nor depends upon the false principle that knowing that the content of *t* is that *p* entails (or is entailed by) knowing that *p*.

[31] It might be argued that *t*'s being conscious is sufficient to ground Immediate knowledge of *t*'s content because what makes *t* conscious is the higher-order thought that *t* has the content it does (cf., e.g., Rosenthal 1986). But it would not follow from this that *t*'s being conscious is sufficient to ground *conscious* knowledge of its content, since *t′* itself could still be (and, according to Higher-Order Thought theories of consciousness, typically is) *not* conscious. It is *t′* whose (properly caused) occurrence is supposed to ground knowledge of the content of *t*; so if *t′* is unconscious, the knowledge is unconscious – even if *t* itself is conscious.

Again, the issue is not how one *knows* what is in one's belief box (the etiology of its occurrence might take care of that), but how one knows *what* is in one's belief box – that is, how one knows what the *content* of one's occurrent conscious thought is. And what I am arguing is that the object of knowledge would not be Immediately identifiable as having *any* property *F* if its *F*ness were not *directly presented* in consciousness. But phenomenal properties are the *only* properties that can be directly presented in consciousness. Hence, Immediate knowledge of the content of an occurrent conscious thought depends on sensitivity to its distinctive phenomenal properties.

2.3.2. Perhaps one could give a deflationary account, not of the claim *that* there is something it is like to think that *p*, but of *what* it is like to think that *p*.[32] On such an account, the direct presentation of a property in consciousness would not require any distinctively mental qualitative character, or quale.

I can think of two ways to develop this objection. The first models itself on the following argument (due to Harman (1990); see also Dretske 1999, Lycan Forthcoming and Tye 2000) that no awareness of phenomenology is involved in being in a *perceptual* state. What one is aware of when one sees something, for example, is not visual phenomenal qualities inhering in one's own mind, but objective qualities of the thing seen (perceptual experience is "transparent"). The blueness we experience when we look at the sky is the blueness *of the sky*, not of our perceptual experience of it. One need not postulate subjective phenomenal characters – purely mental properties, problematically related to the objective properties of extramental objects – in order to explain the qualitative aspects of perceptual experience. So, analogously (one might argue), it is not necessary to postulate phenomenal characters in the case of introspective awareness of thought content either. What one introspects is just the contents of one's thoughts. What distinguishes the contents of my thoughts for me is nothing about me, or my experience, but something about them. Thinking that *p* is distinguishable from thinking that *q* because the contents *that p* and *that q* are distinct – just as seeing the blue of the sky is distinguishable from seeing the green of an apple because the blue *of the sky* and the green *of the apple* are distinct. Thus, one may accept (P), but interpret it in such a way as to block the attribution of phenomenal character to thinking, and locate whatever detectable differences there are between thinking that *p* and thinking that *q* in the objects of the thinkings – i.e., the contents themselves. To quote Worley:

> It might also be claimed that *t'*s being conscious makes only *part* of its content consciously available – i.e., the content that *the content of t is that p* (as opposed to the content that *the content of t' is that the content of t is that p*). But it is difficult to see what the nonquestion-begging justification for this claim could be.

[32] I owe this objection also to Sara Worley and, independently, to Mark van Roojen.

To have a conscious thought is to be consciously aware of the contents of the thought. There is no qualitative experience over and above the awareness of the content of the thought. ... the qualitative difference between thinking that today is Wednesday and thinking that it is Thursday ... is that one is aware of different contents in the two cases. That difference exhausts the difference between the two thoughts.[33]

But this move is ineffective. For one thing, as noted above, there is no phenomenological distinction between a conscious thought and the experience of it, any more than there is a difference between a conscious pain and the experience of it. But, even if there were such a difference, perception is disanalogous to introspection in a very crucial respect. The objects of perception may be mind-independent objects; but what one introspects certainly are not. When we introspect, we turn our attention inward, toward the contents of our minds – which are mental if anything is. Thus, to be a direct realist about introspectable properties *is* to recognize subjective phenomenal characters. To deny qualitative character to conscious experiences would be like granting that there is a perceptible difference between seeing blue and seeing green, but denying that the blue and the green are in *either* the experience *or* the objects of experience. And this hardly seems credible.[34]

The second sort of deflationary move is more radical. It claims that what we are aware of when we are aware of the contents of our thoughts has no qualitative character whatsoever. The first deflationary strategy accepted that one introspects qualitative character, but claimed that this is the character of what is introspected, not of the experience of introspecting. I argued that this is a distinction without the desired difference. The second strategy is to say that thoughts' being consciously distinguishable implies no qualitative character *anywhere* – either in what is experienced or in the experience itself. What distinguishes our thoughts qua objects of introspective awareness is *just* their contents: I am able to tell that a thought is the thought that p, and not that q, because the content of the thought is that p, and not that q. And that's all there is to say about it.

---

[33] This passage is from Worley's comments on my APA presentation. I think there is some vagueness here: Does '[t]here is no qualitative experience over and above the awareness of the content of the thought' involve the claim that thought contents have qualitative character, or not? In the paragraphs to follow I assume the reading on which it does.

[34] Perhaps Worley has in mind the external perception of the objects and properties that constitute the (Russellian) proposition that is the propositional content of the thought. However, since such propositions are *sets* (ordered n-tuples), and sets are not perceivable, the contents of one's thoughts (on this view) are not perceivable either. Moreover, there are obvious, similar problems with this approach as applied to thoughts about abstract objects. Finally, unless the propositions are thought of (à la McDowell 1986) as *part of* the thoughts whose contents they are, this is not introspection. (In any case, as outlined above, there are plenty of reasons for doubting that externalism has these sorts of consequences for self-knowledge of content.)

This is as credible as the claim that there can be perception of, say, the colors of objects, but that perception of these involves neither objective nor subjective qualitative properties. Which is to say, not credible at all.

If my arguments so far are sound, and there *is* something it is uniquely like to think that *p*, then an objector must focus on the other claim embedded in the second premise of the argument from Immediate knowledge of content – viz., that the phenomenology of thinking that *p* is *proprietary* – i.e., that it is neither sensory, nor proprioceptive, nor emotional, etc.

2.3.3. (K2) claims that if each type of conscious thought did not have a proprietary, distinctive, individuative phenomenology, then it would not be possible Immediately to identify and distinguish one's occurrent conscious thoughts. I have so far only considered objections to the claim that phenomenology is necessary for Immediate identification of conscious occurrent thoughts. I shall now consider three objections to the claim that such phenomenology is neither distinctive nor proprietary.

Lormand (1996) claims that conscious attitudes, thoughts and moods do not have phenomenology of a proprietary kind, since "[e]xcluding what it's like to have accompanying [phenomenal] states, ... typically there seems to be nothing left that it's like for one to have a conscious [thought] that snow is white" (247). But he provides no argument that the only sorts of phenomenology that could be assigned to conscious thoughts are those of the familiar sorts he lists (perceptual representations, bodily sensations, images or inner utterances); he simply asserts it. The possibility of a distinctive, cognitively proprietary kind of phenomenology is not even considered.[35]

Lormand offers as an *explanation* of the fact that conscious attitudes, moods and thoughts are not phenomenal that they exhibit neither of a pair of correlative "inner-perceptual illusions" which are necessary for phenomenality. In the "image illusion," mental states seem to have properties of non-mental objects (they are *sharp*, *round*, *red*, ...). In the "appearance illusion," non-mental objects seem to have properties of mental states (they are *sweet*, *cold*, *delicious*, ...). Conscious thoughts are not phenomenal because they are subject to neither of these illusions.

Though these considerations are offered by Lormand as an explanation of – not an argument for – the non-phenomenality of conscious thought, it might be thought that an argument could be based on them, viz.: If a conscious state is phenomenal, then it is subject to one of the inner-perceptual illusions; conscious thoughts are subject to neither; hence, conscious thoughts are not phenomenal. But such an argument would simply beg the question against the proponent of a proprietary cognitive phenomenality of conscious thought.

---

[35] If Lormand is making a *Humean* case against cognitively proprietary phenomenology, my reply can be found in section 3, below.

The inner-perceptual illusions may apply to the familiar kinds of phenomenology; but what is the argument that they apply to *any* sort of phenomenology?

The third sort of objection to (K2) grants that conscious thoughts have qualitative character, and that it is in virtue of Immediate awareness of that character that one knows what one is thinking, but denies that it is *proprietary*. The phenomenology of cognition is, the objection goes, that of either sensory experiences or imagery associated with characteristic linguistic modes of expressing the thought, or emotional experiences triggered by the content of the thought – or, perhaps, a particular kind of proprioceptive or other experience associated with conscious propositional attitude *relations* (fearing, hoping, worrying, etc.). If there is a phenomenology of thinking, it is not a *content*-phenomenology, but merely an associated phenomenology of some familiar kind. And if we do have knowledge of the contents of our thoughts in virtue of introspective awareness of qualitative states, it is just these kinds, none of which are distinctive of cognition *per se*, that explain it.

No doubt the most promising of these familiar sorts of phenomenology is the first. The claim would be that it is either phonological or orthographic imagery (or both) that we introspect when we attend to our thoughts, and that we know what we are thinking in virtue of introspective acquaintance with the phenomenal properties of such imagery. Thinking is just "hearing (or seeing) words in our heads." (Indeed, as is well known, some philosophers have been tempted into *defining* thought as inner speech.[36]) What goes on when one thinks that *p* is, according to this account, among other things a replay of auditory or visual experiences of some sentence- form in one's language that expresses the content that *p*. There need be no proprietary cognitive phenomenology telling you what you are thinking: hearing or seeing a sentence in your head is sufficient.

But sentences in one's head are no more thoughts than sentences on paper. One can host an inwardly tokened (natural-language) sentence without thinking the thought it conventionally expresses – or, indeed, without thinking anything at all: imagine a child of three, or a monolingual speaker of Hindi, replaying an utterance of 'muons and neutrinos are leptons' in her head. And if the tokening of a sentence in one's head is not sufficient for thinking the thought it expresses, then introspective awareness of the phenomenal properties of the sentence-image cannot be sufficient for introspective awareness of what is thought. Thus, even if thoughts are, as a matter of fact, usually or even always accompanied by some kind of auditory or visual imagery (so that such imagery is *part of* the total phenomenology of a conscious thought) – or even if such imagery is necessary for thinking a thought – it is not identical to the thought, and it is not in virtue of introspecting such experiences that one introspectively knows what one is thinking.

---

[36] Carruthers 1996 defends such a view.

A more articulated version of this objection holds that what is going on in the phenomenal mind when one has a conscious thought is production or replay of phonological or orthographic imagery *together with the thoughts that the imagery is imagery of a sentence in a particular language, and that that sentence has a particular meaning*. One knows what thought one is thinking because one knows what sentence one is inwardly tokening, and what it means. There is no proprietary qualitative experience of "grasping" the meaning of that sentence.

But note that all of the work here is being done by knowledge of the meaning of the sentence. Since the mere inward tokening of a sentence-form does not constitute thinking the thought it conventionally expresses, Immediate identification of the phenomenal properties of the sentence-image does not constitute Immediate knowledge that one is thinking that thought. It is only in grasping what the sentence means that one knows what thought one is thinking. But grasping what the sentence means entails thinking the thought it expresses (one thinks, e.g.: *the sentence 'snow is cold' means that snow is cold*). And if such knowledge is to be Immediate, the thought must be known Immediately, as well. (Conscious introspection of the phonological/orthographic imagery alone could not deliver Immediate knowledge of the meaning of the sentence. So if conscious introspection delivers such knowledge, it must be in virtue of conscious introspective identification of the relevant thoughts.) But if this does not simply presuppose the ability to do what the account is supposed to explain, it requires still further inward utterances and content-identifications. Hence, the account is either circular or viciously regressive.

A subtler objection along these lines would claim that introspective knowledge by acquaintance that a thought is the thought that *p* is like perceptual knowledge by acquaintance that an apple is rotten. Rottenness need not be a distinctive perceivable property in order for an apple to be perceived to have it. To perceive that an apple is rotten is to believe that it is rotten on the basis of its appearing in ways (brown, wrinkled, pulpy, malodorous, ...) it would not appear unless it were rotten. In general, perceiving that *O* is *F* may be either *direct* (i.e., based on its appearing *F* in circumstances under which it would not appear *F* unless it were *F*) or *indirect* (i.e., based on its appearing *G* (*H*, *I*, ...) ($\neq$ *F*) in circumstances under which it would not appear *G* (*H*, *I*, ...) unless it were *F*). Properties that are not directly perceivable (including those that are not perceivable *at all*) may nonetheless be indirectly perceivable; and those that are not directly perceivable may be thought of as *manifesting themselves* in those that are.

Likewise, *introspective* knowledge by acquaintance may be either direct or indirect. Hence, one need not assume that the (representational) content of a thought is itself a phenomenal property in order to explain

how one could have such knowledge of it. Immediate knowledge that a thought is the thought that *p* could be based on the phenomenal properties of phonological/orthographic imagery in circumstances under which such imagery would not be occurring unless one were consciously thinking that *p*, even if the phenomenal properties of such inner utterances do not constitute representational content. One could, that is, claim that such phenomenal properties constitute the way thoughts manifest themselves in consciousness.

But this objection does not succeed either. If it is conceded that there is a phenomenology of conscious thought, then such phenomenology *must* be proprietary, distinctive and individuative. Since (as shown above), phenomenal states in general cannot appear other than as they are, there are no conditions under which a conscious thought can appear *as something else*. (A *pain* cannot appear as pleasurable sensation, or an orange after-image, or the thought that snow is cold.) But inner utterances/inscriptions *are* something else, since what appears as an inner utterance or inscription *is* an inner utterance or inscription, and phonological/orthographic phenomenology can occur in the absence of conscious thought. Hence, conscious thoughts cannot manifest themselves as inner utterances or inscriptions in such as way as to ground Immediate knowledge of their content.

Moreover, as I shall now attempt to demonstrate, there is an introspectable difference between running words through one's head and thinking.

3. Getting Acquainted with Cognitive Phenomenology

A natural response to the antirealist about sensory phenomenology is to induce in him a particularly vivid sensory experience (say, a sharp pain), and then to point out that what he was just aware of is what you claim exists. That is, you acquaint him with the phenomenon you wish to defend the existence of, and argue that the thing just experienced exists. This is admittedly a somewhat crude form of argumentation (if in fact it is really argumentation at all). But I think it does have a significant effect on the polemical situation: it forces the skeptic to deny the obvious.

What I want to try to do now is rather deliberately to draw attention to the sort of experience I claim exists, by inducing pairs of experiences that are identical save for the absence or presence of a particular phenomenological aspect. No doubt it would still be possible to deny the existence of this aspect; but I hope to construct a situation in which to do so would be to deny the obvious.

The experience I want to isolate is the proprietary experience of thinking a thought, or entertaining a content. I shall try to do this by conducting a little amateur psycholinguistic experiment – one that will, I hope, display the phenomenological difference between reading or hearing a sentence without understanding and reading or hearing it with understanding. Recall the claim above that there can be orthographic or

phonological phenomenology (I shall call such phenomenology "linguistic phenomenology") without thought. Comparing a linguistic experience without understanding to the same linguistic experience with understanding will isolate what Strawson (1994) calls "understanding experience" – which is what I mean by "what it is like to think that p."[37]

There are several classes of sentences that linguists use to illustrate properties of our language processing abilities – specifically, limitations on our ability spontaneously to understand grammatically well-formed sentences of our language. These are sentences with multiple "center-embeddings," such as

(1) The boy the man the girl saw chased fled,

so-called "garden-path" sentences,[38] such as

(2) The boat sailed down the river sank,

and what I shall call "machine-gun" sentences, which contain a sequence of iterated tokens of an orthographic type, such as

(3) Buffalo buffalo buffalo buffalo buffalo buffalo buffalo

For most people, multiply center-embedded, garden-path and machine-gun sentences are not comprehensible at first reading/hearing. If, however, it is made clear what the sentences mean, one's experience of them ought to change. If one comes to know that 'the boy the man the girl saw chased fled' means the same as

(1') The boy, who was chased by the man that the girl saw, fled

or that 'the boat sailed down the river sank' means the same as

(2') The boat that was sailed [by someone] down the river sank

---

[37] Cf. also Moore 1962: 72-76:

> It is quite plain, I think, that when we understand the meaning of a sentence, something else does happen in our minds *besides* the mere hearing of the words of which the sentence is composed. You can easily satisfy yourselves of this by contrasting what happens when you hear a sentence, which you *do* understand, [with] what happens when you hear a sentence which you do *not* understand .... Certainly in the first case, there occurs, beside the mere hearing of the words, another act of consciousness – an apprehension of their meaning, which is absent in the second case. And it is no less plain that the apprehension of the meaning of one sentence with one meaning, differs in some respect from the apprehension of another sentence with a different meaning ....

[38] Cf. Goldman 1993: 25.

or that 'buffalo buffalo buffalo buffalo buffalo buffalo buffalo' means the same as

(3') Buffalo that are outwitted by buffalo outwit buffalo that are out-witted by buffalo

one ought to experience them differently. What is the difference? Before these sentences were explained, while reading/hearing them one is not in the mental state of consciously thinking that the boy the man the girl saw chased fled, that the ship sailed down the river sank, or that buffalo buffalo buffalo buffalo buffalo buffalo buffalo. But after they are explained, the experience of them should change. The change consists in the presence of cognitive phenomenology: it is the difference made by *thinking*.

The reliabilist of section 2.2.1, above, will no doubt wish to argue that the difference concerns whether or not one has a certain higher-order thought in one's belief box to the effect that one is reading/hearing a particular sentence with a particular meaning, or whether or not one's language faculty is able to assign an interpretation to the utterance. In fact, I do not want to dispute any of this. What I want to do is to call attention to what that difference *was like*. And I want to claim that that sort of thing goes on whenever one consciously entertains a content, where consciously entertaining a content is what goes on when one understands a sentence – whatever else may be going on in your computational or sub-personal mind.

Another objection to this little experiment is the claim that, at best, what it induces is an experiential difference between reading/hearing a sentence with and without knowing how, or being able, to *parse* it. What has been isolated, the objection goes, at best, is an experience of a sentence under a syntactic description, and not an experience of it under a semantic one. Since knowledge of syntactic structure alone is not sufficient for knowledge of meaning (since syntactic well-formedness is not sufficient for semantic well-formedness) these examples do not succeed in isolating the phenomenology of understanding.

In response, I would argue that what is going on when one comes to understand sentences like (1)-(3), even if it does involve an experience of them as parsed, or the experience of realizing or learning how to parse them, is not *just* that. One also comes to understand what the sentences *mean* – that is, one thinks the thoughts the sentences express. So let me factor out the parsing by inducing an analogous experience of a syntactically transparent sentence. Consider (4):

(4) The rhodomontade of ululating funambulists is never idoneous

If one does not know what at least one of 'rhodomontade', 'ululating', 'funambulist' and 'idoneous' means, one does not understand (4) (or at least not completely). Yet, (4) is syntactically transparent and composed of perfectly good, if slightly obscure, English words. Now if one learns that 'rhodomontade' means *rant*, that 'ululate' means *howl*, that funambulists are tightrope walkers, and that 'idoneous' means *appropriate*, one should experience the sentence differently. It will be read or heard as meaning that the rant of howling tightrope walkers is never appropriate. But this time the difference cannot be attributed to parsing.

It might be claimed that what has been isolated, if anything, is just the experiential difference between reading/hearing a sentence with and without *bewilderment* – some sort of general "Aha!" experience – and has nothing to do with any experience of the content of the sentence. If bewilderment is just lack of understanding, this is not an objection. For what one experiences on coming to understand a sentence not previously understood is just what I am trying to isolate. If bewilderment is some kind of positive state, like cognitive disequilibrium, or perhaps some sort of psychosomatic state, I would recur to multiply center-embedded sentences. Such sentences remain hard to understand even after they cease to be puzzling or bewildering (and even after you know how to parse them); so the experience of grasping them can be teased apart from the experience of relief from any bewilderment their novelty may cause. Likewise if one forgets what 'the rhodomontade of ululating funambulists is never idoneous' means. It seems likely that having heard it once, hearing it again, even without understanding, will not be especially bewildering.

Or perhaps what is isolated is an experience of coming to understand, where *what* is understood is not itself experienced. This experience would be the same for all instances of coming to understand a sentence, and so would not be the sort of experience I claim exists. I do not see, however, how the experience of coming to understand a sentence could be separated from an understanding of the sentence itself.[39] Moreover, given that only phenomenal states are consciously introspectable, and that conscious thoughts are Immediately identifiable, it could not be the case that what it is like to think that p is the same as what it is like to think that q, for all contents p and q.

4. Eliminativism About Perceptual Phenomenology

Before concluding, I would like to point out an important consequence this view has for a popular form of cognitivist eliminativism about the more familiar sorts of qualia.

The functionalist/computationalist objection discussed in section 2.2.1, above, has been made to claims that there is something it is like to see red, or feel pain – i.e., that there is sensory phenomenology. Dennett

[39] See Moore, *loc. cit.*

(1991), Harman (1990), Lycan (1990) and Nelkin (1989), for example, have argued that knowledge of such paradigmatically phenomenal states as seeing red and feeling pain should also be understood in a functionalist/computationalist way. They take the state of knowing that one is seeing red or in pain to be (something like) the state of having a token of 'I am seeing red' or 'I am in pain' reliably tokened in one's belief box. One need not postulate qualia in order to explain introspective knowledge.

The existence of a cognitive phenomenology would undermine such eliminativism. For the appearance of a representation of *I am in pain* in my belief box must be recognized for the message it is in order for me consciously to know that I am in pain. But if the relevant recognition is, as I have argued, itself a qualitative experience, the cognitivist elimination of perceptual phenomenology would be at best a limited victory. At best, it would show that there is only one kind of phenomenology, the cognitive sort.

In addition, the existence of cognitive phenomenology would provide for a defense of an argument for the existence of sensory phenomenology due to Raffman. Raffman (1995) cites psychological literature showing that we are capable of discriminating perceptual stimuli that we cannot type-identify. For example, we are capable of making discriminations among colors that we are not able to recognize independently as instances of a particular type. (This is similar to Dretske's reason for postulating simple perception.) Raffman argues that these psychological facts are most plausibly explained by supposing that there are some *percepts* (qualitative states) for which we do not have determinate *concepts*. The computationalist is thus, according to Raffman, denied the resources needed for an eliminative explanation of sensory experience.[40]

Jolley and Watkins (1998) argue the converse – that there could be cases in which colors that cannot be discriminated phenomenologically are independently identifiable as instances of particular types. Presumably, this could occur only if there were some mechanism associated with perception that was able to make finer discriminations than could be detected phenomenologically – a mechanism that, say, could detect differences of wavelength that do cause differences in conscious experience, *and* could make the discriminative information available to the perceiver in a form that would enable conscious identification. (This would amount to a kind of nonpathological blindsight.) Since, *ex hypothesi*, this information is not presented in a

---

[40] Note that Raffman's defense is not plausible in the cognitive case, since it would require claiming that we have more thoughts than we are capable of thinking, or that we have more concepts than we can conceive. And unless this is understood as meaning that we are capable of thinking some thoughts or concepts that in principle cannot get "on line" – which does not seem very plausible (why can't they?) – it is contradictory.

phenomenally discriminable state, it must be presented in a cognitively discriminable state – that is, something like the state of having a token of 'that is color X' (where color X is not phenomenologically distinguishable) in one's belief box.

Jolley and Watkins's story presupposes the experiencer's ability to make conscious distinctions among thoughts with respect to their content. But the argument from Immediate knowledge of content shows that this is not possible without a detectable difference in phenomenology. So, at best, Jolley and Watkins have answered an argument for the existence of perceptual phenomenology. Again, a limited victory. And, given that the point of such arguments is the elimination of phenomenology *tout court*, in relying on states with phenomenal properties to make their point, Jolley and Watkins would defeat their own purposes. Moreover, if we are forced to recognize the existence of phenomenology of one sort, one may wonder, what is the point of trying to eliminate phenomenology of so significantly many other and so much more obvious sorts?

5. Conclusion

I would not characterize cognitive phenomenology as a *feeling*. That would be (in my view) to class it with such experiences as tactile sensations, proprioceptions and emotions. But I do not think it is like any of those. It is a kind of awareness closer, in fact, to visual or auditory experiences – which, I take it, are not literally felt. Yet it is unique; it is not like any other kind of phenomenology.

I do not think this in itself makes the view especially implausible; for neither is any of visual, auditory, olfactory, gustatory, tactile, proprioceptive and emotional phenomenology like any of the others. Each is unique and sui generis. If one were to ask, "Well, what *is* it like to think that the weather is changing?", I could only answer in the way I would if asked what it is like to see orange by someone who had never seen it. I might say that seeing orange is like seeing red, and like seeing yellow, but definitely different from each. So, I might reply that thinking that the weather is changing is kind of like thinking that the weather is pleasant, and kind of like thinking that the weather is the same as it was yesterday, but definitely different from each.

One cannot give an informative answer about seeing orange to the congenitally blind. The best one can *say* is that it is radically dissimilar from any of the other kinds of phenomenology they are capable of experiencing. One cannot say what it is like to see orange – one must *show it*, by somehow inducing the experience. Likewise, if one has not experienced what it is like to think that the weather is changing, or has not noticed that experience before, and asks to be told what it is like, there is no answer – it can only be shown. Phenomenology is not describable in nonphenomenological terms; and different sorts of

phenomenology are not describable in the terms of any other sort.

I also do not think it is especially implausible to claim a distinctive phenomenology for each of the indefinitely many thoughts we are capable of consciously entertaining – because, in particular, there are just too many of them. Consider, for example, how many distinct visual scenes we are capable of distinguishing – an astronomically large number, I would say. The productive capacity of the phenomenal mind is prodigious.[41]

[41] I would like to thank Jonathan Adler, Brad Armour-Garb, Mark Balaguer, Al Casullo, Brie Gertler, Jerry Katz, Joe Levine, Joe Mendola, Lex Newman, Jeff Poland, Mark van Roojen, Galen Strawson, Adam Vinueza, Michael Watkins, Sara Worley and Aaron Zimmerman for helpful discussion of the issues raised in this paper. Special thanks are due Ned Block, whose comments on an earlier version led to significant improvements. Early work on this paper was supported by a fellowship from the National Endowment for the Humanities.

REFERENCES

Armstrong, D.M. 1969. *A Materialist Theory of Mind*, London: Routledge and Kegan Paul.

Baars, B.J. 1988. *A Cognitive Theory of Consciousness*, Cambridge: Cambridge University Press.

Block, N. 1996. "Mental Paint and Mental Latex," *Philosophical Issues 7: Perception*, E. Villanueva, ed., Atascadero: Ridgeview Publishing Company: 19-49.

---------- 1997. "On a Confusion About a Function of Consciousness," Block, Flanagan and Güzeldere 1997: 375-415.

Block, N., Flanagan, O. and Güzeldere, G., eds. 1997. *The Nature of Consciousness: Philosophical Debates*, Cambridge, MA: The MIT Press.

Boghossian, P.A. 1989. "Content and Self-Knowledge," *Philosophical Topics* 17: 5-26. (Reprinted in Ludlow and Martin 1998. Page number references in the text are to this reprint.)

---------- 1994. "The Transparency of Mental Content," *Philosophical Perspectives 8: Logic and Language*, J. Tomberlin, ed., Atascadero: Ridgeview Publishing Company: 33-50.

Brandom, R. 2002. "Non-inferential Knowledge, Perceptual Experience, and Secondary Qualities: Placing McDowell's Empiricism," in Nicholas H. Smith, ed., *Reading McDowell: On Mind and World*, London: Routledge.

Burge, T. 1988. "Individualism and Self-Knowledge," *The Journal of Philosophy* 85: 649-663.

---------- 1997. "Two Kinds of Consciousness," Block, Flanagan and Güzeldere 1997: 427-433.

Carruthers, P. 1996. *Language, Thought and Consciousness*, Cambridge: Cambridge University Press.

Chalmers, D. 1996. *The Conscious Mind*, New York: Oxford University Press.

---------- In preparation. "The Foundations of Two-Dimensional Semantics," http://www.u.arizona.edu/~chalmers/papers/foundations.html.

Conee, E. 1994. "Phenomenal Knowledge," *Australasian Journal of Philosophy* 72: 136-148.

Dennett, D. 1991. *Consciousness Explained*, Boston: Little, Brown and Company.

Davidson, D. 1984. "First Person Authority," *Dialectica* 38: 101-111.

Dretske, F.I. 1969. *Seeing and Knowing*, Chicago: The University of Chicago Press.

---------- 1979. "Simple Seeing," *Body, Mind, and Method*, D.F. Gustafson and B.L.

Tapscott, eds., Dordrecht: Kluwer Academic publishers: 1-15. (Reprinted in F. Dretske, *Perception, Knowledge and Belief: Selected Essays*, Cambridge: Cambridge University Press, 2000: 97-112.)

---------- 1999. "The Mind's Awareness of Itself," *Philosophical Studies* 95: 103-124.

Dummett, M. 1978. *Truth and Other Enigmas*, Cambridge, MA: Harvard University Press.

Flanagan, O. 1992. *Consciousness Reconsidered*, Cambridge, MA: The MIT Press.

Goldman, A. 1993. "The Psychology of Folk Psychology," *Behavioral and Brian Sciences* 16: 15-28.

Harman, G. 1987. "(Nonsolipsistic) Conceptual Role Semantics," *New Directions in Semantics*, E. Lepore, ed., Academic Press, Inc. London Ltd., London: 55-81.

---------- 1990. "The Intrinsic Quality of Experience," *Philosophical Perspectives 4: Action Theory and Philosophy of Mind*, J. Tomberlin, ed., Atascadero: Ridgeview Publishing Company: 31-52.

Heil, J. 1988. "Privileged Access," *Mind* 97: 238-251.

Horgan, T. and Tienson, J. 2002 "The Intentionality of Phenomenology and the Phenomenology of Intentionality," in D. Chalmers, ed., *Philosophy of Mind: Classical and Contemporary Readings*, Oxford: Oxford University Press

Husserl, E. 1900/1970. *Logical Investigations*, trans. J.N. Findlay, London: Routledge and Kegan Paul.

---------- 1913/1962. *Ideas: General Introduction to Pure Phenomenology*, W.R. Boyce Gibson, trans., New York: Collier Books.

Jackendoff, R. 1987. *Consciousness and the Computational Mind*, Cambridge, MA: The MIT Press.

Jolley, K. D. and Watkins, M. 1998. "What Is It Like to be a Phenomenologist?" *The Philosophical Quarterly* 48: 204-209.

Kant, I. 1781/1929. *Critique of Pure Reason*, N. K. Smith, trans., New York: St Martin's Press.

Kobes, B. 1995. "Access and What it is Like," *Behavioral and Brain Sciences* 18.

Kripke, S. 1980. *Naming and Necessity*, Cambridge, MA: Harvard University Press.

Langsam, H. 2000. "Experiences, Thoughts and Qualia," *Philosophical Studies* 99: 269-295.

Levine, J. 1993. "On Leaving Out What It's Like," *Consciousness*, M. Davies and G. Humphreys, eds., Oxford: Basil Blackwell Publishers: 121-136.

---------- 1995. "On What It Is Like to Grasp a Concept," *Philosophical Issues 6: Content*, Enrique Villanueva, ed., Atascadero: Ridgeview Publishing Company: 38-43.

Loar, B. 1987. "Subjective Intentionality," *Philosophical Topics* 15: 89-124.

--------- 1998. "Phenomenal States," Block, Flanagan and Güzeldere 1997: 597-616.

--------- 2001. "Phenomenal Intentionality as the Basis for Mental Content," *Objections and Replies*, M. Hahn and B. Ramsberg, eds., Cambridge, MA: The MIT Press

Locke, J. 1975. *An Essay Concerning Human Understanding*, P.H. Nidditch, ed., Oxford: Clarendon Press.

Lormand, E. 1996. "Nonphenomenal Consciousness," *Noûs* 30: 242-261.

Ludlow, P. and Martin, N., eds. (1998). *Externalism and Self-Knowledge*, Stanford: CSLI Publications.

Lycan, W.G. 1990. "What is the 'Subjectivity of the Mental?," *Philosophical Perspectives 4: Action Theory and Philosophy of Mind*, J. Tomberlin, ed., Atascadero: Ridgeview Publishing Company: 109-130.

---------- 1996. *Consciousness and Experience*, Cambridge, MA: The MIT Press.

---------- Forthcoming. "The Case for Phenomenal Externalism," *Philosophical Perspectives*.

McDowell, J. 1986. "Singular Thought and the Extent of Inner Space," *Subject, Thought and Content*, P. Pettit and J. McDowell, eds., Oxford: Clarendon Press: 137-68.

McGinn, C. 1992. "Content and Consciousness," *The Problem of Consciousness*, Cambridge: Blackwell Publishers: 23-43.

McCulloch, G. 1999. "Bipartism and the Phenomenology of Content," *Philosophical Quarterly* 49: 18-32.

Moore, G.E. 1962. "Propositions," *Some Main Problems of Philosophy*, New York: Collier Books: 66-87.

Nelkin, N. 1989. "Propositional Attitudes and Consciousness," *Philosophy and Phenomenological Research* 49: 413-430.

---------- 1996. *Consciousness and the Origins of Thought*, Cambridge: Cambridge University Press.

Peacocke, C. 1992. "Scenarios, Concepts and Perception," *The Contents of Experience: Essays on Perception*, T. Crane, ed., New York: Cambridge University Press: 105-135.

---------- 1998. "Conscious Attitudes, Attention, and Self-Knowledge," Smith and Macdonald 1998: 63-98.

---------- 2001. "Does Perception have a Nonconceptual Content?" *The Journal of Philosophy* 98: 239-264.

Pitt, D. 2000. "Nativism and the Theory of Content,"*ProtoSociology* 14, 2000: 222-239.

--------- In Preparation (a). "Unconscious Intentionality."

--------- In Preparation (b). "The Burgean Intuitions."

Raffman, D. 1995. "On the Persistence of Phenomenology," *Conscious Experience*, T. Metzinger, ed., Paderborn: Schönigh/Imprint Academic: 293-308.

Rosenthal, D. 1986. "Two Concepts of Consciousness," *Philosophical Studies* 49.

---------- 1991. "The Independence of Consciousness and Sensory Quality," *Philosophical Issues 1: Consciousness*, Enrique Villanueva, ed., Atascadero: Ridgeview Publishing Company: 15-36.

Russell, B. 1910-1911. "Knowledge by Acquaintance and Knowledge by Description," *Proceedings of the Aristotelian Society* 11.

Schweizer, P. 1994. "Intentionality, Qualia and Mind/Brain Identity," *Minds and Machines* 4: 259-282.

Searle, J. R. 1990. "Consciousness, Unconsciousness and Intentionality," *Propositional Attitudes* (CSLI Lecture Notes No. 20), C. A. Anderson and J. Owens, eds., Stanford: Center for the Study of Language and Information: 269-284.

Siewert, C. P. 1998. *The Significance of Consciousness*, Princeton: Princeton University Press.

Strawson, G. 1994. *Mental Reality*, Cambridge: The MIT Press.

Tye, M. 1995. *Ten Problems of Consciousness*, Cambridge: The MIT Press.

--------- 2000. *Consciousness, Color, and Content*, Cambridge: The MIT Press.

Wright, C., Smith, B.C. and Macdonald, C., eds. 1998. *Knowing Our Own Minds*, Oxford: Clarendon Press.