

Loar's Compromised Internalism

David Pitt

In a series of subtle and penetrating papers spanning four decades, Brian Loar developed a distinctive two-factor theory of mental content.¹ On Brian's account, an individual's conceptual mental states have a kind of content that is determined entirely by factors internal to the individual—a *narrow* ("psychological") content, as well as a kind of content that is determined by (natural and) social factors external to the individual—a *wide* ("social") content. A distinguishing feature of Brian's theory is the foundational role it gives to qualitative experience—*phenomenology*—in the determination of narrow content. His guiding insight was that an adequate theory of mental intentionality (including mental reference) must take into account the individual's *experience* of his or her intentional states, from the *subjective* point of view. This is essential to capturing psychological content, which is constituted by how an individual "conceives the facts, how he conceives the world as being" (Loar 1987, 2017, 165).²

Brian was thus a pioneer in the research program in analytic philosophy of mind that has come to be called "Phenomenal Intentionality". (Brian, (Loar 2003b) was one of the first to use this term. See also Horgan and Tienson 2002.) Though a given in the Phenomenological tradition, the idea that conscious qualitative experience has something to do with conceptual representational capacities borders on heretical for most analytic philosophers of mind. If Brian (and the Phenomenologists, Horgan and Tienson, and Searle, among others) are right, the project of *naturalizing* intentionality cannot proceed without naturalizing consciousness and qualitative experience—which is, needless to say, not something that will be accomplished by the middle of next week (if it can be done at all). (Jerry Fodor proclaimed acceptance of an essential link between intentionality and consciousness "intellectual suicide".³) Though Brian was no dualist, he saw clearly that the impossibility of giving a naturalistic explanation of consciousness and phenomenology with presently available resources did not justify ignoring it or its foundational role in the determination of

intentional content. We should, he thought, “give[] up the idea that naturalism requires that intentionality be explicated in objective and externalist terms” (1991, 236). Such an explanation may arrive at some point in the future; but the theory of content cannot wait for it. The facts about intentionality *as we experience it* are there before us (“there for the noticing” (Loar 2003b, 2017, 293)), to be discounted at one’s theoretical peril.

As bold as his insistence on the importance of phenomenology and the subjective point of view was, in the face of mainstream analytic philosophy of mind and its central commitment to naturalizability as a condition of acceptable theorizing, I think Brian did not go far enough. I think he was prevented from following his internalist insights where they inevitably lead by his acceptance of widely-held views on the semantics of singular terms, the sociolinguistic determination of conceptual content, and the scope of phenomenality. I think these views are either not inevitable, not well motivated, or short-sighted, and, in any case, ought to be rejected by anyone who, like Brian, recognizes the central role phenomenology plays in the constitution of intentionality. In accommodating these views, Brian’s internalism is compromised.

In this paper, I will focus on the second and third of Brian’s compromises, concerning sociolinguistic determination of conceptual content, and the restriction of phenomenality to the familiar sensory kinds. I will argue that the thesis that an individual’s concepts have social content in addition to psychological content is not supported by the thought experiments that Brian (like so many others) relies on, in “Social Content and Psychological Content” (1988), that it has unpalatable consequences, and that, in any case, Brian’s conclusion about the role of that-clauses in capturing psychological content is too hasty. Then I will argue that the perception-based theory of conceptual content that Brian advocates in his paper “Phenomenal Intentionality as the Basis of Mental Content” (2003b) cannot account for conceptual content, and that an internalist must embrace a broader conception of phenomenology that includes a *sui generis* kind of phenomenology proprietary to conceptual thought.⁴

Though this paper is critical, I want to emphasize how much I admire Brian’s work, and how much I have learned from studying it. Brian was a deep and subtle thinker.

1. Social Content and Psychological Content

In “Social Content and Psychological Content”, Brian argues that the following principle is false:

(A) Sameness of *de dicto* or oblique occurrence of a general term in two belief ascriptions implies, if everything else is the same, sameness of the psychological content of the two beliefs thus ascribed.

(Loar 1988, 2017, 153)

and he concludes from this that there are two different kinds of contents that beliefs have, *social* content, which is determined by the socially determined meanings of the terms in the ascriptions, and *psychological* content, which is determined by the way an individual subjectively conceives things as being (which involves the psychological, including inferential, role the state has for the individual). Psychological content is, according to Brian, entirely subjective (and grounded in phenomenology), and is not determined by sociolinguistic facts external to the thinker. Social content, which is what is captured by that-clauses, on the other hand, is sociolinguistically determined. Though only the former kind of content, which captures how things are from the thinker's point of view, is relevant to psychological explanation, Brian maintains that beliefs (and their constituent concepts) *also* have socially determined contents.

Brian's rejection of (A) is based on cases in which univocal ascriptions fail to capture personal psychological differences. Some of these cases involve translation and proper names, which I would reject for tendentious reasons that I will not defend here. (I do so in *The Quality of Thought*.) So I will focus on the cases involving general terms.

There is an initial question about how (A) is to be interpreted. There is a reading of it on which it means that for an individual to apply the same belief ascription to different individuals is for him to attribute the same belief to them. Use of the same of that-clause implies that the ascriber means to be attributing the same belief. In the ordinary case, this is true—perhaps even trivially so. Without some special reason to the contrary, we assume that the same words mean the same thing when one uses them on different occasions. If I say that two individuals believe that cats are furry, I am ascribing the same belief to them. Brian's intended interpretation, I take it, is, rather, that if two ascriptions using the same that-clause are *true*, it follows that the ascribers believe the same thing. And this is what he claims is not the case. The focus, then, is on the truth of the ascriptions.

Brian is here (and throughout his work on the philosophy of content) relying heavily on the Burgean idea that there are cases in which it is *appropriate* to use an ascription that does not capture subjective content, but in which we should nonetheless take the ascriptions to be *literally true*. He imagines a case in which an individual, Paul, associates two

different concepts with the word “cat”, wrongly thinking that there are two different kinds of animals that go by the same general name (1988, 2017, 157–158). No one knows this, however; so when Paul says, for example, “I like cats”, he is interpreted by his audience as meaning by “cat” what they mean by “cat”—i.e., what “cat” means in the language they speak. And when they attribute this belief to him, on the basis of what he has said, they use the word “cat” as they understand it (with its socially determined meaning).⁵ But while the ascription may be appropriately made and true, its that-clause content need not match the particular conception of cats that Paul had in mind when he said that he likes cats. In a case where it does not, Paul’s belief has a social content (the content of the that-clause used to report what he believes) and a different psychological content (Paul’s personal conception).

I think there are two serious problems for Brian’s argument, and, hence, for the two-factor view of intentionality he constructs as a compromise between his internalism and the falsehood of (A). The first is that the Burgean principle it relies on—that ascriptions it is appropriate (intuitively natural) to make are, all things equal, literally true—does not apply in *any* of the cases that are supposed to show that content (for Brian, a *kind* of content) is determined by factors extrinsic to an individual. The second is that in allowing that a psychological state can have two different kinds of content, Brian condones a practice that allows that we may all be massively incoherent, or at least duplicitous, in what we believe—that our beliefs have inconsistent socially and psychologically determined contents. In the next two sections I develop these complaints, beginning with the second.

Before doing that, however, I want briefly to take issue with Brian’s claim that that-clauses *cannot* capture subjective psychological content. He maintains that it is “*impossible*” . . . “to explain that-clause ascriptions in internalist terms”, (1987, 2017, 178; my emphasis), and that “we lack forms of words conventionally designed for reporting personal conceivings” (*id.* 179).

This seems very implausible. If the contents of attitudes are propositions, and that-clauses express (or denote) propositions, then how could there be *no* that-clause that captures psychological content? How could an individual have a *conception* of things without employing *concepts*? And how could a conception not be constituted by what an individual *thinks*—i.e., which propositions he entertains?

In taking up a first-person perspective on intentionality, I think Brian should have been more careful in distinguishing the various functions propositional attitude ascriptions have, and avoided their role in third-person explanation and prediction of behavior. That-clauses *can* capture subjective intentional content, even if that is not their only, or

even primary, function in our ascription of propositional attitudes to each other. The practical purposes typically served by such ascriptions do not require that we specify precisely what someone is thinking. Often enough, information about what an ascriber is referring to, what properties are being attributed to it, and what the ascriber intends to do with it or to it, are sufficient. This is the basic information we need in order to anticipate others' behavior, especially as it might affect ourselves, and to communicate our findings to relevant parties. Details about how an individual is conceiving of objects and properties may not matter if we are just trying to get a bead on what he is likely to do, to what, or to whom. Inexact attributions can count as *true for practical purposes*, even if they are not literally true.

But it does not follow from this that a that-clause *cannot* capture what someone is thinking, the way he is thinking it. For an individual, from the subjective perspective, there is no thinking of a thing and attributing a property to it that does not involve some specific *ways* of thinking of the thing and the property. Unless one supposes that individuals never accurately grasp the socially-determined meanings of terms in the language they speak, and which they use to express their thoughts, then it ought to be possible for a that-clause ascription—especially a *self*-ascription—to express (or denote) the very same proposition that an individual grasps from the subjective point of view. (Indeed, this seems presupposed by the case presented above in which Paul associates two concepts with “cat”.) Brian does discuss situations in which our primary concern *is* precisely describing individuals' mental states *as they have them*; but he says that in such cases we *project* from our own case, and what we are projecting cannot be captured by that-clauses. I do not think he has provided sufficient support for this claim.

1.1. Epidemic Duplicity

Suppose that one of the concepts Paul associates with “cat” is the standard one—say, *feline mammal*—while the other is very idiosyncratic—say, *feline amphibian*. He thinks that some cat-like creatures are in fact cat-mimicking giant salamanders (*Andrias catus*) (he once watched *Planet Earth* while tripping). Though he thinks these are native to East Asia, he suspects that some have been smuggled into the US and sold as pets. In fact, he thinks his neighbor has one. (It spends a lot of time near her bird bath, which Paul finds suspicious.) He knows the difference between mammals and amphibians, and that the classes *Mammalia* and *Amphibia* are, logically, mutually exclusive (nothing can be both a mammal and an amphibian).

One day, Paul's aunt Lydia overhears him muttering “There's that cat again; stalking crickets, no doubt!” and attributes to him the belief that a

cat is stalking crickets in the neighbors' yard. If we accept this attribution as literally true (by Lydia's lights), we are committed to Paul's believing that there is a feline mammal in the neighbors' yard. From a subjective perspective, however, what Paul believes is that there is a feline amphibian in the neighbors' yard. Moreover, for Brian, *it is one and the same mental state that has these two contents*, the first socially determined and the second personally determined. So we end up having to say that Paul has a belief with two distinct, contradictory contents. Given the likelihood that the concepts individuals associate with terms in the language they speak are quite often out of sync with those terms' socially determined contents, we seem committed to a practice that to a large extent makes a mess of our fellows' mental lives. (Nor need the double contents we count an individual as having be contradictory. The very idea that a given belief typically has two distinct contents is suspect enough.)

With respect to his (less fanciful) Paul case, Brian argues that we should say that Paul's beliefs "are not inconsistent in their psychological contents" (1988, 2017, 158)—because they are, for Paul, inferentially isolated. Paul might have distinct beliefs univocally attributable to him by "Paul believes that cats have tails" and "Paul believes that cats do not have tails" (where "cat" has its social meaning in both), but his rationality is not impugned if these beliefs have different psychological roles for him. But what I am concerned with here is cases in which we accept a third-person attribution to a single occurrent belief state that has a different psychological (in Brian's sense) content. It seems that Brian is committed to saying that a belief can have two distinct, even contradictory, contents. From Aunt Lydia's (social) point of view, Paul believes there is a feline mammal in the yard. And her attribution of this belief to Paul is true. From Paul's (subjective) point of view, Paul believes that there is a feline amphibian in the yard. And his self-attribution of this belief (were he to make it) is true. This is not a case in which two distinct beliefs having *prima facie* contradictory content are assigned on the basis of distinct utterances, but of assigning a single, occurrent belief different (incompatible) contents. There is nothing for Paul to compartmentalize. The issue concerns what we ought to say Paul believes at the moment Aunt Lydia says that Paul believes there is a cat in the neighbors' yard. It is not at all plausible that Paul's belief has, simultaneously, the content *there is a feline mammal in the yard* and the content *there is an amphibian in the yard*. And we should be suspicious of a theory that implies that this is possible.

1.2. The Burgean Intuitions⁶

Tyler Burge (1979, 1982, 1986) presents cases that are meant to elicit intuitions that, in combination with a plausible general principle about belief ascription, provide a strong reason to reject any individualist conception of mind. The intuitions (the “Intuitions”) concern *what it is natural to say* about what the individuals described in the thought experiments believe, and the general principle (the “Principle”) is that, *all things being equal*, belief ascriptions it is natural to make are literally true. The Intuitions are supposed to override any sense that the *ceteris paribus* clause of the Principle is sprung because of the conceptual idiosyncracies of Burge’s imagined subjects: they believe what they say *in spite of* their deviance from the communal norms governing the usage of their words.

The Intuitions about what it is natural to say are powerful. Burge is surely right that in the situations he describes it is natural to attribute beliefs to his subjects using the very words they utter, even though their understanding of those words is at odds with their socially determined meanings. Nonetheless, I think the Intuitions are not sufficient to establish Burge’s thesis. What it is natural to say in a given case need not be literally true; and there may be compelling reasons for thinking that it is not. I will argue that in *any* case of the kind required by a Burgean thought experiment there is a reason for not taking the homophonic ascriptions at face value, which is not neutralized by the intuitive naturalness of making them.

In the arthritis example, we have an individual (I will call her Anna) who, according to Burge, believes that she has arthritis in her thigh, in spite of the fact that her understanding of the term “arthritis” deviates from its meaning in the language she speaks (“arthritis” “as used in [her] linguistic community, does not apply to ailments outside joints. Indeed, it fails to do so by a standard, non-technical dictionary definition” (Burge 1979, 78)). Though she does not have a proper grasp of the concept *arthritis*, Anna nonetheless *has* this concept, and it partly constitutes the content of her belief. Counterfactual Anna, who has the same physical history as Anna up until the moment she expresses her worry about her thigh, is a member of a different linguistic community, in which “arthritis” applies by definition to rheumatoid diseases that can occur in muscle as well as joints. In this counterfactual context, none of Anna’s “arthritis” utterances mean what they do in her actual context, and none of her beliefs have their actual content. Since by hypothesis the only difference between the actual and counterfactual situations is the practices of the linguistic communities Anna is a member of, we seem compelled to conclude that it is these factors, and not something internal to Anna, that determine the contents of her “arthritis” beliefs.

It is essential to Burge’s argument that it be literally true in the actual context that, in spite of her confusion, Anna believes that she has

arthritis in her thigh. If she does not believe this in the actual case, then the counterfactual case is irrelevant, and no conclusions about social content determination are established by the thought experiment. (In the counterfactual case, there is no discrepancy between what “arthritis” means and what Anna thinks it means.) It is only on the assumption that Anna believes she has arthritis in her thigh that the content-determining role of sociolinguistic context is made salient by a comparison with the counterfactual case.

One might be tempted to think that Anna’s statement to her doctor is evidence that she does *not* have the concept *arthritis*, and, hence, that she cannot have the belief that she has arthritis in her thigh. After all, counterfactual Anna does not have the concept *arthritis* because “arthritis” in the counterfactual community applies by definition to a disease that can occur in muscle (or otherwise outside of joints). So why should actual Anna’s belief that “arthritis” applies to (we may suppose) those very muscle ailments not preclude her from having the concept in the actual case? The reason is, Burge argues, that Anna is not, to our knowledge, an abnormal speaker of English, or obviously deceitful, dissembling or deranged; and there is nothing odd about any of her previous “arthritis” utterances. So we should take her utterances at face value: what she says is what she believes; and we should take our own ascriptions at face value as well: what she believes is what we say she believes. There is, as Burge says, “a methodological bias in favor of taking natural discourse literally, other things being equal. . . . Literal interpretation is *ceteris paribus* preferred” (Burge 1979, 88). And we should continue to do this even after we become aware of Anna’s misconception, because there is a very strong intuition that the natural and correct way to describe her mistake is to say that she (falsely) believes that she has arthritis in her thigh.

These considerations show that comparison with the counterfactual case is not the crux of Burge’s argument. The reasons for thinking that it is literally true that Anna believes that she has arthritis in her thigh, despite her confusion, are themselves reasons for thinking that psychological content can be (and typically is) socially determined. Given that literal interpretation is preferred even in cases of partial understanding, or misunderstanding like Anna’s, and that literal meaning is determined socially, it follows that the contents of Anna’s thoughts are determined socially. The counterfactual case is meant to reinforce this conclusion by showing that thought content can change with change of linguistic context, without any difference in intrinsic properties of thinkers. But the contrast depends upon having established that actual Anna’s thought contents are sociolinguistically determined. The heart of the argument is the explanation of how this works.

A different kind of example is developed in “Intellectual Norms and Foundations of Mind” (Burge 1986). Here we have a subject (I will call him “Andy”) who is the victim of neither partial understanding nor misunderstanding with respect to the specimen term (“sofa”). Andy knows the meaning of the word “sofa” in his language. He knows that sofas are by definition pieces of furniture, and he is fully competent in the use of the word. Nonetheless (according to Burge), he doubts that sofas are furniture. Andy has become convinced that they are not furniture, but some sort of artworks or religious artifacts. He expresses his suspicion by saying “Sofas are not pieces of furniture!” He is wrong, however. His belief that sofas are not furniture was formed on the basis of systematically misleading evidence.

In a counterfactual case, in contrast, we are to suppose that a physically identical Andy says something *true* when he utters the sentence “sofas are not pieces of furniture”. In the counterfactual language community, the word “sofa” is applied by definition not to pieces of furniture, but to works of art or religious artifacts, just as Andy suspects. When counterfactual Andy says “Sofas are not pieces of furniture!” his utterance, and the belief it expresses, are both true. They do not, however, have the same content as Andy’s actual utterances. He counterfactually neither says nor believes that sofas are not pieces of furniture, because “sofa” in the counterfactual community does not mean *sofa*.

The pattern of reasoning in the Andy case is the same as in Anna’s: the difference between the contents of actual and counterfactual beliefs is due to the actual and counterfactual meanings of the words uttered, which are determined by factors extrinsic to the speaker. Moreover, the points made above with respect to the role of the intuition about what Anna believes in the actual case, and the comparison to the counterfactual case, apply to Andy’s case as well. The comparison of actual Andy and counterfactual Andy is not what establishes the anti-individualist conclusion; it merely dramatizes it.

Further, in Andy’s case it is not misunderstanding that is overridden by communal meaning, and which provides the essential contrast to the ordinary case. Andy understands perfectly well what his community takes “sofa” to mean. It is, rather, as Burge says, “nonstandard theory” (1986, 709). Nonetheless, there is an important parallel to Anna’s case. Both Anna’s and Andy’s utterances are *prima facie* conceptually problematic: neither involves a merely empirical divergence from communal standards and belief. Anna’s utterance indicates conceptual *confusion*, and Andy’s indicates an attempt at conceptual *subversion*. I will characterize what the cases have in common as conceptual *dissonance*. In all of Burge’s examples it is the fact that the individuals in question think what they say, in spite of the conceptual dissonance of

their utterances, that provides the reason for thinking that the Principle has general application, and that communal linguistic meaning determines mental content.

The inclusion of the *ceteris paribus* clause in the Principle indicates that Burge thinks there *are* conditions under which literal interpretation is not preferred. This might mean that the individuals' words are themselves to be reinterpreted. If there can be cases of mismatch between thought content and linguistic content, then, it remains to be seen what sorts of criteria there are for identifying them. Burge (1978, 134, 1979, 90–91) mentions utterances involving slips of the tongue, spoonerisms, malaprops, and radical misunderstandings as appropriate candidates for reinterpretation. In a parenthetical remark (1979, 91), however, Burge indicates that he thinks utterances involving malaprops or radical misunderstandings are *not* exempt from the Principle. He says that he is “not convinced” that someone who believes that “orangutan” is a word for a fruit drink and says, “An orangutan is a fruit drink” should not therefore be taken to *mean that* an orangutan is a fruit drink and *believe that* an orangutan is a fruit drink.

But it seems that Burge could reject such radical cases and still have an important thesis about the determination of mental content. (He has in fact, in conversation, disavowed the intuition about the orangutan case.) He says (*id.*, 92): “The thought experiment depends only on there being some cases in which a person's incomplete understanding does not force reinterpretation of his expressions in describing his mental contents. Such cases appear to be legion”. I take it that Burge means that his argument does not require commitment to the Principle for *every* utterance. Standard practice can call for reinterpretation in cases of linguistic incompetence or momentary lapses without derailing the anti-individualist argument. What the argument requires is cases in which the Principle overrides conceptual dissonance of one kind or another. Conceptual dissonance is a *prima facie* reason for thinking that all things are not equal. This is why Burge stresses that the deviance from community norms is not merely empirical. Mere empirical deviance would not raise the issue of non-literal interpretation, and would therefore not provide the resistance needed to reveal the determinative role of the Principle. Since homophonic interpretation is still preferable in these cases, they do show something interesting and important about psychological content.

But *why* should we think that homophonic interpretation is preferable in these cases? Why is the Principle's *ceteris paribus* clause not sprung by the conceptual oddity of the utterances? The answer constitutes, I think, the very heart of Burge's argument. We do not reinterpret in these cases because of the Intuition about the correct way to describe them is

so powerful. Burge (*id.*, fn. 4) says “I used to believe that a fortnight is a period of ten days”. He knows that what he used to believe cannot be true; he knows he made a conceptual error. But he describes what he used to think using the very word he did not understand. This is, Burge stresses, the natural and intuitively correct thing to say. But it shows that communal meaning trumps individual misconception in determining the contents of a generally competent speaker’s thoughts.

The role of the Burgean Intuitions should now be clear. The Principle applies in these cases in spite of *prima facie* evidence that all things are not equal because of the clear and firm Intuition about the naturalness of our homophonic ascriptions.

The Intuition is undeniable. Nevertheless, I think it is misleading. It cannot be uncritically accepted as evidence that all things are equal in the situations Burge describes. I think there are features of our ordinary belief-ascribing practices that militate against Burge’s anti-individualist conclusion. These practices are governed by a principle of charity enjoining ascription of contradictory or incoherent beliefs except, perhaps, in extraordinary circumstances, and also by the principle that utterances betraying linguistic incompetence should not be interpreted literally.

It has been objected to Burge (beginning with Fodor 1982) that all things are not equal in Anna’s case because the sentence “she has arthritis in her thigh” is, *by hypothesis*, self-contradictory or conceptually incoherent, and there is a competing principle governing ordinary discourse that says that one’s attributions ought to be charitable. The question whether or not Anna has such a belief should be distinguished from the question whether or not we attribute one to her in using a homophonic ascription. If Anna utters the sentence “my arthritis has spread to my thigh”, it might be open to question what she actually meant (what she actually thought). But if *we* utter the sentence “Anna believes that she has arthritis in her thigh”, then there is no question that the embedded sentence “she has arthritis in her thigh” is *for us* conceptually incoherent. For, again as Burge stresses, Anna’s error is not an ordinary empirical one; it is a *conceptual* matter that arthritis cannot occur outside of joints. But given that the kinds of errors ~~that~~ Burge describes are common, to suppose that individuals’ words should be taken at face value in such cases is to turn a practice of intending to make sense of their behavior (attitude attribution) into a practice that routinely makes nonsense of their mental lives. If this is not our practice, then the kinds of cases Burge describes in “Individualism and the Mental” are ones in which literal interpretation is not preferable, and the conclusion he wants about the contents of thought does not follow from the naturalness of our attributions. The fact that we find it natural to make them in spite of their incoherence no doubt requires an

explanation. But if the principle of charity is upheld, it cannot be maintained that we ought to take them to be literally true.⁷ It can also be argued that the principle of charity prevents attribution of the belief that sofas are not furniture to Andy since, *by the ascriber's lights*, that belief is conceptually incoherent.

Nonetheless, as strong as the charity objection is, I think it does not really get to the bottom of Burge's argument for anti-individualism. Given the role the examples are supposed to play in the thought experiments, it really is not necessary that they involve the sort of *prima facie* conceptual incoherence Burge's actual examples exhibit. The examples are supposed to illustrate the domineering role of the Principle by presenting it with a serious obstacle to overcome. They are supposed to show that even when an individual's errors raise the specter of reinterpretation, homophonic attribution is completely natural and compelling, and therefore the Principle still applies. I have argued that the Anna and Andy cases run afoul of the principle of charity. But perhaps Burge can accept (because he can explain away) cases in which reinterpretation is forced, because there remain so many cases in which it is not. Giving up the original examples need not spoil the argument for anti-individualism, if there are others that can play the same role. Here is one.

Consider Marge. Marge fancies herself a sophisticate, and often opines upon the finer things in life. On one occasion, she declares "You can tell the quality of a diamond by how many faucets it has". (It is not a slip. When pressed ("*Faucets?*"), Marge does not retract her statement.) If we say that Marge *said that* the quality of a diamond depends on how many faucets it has, it seems clear that what we are saying is literally true. That *is* what she said—what the words she used mean in her language. Moreover, it seems entirely natural to *say that* Marge believes that the quality of a diamond depends on how many faucets it has. Marge is, of course, confused. But we describe her confusion using the very word she misuses, just as we do in the case of Anna. And after she has been corrected Marge might describe her misstatement using the words she uttered (just as Burge did in the "fortnight" case).

Marge's utterance is sufficiently dissonant to raise suspicions about her competence with the term "faucet", and, hence, to present a *prima facie* obstacle to literal interpretation. But it does not run afoul of the principle of charity. It is certainly odd to say that diamonds have faucets; but it is not conceptually or logically impossible that it be true. (If you look closely and carefully enough, as at Albritton's pencil, you can see the tiny things.) So perhaps this is just the kind of example that Burge needs. It does not involve a contradiction, but neither is it an ordinary empirical error. It is a rather extraordinary one: faucets are bits

of plumbing. It is bizarre to think that diamonds have plumbing. Anyone who seriously says so thereby provides *prima facie* evidence of not knowing what a faucet is—of not being competent with the concept *faucet*. However, since the Intuition applies in Marge's case, the Principle ought to apply as well, and Marge ought to have the (our) concept *faucet*, and think that diamonds have them, in spite of her confusion.

But there are strong counterintuitions in cases of this kind. If we reflect, and ask ourselves if Marge really was thinking that diamonds have *faucets*, I think the answer has to be "No". This really is not what she thought. She simply misexpressed herself. She had a false belief about which word to use to say what she was thinking, and that belief led to her anomalous utterance. (Cf. Bach 1988; Crane 1991.) In spite of the fact that the Intuition holds, the *ceteris paribus* clause of the Principle is sprung: all things are not equal in Marge's case, because she is not fully competent with the words she uses.

Since Burge accepts that malaprops are evidence of a lack of the minimal competence required for an individual to be subject to the Principle (1979, 90), he ought to agree that it is not literally true that Marge believes that diamonds have faucets. The crucial point here, however, is that *it is nonetheless intuitively completely natural* to say that this is what she believes. But it is clear that we do not think that our homophonic ascription to Marge is literally true. We know that she meant to say "facets", and that this is what she was actually thinking. So the presence of the Intuition provides no assurance that we should apply the Principle in spite of a speaker's errors. But Burge's argument for anti-individualism depends on this: naturalness of homophonic ascription is supposed to overpower any inclination to reinterpret. Cases of this type show that the Intuitions concern only what it is natural to say, since, on reflection, we do not take our ascriptions to be literally true. The Intuitions do not prevent the *ceteris paribus* clause of the Principle from being sprung.

Insofar, then, as Brian's case for a socially determined kind of belief content relies on Burgean Intuitions, his two-factor theory of content is unmotivated. An internalist need not (and should not) recognize two kinds of intentional content.

2. Phenomenal Intentionality

In "Phenomenal Intentionality as the Basis of Mental Content", Brian says that

Internal intentionality is to be located primitively in perceptually based concepts. It will be derivatively located in nonperceptual

concepts via their conceptual connections with perceptual concepts. The subjective intentional properties of nonperceptual concepts are always a matter of, as it were, looking sideways via their connections with perceptual concepts.

(Loar 2003b, 2017, 310; page number references in this section are to this work.)

These foundational, perceptually based concepts include *singular perceptual demonstratives*, which pick out individuals, *perceptual recognitional concepts*, which are “demonstrative concepts that purport to pick out, perceptually, kinds and properties” (*ibid.*), “certain general concepts of approximate spatial relations, shapes, and the like, and a certain conception of a three-dimensional object as it persists over time” (313). The intentional content of these concepts, their “directedness” toward particular objects, properties and relations (whether or not they refer), is derived from the directedness of perceptual experience—the experience *as of* objects, properties, and relations, which is itself determined by perceptual phenomenology. The intentional content of perceptual experience is an intrinsic feature of its phenomenology, and is determined entirely independently of relations to external objects and properties. Perceptual concepts inherit their intentional contents from perceptions, and nonperceptual (“systematic” or “theoretical”) concepts acquire theirs in virtue of their inferential relations with them—their conceptual roles. All internal mental intentionality, therefore, is derived from the intentionality of perception.

I have two problems with Brian’s view. The first is with the notion of a *perceptual concept*. I do not think this is workable. Perceptual concepts, as Brian understands them, *incorporate* perceptions:

The feature [the directedness of visual demonstrative concepts] presumably belongs primarily to a visual perception, and derivatively to a visual demonstrative concept that incorporates the perception.

(302)⁸

I do not think it is possible for concepts to incorporate percepts. Concepts and percepts are fundamentally distinct and incompatible kinds of mental entities, and cannot be combined in such a way that their amalgamation constitutes a *concept*. We can have concepts *of* (*referring to*) perceptual experiences; but the internal intentional contents of these concepts are not individuated with respect to the percepts they are of, by incorporation or otherwise.⁹

I also want to take issue with Brian's account of a particular kind of perceptual concept, *recognitional* concepts. I do not think these kinds of concepts exist either, and not just because concepts cannot contain perceptions. There are no, in my view, and cannot be, concepts of the kind that Brian places at the foundation of mental intentionality.

My second problem with Brian's view is that it is too conservative about what kinds of phenomenology there are. In recognizing only familiar sensory kinds of phenomenology, Brian limits the possibilities of his internalism, and its effectiveness in giving a general account of conceptual content. An experiential theory of conceptual content that restricts phenomenality to sensory experience inherits the notorious difficulties of the Empiricist program in coping with non-perceptual concepts.¹⁰ It does not have the resources to explain non-sensory intentionality. I do not think the inclusion of conceptual roles can remedy the situation, unless non-sensory phenomenology is somehow thereby smuggled in.

I develop these objections in the next two sections.

2.1. Perceptual Concepts

It is traditional to make a sharp distinction between concepts and percepts, thought and perception, sapience and sentience. I think there are powerful intuitions supporting this distinction, and that these two kinds of mentality must always be kept strictly separated in theorizing about the mind. Though there may be total conscious occurrent states containing concepts and percepts, percepts cannot be parts of concepts, and concepts cannot be parts of percepts. The tendency to blur the distinction comes, it seems to me, from a commitment (tacit or explicit) to the idea that the content of a concept must determine what, in general or in particular, it purports to refer to. This is especially so in the case of indexicals and demonstratives.

Intuitively, my thoughts *this is an unpleasant sensation* directed at two different sensory experiences on two separate occasions, are different thoughts. They are about different sensations, differ in their truth conditions, and may differ in truth value; so, the standard argument goes, they must have different contents. And it seems the only way to explain this is to assign different contents to the two tokens of *this*. One way to do this is to incorporate sensations into the token demonstrative concepts. It is also often remarked that in order to understand an experiential concept, one must have had the relevant experience. For example, a congenitally blind person could not have the concept *red*. To understand red, to cognitively grasp the concept *red*, one must experience red. This suggests that the experience is somehow intimately involved with the concept—perhaps as a constituent of it.

Another source of this tendency, it seems to me, is a bias toward approaching the study of mentality from the third-person perspective—particularly from the point of view of our practices of attribution and communication. If you want to tell someone what I said when I said, “This is an unpleasant sensation”—if you want them to understand what I said—you cannot just leave it at repeating my words. Many have the intuition that if the third party has no idea what sort of sensation I was referring to, she will not really (or not completely) understand what I said. The attributor/communicator must supplement the content of the demonstrative “this sensation” in such a way as to direct the hearer’s attention on the relevant kind of sensation.¹¹ But these practical requirements of third-person attribution and communication can be met without shoehorning identifying content into demonstrative concepts. The relevant information can be communicated separately.

Moreover, one need not accept that *grasping concepts of* phenomenal properties requires experiencing them. I once knew someone who was born without a sense of smell. After many conversations with her, I realized that her grasp of the *concept* of olfactory experience was no different from mine. We would say the same thing if we were asked to explain what it *is*: it is a kind of sensory experience, as different from others as they are from each other, involving the nose, often confused with taste (she did not fail to notice that other people enjoyed food a lot more than she did), etc., etc. In contrast, knowing *what smelling is like* in general (knowing *what it is like* to have olfactory experience), or what, in particular, *lavender* smells like—does require experience. However, I maintain that this kind of knowledge is not *conceptual* at all.¹²

But why (people keep asking me) can we not just say that perceptual (phenomenal) concepts are “hybrids”, like (on some views) emotions, consisting of sensory and cognitive components? As I said above, nothing prevents the occurrence of complex occurrent states of consciousness comprising sensory and cognitive components; but I insist that there cannot be *cognitive* components that themselves incorporate perceptual ones, nor *perceptual* components that incorporate cognitive ones. This is not a quibble about how to use the word “concept”. It is an insistence that we recognize the fundamental differences among different kinds of mental states, and the impossibility of their combining in such a way as to produce a state that is of one kind, but not the other.

It is clear that there cannot be cross-modal *sensory* incorporations. One’s total conscious state may include simultaneous auditory and olfactory experiences; but there can be no sounds parts of which are

smells, or smells parts of which are sounds. This is simply impossible. Phenomenologies do not mix *in this way*.¹³

I think this is true of cross-categorical (sensory/cognitive) incorporations as well.

My tendentious reason for believing this, which I will defend shortly, is that occurrent concepts (and thoughts) are themselves experiences, of a unique *sui generis*, cognitive kind, and that, therefore, the Principle of Phenomenal Immiscibility applies to them, with respect to any other kind of experience. Sounds cannot have smells in them, and neither can concepts—for the same reason.

But there are also good non-tendentious, intuitive reasons for resisting cross-categorical incorporation even if one denies that there is conceptual phenomenology. The traditional distinction between the two kinds of mind—sentient and sapient—is, I think, based on our own experience. It seems clear, from the subjective point of view, that thinking is different from sensing. We recognize that they are different kinds of mental activity, and that we can engage in one without engaging in the other. Further, we find it conceivable that there are creatures who can sense but cannot think, and creatures who can think but cannot sense. But if one can sense without thinking, and *vice versa*, it must be that percepts and concepts are, metaphysically, different kinds of things.

Another strong intuitive reason for resisting conceptual-sensory assimilation is this. Concepts (conceptual contents) must be *thinkable*; but percepts (and images) are not. It is nonsense, a category mistake, to say that what you were *thinking* (or part of what you were thinking) was, e.g., *the taste of cinnamon*, or *the sound of a piano*. It is true that you can think *about* these experiences, but only in the sense that you can have otherwise-content-individuated concepts that refer to them. Concepts must be capable of being *thought*, in the course of thinking a complete thought of which they are constituents. But sensory experiences cannot be thought. One can no more think the taste of cinnamon than one can hear it. (Nor, for that matter, can one think physical objects and properties—trendy Continentalisms like “How do you think Derrida?” notwithstanding. This is a reason for rejecting certain widespread accounts of the contents of singular thoughts.) If these intuitions are sound, then there can be no perceptual concepts, as Brian understands them.

I also think there are no recognitional concepts, as Brian understands them, though for different reasons. Recognition, in my view, essentially involves a particular kind of nonconceptual experience—a feeling of familiarity. The application of a demonstrative concept (*that feeling*) to a familiar experience is not what makes it familiar; nor is a demonstrative concept individuated by the experience it is applied to. A

person with perfect pitch does not have a different concept of middle C from me. The fact that she instantly identifies the pitch when she hears it does not make her concept different from mine. What is different between us is her infallibly correct application of it on the basis of her experience of it. Her experience of middle C is like my experience of red: I recognize red on sight; she recognizes middle C on hearing. But I do not think I have a different concept of red from someone who is color blind, or totally blind.

Now, it might be wondered why it should matter to Brian if there are no perceptual concepts (as he understands them). Perhaps conceptual incorporation of percepts is not a necessary feature of his view: perceptual experience could be a foundation for conceptual intentionality without incorporation. I agree that incorporation is not necessary (for me, because impossible) for a phenomenal theory of intentionality. But I think it is essential to Brian's account, on which the intentionality of all concepts is derived from the intentionality of perceptual experience.

For Brian, content cannot arise from inferential relations among concepts, since, for him, thinking is "something lively—there is something it is like to engage in it", whereas conceptual roles on their own are "too *blank* to constitute internal mental content as we perceive it"—because they are "purely dispositional" (296). There is a phenomenology of thinking; but there is no phenomenology of dispositional states. Conceptual roles are "central in individuating thought contents" (*ibid.*), but they cannot be *sources* of contents, since they are not phenomenal.

Furthermore, a view on which non-intentional conceptual roles are "phenomenalized" by inclusion of pure, non-intentional perceptual states—which are, in their turn, thereby "conceptualized" (i.e., *intentionalized*)—is, according to Brian, not available. Sensory phenomenology is inseparable from its intentionality ("We can hardly peel the phenomenal aspects of [e.g.] vision away from its intentionality" (297)). The phenomenalization of conceptual roles would, necessarily, simultaneously be their intentionalization. And since (as it seems Brian assumes) there is only sensory phenomenology, sensory (perceptual) states are the only source of intentionality.

Finally, given that phenomenal properties are intrinsic, phenomenal concepts cannot get their content from extrinsic relations to sensory states, including reference. Brian thinks psychological contents are not individuated by external referential relations. But if he thinks this because he thinks that content is phenomenal, and phenomenology is intrinsic, then he ought to accept it for internal reference as well. Phenomenal concepts (including recognitional concepts) are not individuated by phenomenal states in virtue of referring to them. They

must incorporate them. This is, on Brian's theory, the only way to *conceptualize* sensory phenomenal intentionality. The content of a concept is internally determined. This does not just mean that it is determined by features internal to thinkers: it is determined by features internal to concepts themselves.

To close this section, I want to identify what seems to me a significant tension between Brian's claims that conceptual roles are "too blank to constitute internal mental content", and that they are "central in individuating thought contents" (296). I think what Brian has in mind is that position in a conceptual network is not in itself sufficient to confer content: the interconnected concepts must have independently derived content. It cannot be that *every* concept derives its contents from its inferential connections to other concepts. There must be a foundation of *intrinsically* contentful concepts.

However, it is not clear that relatedness to other concepts can *constitute* the content of a concept if phenomenology is both intrinsic and conceptually foundational. In general, the phenomenal character of a state will not be altered by causal connections (dispositional or manifest) with other states. Nor, therefore, will its content. (Phenomenal context effects are not to the point here.) At best, then, it seems that tracing connections to other concepts will enable individuation only in an *epistemic* sense: the intrinsic content of a concept can be revealed, if it is not immediately available to introspection, in its inferential connections to other concepts.

2.2. Conceptual Phenomenology

My tendentious explanation for the intuitive facts appealed to above about what is thinkable and not, is that thinking is a distinctive kind of experience, and that different kinds of experience are individuated by their different kinds of phenomenology. Thinking is not the same kind of experience as seeing, hearing, smelling, tasting, and so on. So thinking must have its very own brand of phenomenology—a phenomenology which I have elsewhere (Pitt 2004) characterized as *proprietary*, *distinctive*, and *individuating*. Experiential modes in general are distinguished phenomenally. Vision, audition, olfaction, gustation, etc. are modes of experiencing, and each is constituted by its own proprietary kind of determinable phenomenology (visual, auditory, olfactory, gustatory, . . .). Hence, there is a proprietary determinable phenomenology for thought—what I have called (Pitt 2004) "cognitive" phenomenology ("conceptual" and "propositional" phenomenology will do as well). And just as there are differences *within* each kind of sensory phenomenology that distinguish them from others

of the same general kind (e.g., green percepts from yellow percepts), there are differences within conceptual phenomenology that distinguish conceptual contents from each other. There are cognitive, as well as sensory, determinables and determinates. Moreover, on this view, the various determinate cognitive phenomenologies *are* intentional contents (Pitt 2009).

So, for the same reason that there cannot be sounds with smells as constituents, there cannot be concepts with smells as constituents. Furthermore, just as sounds cannot be individuated with respect to smells, phenomenal concepts cannot be individuated with respect to their nonconceptual referents. The contents of my concept *that smell* do not change depending upon which smell I am attending (referring) to.¹⁴

I have both epistemic and metaphysical reasons for believing that there is a proprietary, distinctive, individuating phenomenology of thought content.¹⁵ Briefly, the epistemic reason (developed in Pitt 2004, 2009) is that there is available to us a mode of access to the contents of our conscious occurrent thoughts—viz., introspective and non-inferential—that would not be available to us if occurrent conscious thought contents were not distinctively presented to us in conscious experience. In general, discriminatory non-inferential introspective awareness of occurrent conscious states requires that the states accessed be differentiated *in consciousness*—that is, they must be introspectively distinguishable and identifiable as the states they are. (This is analogous to the role that distinguishability and identifiability of objective properties plays in purely perceptual discrimination.) But differentiation in consciousness is entirely a matter of difference of phenomenology. Thus, we can be non-inferentially introspectively aware that we are experiencing a pain, and not a smell or a sound, that the pain is burning, and not achy or slashing, or that it has gotten worse, because pain experiences have proprietary, distinctive, individuating phenomenologies. And we can be non-inferentially introspectively aware *that* we are thinking, and not feeling pain or hearing a sound or smelling a smell, and that we are thinking about covalent bonding, and not about beer, or the end of American democracy. But if conscious thoughts can be thus discriminated from other conscious states (or events) and identified as the thoughts they are, introspectively and non-inferentially, they too must have proprietary, distinctive, individuating phenomenologies, which constitute their intentional contents.

This is a transcendental argument: a certain kind of access to conscious occurrent thought content is possible (which is not to say either that it is infallible or that it is the only kind of access we can have to our thoughts); but it would not be possible if there were no

proprietary, distinctive, individuating phenomenology of thought; hence, there is such phenomenology.

The metaphysical reason (developed in Pitt 2011) focuses on the fact that conscious states in general are, *qua* conscious, phenomenally individuated. What distinguishes conscious smells from conscious sounds is their distinctive kinds of phenomenologies (olfactory and auditory). Hence, if conscious thoughts are not conscious sights, smells, sounds, . . . , then they must have their own kind of phenomenology that constitutes their determinable phenomenal kind. And if the thought that yttrium conducts electricity is a different thought than the thought that robots are stealing our luggage, then, like the taste of honey and the taste of ashes, they must have different determinate phenomenologies.

Now, Brian claims that not all contents of occurrent conscious concepts (hence, thoughts) are immediately accessible in introspection. This is, for example, true of what he calls “systematic” or “personal theoretical” concepts, such as *mother*, *female*, and *child*:

We do not take in the intentional properties of a systematic concept all at once. . . . We do so rather by finding our way about among a systematic concept’s lateral interconceptual connections. . . . The phenomenological world-directedness of a personal theoretical concept, I want to propose here, derives from its intimate conceptual connections with perceptual intentionality. . . . So the idea that every concept can be revealed in an introspective glance, or even in an introspective stare, is not essential to the defense of internal, phenomenological, intentionality.

(Loar 2003b, 2017, 315)

But it cannot be the case that *all* concepts are systematic in this way. Some must wear their contents on their face. Otherwise, given Brian’s rejection of the idea that sheer relatedness—blank conceptual roles—can confer content, the search for intentional properties would be viciously circular. And the question is whether these “facial” contents can all be sensory.

The history of failed attempts to construct an empiricist psychosemantics strongly suggests that they cannot. It is highly doubtful that a theory that relies on sensory phenomenology as the ultimate source of all intentional content can account for the contents of, for example, philosophical, mathematical and logical concepts and beliefs. The concept of objective causal necessity cannot (as Hume showed) be derived from experience, either introspective or extrospective, which only presents us with the way things are, and not how they must or cannot be. (This led Hume to deny that we have any such concept,

redirecting the term “necessity” to experiences of constant conjunction and expectation.) Nor do concepts like *transfinite ordinal*, *least convergent series*, *infinitesimal*, *entailment*, *logical necessity*, and *possible world* have sensory contents. And it is not obvious that how the contents of such concepts can be generated by inferential connections to sensory concepts.

Moreover, even if sensory experience provides occasions for the formation of such concepts, it cannot confer their contents upon them. One might form a general geometrical concept like *triangle* in response to sensory experiences of a variety of physical representations; but, as has been long observed, general triangularity is not represented by any particular sensory experience (and not least because there are no experienceable triangles). The contents of concepts (the conceptual machinery) needed to form a general concept on the basis of sensory experiences are not themselves to be found in those experiences, all of which are particular. Nor do they arise from the connections particular concepts bear to each other. What is *understood* when we grasp abstract philosophical, logical and mathematical concepts is not sensory, and cannot be reduced to the sensory.

To accommodate these facts, a phenomenally-based theory of content must recognize a non-sensory cognitive phenomenology, in addition to sensory phenomenology. Moreover, given the impossibility of sensory incorporation, cognitive phenomenology must be taken to constitute *all* distinctively *conceptual* intentional content. (This is not to deny that sensory experiences have *nonconceptual* intentionality.) Insofar as Brian’s theory does not recognize this kind of phenomenology, it is, in my view, compromised. It is not yet a theory that a Phenomenal Intentionalist should be satisfied with.

Notes

-
- 1 Loar (1976, 1987, 1988, 1991, 1995, 1997, 2003a, 2003b).
 - 2 Where a cited paper is reprinted in Loar 2017, page number references are to that volume.
 - 3 I have not seen this remark in print. I heard Fodor make it, and Strawson reports it in his 2008.
 - 4 I challenge the standard semantics of singular terms (proper names and indexicals), as applied to the theory of mental intentionality, in book *The Quality of Thought* (under contract with Oxford University Press).

-
- 5 I set aside here the complication that an ascriber's concept need not match the concept expressed in his language by the term he uses.
- 6 This section is adapted from a much longer and more detailed discussion of Burge in *The Quality of Thought*.
- 7 I discuss Burge's response to this objection, which I think is inadequate, in *The Quality of Thought*.
- 8 See also Loar 1997.
- 9 I have in mind here accounts on which phenomenal concepts are individuated referentially. Referential individuation is as problematic for internally directed concepts as it is for externally directed ones. Externalism is just one form of referentialism (extensionalism). (Cf. Farkas 2008, 77–79.)
- 10 I also believe that Brian's account of *reference* as an external, causal (as opposed to descriptive) relation will run into problems with (at least some) concepts of abstract objects. I develop this objection in *The Quality of Thought*.
- 11 Analogous considerations are often appealed to in the case of proper names (and indexicals). For example, I will not understand what you said when you said "Don is totally bogus" if I do not know who Don is—which Don you are referring to; and you must somehow supplement the name in such a way as to get me on to the right one. And neither will *you* understand what you said, or succeed in thinking a thought, if you cannot supply a referent (Evans 1982, 74).
- 12 I elaborate and defend this nonconceptual account of knowing what it is like, what Earl Conee (1994) calls "phenomenal knowledge" (I call it "acquaintance knowledge") in my forthcoming paper "Acquaintance and Phenomenal Concepts".
- 13 This is what I call the *Principle of Phenomenal Immiscibility*. To say that oil and water are immiscible is not to say that there cannot be a collection of water molecules and oil droplets in suspension. Oil and water can be mixed in *that* sense. But oil droplets cannot be partially composed of H₂O molecules, and H₂O molecules cannot have oil droplets as constituents.
- 14 I present and defend a cognitive-phenomenal account of the contents of indexical and demonstrative concepts in chapter four of *The Quality of Thought*.
- 15 This should be carefully distinguished from the thesis that there are proprietary phenomenologies for the various propositional *attitudes*, which may in fact be individuated by psychological roles. Cognitive phenomenology is the phenomenology of *content*—a *conceptual/propositional* phenomenology.

References

- Bach, K. 1988. "Burge's New Thought Experiment: Back to the Drawing Room," *Journal of Philosophy* 85, 88–97.
- Burge, T. 1978. "Belief and Synonymy," *Journal of Philosophy* 75, 119–138.
- Burge, T. 1979. "Individualism and the Mental," *Midwest Studies in Philosophy* 4, 73–121.
- Burge, T. 1982. "Other Bodies," in A. Woodfield, ed., *Thought and Object*, Oxford University Press, 97–121.
- Burge, T. 1986. "Intellectual Norms and Foundations of Mind," *Journal of Philosophy* 83, 697–720.
- Conee, E. 1994. "Phenomenal Knowledge," *Australasian Journal of Philosophy* 72, 136–150.
- Crane, T. 1991. "All the Difference in the World," *Philosophical Quarterly* 41, 1–25.
- Evans, G. 1982. *The Varieties of Reference*, Oxford University Press.
- Farkas, K. 2008. *The Subject's Point of View*, Oxford University Press.
- Fodor, J. A. 1982. "Cognitive Science and the Twin-Earth Problem," *Notre Dame Journal of Formal Logic* 23, 98–118.
- Horgan, T. and J. Tienson 2002. "The Intentionality of Phenomenology and the Phenomenology of Intentionality," in D. J. Chalmers, ed., *Philosophy of Mind: Classical and Contemporary Readings* Oxford University Press, 520–533.
- Loar, B. 1976. "The Semantics of Singular Terms," *Philosophical Studies* 30, 353–77.
- Loar, B. 1987. "Subjective Intentionality," *Philosophical Topics* 1, 89–124.
- Loar, B. 1988. "Social Content and Psychological Content," in R. H. Grimm and D. D. Merrill, eds., *Contents of Thought: Proceedings of the 1985 Oberlin Colloquium in Philosophy*, University of Arizona Press, 99–110.
- Loar, B. 1991. "Can We Explain Intentionality?," in B. Loewer and G. Rey, eds., *Meaning in Mind: Fodor and His Critics*, Blackwell, 199–136.
- Loar, B. 1995. "Reference from the First-Person Perspective," *Philosophical Issues: Content* 6, Ridgeview Publishing Company, 53–72.
- Loar, B. 1997. "Phenomenal States: Second Version," in N. Block, O. Flanagan and G. Güzeldere, eds., *The Nature of Consciousness: Philosophical Debates*, MIT Press, 597–616.
- Loar, B. 2003a. "Transparent Experience and the Availability of Qualia," in Q. Smith and A. Jokic, eds., *Consciousness: New Philosophical Perspectives*, Oxford University Press, 77–96.

-
- Loar, B. 2003b. "Phenomenal Intentionality as the Basis of Mental Content," in M. Hahn and B. Ramberg, eds., *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, MIT Press/Bradford Books, 229–257.
- Loar, B. 2017. *Consciousness and Meaning: Selected Essays*, K. Balog and S. Beardman, eds., Oxford University Press.
- Pitt, D. 2004. "The Phenomenology of Cognition, Or, *What Is It Like to Think That P?*," *Philosophy and Phenomenological Research* 1, 1–36.
- Pitt, D. 2009. "Intentional Psychologism," *Philosophical Studies* 146, 117–138.
- Pitt, D. 2011. "Introspection, Phenomenality, and the Availability of Intentional Content," in T. Bayne and M. Montague, eds., *Cognitive Phenomenology*, Oxford University Press, 141–73.
- Pitt, D. Forthcoming. "Acquaintance and Phenomenal Concepts," in S. Coleman, ed., *Cambridge Classic Arguments Series: The Knowledge Argument*, Cambridge University Press.