

## Is Phenomenology the Basis of Mental Content?

---

Adam Pautz

Fortunately for the determinate character of intentional content, content determinacy is fixed phenomenally.

--Graham, Horgan and Tienson (2007)

Many go in for the *reductive externalist program* concerning the mind and its intentionality. It was first developed for cognitive side of the mind. The goal of the program is to “naturalize” intentionality in externalist terms. For instance, many have said that at least some beliefs have their contents thanks to the fact that their neural realizers bear an appropriate wide physical relation to certain external conditions. We might call the relevant relation the *tracking relation*, leaving it open whether it is to be explained in terms of causal covariation under optimal conditions (Stalnaker, Tye), asymmetric dependence (Fodor), indicator function (Dretske), or normal conditions for the proper function of output systems (Millikan).

Many have extended reductive externalism to the phenomenal side of the mind. They accept *intentionalism* about experience, according to which the phenomenology of experiences is fully determined by their intentional contents. For instance, an experience of a tomato has a rich built-in intentional content which determines what it is like. So for them the hard problem of experience becomes a special case of the hard problem of intentionality. To solve this problem, they invoke their usual view, claiming that the intentional contents of experiences are fixed by mind-world tracking relations. In this way they are led to a radically externalist view of phenomenology that I call *tracking intentionalism*. On this view, an accidental, life-long brain-in-a-vat duplicate of your brain would not support phenomenology.<sup>1</sup>

The reductive externalist program is vague and programmatic. It faces many longstanding problems of detail. Among them are the disjunction problem, the distance (depth) problem, and problems about content determinacy due to Quine and Kripkenstein.<sup>2</sup> Indeed Lycan (2009: note 1) has recently spoken of the “dismal failure” of all existing proposals within the reductive externalist program, suggesting that this provides the best argument for a non-reductive or “primitivist” approach to intentionality. I tend to agree. It is fair to say that the reductive externalist program is in a state of stagnation.

But recently a new approach has come to the fore, the *phenomenal intentionality program*. This program gives a sense of revolution, of upsetting the applecart. Proponents generally insist that the reductive externalist program ignores the role of phenomenology and consciousness in grounding intentionality. We can finally solve the problems of intentionality by maintaining that phenomenology is in some sense the

---

<sup>1</sup> Defenders include Armstrong (1968), Byrne and Hilbert (2003), Dretske (1995), Harman (1990), Hill (2009), Lycan (2001), and Tye (2000). In Pautz (2006), (2010a) and (2010b) I critically discuss “tracking intentionalism” and characterize the reductive externalism program in terms of a generic “tracking relation” which can be spelled out in various ways. See also Tye and Cutter (2011). Kriegel (in his 2011 and in this volume) also characterizes the reductive externalist program in terms of a generic “tracking relation”.

<sup>2</sup> See Quine (1960) and Kripke (1982). ‘Kripkenstein’ refers to an imaginary proponent of the views that Kripke attributes (some think wrongly) to Wittgenstein.

source of all intentionality. A consequence of the approach is that a life-long zombie could not have “genuine” intentional states.

More specifically, we may identify a few quite extreme theses associated with the phenomenal intentionality program. First, they often accept “prioritism”: they hold that, in all cases where phenomenology and intentionality are intertwined, phenomenology is somehow *metaphysically prior* to intentionality, whereas the aforementioned intentionalists about phenomenology are supposed to hold that intentionality is prior to phenomenology. Some, for instance Siegel (2010: 8.3) and Chalmers (2004: 154), take prioritism to be the defining thesis of the phenomenal intentionality program. Indeed, Horgan and Tienson (2002: 520) build this somewhat obscure thesis into the very definition of ‘phenomenal intentionality’: they define it as intentionality that is possessed “in virtue of” phenomenology, where this stands for an antisymmetric relation of explanatory priority. Their argument for the thesis that there exists phenomenal intentionality (which for them presupposes prioritism) is based on the assumption that it simply amounts to the plausible modal thesis that “there is a kind of intentional content, pervasive in human mental life, such that any two possible phenomenal duplicates have exactly similar intentional states vis-à-vis such content” (2002: 524).<sup>3</sup>

Second, proponents of the phenomenal intentionality program often accept a radical version of phenomenal internalism. For instance, Graham and Horgan (2004) conjecture that, for *any* arbitrary metaphysically possible *physical* duplicate of one’s brain, most of us will have an “intuition” that it is also a *phenomenal* duplicate of oneself. Given the thesis of phenomenal intentionality, this would mean that it is also shares many of one’s intentional states. They think the intuition even applies to a *lone, life-long, accidental* “brain-in-a-vat” duplicate of one’s brain, which has always been totally disconnected from the environment (see also Kriegel this volume: sect. 2.1). This yields a quite radical form of brain-based phenomenal internalism. *If* their radical internalism could be established simply by consulting our intuitions, it would provide a quick and easy argument against all possible versions of the reductive externalist program which seek to ground all intentionality and even phenomenology in tracking relations to the environment.<sup>4</sup>

Third, proponents of the phenomenal intentionality program are often phenomenal liberals as opposed to phenomenal conservatives. For instance, they hold that occurrent beliefs have a special “cognitive phenomenology” *distinct from* associated sensory phenomenology, understood broadly to include perceptual, imagistic, emotional and agentic phenomenology. Indeed, just as they say that the sensory phenomenology of perceptual experiences determines their intentional contents, they often say that the non-sensory cognitive phenomenology of occurrent beliefs *determines* their (nar-

---

<sup>3</sup> In Pautz (2011b) I show that this argument for prioritism fails because such modal theses *leave open* issues about priority or grounding (Schaffer 2009: 364). Indeed standard intentionalists like Dretske (1995), Lycan (2001), and Tye (2000) who reduce phenomenology to intentionality would accept Horgan and Tienson’s modal thesis but would certainly reject their priority thesis. No convincing argument for (or even explanation of) phenomenal intentionality has yet been provided, if it presupposes the obscure thesis of prioritism (Pautz 2008).

<sup>4</sup> Against this, in view of the intuitive possibility of absent qualia and inverted qualia in physical duplicates, if we attempt to decide the matter merely on the basis of “intuition”, then we should arguably assign equal credence to the various possible hypotheses about the phenomenal lives of particular hypothetical accidental internal physical duplicates of us in other metaphysically possible worlds. It is not even intuitive that phenomenology supervenes with metaphysical necessity on an individual’s *total* (physical and non-physical) intrinsic character (*pace* Hawthorne 2004), as witness the coherence of radically externalist act-object (Jackson 1977: 77-8, 102-3) and naïve realist theories (Alston 1999: 191) on which all perceptual phenomenology (even hallucinatory) essentially involves a relation to contingently-existing items (sense data, regions of space) located before the subject. Horgan and Tienson (2002: 526) do provide a different, quasi-empirical argument for their brain-based internalism; but Lycan (2008) effectively criticizes it.

row) contents. And they declare that this finally solves the longstanding problems raised by Quine, Kripkenstein and others concerning content determinacy.

If one wanted a single slogan for the phenomenal intentionality program, it might be *phenomenology first*, since at its heart is the idea that phenomenology (or consciousness) plays a foundational role in grounding all intentionality. As for the nature of phenomenology itself, proponents of the phenomenal intentionality program have had very little to say.<sup>5</sup> Of course, provided they accept radical internalism about sensory and cognitive phenomenology, they must reject all reductive externalist theories of such forms of phenomenology, for instance “tracking intentionalism”. Surprisingly, they have simply not addressed the question of whether they might put a reductive *internalist* theory of phenomenal intentionality in their place. There is reason to doubt that such a theory could be provided. For instance, Graham and Horgan (2004: 305) claim that a life-long a life-long accidental brain-in-vat might stand in a phenomenologically-constituted “acquaintance-relation” to various *shapes* and other spatial properties and have beliefs about such properties. How might they explain the brain’s intentional relations to shapes in physical terms, given that its internal states do not have the function of indicating any spatial properties (or indeed properties of *any* sort) or producing any behavior to move in real space (Evans 1982: chap. 6)? They must apparently take such intentional relations to be primitive but somehow mysteriously supervenient on purely internal factors.<sup>6</sup>

My own view is that the specific theses listed above are very underdeveloped and poorly motivated, but there are defensible theses in the vicinity. So my own approach falls within vague boundaries of the phenomenal intentionality program. For instance, instead of the general thesis that phenomenology is “prior to” intentionality, I have elsewhere argued for what we might call (following David Chalmers) an *integrativist* or *no-priority* view concerning the relationship between the intentionality and phenomenology of *sensory* experience. I also believe that proponents of the phenomenal intentionality program are quite wrong to put any stock in intuition-based arguments for phenomenal internalism. But I have argued that a modest version of phenomenal internalism can be established *empirically*, and that it is enough to rule out externalist “tracking” theories of sensory intentionality.<sup>7</sup>

In this paper, I will be mainly concerned with “cognitive phenomenology”, although these issues concerning prioritism and phenomenal internalism will also come up. In particular, I will focus on what is perhaps the least discussed thesis associated

---

<sup>5</sup> My “phenomenology first” slogan is inspired by Williamson’s (2000) “knowledge first” slogan. Williamson holds that knowledge plays a foundational role in the epistemic domain and that it is (at least conceptually) irreducible. Likewise, proponents of the phenomenal intentionality program hold that phenomenology plays a foundational role in the mental domain and appear at least open to the view that it is (ontologically) irreducible.

<sup>6</sup> Pautz (2010a: sect. 7) argues in detail that the radical internalism defend by Horgan, Graham and Tienson requires a primitivist, “no-theory” theory of phenomenal intentionality, a theory that I happen to favor myself. For instance, their radical internalism is not consistent with the more modest internalist theory of intentionality formulated by Lewis (1994: 425) and Braddon-Mitchell and Jackson (2007: 240). Nor is it consistent with Kriegel’s response-dependent tracking intentionalism (for critical discussion see Pautz 2010b: sect. 5); for one thing, the internal states of an accidental brain-in-the-vat do not “track” *any* properties, including Kriegel’s response-dependent properties, in the sense in which this notion explained by Dretske, Tye and others (*pace* Kriegel 2011: 3.2.3).

<sup>7</sup> Pautz (2011b) develops my general take on the phenomenal intentionality program. For empirical arguments in favor a modest version of phenomenal internalism and against externalist tracking theories of sensory intentionality, see Pautz (2006), (2010a) and (2010b). For discussion of these empirical arguments, see Chalmers (2005), Cohen (2009: 81-88), Hill (2012), and Tye and Cutter (2011). For integrativism as against prioritism, see Pautz (2008) and Chalmers (2008). Kriegel (this volume: 1.3) backs off from the claim of prioritism that there is an “antisymmetric” explanatory relation between phenomenology and intentionality; he also says (2011: 63) that he is attracted to the kind of “integrativist”, identity view that I defended in my (2008).

with the phenomenal intentionality program: the thesis that the “cognitive phenomenology” of our occurrent beliefs *determines* their (narrow) contents, thus answering long standing problems about content determinacy due to Quine and Kripkenstein.

My plan is as follows. First (§1-2) I will argue that the thesis that cognitive phenomenology is the basis of mental content is very underdeveloped and poorly motivated. Then (§3) I will develop several arguments against it. Finally (§4), I will briefly argue in favor of what I regard as a much more promising approach in the same general vicinity. Drawing on David Lewis’s important and influential work on intentionality, I will sketch a view I call *phenomenal functionalism* which entails that *sensory* phenomenology is the source of all determinate intentionality. (Phenomenal functionalism has similarities to views recently proposed by Schwitzgebel and Chalmers.) Phenomenal functionalism avoids an overlooked problem for Lewis’s theory involving the notion of sensory evidence, which is crucial in Lewis’s approach to ruling out “deviant interpretations”.

## 1. Cognitive Phenomenology as the Basis of Mental Content?

I begin by describing my target in greater detail. As I said, proponents of the phenomenal intentionality program typically say that occurrent beliefs (as opposed to your unconscious standing beliefs) have *cognitive phenomenology*. This is a kind of phenomenology attaching to occurrent beliefs that is *distinct from* associated sensory phenomenology, understood broadly to include perceptual, imagistic, emotional and agentive phenomenology. They also say that occurrent desires have a non-sensory conative phenomenology, although in what follows I will focus mostly on occurrent beliefs. (I will use ‘cognitive phenomenology’ broadly to cover the alleged non-sensory phenomenology of desire as well as the alleged non-sensory phenomenology of belief.) Call this rough thesis the *CP-existence thesis*.<sup>8</sup>

Consider an example due to Horgan and Tienson (2002). You hear ‘Visiting relatives can be boring’ first as a remark about the people who are visiting and then as a remark about visiting certain people oneself. Horgan and Tienson say that the actual sound or auditory imagery may be the same, but the total experiences are phenomenally quite different. They explain this by saying that the occurrent beliefs you have on the two occasions have different cognitive phenomenologies. Indeed, they say (2002: 522) that cognitive phenomenology is quite rich: “Change either the attitude-type (believing, desiring, wondering, hoping, etc.) or the particular intentional content, and the phenomenal character thereby changes too”.

I myself find the CP-existence thesis alone to be not very interesting, even though it has been much discussed. Even those who think the CP-existence thesis is false in our world would presumably grant that it is true in other possible worlds. In some worlds a special phenomenology over and above sensory phenomenology attaches to believing and desiring. Why then does it matter so much if our own world is one of those in which the CP-existence thesis is true? Here I will not be primarily focused on the CP-

---

<sup>8</sup> Traditional empiricists held that sensory phenomenology is quite thin and made up of momentary “raw feels” or “givens”. But those who deny the CP-determination thesis are not committed to this kind of traditional empiricist conception of sensory phenomenology. They might plausibly maintain that sensory phenomenology is rich and itself imbued with intentionality. For instance, there is a genuine sensory-perceptual difference between seeing the duck-rabbit figure as a duck and seeing it as a rabbit; and this essentially involves some difference in the intentional contents of the experience. Those who accept CP-determination hold that, even on such a rich conception of sensory phenomenology, there is more phenomenology than sensory phenomenology. There is, for instance, also a purely non-sensory, non-perceptual difference between entertaining different mathematical thoughts.

existence thesis. I will assume the reader is familiar with the thesis and the arguments in its favor.

What I think is more interesting, and what I will primarily focus on here, is a stronger thesis that has not been much discussed. Here are some representative passages:

[H]ow can [cognitive] experience ever deliver determinateness? It just can. Cognitive experience in causal context can do just this. Such is its power. (When it comes to [thinking of] the number 2, it doesn't even require causal context.) . . . If God could look into my mind and apprehend the cognitive [phenomenology] of my experience he would certainly know what I was thinking about, given that he also knew—and how could he not—about my causal circumstances. It is the same power that makes it the case that I can think determinately about the number 2 although there is no relevant causal context. *Pffff!* This is the correct account of how it is that content can be determinate in spite of all the problems raised for this idea by Kripke in his book *Wittgenstein on Rules and Private Language*. (Strawson 2010: 351, 354)

The part of what is thought that is fully determined by [cognitive] phenomenal character [is] a kind of thought content. (Siewert 2011)

[E]ach specific occurrent intentional state with phenomenal intentional content is constitutively determined by its own distinct phenomenal character – viz., the what-it's-like of undergoing that particular attitude-type vis-à-vis that particular phenomenal intentional content. . . Suppose that you are now undergoing a psychological state with the distinctive [non-sensory] phenomenal what-it's-like of believing that a picture is hanging crooked on a wall directly behind you. Then you thereby believe that there is a picture hanging crooked on a wall directly behind you; undergoing this phenomenology constitutively determines that you are instantiating that belief-state. Any experiencing creature [e. g. a brain in a vat] undergoing this phenomenology would thereby instantiate the belief-state, even if its overall phenomenology is otherwise quite different from your own. (Horgan and Tienson 2002: 526)

Physically and apart from phenomenology, there is no “one, determinate, right answer” to the question of what is the content of an intentional state. For . . . the content of each mental state is not determinately fixed once the physical facts (including perhaps physical facts about the internal-physical linkages) are fixed. Fortunately, however, for the identity or determinate character of intentional content, content identity or determinacy is fixed phenomenally. For example, the what-it's-like of thinking “Lo, a rabbit” is different from the what-it's-like of thinking “Lo, a collection of undetached rabbit parts”. . . [This] commitment to phenomenal individuation of intentional content, combined with rejection of physical individuation, [might be] tantamount to dualism. (Graham, Horgan and Tienson 2007: 476, 481)

You know what you are thinking and what you mean by your utterance, and there is a determinate fact of the matter about what you are thinking and what you mean by your utterance, because there is something it is like to think a determinate thought and to make an utterance that expresses that thought. Developing in further detail our proposed account of content determinacy is a task beyond the scope of the present paper—and an agenda item for the future. (Horgan and Graham 2010)

[My view] maintains that the intentional content of a thought is determined by its *intrinsic* phenomenal properties, *not its relational properties*. My teachers will be very disappointed in me. (Pitt 2009: note 5, my italics)

Let a *cognitive phenomenal property* be a property of individuals of the form *having a state with such-and-such cognitive phenomenology*. All of the quoted philosophers apparently endorse the following *CP-determination thesis*: for at least some cognitive phenomenal properties *P*, there is a unique content *c* such that it is metaphysically necessary that, if an individual has *P*, then he has an occurrent belief (or desire) with content *c*.<sup>9</sup> On this view, “cognitive experiences” are not mere “raw feels” or bits of “mental paint”, akin to a state of undirected depression. Instead, they have built-in intentionality, just as perceptual experiences arguably have built-in intentionality.

Let me make some clarifications. First, CP-determination of course cannot be applied to our unconscious, standing beliefs and desires which certainly lack phenomenology. To explain such beliefs and desires, the proponent of CP-determination needs a different account (for some options see Graham, Horgan and Tienson 2007). This will not concern us here. Second, on the assumption that it is “intrinsic”, cognitive phenomenology of course cannot determine the “wide contents” of our beliefs and desires which can differ between internal duplicates: natural kind contents (e. g. about water or rabbits), singular contents (e. g. contents involving a particular river), and so on. Consequently, proponents of CP-determination only claim that cognitive phenomenology determines the *narrow* (and perhaps “centered”, *de se*) contents of beliefs and desires that do not differ between such duplicates. In this category they include mathematical contents (Strawson mentions contents about the number 2), certain general descriptive *de se* contents (the watery stuff of my acquaintance is wet), contents about perceptible properties (for instance, *that there is a red and round thing there*), artifactual contents (Horgan and Tienson mention a content about a picture), and so on. Some such narrow contents might be character-like entities that only determine a truth-value relative to a “centered world”.<sup>10</sup> To explain *wide* content, proponents of

<sup>9</sup> I will ignore *degrees* of belief and desire. (Would advocates of CP-determination say that in some cases these too are fixed by non-sensory feel, rather than functional role?)

<sup>10</sup> It might be wondered whether Strawson actually commits to CP-determination in the passage quoted above, since he suggests that in fixing content “causal context” (a subject’s causal relations to the outside world) is relevant in addition to cognitive phenomenology. I think he does so commit, since he says that a certain cognitive phenomenal property necessitates thinking of the number two, and that in this special case causal context is not relevant. Further, he says that cognitive phenomenology solves the Kripkenstein problem about content determinacy. So he presumably thinks that some cognitive phenomenal properties alone determine entertaining a *complete* mathematical content. It is also worth stressing that while Strawson is correct that in typical cases causal context and not just cognitive phenomenology might be relevant to *wide* content (e. g. believing a singular content to the effect that a particular tree is tall), it is open to him to say that even in these cases cognitive phenomenal alone determines believing a certain underlying descriptive, narrow content (e. g. the tree-like thing before me is tall), without help from causal context.

CP-determination need some other account. For instance, Horgan and Tienson adopt David Lewis's (1994) view that wide content is derivative from narrow content and relations to the environment.

The CP-determination thesis is interesting for two reasons. First, it is extremely unorthodox. Let us assume the standard relational view of belief and desire: to have a particular belief or desire is to stand in the belief relation or the desire relation to a particular proposition or "content". Then the puzzle of intentionality can be put like this: how is it that one manages to stand in the belief relation and the desire relation to various contents? Here is a very rough statement of the orthodox view:

- The contents of (occurrent and non-occurrent) beliefs and desires are always determined by *non-intrinsic factors*: factors such as behavioral dispositions, wide relations to the environment, causal or inferential relations among internal states, what sentences one accepts and their contents as determined by their overall pattern of use in the language, and so on.

Granted, philosophers like Lewis and Jackson recognize a notion of "narrow content". But even they accept the non-intrinsic claim. For, on their notion of "narrow content", narrow content is *not* entirely intrinsic to you, because it is determined by the typical world-involving functional roles of your internal states in your population, which are not intrinsic.<sup>11</sup>

The CP-determination thesis could not be more different. According to it, "cognitive phenomenology" determines that our beliefs and desires have certain contents. Further, Horgan and Tienson (2002) and Graham and Horgan (2004) say that all phenomenology is "intrinsic" and constitutively independent of anything outside the brain.<sup>12</sup> Thus an accidental, life-long brain-in-vat duplicate of yourself would have all the same cognitive phenomenal properties. Likewise in the above quotation Pitt says that phenomenal properties are "intrinsic" and "non-relational". Putting this together with CP-determination, we get the following very unorthodox idea:

- Some of a subject's occurrent beliefs and desires have contents that are fully determined by certain of his *intrinsic properties* at

---

<sup>11</sup> For this point, see Lewis (1994: 425) and Braddon-Mitchell and Jackson (2007: 240).

<sup>12</sup> What exactly do Horgan and coauthors mean by the claim that phenomenology is intrinsic? I find their own often-repeated gloss difficult to interpret: "phenomenology is not constitutively dependent on anything outside *phenomenology itself*". By 'anything outside phenomenology itself' do they mean any state "whose nature is describable in non-phenomenological [e. g. physical] language", in the words of Horgan and Tienson (2002: note 23)? But then the intrinsicness thesis entails *dualism*. Alternatively, do they perhaps mean anything that phenomenology does not constitutively depend on? But on this interpretation the intrinsicness thesis is equivalent to the utterly trivial thesis that phenomenology does not constitutively depend on anything it does not constitutively depend on. There are also problems with their often-repeated thesis that "phenomenology does not depend constitutively on factors outside the *brain*" (e. g. Horgan and Tienson 2002: 526), which they assert to be intuitively certified. If they are right, then we can rule out virtually *a priori* both Humeanism about laws (as Hawthorne 2004 explains) and even all versions of *substance dualism* on which phenomenology constitutively depends on the properties of souls which are not literally located in the brain. Here is the gloss of the intrinsicness thesis I will assume here: all phenomenal properties of the form *having an experience with so-and-so phenomenal character* are intrinsic properties of subjects, where an intrinsic property is intuitively one whose instantiation depends only on the character of its *bearers* (whether physical or non-physical) and does not constitutively depend on anything outside its *bearers*. Contrary to Horgan and Tienson (2002: note 23), I do not believe that this general thesis "is self-evident to reflective introspection" (see note 4 of the present paper). In any case, my aim here is not to evaluate the thesis that phenomenology is intrinsic, but to evaluate CP-determination on the assumption that it is true.

the time he forms the beliefs and desires (namely, cognitive phenomenal properties), where the relevant intrinsic properties are *distinct from* all of his sensory and functional properties (past, present and future).

There is a second reason why the CP-determination thesis is interesting besides its unorthodox character. Many of those quoted above declare that it finally solves longstanding content determinacy worries due to Quine and Kripkenstein. Here is a rough formulation of the Kripkenstein problem as it arises for physicalists. Let  $quus_1$ ,  $quus_2$ ,  $quus_3$  denote different functions defined over numbers that are just like the plus function but that differ from the plus function only when it comes to numbers that are too large for us to compute. Then, at least if we set aside the widely rejected view that necessarily equivalent propositions are identical, we have infinitely many distinct contents up in Plato's heaven: that two plus two equals four, that two  $quus_1$  two equals four, that two  $quus_2$  two equals four, that two  $quus_3$  two equals four, and so on. If we say that the non-intentionally characterized functional and behavioral facts determine (in the sense of *metaphysically necessitate*) that one believes one of these contents (in particular, *that two plus two equals four*) to the exclusion of all the others, then we want some kind of *explanation* for this. (Likewise we want an explanation for how they determine that the expression '+' refers to one of them to the exclusion of the others.) *How* do these facts *select* or *point to* that content to be what you believe? Is there a physicalist-functionalist (perhaps *a posteriori*) account of the belief relation which *explains* how the physical facts determine that this is the content you believe, as opposed to the other candidates? In my view, the main problem here is that our dispositions are in some clear sense *finite*. We have (non-intentionally characterizable) dispositions involving a certain finite set of numerals in language (and perhaps in the language of thought) but we do not have dispositions involving numerals that are too long for us to compute. Quine's well-known problem about rabbits and undetached rabbit-parts is related to the Kripkenstein problem. In the above quotations, Strawson as well as Graham, Horgan and Tienson suggest that in these cases cognitive phenomenology somehow manages to determine what you believe and to rule out deviant interpretations.

The CP-determination thesis is underdeveloped. One does not get much more than the bare assertion that cognitive phenomenology determines content. This leaves many obvious questions unanswered. What exactly is the relationship between cognitive phenomenal properties and intentional properties? What is the relationship between cognitive phenomenal properties and physical-functional properties? I will be looking at these questions.

Let me lay some cards out on the table. The CP-determination thesis presupposes the CP-existence thesis. I find even the CP-existence thesis very hard to believe. In fact, I think this issue *may* be irresolvable (or merely verbal) and so not worth discussing. However, my main aim is to argue that, even if the CP-existence thesis is true, the stronger and more interesting CP-determination thesis is not. Along the way I will present some reasons to doubt the CP-existence thesis as well.

Even those who reject the CP-existence thesis should find the discussion interesting. As already mentioned, given that cognitive phenomenal properties are *intrinsic* and *distinct from* all sensory and functional properties, the CP-determination thesis entails a more general thesis: that some occurrent *belief-desire* properties are likewise intrinsic and distinct from all sensory and functional properties. Belief and desire

form a level of mental reality distinct from the sensory and functional levels. As we shall see, this general irreducibility thesis arguably in turn entails that the belief and desire facts do not even completely *supervene on* the totality of sensory and functional facts. I think that this general claim is what is really interesting about the CP-determination thesis. And it is actually neutral on the CP-existence thesis. For instance, Searle (1992) holds that some occurrent *belief-desire* properties are intrinsic and distinct from all sensory and functional properties; but he does not explicitly endorse the CP-existence thesis. This general idea is my chief target in what follows.

## 2. The CP-Determination Thesis as Unmotivated

Let us start by looking at arguments *for* the CP-determination thesis. Contrary to some proponents, it is certainly not introspectively obvious. I will consider two more elaborate arguments for CP-determination, suggesting that they fall well short. No persuasive argument has yet been given for CP-determination.

### 2.1 First argument: the determinacy argument.

As we saw, many declare that CP-determination solves determinacy problems due Quine and Kripkenstein. In several places, Graham, Horgan and Tienson have briefly argued for CP-determination on the grounds that it provides the *only* adequate solution to these problems, as follows:

- 1 Against Quine and Kripkenstein, there generally are determinate facts about what we believe and mean.
- 2 But Quine and Kripkenstein are right that such facts could not be “determinately fixed by” the physical facts
- 3 If premise 2 is true, then there could be determinate facts about what we believe and mean only if they are fixed by *non-sensory, cognitive* phenomenology: this is the only alternative to physical determination.
- 4 So, there could be determinate facts about what we believe and mean only if they fixed by cognitive phenomenology. (2, 3)
- 5 So, what we believe and mean is fixed by cognitive phenomenology. (1, 4)<sup>13</sup>

There are several problems with this argument.

(i) Against premise 2, there are numerous important physicalist answers to Quine and Kripkenstein’s arguments in the literature on “naturalizing intentionality”. But Graham, Horgan and Tienson do not provide an in-principle reason to think that all they fail, nor do they eliminate them one by one.

(ii) In support of premise 3, Graham, Horgan and Tienson (2009: 531) say that if we “[agree] with Quinean misgivings about the physical determination of content, then all that’s left as a mode of individuation is phenomenological individuation”. From the examples they discuss (they speak of the what-it’s-like of thinking “Lo, a rabbit”) it is clear they mean *cognitive* phenomenology and not mere *sensory* phenomenology; and Horgan and Graham (2010) explicitly say that cognitive phenomenology over and above sensory phenomenology is needed to secure our actual level of content

<sup>13</sup> See Graham, Horgan and Tienson (2007) and (2009) and Horgan and Graham (2010).

determinacy. (Otherwise they do not have an argument from determinacy for *cognitive* phenomenology.) But they miss *phenomenal functionalism*: the modified, non-reductive version of Lewis's functionalism which in §4 I will recommend as an alternative to CP-determination. Briefly, phenomenal functionalism "agrees Quinean misgivings about the physical determination of content" in a sense: facts about the determinate content do not reduce all the way down to purely physical facts, for instance non-intentionally characterizable behavioral and functional facts. But they do reduce to clusters of narrow and wide functional facts *and first-person sensory* facts, where sensory facts are themselves *richly intentional* (e. g. essentially involve the representation of observational properties) and *not* reducible to purely physical facts. To show that this general type of view is wrong and that they are right in maintaining that *cognitive* phenomenology over and above sensory phenomenology is needed to secure the level of content determinacy that our intentional states actually possess, Graham, Horgan and Tienson would need to show that a community which is *exactly like* us in all physical respects *and* all actual and dispositional *sensory* respects (same visual experiences, same language, etc.), but which lacks the occasional extra bits of "cognitive phenomenology" that we allegedly enjoy (an *absent cognitive qualia case*), would differ profoundly from us in having intentional states whose contents are much less determinate than our intentional states. In short, to make good on their determinacy argument for cognitive phenomenology as distinct from sensory phenomenology, they would have to show that the intentional facts about a community *fail to supervene* on the total sensory and functional facts of a community. But Graham, Horgan and Tienson have certainly not shown this or even addressed the issue. How *could* they show this? Later on I will argue that this supervenience claim is actually correct (§§3.2-3.3).

(iii) In maintaining that the only alternative to physical determination is determination by cognitive phenomenology (premise 3), Horgan and Tienson ignore yet another option, one that is historically well-known. As Boghossian (1989: 542) notes, Brentano's thesis of the irreducibility of intentional idioms answers Kripkenstein and gives you content determinacy but does not require cognitive phenomenology. So considerations concerning content determinacy alone simply do not justify acceptance of CP-determination at all. Horgan and Tienson cannot object to this rival position on the grounds that primitivism is intolerable. As we will see (§3.1), their own brief remarks in response to Quine and Kripkenstein suggest a primitivist (or in their words "dualist") view of content determinacy.

(iv) Finally, against premise 1, one might say that, while perhaps some of our beliefs clearly determinately possess certain contents (e. g. immediate perceptual beliefs) thanks to having an especially close connection to sensory experience, other beliefs (beliefs about electrons, arithmetical beliefs) must be admitted to be radically indeterminate in content. In response, Horgan and Graham (2010) would rely on their often-repeated claim that in general content determinacy "is just obvious", contrary to Quine and others. But, according to them, in the *absent cognitive qualia case* mentioned above, when Horgan and Graham's counterparts (sensory-functional duplicates) say (just as they do) that "in general it is obvious that there is content determinacy", their counterparts are saying something wrong, for their counterparts in this case lack the cognitive qualia which, according to them, are necessary for securing general content determinacy. So by their own lights intuitions about content determinacy are fallible.

### 2.2 Second argument: the access argument

Very roughly, the first premise of this argument (e. g. Pitt 2011) is that one has some kind of “special access” to one’s occurrently believing or desiring that *p*. The second premise is that this requires that occurrently believing or desiring that *p* have an “individuating” non-sensory, cognitive phenomenology, in a sense that entails CP-determination.

But there are two problems with the second premise of this argument. First, there are many models of how one has special access to one’s occurrent and standing beliefs and desires that do not require CP-determination. Indeed, even Pitt’s own *acquaintance theory* does not require that our occurrent thoughts be phenomenal states, since Russell (1912: 105) and more contemporary acquaintance theorists allow for immediate, acquaintance-based knowledge of *non-phenomenal* states (e. g. two plus two equaling four).<sup>14</sup> Proponents of the access argument have not provided a principled argument for thinking that all such alternatives fail. Nor have they exhaustively eliminated all of them one by one.

A second problem is this. Consider once again the *absent cognitive qualia case*. It is exactly like the actual case in all respects (sensory, perceptual, functional, physical) except that we lack cognitive qualia. The second premise of the access argument has a very radical consequence: that in this case our counterparts *could not* enjoy the same kind of special access to their occurrent beliefs and desires that we enjoy. This radical claim strikes me as very implausible. By hypothesis, in this case the same inner sentences would be running through their heads; they would manifest understanding these sentences just as we do; they would have the same sensory images and feelings; they would say ‘I certainly have conscious beliefs and desires and know what they are’ and would just as easily provide generally right answers about what they occurrently believe and desire. All of this provides very strong evidence that they have “special access” to their occurrent beliefs and desires, despite their lack of “cognitive qualia”. In any case, proponents of the access argument have never argued otherwise or even considered absent cognitive qualia cases, even though this is precisely what the argument requires.

## 3. Arguments Against CP-Determination

Hence no convincing argument for the CP-determination thesis has yet been provided. I will now develop several arguments against it. Since there are no arguments for it and several arguments against it, it does not merit belief.

### 3.1 First argument: the danglers argument

My first argument is that there are reasons to interpret proponents of CP-determination as “anti-reductionists” or “primitivists” about the intentionality of occurrent belief and desire. In a sense to be explained, they accept what is sometimes called “Brentano’s thesis”. But their primitivism requires an especially objectionable form of “danglers”, in Smart’s (1959) sense.

Let me first say why I think that proponents of CP-determination are best understood as “primitivists”. What in their view is the nature of cognitive phenomenology, such that it delivers determinate content? They never say. They *might* take a reductive physicalist view of cognitive phenomenology. Here I understand reductive views

<sup>14</sup> The acquaintance model is in any case problematic (Pautz 2011).

broadly to include views that identify mental properties with physical properties and physically-realized functional properties. There are two types of reductive physicalist theories of *sensory* phenomenology: *brain-based* or *biological theories* and *functional theories*. Proponents of CP-determination might apply the same theories to cognitive phenomenology. But there is reason to think that they would not.

To see this, consider first the *biological theory* of cognitive phenomenology, akin to the familiar biological (type-type identity) theory of sensory phenomenology. To illustrate, let  $P$  be the non-sensory phenomenal property you allegedly possess when you occurrently believe that two plus two equals four on a certain occasion. On the biological theory of cognitive phenomenology,  $P$  is *necessarily* identical with neuro-computational property  $M$ , which has its connections with the external world and behavioral outputs only *contingently*, so that a brain in a vat could have  $M$ . Maybe  $M$  is of the form: *tokening mentalesse sentence  $s$* . The biological theory of phenomenology might be the *only* reductive physicalist theory compatible with the claim of Horgan and his coauthors that phenomenology in general is *intrinsic* (Pitt says “non-relational”). Recently, Levine (2011) has actually advocated a theory of cognitive phenomenology in terms of awareness of inner brain-writing that is similar to the biological theory, although he is not a proponent of CP-determination.

Is the biological theory a plausible model for “naturalizing” phenomenal intentionality? Might proponents of CP-determination accept the biological theory of cognitive phenomenology? They cannot, because this combination of views would yield a radically implausible view of the ground of (narrow) intentionality. By CP-determination, having cognitive phenomenal property  $P$  determines *as a matter of metaphysical necessity* believing two plus two equals four. Putting this together with the claim of the biological theory that  $P$  is necessarily identical with neuro-computational property  $M$ , we get the result that having  $M$  determines as a matter of metaphysical necessity believing two plus two equals four. In general, the conjunction of CP-determination and the biological theory entails that, for every narrow (non-Twin-Earthable) belief and desire, there is a single neural state that metaphysically necessitates having that belief or desire. Call this *biological CP-determination*. It is analogous to what I have elsewhere called *biological intentionalism* about sensory intentionality (Pautz 2010b; 2011b).

Biological CP-determination is unacceptable for two reasons. First, why should  $M$  metaphysically necessitate occurrently believing the plus-content as opposed to a bent, quus-like content or indeed any other content? There is no credible general account of intentionality from which the modal claims like “necessarily, neural duplicates believe the same narrow contents” might be derived (Pautz 2010a: sect. 7). They would apparently just have to be taken as *brute* metaphysical necessities (in the sense of Dorr 2007). Biological CP-determination can fairly be described as a *magical theory* of narrow intentionality. Second, biological CP-determination is open to what we might call the “separation argument”. (Pautz (2010c: sect. 4) develops this argument against biological theories of *sensory* intentionality.) According to biological CP-determination, having the mere neuro-computational property  $M$  metaphysically necessitates believing that two plus two equals four. Against this, there are possible “separation cases” in which  $M$  (perhaps just a bit of brain-writing) plays no interesting functional role with respect to collections of objects, counting, imagery and so on, and also worlds in which it plays a functional role appropriate to quite different arithmetical beliefs. In these cases  $M$  is “separated from” its actual functional role. It would be absurd to sug-

gest that in these possible cases *M* realizes believing that two plus two equals four. So biological CP-determination is not credible.

In short, the biological theory of cognitive phenomenal properties (Levine 2011) entails that cognitive phenomenal properties are raw feels like a feeling of undirected depression that do not necessarily determine any intentional features, contrary to CP-determination.

Most philosophers think that intentionality is grounded in functional role, not mere neural states. Accordingly, the friend of CP-determination might consider providing a *reductive functionalist theory* of cognitive phenomenology. To illustrate, consider a non-arithmetical example. In the passage quoted in §3.1 Horgan and Tienson say that there is a unique cognitive phenomenal property that metaphysically necessitates believing *that there is a picture hanging on the wall behind one*. Call it *P* for “picture”. Reductive functionalist theories of properties like *P* come in narrow and wide versions.

On a *narrow version* of reductive functionalism about cognitive phenomenology, the cognitive phenomenal property *P* that determines the picture-content is identical with some narrow functional property, where a narrow property is (roughly) one that is shared by suitable total physical duplicates. To choose a specific narrow functional property, the proponent of CP-determination could draw on his favorite narrow functionalist or “conceptual role” theory of content and concept-possession. Maybe he thinks, as I do, that only Lewis (e. g. 1972, 1974) and Loar (1981) have actually begun to work out the details of such a theory, explaining how functional roles might be systematically paired with contents. Then he might simply identify the cognitive phenomenal property *P* with the very same holistic, world-involving functional property that Lewis would say determines occurrent belief that there is a picture hanging on the wall. In general, his view of the basis of intentionality might be identical with Lewis’s view, except he adds that Lewisian functional properties that ground intentionality have a special non-sensory “feel” to them. Now on Lewis’s view there is a sense in which some belief and desire properties are *narrow*, which is certainly something defenders of CP-determination would agree with. But, as I mentioned above, even Lewis’s narrow theory is incompatible with their much more unorthodox claim that some belief and desire properties are entirely *intrinsic* (Pitt says “non-relational”), for reasons explained by Lewis (1994: 425) and Braddon-Mitchell and Jackson (2007: 240). But it is a possible view for fans of CP-determination willing to forsake the intrinsicness claim.

On a *wide version* of reductive functionalism about cognitive phenomenology, the cognitive phenomenal property *P* is instead identical with a wide functional property. To settle on a specific wide functional property, the proponent of CP-determination might look to his favorite wide theory of intentionality. For instance, if he likes Fodor’s (1994) theory of content determination, he might say that the cognitive phenomenal property *P* is identical with the complex functional property involving the tokening of a mentalese sentence, functional role, and tracking (asymmetric dependence) relations to the external world that Fodor would say constitutes the occurrent belief that there is a picture hanging on the wall. This would yield a tracking theory of cognitive phenomenology very similar to the familiar kind of tracking intentionalism about sensory phenomenology (defended by Tye and Dretske and others) that I mentioned in the introduction. The idea is that cognitive phenomenology and hence cognitive intentionality is not in the head, just as tracking intentionalists insist that sensory phenomenology is not in the head. Evidently, this radically externalist theory of

cognitive phenomenal properties is incompatible with the claim of Horgan and his coauthors that phenomenal properties in general are *intrinsic* (Pitt says “non-relational”). Nevertheless, it is a possible view. In the introduction, I suggested that the phenomenal intentionality program is in general opposed to the reductive externalist program, in part because it typically is associated with the claim that phenomenal intentionality is intrinsic and fully internally-determined. But, as I have pointed out elsewhere (Pautz 2008: 272-273), in principle the core thesis of the phenomenal intentionality program that phenomenology is the source of intentionality is quite compatible with the program of naturalizing phenomenology and intentionality in terms of “tracking relations” (a point Kriegel (2011: chap. 2) develops into an entire theory).

Let *functional CP-determination* be the conjunction of the CP-determination thesis and some such a wide or narrow functionalist (or “conceptual role”) theory of cognitive phenomenal properties. Might friends of CP-determination then accept functional CP-determination, thereby providing a “naturalistic” theory of how cognitive phenomenal properties ground intentionality? The trouble is that their other commitments rule out functional CP-determination.

(i) As I have already said, Horgan and coauthors as well as Pitt say that cognitive phenomenal properties are intrinsic and “non-relational”, which rules out standard functionalist and tracking theories of such properties (see also note 6).

(ii) As we saw (§2.1), Horgan and coauthors reject all *orthodox physicalist solutions* to the determinacy problems due to Quine and Kripkenstein. And the quote cited in §1 suggests that Strawson rejects orthodox physicalist solutions to the Kripkenstein problem. But functional CP-determination provides the *very same problematic orthodox physicalist solutions* to these problems developed by Lewis and Fodor in terms of functional role, inferential-conceptual role, tracking, and so on. It is at bottom reductive physicalist account of content determinacy. The only difference is that it adds that the physical-functional properties that pin down content have a special non-sensory, cognitive phenomenology.

(iii) Presumably, fans of the phenomenal intentionality program would say that their claim that phenomenology grounds intentionality plays an *essential role* in solving the well-known problems of detail plaguing the reductive externalist program and the naturalization program more generally. But, on functional CP-determination, this is not so. To repeat, functional CP-determination provides the very same problematic orthodox physicalist solutions to these problems developed by Lewis, Fodor and others in terms of functional role, inferential role, tracking, and so on. The only difference is that it adds that the physical-functional properties that pin down content have a special non-sensory, cognitive phenomenology. But this claim is not essential to these solutions. It is not doing any extra work.

Since I can think of no plausible reductive theory of cognitive phenomenology consistent with the things proponents of CP-determination have said, I provisionally conclude that they must accept *primitivist CP-determination*. Indeed, the passages from Horgan and Tienson and Strawson quoted in §3.1 clearly suggest some kind of “primitivist” view of cognitive phenomenology.

Primitivist CP-determination comes in two possible versions. On a quite extreme version, there is *no real definition* of a cognitive phenomenal property. In other words, there is no (*a priori* or *a posteriori*) answer, in neural or functional or other terms, to the question: *what is it* to have a given cognitive phenomenal property? In this sense, properties are utterly simple, just as G. E. Moore said goodness is utterly simple. Yet they somehow “ground” determinate belief and desire properties, in accordance with

CP-determination. For instance, it is just a brute fact that, if you have a certain utterly simple, unstructured cognitive phenomenal property, you stand in the belief relation to a plus-content, as opposed to a bent, quus-like content. On another version of primitivist CP-determination, there *is* a partial answer to the question: what is it to have a particular cognitive phenomenal property? The idea is that having a particular cognitive phenomenal property simply *consists in* (and so trivially determines) occurrently standing in a particular intentional relation to a particular narrow content. So cognitive phenomenal properties have a relational structure. Of course this *identity view* would explain the tight relationship between cognitive phenomenology and intentionality. It would explain why it is necessary that if you “change either the attitude-type (believing, desiring, wondering, hoping, etc.) or the particular [narrow] intentional content” then “the phenomenal character thereby changes too” (Horgan and Tienson 2002: 522).<sup>15</sup> What makes the view I have in mind *primitivist* is that it adds that there is in turn no reductive account of the relevant phenomenally-constituted intentional relations towards narrow contents. For instance, there is no answer to the question: what *is it* for you to occurrently *believe* that two plus two equals four, as opposed to two quus two equals four? This kind of “constitutive” question cannot be answered by appealing to functional role or anything else. (As I am about to explain, this is compatible with our beliefs having physical *supervenience-bases*.) This view of cognitive phenomenal properties is analogous to *non-reductive, primitivist intentionalism* about sensory phenomenal properties (Chalmers forthcoming, Pautz 2010a).

Now I can finally state my “danglers” argument. It targets primitivist CP-determination. It is nothing more than a generalization of Smart’s (1959) famous “danglers” argument against dualism based on its apparent commitment to complex physical-to-mental “bridge principles” (which he called “danglers”). Let *P* be the phenomenal property which allegedly determines (perhaps just is) your occurrently believing that two plus two equals four, as opposed to two quus two equals four. On primitivist CP-determination, *P* is not a physical property, at least in a narrow sense where physical properties just include mainstream physical properties like neural properties, functional properties, tracking properties, and so on. Yet even on this view your having this primitive cognitive phenomenal property is not a fluke; it is in *some* sense determined by one of your mainstream physical properties. Thus Strawson (2010: note 54) speculates that cognitive phenomenal properties are dependent on neural properties instantiated somewhere near the sensory regions of the brain. Alternatively, perhaps your having *P* is connected to (but distinct from) your having a more complex functional property involving your inferential dispositions. In any case, there is a true conditional of the form: necessarily, if any individual has physical property *F*, then that individual has *distinct* primitive cognitive phenomenal property *P*, which in turn determines (and perhaps just is) his occurrently believing the content *two plus two equals four* as opposed to the content *two quus two equals four*. These conditionals

---

<sup>15</sup> For a discussion of the identity view, see Pautz (2008: 269). Pitt (2009) also develops a version of the identity view. As we saw in the introduction, many fans of the phenomenal intentionality program endorse “prioritism”, the thesis that phenomenal properties are in general metaphysically prior to intentional properties. Since the priority relation is presumably irreflexive, the view that certain cognitive intentional properties *just are* cognitive phenomenal properties would entail that prioritism is mistaken (Pautz 2008: 269; Kriegel 2011: 63). Another consequence of the identity view would be that slight differences in cognitive phenomenology between individuals (as it might be, between two people who say “Aristotle was smart”) would entail that they entertain slightly different “narrow” contents, even if they might entertain the same “wide content”. It entails a difference in content at least at the level of modes of presentation, even if does not entail a difference in truth-conditions.

link the physical properties of individuals with their determinately believing (or desiring or wondering) some narrow contents rather than others.

Now there are two possible views available to defenders of primitivist CP-determination concerning the modal strength of these conditionals. Both are open to the danglers arguments.

On one view, the conditionals are merely *nomically* necessary. They are fundamental laws of nature distinct from the laws of physics that Smart (1959) called *nomological danglers*. For instance, maybe there is a basic contingent law of dualistic psychophysics according to which, if someone has specific neuro-computational property *M* with the right cognitive background, then they have cognitive phenomenal property *P* and hence believe the content *two plus two equals four* rather than the content *two quus two equals four*. In general, for each narrow content, we apparently need a separate nomological dangler connecting having an occurrent belief with that content rather than some other content with some underlying physical condition. In this version, primitivist CP-determination amounts dualism. There are physical duplicates of us (zombies) who have no cognitive phenomenology and hence no determinate beliefs or desires. John Searle (1992: 158-159), an anti-functionalism about intentionality who has also stressed the bearing of the first person on the issue of content determinacy, speaks of such neural-intentional laws, although he disavows the dualist label. The “danglers” argument against the present dualist view is simply that it solves the determinacy problems due to Quine and Kripkenstein at the cost of providing an implausibly complicated account of the place of intentionality in the natural world. We have general Occamist reasons to keep our basic laws of nature to a minimum, but the present dualist view multiples them. This is not a knock-down argument; but it gives us a reason to disbelieve the view. Smart (1959) objected to dualism about sensations on the grounds that it requires fundamental laws connecting physical states with sensory experiences. The present dualist view requires *in addition* an indefinite number of fundamental laws connecting physical states with the enormous variety of basic cognitive states.

One another possible version of primitivist CP-determination, the conditionals linking physical conditions with distinct intentional conditions are *brute metaphysical necessities*. They are not just epistemically brute because of the “explanatory gap”. They are metaphysically brute in the sense of Dorr (2007): roughly, they cannot be derived real definitions (e. g. *a posteriori* identities) together with logic. Just as nomological danglers would dangle from the rest of the body of nomological truths, so these *modal danglers* would dangle from the rest of the body of modal truths. In this version, primitivist CP-determination amounts to what I have called *primitivist physicalism* (Pautz 2010a). Horgan (2010) has called this kind of view *Moorean emergentism* and says he takes it seriously. (Indeed, I have argued that his remarks on CP-determination suggest *some* kind of emergentism.) Of course, the view that cognitive phenomenology is primitive but supervenient with metaphysical necessity on the physical is analogous to Moore’s view that goodness is primitive but supervenient on the natural. Ontologically, it is just like the dualist view described above; it differs from dualism only modally, claiming zombie worlds are metaphysically impossible. My objection to this view is analogous to my objection to dualism about cognitive intentionality. Even though we must admit some brute necessities (e. g. certain basic metaphysical principles), we have general Occamist reasons to keep them to a minimum. But

the present account of the intentionality of occurrent beliefs and desires multiplies them.<sup>16</sup>

The argument I have been developing is best taken as a dilemma for proponents of CP-determination. Either they say that cognitive phenomenal properties are identical with mainstream neural or (wide or narrow) functional properties (Kriegel 2011), or else they accept one of the versions of primitivism I described above. If they accept the first horn, then they are in effect accepting one or another of the familiar attempts to “naturalize” intentionality (Fodor, Dretske, Lewis, etc.), without doing anything to answer the familiar problems due to Quine and Kripkenstein and others. They are only adding that the physical-functional properties that determine content have a special non-sensory “feel”. In that case, their view is just “old wine in a new bottle”. Given the things they have said, most proponents of CP-determination probably will not opt for this horn. But if they instead accept the second horn, then again their view is old wine in a new bottle. They are simply accepting Brentano’s thesis of the irreducibility of intentionality. This has been on the table from the start (Quine 1960: 221). This is what is doing the work in their answer to the determinacy problems due to Quine and Kripkenstein. They are just adding that the relevant irreducible properties have a special non-sensory “feel”. Moreover, in that case their view is open to a generalization of the danglers argument (or “simplicity argument”) due to Smart (1959).

### 3.2 *The separation argument*

My next few arguments apply against all versions of CP-determination, not just “primitivist” versions. They have a common form. All versions of CP-determination accept a certain irreducibility thesis. They hold that cognitive phenomenal properties, and hence belief and desire properties, are distinct from all sensory properties. Indeed, Strawson (2010: note 54) says that they constitute an experiential modality over and above the usual sensory modalities. In addition, Pitt as well as Graham, Horgan and Tienson also hold that cognitive phenomenal properties, and hence some belief and desire properties, are “intrinsic” and “non-relational” properties. This presumably entails that they are distinct from all functional properties. In short, CP-determination entails that cognitive phenomenal properties, and hence some belief and desire properties, form an autonomous level of mental reality, distinct from the sensory and functional levels.

Now irreducibility entails *modal independence*. Therefore the CP-determination thesis entails that cognitive phenomenal properties, and hence some belief and desire properties, are radically *modally independent* of sensory and functional properties

---

<sup>16</sup> On the kind of view I will sketch in §4, all facts about beliefs and desires reduce to clusters of functional facts and sensory facts. However, sensory facts are not in turn reducible to physical-functional facts. So my own view requires *physical-sensory danglers*: brute bridge principles connecting the physical with the sensory. However, since it holds that beliefs and desires reduce to sensory-cum-functional facts, it does not require *in addition* the kind of *physical-cognitive danglers* that primitivist CP-determination is committed to: additional brute bridge principles connecting physical conditions with beliefs and desires. However, one might wonder, if my own view requires physical-sensory danglers, how I can consistently object to physical-cognitive danglers? I have several responses. (i) Physical-cognitive danglers are more objectionable than physical-sensory danglers. We can hope that danglers connecting the physical and the sensory are systematizable (for instance, average firing rates might determine perceived intensities). In that case, my view might only need a handful of brute physical-sensory danglers. By contrast, according to primitivist CP-determination, for each basic cognitive phenomenal property, we apparently need a separate brute physical-cognitive “dangler”. (ii) In any case, all else being equal, we have Occamist reasons to keep danglers to a minimum. My view only requires physical-sensory danglers; so it is superior to primitivist CP-determination, which requires both physical-sensory danglers and physical-cognitive danglers. (iii) As we have seen, there is no reason to accept CP-determination, and so no reason to accept this added complexity. (iv) As we shall see, I have additional objections to CP-determination, which are independent of my present appeal to simplicity.

(Chalmers 2008). *They can entirely “float free” from sensory and functional properties.* (Consider an analogy: a strong thesis of the irreducibility of sensory qualia to physical properties entails the possibility of qualia inversion in physical duplicates, zombies in which qualia are entirely absent, disembodiment, and so on.) But this strong modal independence thesis is simply not true. My next three arguments are meant to establish that it is not true, by considering whether we can make sense of three type of scenarios in which cognitive phenomenal properties, and hence some belief and desire properties, float free from sensory and functional properties. I will argue that we cannot make sense of these scenarios.

My first argument, the *separation argument*, generalizes my “separation argument” against biological CP-determination (§3.1) to any version of CP-determination. Consider an example. According to the CP-determination thesis defended by Horgan and Tienson, there is a cognitive phenomenal property *P* that metaphysically necessitates occurrently believing the narrow, *de se* content *there is a picture on the wall behind oneself*.

The first step of my separation argument asserts that, *if* there is such a property as *P* (I myself am doubtful), then there should be possible “separation cases” in which it is separated from its actual associations with sensory and functional conditions and is associated with quite different such conditions. In one such case, while Charlie has *P*, the sentence ‘there is a *clock* on the wall’ (with its actual meaning) runs through his interior monologue, he forms a sensory image of a *clock* on a wall, and behaves (and has the first-person experience of behaving) exactly as if he believes a *clock* is on the wall. My case for asserting this conditional possibility claim is based on *what proponents of CP-determination themselves say about cognitive phenomenal properties*. They say that they are *distinct from* all sensory properties and indeed constitute a separate experiential modality. Given this, if there is such a cognitive phenomenal property as *P*, it might be typically associated with different sensory-functional properties than those which it is in fact associated with. Consider an analogy: since auditory qualia are distinct from visual qualia, they can be combined with any series of visual qualia. Here I am provisionally assuming a kind of “principle of recombination” as regards distinct and mutually irreducible ranges of phenomenal properties. To repeat, irreducibility naturally goes with modal independence.

The second step of my separation argument asserts that, in this case, even though Charlie has *P*, he does *not* occurrently believe that there is a *picture* on the wall behind him, contrary to Horgan and Tienson’s CP-determination thesis. The *clock*-appropriate sensory and functional facts present in the case are incompatible with his having this occurrent belief.<sup>17</sup> Even if we reject crude behaviorism or a functionalist theory according to which beliefs *reduce to* functional states entirely characterizable in non-mental terms, we should recognize that, for any belief and desire, some combinations of sensory and functional facts are *a priori* incompatible with having that belief or desire.<sup>18</sup> In this sense, there are *functional constraints on belief and desire*.<sup>19</sup>

---

<sup>17</sup> The proponent of CP-determination might grant that a wide range of bizarre separation cases involving a given cognitive phenomenal property are possible (not just the Charlie case, but the baby case discussed below), but insist that across these cases it determines the same narrow content. This what Siewert (2011) calls “the part of what is thought that is fully determined by [cognitive] phenomenal character”. The proponent of CP-determination might suggest that my separation argument only shows that this is not a content of the sort proposed by Horgan and Tienson and Strawson involving ordinary concepts like *picture* or *two* or *plus*. (Thanks to David Chalmers here.) But if he cannot even *gesture at* some specification of them, we have no reason to believe there are such contents.

<sup>18</sup> Here I use “functional facts” broadly. Thus, suppose you might have life-long brain-in-a-vat phenomenal duplicate. Even though it does not have states caused by ordinary worldly conditions or that cause actual bodily movements, it

Here is another illustration. Let  $P$  be the cognitive phenomenal property which, according to proponents of CP-determination like Strawson, metaphysically necessitates occurrently believing the narrow content *two plus two equals four*, thus solving the Kripkenstein problem. Again, my separation argument against the CP-determination thesis has two steps. First, if there is such a non-sensory, intrinsic property as  $P$ , then a baby (say) presumably might have  $P$  for a few seconds while otherwise remaining the same. This might involve temporarily modifying its neural state in the “cognitive phenomenology center” of its brain (Strawson 2010: note 54); but I see no reason why it should not be metaphysically possible. Second, even though the baby has cognitive phenomenal property  $P$ , it does not occurrently think that two plus two equals four, contrary to CP-determination. That would require that it have certain arithmetical concepts (e. g. *plus*), which in turn would require that it have certain arithmetical *abilities*. But by stipulation it lacks such abilities.<sup>20</sup>

Such cases could be multiplied indefinitely. In fact, there are presumably possible cases in which the alleged cognitive *quale* of an occurrent *belief* plays the functional role of a *desire* (for instance, the desire for food) and occurs with certain types of *affective phenomenology* (for instance, hunger pangs). Surely in such cases the relevant quale does not realize a *belief*: the sensory-functional facts trump “cognitive qualia” in determining what kind of intentional states we have, contrary to CP-determination.

Since cognitive phenomenal properties are modally independent of sensory properties, CP-determination thesis entails quite radical violations of the intuitively plausible claim that the intentional facts about a population supervene merely on the sum total of sensory and functional facts about that population. For instance, suppose that in the actual world Elmer attends to a whole rabbit and says ‘there’s a rabbit’. According to Graham, Horgan and Tienson, he has the cognitive phenomenal property  $R$ , which somehow determines his occurrently believing that there is a *rabbit*, thus answering Quine’s challenge.<sup>21</sup> Now consider Twin Elmer who belongs to a community of sensory and functional duplicates of us. Since he is a sensory-functional duplicate of Elmer, he also attends to a whole rabbit and says ‘there’s a rabbit’, he is disposed to answer questions about rabbits and his rabbit-beliefs in the same way as Elmer, and so on. But suppose that, when says ‘there’s a rabbit’, instead of the cognitive phenomenal property  $R$ , he has the cognitive phenomenal property  $U$ , which according to Graham,

---

has *experiences* of worldly conditions, which in turn lead to *conscious attempts* to engage in bodily movements, and *experiences* as of successfully engaging in such movements. These are functional facts about its patterns of sensory experiences. Functional facts also include facts about conceptual or inferential role. My claim that there are functional constraints on belief and desire is quite compatible with the claim that Strawson’s (2010) weather watchers could have beliefs and desires.

<sup>19</sup> While I think that there are *a priori* functional constraints on cognitive intentionality, I do not think that there are *a priori* functional constraints on *sensory* intentionality (or at least those aspects of sensory intentionality fixed by sensory phenomenology). See Pautz (2010c: 271). In my view, this is an important difference between cognitive intentionality and sensory intentionality.

<sup>20</sup> Horgan and Tienson in one place (2002: 526) appear to accept *qualified CP-determination*: necessarily, if an individual has a certain cognitive phenomenal property and “has a sufficiently rich network of actual and possible phenomenal/intentional states”, then the individual has a thought with a certain narrow content (e. g. *there is a picture on the wall*, or *two plus two equals four*). Unfortunately, they do not explain what this means or why they say it. In any case, my Charlie case and my baby case are counterexamples even to qualified CP-determination. For Charlie and the baby *do* have a rich network of phenomenal/intentional states in addition to the relevant cognitive phenomenal properties (they both have many experiences and thoughts). Yet they do not have the occurrent thoughts which are allegedly necessitated by those cognitive phenomenal properties, provided that they occur in the context of a rich network of phenomenal/intentional states.

<sup>21</sup> Actually, since contents involving the natural kind *rabbits* are wide, their view must be that intrinsic cognitive phenomenology somehow determines that one believes some underlying narrow content, and that believing this narrow content, together with standing in certain causal relations to the world, somehow determines believing a content involving rabbits. But they do not fill in the details and in any case this point does not matter to my example.

Horgan and Tienson determines believing that *there is an undetached rabbit-part* (see the quote in §1). Maybe this is due to a neural difference in the cognitive phenomenology center of his brain (Strawson 2010: note 54). In short, Elmer and Twin Elmer exhibit a case of *altered cognitive qualia* among sensory-functional duplicates. Auditory qualia can differ in just about any possible way between total *visual* duplicates. So if there cognitive phenomenal properties that are *totally distinct from* all sensory properties, forming a distinct experiential modality, why couldn't they likewise vary across total sensory duplicates? This case should be possible.<sup>22</sup> Of course, by CP-determination, whereas Elmer believes that there is a *rabbit* present, Twin Elmer believes that there is an *undetached rabbit-part* present. This is so despite the fact that they are totally sensory-functional duplicates, so that Twin Elmer, like Elmer, will insist 'my belief is about a *rabbit*, not an *undetached rabbit part*!' Intuitively, this consequence of CP-determination is quite implausible.

Proponents of CP-determination need to address the separation argument, because it only relies on the widely-accepted intuition that there are functional constraints on belief and desire. Siewert (1998: 285) at least addresses a case similar to one of my cases, the baby case. He simply asserts that such cases are impossible, saying that he accepts "holism" about cognitive phenomenal properties and rejects "atomism". He does not address cases like my Charlie case or my Elmer case in which cognitive phenomenal properties are associated with different sensory and functional conditions than those with which they are in fact associated. But he might just declare that they too are impossible. Other proponents of CP-determination (Graham, Horgan and Tienson, Pitt, Strawson) might likewise assert that "separation cases" are impossible. They might say that they cannot imagine what they would be like from the inside. In short, they might say that, since there are necessary functional constraints on belief and desire, and since cognitive phenomenal properties determine belief and desire, there are necessary functional constraints on cognitive phenomenal properties.

But this is not enough. Remember: I gave an *argument* for claiming that, if there are cognitive phenomenal properties, then separation cases are possible. Graham, Horgan and Tienson as well as Pitt themselves say that cognitive phenomenal properties are *distinct from* all sensory properties, and that they are *intrinsic, non-relational, non-functional* properties. Given this, they should be *modally independent* of sensory-functional properties: there is no apparent reason why they might not be associated with different combinations of sensory and functional facts than they are in fact associated with. If proponents of CP-determination say that my separation cases are not possible, they need to answer this argument. When something is impossible, we generally expect an explanation *why* it is impossible. If there are necessary functional constraints on cognitive phenomenal properties, we need an explanation. We should keep "brute necessities" to a minimum (Dorr 2007).<sup>23</sup>

<sup>22</sup> Note that my claim here is conditional: *if* we enjoy cognitive phenomenal properties over and above sensory properties, then such radical phenomenal variation among sensory-functional duplicates should be possible, and indeed imaginable. I myself think that we have no grip on these cases, and that this casts doubt on the claim that we do enjoy cognitive phenomenal properties over and above sensory properties. I will discuss this point in §3.3 in connection with "absent cognitive qualia cases".

<sup>23</sup> In discussing what determines narrow content, at one point Horgan and Tienson (2002: 525) say that some *perceptual beliefs* are "experienced as having [a] belief-specific [functional] role" involving perceptual inputs and bodily outputs. Could this help them answer my separation argument, at least as applied to immediate perceptual beliefs? The trouble is that Horgan and Tienson's remarks on this score are not elaborated and admit of multiple interpretations. Suppose I just experienced a ball go behind a tree and now believe it is behind the tree. Do they think that, *at the present time*, I have some kind sophisticated non-sensory cognitive experience (not just a potential sensory memory or anticipation) of *my belief* as having been caused a second ago by a certain visual experience and about to cause me to go

In response, the proponent of CP-determination might defend a functionalist theory of such properties on the model of (say) Lewis's functionalism about belief properties. According to Lewis's theory of belief, occurrently believing that two plus two equals four is a complex functional property  $F$  involving relations to experiences, other beliefs and behavior. This view explains why belief properties are holistic and not atomistic, so that a baby (for instance) cannot have arithmetical beliefs. It explains why there are necessary functional constraints on belief and desire. Now let  $P$  be the phenomenal property which, according to CP-determination, determines occurrently believing that two plus two equals four. The proponent of CP-determination might answer my separation argument by identifying the cognitive phenomenal property with the Lewisian functional property  $F$ . In general, cognitive phenomenal properties reduce to functional properties, he might say. This would *explain* why cognitive phenomenal properties are holistic and not atomistic (why there are necessary functional constraints on having such properties), something Siewert asserts without explaining. It would explain why my separation cases are not possible for cognitive phenomenal properties, thus answering my separation argument against CP-determination. The trouble is that CP-determination would now just amount to what I previously (§3.1) called functional CP-determination. It would just be a kind of functionalism. As we saw, this is not compatible with what proponents of CP-determination have said.

Let me suggest a different reply to my separation argument on behalf of proponents of CP-determination. I will introduce the reply by analogy. Swyer (1982) accepts an interesting view about fundamental physical properties, like masses and charges. He does not say that they reduce to clusters of dispositions or functional properties. But he says that such properties essentially possess certain (second-order) nomic properties and essentially stand in certain nomic relations; these are given by the fundamental physical laws. This is just a basic modal fact. So they are necessarily connected with certain dispositional and functional properties. And, even though they are intrinsic, "separation cases" in which they are recombined in ways that violate their nomic profiles are metaphysically impossible.<sup>24</sup> There are unHumean metaphysically necessary connections between distinct existences. Likewise, proponents of CP-determination might say that cognitive phenomenal properties (which on his view might just themselves be belief and desire properties) are fundamental, intrinsic, non-functional properties that are distinct from sensory phenomenal properties. But they have certain functional profiles essentially. They essentially bear certain systematic

---

behind the tree in second later, so that my current cognitive phenomenology encompasses a quite large temporal frame (larger than the usual estimates for the specious present)? This is just not plausible. On another interpretation, Horgan and Tienson are saying that the narrow contents of perceptual beliefs are *not* fully determined by merely their "cognitive phenomenology" *at the time those beliefs are entertained*: they are somehow determined by a complex cluster of *sensory* experiences, including past sensory experiences of the world and actual and potential future sensory experiences of acting on the world (e. g. looking behind the tree and feeling surprised when one doesn't see it there). In that case, their account of the ground of perceptual beliefs is identical to the account given by the "phenomenal functionalism" I will sketch in §3.4, which resembles orthodox functionalism. But it apparently goes against Horgan and Tienson's unorthodox claim that each occurrent belief has its *own non-sensory* phenomenal character which fully determines its narrow content. In any case, as Horgan and Tienson point out, in some cases your beliefs (e. g. beliefs about pictures behind you, or belief that two plus two equals four) are not simple perceptual beliefs that are produced by perceptual inputs and go on to produce bodily outputs. Indeed, sometimes we have occurrent beliefs out of the blue, which do not actually affect behavior at all. In these cases, the alleged intrinsic cognitive phenomenal properties of our beliefs (if such there be) certainly do not involve "experiencing their inputs and outputs" (even if we could make sense of that notion). By my argument, "separation cases" should be possible for such intrinsic cognitive phenomenal properties, showing that CP-determination is false.

<sup>24</sup> For how fundamental physical properties might count as "intrinsic" even under Swyer's view, see Langton and Lewis (1998: 337-338).

relations to various other states and conditions: other cognitive phenomenal properties, behavior or attempts at behavior, mental dispositions and abilities (e. g. inferential dispositions and dispositions to form certain images), sensory experiences, and so on. Hence, even though it is intrinsic, any given cognitive phenomenal property (realizing a particular belief or desire property) necessarily brings with it a *system* of states and abilities. This is not because cognitive phenomenal properties *reduce to* abstract functional properties or clusters of dispositions, as on functionalism or behaviorism. Rather, it is just a brute modal fact, brute in the sense that it cannot be derived from real definitions and logic (Dorr 2007). Hence, under this view, the functional constraints on belief and desire are just brute modal facts. Call this version of CP-determination *systemic CP-determination*. Under *systemic CP-determination*, my separation cases involving cognitive phenomenal properties are impossible, for instance my baby case and Charlie case. So my separation argument against CP-determination fails at the first step.

But, since we have *no reason* to accept CP-determination in general (§2), we have no reason to accept systemic CP-determination. On the other hand, we have strong reasons to reject it. (i) Previously, we saw that proponents of CP-determination apparently must accept primitivist CP-determination, which requires brute modal principles connecting cognitive phenomenal properties with underlying physical states. Now we have seen that, to accommodate the functional constraints on belief and desire and avoid my separation argument, they must accept additional brute “systemic” principles connecting cognitive phenomenal properties with other mental conditions and background abilities. But we ought to keep brute modal principles to a minimum.<sup>25</sup> (ii) What are the essential functional profiles of cognitive phenomenal properties? Presumably, there will be a lot of vagueness here. So systemic CP-determination requires objectionable vagueness-in-the-world concerning the fundamental modal facts. (iii) How exactly do the essential functional profiles of cognitive phenomenal properties rule out the separation cases I described above involving Charlie and the baby and Elmer?<sup>26</sup>

### 3.3 *Absent cognitive qualia cases*

Before (§2) I used absent qualia cases to undermine the determinacy and access arguments for CP-determination. Now I will use such cases in arguments against cognitive phenomenology.

---

<sup>25</sup> In reply, the proponent of CP-determination might say that the functional constraints on cognitive phenomenal properties are not brute. He might say that they follow from the simple fact that having cognitive phenomenal properties amounts to standing in certain intentional relations (occurrently believing or occurrently desiring) to certain contents, together with the fact that there are functional constraints on standing in such relations to contents. But this just passes the buck. The question then becomes ‘Why are there functional constraints on standing in such relations to contents – in other words, why are such intentional relations necessarily constrained by the sensory-functional facts?’ We can avoid treating these constraints as brute only by providing a functionalist account of *what it is to stand* in such relations to contents (for instance, an orthodox functionalist theory or else the kind of phenomenal functionalism I will sketch in §4). But, as we have seen (§3.1), proponents of CP-determination typically reject functionalist theories of intentionality.

<sup>26</sup> Siewert might offer a different reply to my separation argument. (1998: 286) On one formulation of his claim about cognitive phenomenology determining intentionality, it just amounts to the claim that someone who “has, has had, and is disposed to have, experience with phenomenal character indistinguishable from my own” would have many thoughts. But this is compatible with the falsity of CP-determination, and indeed with the non-existence of cognitive phenomenology. For Siewert’s total phenomenal twin has the same sensory experiences of conducting inferences in inner speech, the same sensory experiences as of his environment, and the same sensory experiences of acting on the world. Maybe it is such clusters of past and future *sensory* facts, and not his alleged “cognitive phenomenology”, which determine the contents of his beliefs and desires, in accordance with the kind of phenomenal functionalism I will sketch in §3.4.

My first argument casts doubt on the basic *CP-existence* thesis and has two steps. The first step asserts that, if there are cognitive phenomenal properties, then cases of the following kind are possible and indeed we should be able to *positively imagine* them: (i) you have *exactly* the same sensory properties and functional properties, (ii) your phenomenal life is profoundly different from your actual phenomenal life in that you lack phenomenal properties (namely, cognitive phenomenal properties) you actually enjoy. Indeed, presumably, if the *CP-existence* thesis is true, and cognitive phenomenal properties form a distinctive experiential modality (Strawson 2010: note 54), then there are among *actual humans* individuals who lack cognitive qualia, but otherwise have phenomenal lives very similar to normal individuals. Analogy: our auditory phenomenal properties are distinct from our visual phenomenal properties. So we can easily imagine a condition (namely, deafness) in which we have the same visual phenomenal properties, but our phenomenal lives profoundly differ because we lack auditory phenomenal properties.

The second step of the argument is that we cannot positively imagine such cases. Just try. Suppose you hear a friend say ‘Let’s go to the bar later’ and you quickly form an image of the local bar and follow up with a question as to time. Now try imagine a case that is completely identical to the actual case in all sensory and functional respects, and yet profoundly phenomenally different in that you lack the cognitive phenomenal properties that you allegedly actually possess. I honestly cannot do it. (Similarly, I cannot imagine holding fixed all aspects of sensory, affective phenomenology while removing the additional non-sensory, conative qualia that allegedly characterize our occurrent desires.) Perhaps the proponent of the *CP-existence* thesis will reply: just imagine hearing ‘Let’s go to the bar’ while not being a speaker of English. But that case would differ in *sensory* and *functional* respects from the actual case. If I did not understand English my auditory system would not parse the auditory stream ‘Let’s go to the bar’ in the same way it actually does. So I would have a different auditory experience. I would also not form a quick image of the local bar. Not understanding what was said, I might also feel stressed. And I would not be disposed to follow up with certain questions. What the proponent of the *CP-existence* thesis is committed to is the possibility of a case in which *all* the past and present sensory and functional facts are held *constant* (the very *same* auditory experience of ‘Let’s go to the bar’, the same imagery, the very same feelings, and so on), and yet my phenomenal life is *profoundly different* from my actual phenomenal life in myriad ways. It is this I cannot imagine. Maybe he could insist that it is possible, even if I cannot imagine it. But why should this be? It is not only possible that my experiential life should differ because I lack auditory phenomenal properties while the rest of my phenomenal life is held constant; I can easily imagine this case. So why cannot I easily imagine occupying a profoundly different phenomenal world because all of my actual cognitive phenomenal properties are absent, while *all* of the sensory and functional aspects of my life stay exactly the same?<sup>27</sup>

---

<sup>27</sup> Terry Horgan tells me that in forthcoming work he (independently) uses absent cognitive qualia cases to argue for the *CP-existence* thesis. It is worth mentioning that the allegedly “diaphanous” character of experience cannot explain why we cannot imagine having profoundly different phenomenal lives due merely to the absence of “cognitive experience”. After all, our *visual* experiences are arguably diaphanous in the sense that we cannot have *object-awareness* of them (Tye 2000: 51-52), but this does not prevent us from being able to imagine cases in which they are absent but our other types of experiences are the same. For, as Tye points out, even if we are not aware of experiences in the sense that we do not have object-awareness of them (we cannot attend to our experiences in addition to attending to the objects and properties represented by experiences), we still are certainly quite aware of them in the sense that we effortlessly have *fact-awareness* to the effect that they are present.

Now I turn to my second argument involving absent cognitive qualia cases. While my first argument was directed against proponents of the CP-existence thesis, my second argument is directed against proponents of the CP-determination thesis like Strawson and Graham, Horgan and Tienson. Typically, they accept something stronger than CP-determination. They hold that non-sensory cognitive and conative qualia are *necessary* in order to have the level of content determinacy we actually enjoy. Call this the *necessity thesis*. We have seen that Horgan and Tienson implicitly make this assumption in their determinacy argument for CP-determination (it corresponds roughly to premise 3 of that argument), and it is something they explicitly endorse. Strawson (in the quote in §3.1) likewise says that cognitive qualia are the key to solving the Kripkenstein problem. This is the target of my second argument.

Let an *absent cognitive qualia case* be one in which the entire human population is identical to the actual population in all *sensory* and *functional* respects, but no one possesses any non-sensory cognitive and conative phenomenal properties.<sup>28</sup> For example, suppose that in the *actual* case while hunting Elmer has a visual experience of a rabbit and has the sensory phenomenology of attending to the *whole* rabbit, as opposed to a rabbit part. He says, ‘Lo, a rabbit’. He also has hunger pangs, and imagines having rabbit for dinner. He is disposed to behave exactly as if he believes that a rabbit is present and desires to have it for dinner. Now consider an absent cognitive qualia case in which *all these same sensory and functional facts are present*, but Elmer’s alleged cognitive and conative qualia are absent. Surely in this absent cognitive qualia case as in the actual case Elmer has *some* kind of rabbit beliefs and desires. However, if the necessity thesis is true, then in this case Elmer’s rabbit beliefs and desires are much more *indeterminate* in content than they are in the actual case. (Indeed, a radical version of the necessity thesis according to which cognitive qualia and conative qualia are *essential* to occurrent belief and desire has the even more absurd consequence that he has no genuine occurrent beliefs and desires!) For instance, given what they say in the passage quoted in §3.1, Graham, Horgan and Tienson appear committed to saying that in the actual case Elmer determinately believes that a *rabbit* is present and desires to have a *rabbit* for dinner, whereas in the absent cognitive qualia case it is indeterminate whether Elmer’s belief and desire are about rabbits as opposed to undetached rabbit parts. Therefore there is a difference in intentional facts across the cases. Hence a consequence of the necessity thesis is the failure of the thesis that intentional facts about a population supervene on the totality of sensory and functional facts about that population. My argument against the necessity thesis is simply that this consequence is implausible and proponents of the necessity thesis have said nothing to make it plausible. After all, in the absent cognitive qualia case as in the actual case, Elmer enjoys the sensory phenomenology of focusing on the whole rabbit, he is disposed to respond to Quinean indeterminacy arguments by insisting that it is introspectively obvious to him that he has beliefs and desires about “rabbits” rather than “undetached rabbit parts”, and so on. All the first-person and third-person evidence suggests that across these cases his belief and desire enjoy the same level of determinacy. This supports the view I will sketch in §4 on which intentional facts about a population entirely

---

<sup>28</sup> Everyone should believe in the possibility of absent cognitive qualia cases so understood. If you think we *actually* fail to possess non-sensory cognitive and conative phenomenal “qualia” over and above our sensory phenomenal properties, as I am inclined to do, then you will think that the absent cognitive qualia case is phenomenally identical with the actual case. On the other hand, if you believe we actually possess such “qualia”, then you will also say that the absent cognitive qualia case is possible, but that it would differ phenomenally from the actual case in that we would lack qualia we actually enjoy. My argument will be neutral on this issue.

supervene on (and indeed reduce to) nothing but the sensory and (wide and narrow) functional facts about that population. Or consider arithmetical beliefs. In an absent cognitive qualia case identical to the actual case in all functional and sensory respects (including experiences of counting objects and so on), all the first-person and third-person evidence suggest that our arithmetical beliefs have the same level of content determinacy that they have in the actual case. Then, contrary to Strawson, cognitive qualia are not the essential key to answering the Kripkenstein problem.

### 3.4 Disembodied cognitive qualia cases

I have two related arguments against cognitive phenomenology involving “disembodied cognitive qualia cases”.

My first argument casts doubt on the CP-existence thesis and has two steps. If there are cognitive qualia, then in the actual world they are *embodied* in the sense that they are accompanied by sensory properties, including experiences of having a body and acting on the world. The first step of my argument asserts that, if there are cognitive phenomenal properties distinct from all sensory properties, then *disembodied cognitive qualia cases* ought to be possible and indeed we should be able to positively *imagine* them. In such a case, we allegedly have a rich phenomenal life that *overlaps* with our actual phenomenal life because we have all the same cognitive phenomenal properties we have in the actual world; but our cognitive phenomenal properties are “disembodied” in the sense that they are not accompanied by any *sensory* properties. In other words, we have no visual experiences (including no experiences of having a body or engaging in apparent behaviors), auditory experiences, no mental imagery, no “inner speech”, and so on. Consider an analogy: auditory qualia are distinct from other kinds of qualia. So we can imagine cases in which such qualia occur in the absence of all other sorts of qualia.<sup>29</sup>

The second step of my disembodied cognitive qualia argument against the CP-existence thesis is that we *cannot* positively imagine disembodied cognitive qualia cases. Just try. If the CP-existence thesis is true, then in such cases we have a rich phenomenal life which *overlaps* with our actual phenomenal life, only it is totally non-sensory. But what would it be like? If you try to imagine what it would be like, you might imagine seeing all black, having an experience of inner speech (“nothing much is happening”), and so on. But then you will not be imagining a case in which you have cognitive phenomenal properties but *no* sensory properties. In reply, the proponent of the CP-existence thesis might say that such cases are possible, but that for some reason we cannot positively imagine them. But why not? When it comes to other sorts of qualia, we can easily imagine cases in which they occur in the absence of distinct kinds of qualia: for instance, we can imagine only having auditory qualia.<sup>30</sup>

<sup>29</sup> By a ‘disembodied cognitive qualia case’, I do not mean a case in which an individual actually has no body or brain. Maybe an individual can have cognitive qualia only if he undergoes the right physical processing in the alleged cognitive qualia center of his brain. Even on this view, if there are cognitive qualia distinct from all sensory qualia, disembodied cognitive qualia cases should be possible in my sense. For there might be a very inactive brain which only ever undergoes activity in the cognitive qualia center but no other neural area. According to proponents of CP-determination, such an individual would have cognitive experiences but no other experiences, including visual experiences of the world or his body. In the same way, a subject might only have auditory qualia throughout his lifetime.

<sup>30</sup> Let me guard against a misunderstanding of my argument against the CP-existence thesis involving disembodied qualia cases. What I claim to be impossible is not simply a case in which (i) an individual has phenomenal properties but (ii) that individual has no sensory properties (understood broadly to include perceptual, somatic, and emotional phenomenology). Certainly, *that* case is possible. For, on anyone’s view, there are worlds in which there are phenomenal properties distinct from the sensory properties instantiated in our world and an individual enjoys those phenomenal properties and no other phenomenal properties. (This does *not* entail that we ourselves actually enjoy phenomenal

Now I turn to my second argument involving disembodied cognitive qualia cases. While my first argument was directed against the CP-existence thesis, my second argument is directed against the CP-determination thesis. As we have just seen, the CP-existence thesis arguably entails the possibility of disembodied cognitive qualia cases. For the sake of argument, let us just grant that the truth of the CP-existence thesis and the possibility of such cases, against the skepticism I expressed above. Let Nemo be someone who enjoys a rich series of cognitive phenomenal properties of the kind we actually undergo but who has no sensory phenomenal properties at all. Now, according to CP-determination, Nemo's having such a rich series of cognitive phenomenal properties alone determines (perhaps just consists in) his having a series of our actual sophisticated narrow beliefs and desires: for instance, the belief that two plus two equals four (Strawson), the belief that a picture is hanging on the wall behind himself (Horgan and Tienson), and so on. Hence, according to CP-determination, Nemo has a rich series of occurrent narrow beliefs and desires, even though he has never had any sensory experiences whatever.

But this can be ruled out *a priori*. Even if Siewert (1998: 277) is right that there are occasional cases of simple, purely non-sensory conscious thoughts in us (I would deny even that), the extreme case of Nemo is impossible. As I mentioned in connection with my separation argument, there are *a priori* functional constraints on belief and desire. Like the baby discussed earlier, Nemo does not satisfy these constraints. He has never had any visual or other experiences of walls or pictures, or sets of two or four objects. He does not use a language like English, or even have experiences of using a language like English. He does not have any dispositions to engage in physical behavior. In fact, he does not even have any *experiences* of behaving on the world, so he does not even have dispositions to engage in *apparent* behavior (e. g. counting or adding). He does not even have any experiences *remotely like* these experiences, even if he has "cognitive experiences". Intuitively, all this means that we cannot credit him with the belief that two plus two equals four, the belief that a picture is hanging on the wall behind oneself, and so on. Hence, even if the CP-existence thesis is true, consideration of disembodied cognitive qualia cases show that the CP-determination thesis is false.

The proponent of CP-determination can avoid this argument only by accepting what I above called systemic CP-determination. The idea would be that, even though cognitive phenomenal properties are *distinct from* all sensory phenomenal properties, they are somehow necessarily connected with sensory phenomenal properties, so that entirely "disembodied cognitive qualia" cases are impossible. On this view, cognitive phenomenal properties necessarily bring with them a whole system of mental states and abilities. But, as I said, there is no reason to accept this view and it requires brute modal connections. It is a bit like saying that, even though auditory qualia are distinct from visual qualia, they are necessarily connected with visual qualia.

In sum, the CP-determination thesis holds that cognitive phenomenal properties, and hence some belief and desire properties, are intrinsic properties distinct from sensory and functional properties. They form a distinct level of mental reality. But then they ought to be modally independent from sensory and functional properties. But consideration of separation cases, altered cognitive qualia, absent cognitive qualia, and

---

properties over and above the familiar sensory properties, in accordance with the CP-existence thesis, because it might be that relevant non-sensory phenomenal properties are *alien* phenomenal properties that we humans do not possess.) My argument is instead this: if the CP-existence thesis is true, then the following case is possible and indeed we ought to be able to positively imagine it: (i) individuals have phenomenal lives that *overlap* significantly with our phenomenal lives but (ii) they have no sensory properties. But we (or at least I) cannot positively imagine it.

disembodied cognitive qualia suggest that this is not so. They suggest a view on which the beliefs and desires of a population are fully modally dependent on (and indeed perhaps reducible to) nothing more than the sum total of *sensory* facts and (wide and narrow) functional facts about that population. I will now sketch a view of this kind.

#### 4. An Alternative: Phenomenal Functionalism

I conclude that the CP-determination thesis is false. But the CP-determination thesis does not represent the only way of developing the view that phenomenology is the basis of mental content.

I will now sketch a modified version of Lewis's influential functionalist theory of intentionality that I call *phenomenal functionalism*. Interestingly, it entails that *sensory* phenomenology is in a sense the basis of the content of belief and desire. I will not describe the case for Lewis's basic functionalism here; Lewis (1994) and others have already explained the virtues of his kind of holistic approach. Instead, after briefly stating Lewis's own functionalism, I will raise an overlooked problem for it involving the notion of *sensory evidence*. Then I will explain phenomenal functionalism and how it solves the problem. Finally, I will explain how it accommodates the main ideas of the phenomenal intentionality program while avoiding the problems I have raised for CP-determination.

Roughly, on Lewis's own functionalism (1974: 120), a subject has an (occurrent or standing) intentional state (e. g. belief or desire) with content *p* iff it is part of "best interpretation" that the agent has that intentional state. The same kind of account could be extended to the content of language. Lewis holds that the *best interpretation* is the one that best satisfies a handful of general constraints on interpretation, given the functional facts about an individual and others of his kind.<sup>31</sup>

One constraint on interpretation is the *behavior-rationalization constraint*: all else being equal, a subject has beliefs and desires that make his *behavior* largely rational. But this constraint alone does not rule out deviant interpretations. Suppose Karl is in front on a round thing and reaches for it. One interpretation is that he wants a round thing, believes that a round thing is before him, and so believes that by reaching he will get it. But another is that he has a basic desire for a saucer of mud, believes that one is before him, and so believes that by reaching he will get it.

To rule out such deviant interpretations, Lewis invokes a *humanity constraint*. It has two parts. First, some basic beliefs and desires are objectively reasonable, and people tend to have these. Second, some beliefs (and perhaps even some desires) are objectively reasonable, *given one's sensory evidence*.<sup>32</sup> So for Lewis *evidence* plays a crucial but somewhat overlooked role in ruling out deviant interpretations. Given Karl's evi-

---

<sup>31</sup> In my view, *two* functionalist theories of mental intentionality can be extracted from Lewis. First, there is the *best interpretation account* described in the text. It uses only a handful of general constraints. And it does not necessarily provide functional definitions of our names for the belief and desire relation (or of the infinitely many names of particular beliefs), anymore than Lewis's analogous *best system theory* of laws of nature provides functional definitions of names for laws. Second, there is the more familiar *functional definition account* (Lewis 1972: note 13), which *does* provide such functional definitions, based on a folk theory containing a *huge number* of specific platitudes (e. g. "if someone tastes something bitter, they have the desire to spit it out"). Here I focus on the best interpretation account, because the functional definition account faces serious problems of detail pointed out by Loar (1981: 60). Loar himself (1981: 62) proposes to solve them by explaining the belief relation and the desire relation in terms of functions-in-extension, but I think his proposal fails. (Briefly, there will be infinitely many "bent" functions that satisfy the functional definitions; and the proposal also faces the usual modal problems with trying to explain relations in terms of sets of their actual instances.)

<sup>32</sup> See Lewis (1983) and (1994). Lewis (1974) calls it a 'charity' constraint. Lewis clearly would reject a weak subjective Bayesian interpretation of the humanity constraint.

dence, the belief that there is a round thing before him is more reasonable than the belief that a saucer of mud is before him. And a basic desire for a saucer of mud would be unreasonable. So the humanity constraint will favor the first interpretation over the second, deviant interpretation. Likewise, given the humanity constraint, Lewis could easily rule out Williams' (2007) recent deviant arithmetical reinterpretation of all thought and language, by noting that forming arithmetical beliefs all the time is not a reasonable response to the available *evidence*.

Lewis has yet another constraint on thought interpretation: his famous *naturalness constraint*. This says that the propositions we believe are relatively natural or non-gruesome. Given our history of evidence, we believe that all emeralds are blue, as opposed to believing that all emeralds are *grue*. Although the point has gone unappreciated, Lewis clearly believes that the naturalness constraint *follows from* the more general humanity constraint.<sup>33</sup> It is the naturalness constraint which enables Lewis to answer the specific determinacy problem due to Kripkenstein, without having to invoke cognitive phenomenology.

Lewis's functionalism is not a view on which "there is really no such thing as intentionality, there are no matters of fact about intentionality, all attributions of intentionality are just a matter of theoretical convenience" (Strawson 2010: 348). Rather, there are objective, interpreter-independent facts about beliefs and desires, because there are objective, interpreter-independent facts about best interpretations. Nor is it a view on which we never consciously entertain content. It is an account of what it is to consciously entertain a certain content on which consciously entertaining a certain content supervenes entirely on the sensory-functional facts.

Now for the serious overlooked problem that I think afflicts Lewis's view, which I will call the *problem of evidence*. Given Karl's evidence, the humanity constraint helps to rule out deviant interpretations. But what determines his evidence? Lewis (1974: 112) speaks of "Karl's life history of evidence according to [the physical facts about him]". But arguably a subject's evidence necessarily depends somehow on his history of experiences. Further, there is reason to accept *intentionalism* about experience, according to which Karl's experiences are themselves intentional states with very detailed contents (Chalmers 2004, Pautz 2010c, Tye 2000). Under these assumptions, the problem of determining Karl's evidence is an extremely difficult special case of the problem of intentionality. Since Lewis is a *physicalist* and a *global* functionalist who applies his functionalism to all intentional states, in order to solve this problem he needs constraints on interpretation which go from the purely physical facts about Karl to the rich contents of his *experiences*. But Lewis only formulated constraints on the assignment of *beliefs* and *desires*.

How might the Lewisian functionalist solve the problem of evidence? I know of no good solution available to him. (i) Jackson is a Lewisian functionalist (Jackson and Braddon-Mitchell 2007). He also expresses sympathy for Armstrong's radical doxastic form of intentionalism according to which experience *just is* a special "experiential" form of *belief* (Jackson 2004). In that case the Lewisian functionalist could simply use the *behavior-rationalization constraint* to determine Karl's rich experiential beliefs. This would determine his experiences and hence evidence. But in fact Jackson's doxastic intentionalism is incompatible with Lewisian functionalism. According to Lewis's

---

<sup>33</sup> See Lewis (1983: 375) and (1994: 428). Since Lewis derives the naturalness constraint about thought content from his humanity constraint, I do not think Lewis could accept Williams' (2007: sect. 2) interesting suggestion that the naturalness constraint about thought content can be derived from a general principle about simplicity as a theoretical virtue used to decide between theories that fit the "data" equally well.

charity constraint, Karl's history of experiences play a crucial role in constitutively determining his beliefs, by ruling out deviant interpretations. Given this, his experiences cannot in turn be explained in terms of his beliefs, as doxastic intentionalism maintains. That would be circular. His experiences must be determined in a belief-independent way. (ii) The Lewisian functionalist might simply add to Lewis's constraints on belief and desire mentioned above a separate externalist causal constraint on experiential content (vaguely suggested by Lewis 1983: 374), one according to which suitable causal connections to the world fix the contents of Karl's experiences and hence (given intentionalism) what his experiences are like and what evidence is available to him. In turn, by the humanity constraint, Karl's experiential evidence helps to pin down the contents of his beliefs and desires and ultimately language. This solution to the problem of evidence avoids circularity, since it fixes Karl's experiences independently of his beliefs. The problem is that it amounts to something like the kind of *tracking intentionalism* about sensory intentionality and phenomenology defended by Tye (2000) and others. It is a way of developing the reductive externalist program described in the introduction according to which tracking relations are the ultimate source of all determinate intentionality. It would make Lewis's view a version of *wide functionalism*. (It would be too externalist for Lewis himself.) But elsewhere I have argued that we have empirical reasons to accept a kind of internalism about the sensory domain, which rules out tracking intentionalism (Pautz 2010b).

This brings me to phenomenal functionalism. It relies on a kind of anti-reductive approach about the sensory experiences and their intentionality that I favor along with Alston (1999), Chalmers (2006), Johnston (2007) and Levine (2008). It comes in different versions. My own view, *primitivist intentionalism*, is closest to the view of Chalmers (2006), although my arguments for it are more empirical than his. On primitivist intentionalism, all experience (with the possible exception of undirected moods) is essentially intentional. For instance, having an experience with the distinctive phenomenology of seeing a tomato just is a matter of standing in a special "experiential" intentional relation to a detailed intentional content involving simple perceptible properties like colors and shapes. I call this relation the "conscious-of" relation. (This is a technical term; I do not mean we are literally conscious of contents, which are abstract objects.) What makes the view non-reductive is that it also holds that this relation is primitive. Although what contents we are conscious of in some sense supervenes on physical conditions (on my internalist view, internal physical conditions), the conscious-of relation is not identifiable with a relation characterizable in non-phenomenal, physical terms (e. g. a tracking relation). This anti-reductive view of sensory intentionality and my own case for it are very much within the spirit of the phenomenal intentionality program. As noted, I think we have empirical reasons to hold that sensory experience, with its rich built-in world-directed intentionality, is determined by *internal factors*, not tracking relations to the environment. But, in my view, we simply have no good *reductive* theory of how sensory consciousness might be both essentially externally-directed and yet internally-dependent (Pautz 2010a).

Phenomenal functionalism combines anti-reductionism about Karl's sensory experiences with a theory of all of Karl's intentional states *other than* his sensory experiences, for instance his occurrent and standing beliefs and desires. In particular, Karl has the belief or desire that *p* iff it is part of "best interpretation" that the agent believes or desires that *p*. Here the *best interpretation* is the one that best satisfies the constraints on interpretation, given two sorts of facts about Karl: (i) his history of experiences, which according to primitivist intentionalism are themselves fundamental in-

tentional states; and (ii) the wide and narrow functional facts about Karl and others of his kind. The wide functional facts help determine the “wide” contents of his beliefs and desires.

Lewis’s functionalism is global and reductive: *all* of Karl’s intentional states – his beliefs, desires and experiences – reduce to purely third-personal, physical and functional facts about him. By contrast, phenomenal functionalism is local and non-reductive. On phenomenal functionalism, a holistic functionalist theory applies to all of Karl’s *beliefs* and *desires*, according to which they entirely supervene on, and indeed reduce to, clusters of experiential and functional facts. But no functionalist or other theory applies to Karl’s sensory experiences, which are irreducible intentional states. Hence Quine and Brentano are right that intentionality is not reducible all the way down to the fundamental physical level. On phenomenal functionalism, when it comes to the problem of intentionality, we start with a giant leg up, because Karl’s experiences, with their determinate world-directed intentionality, are among the fundamental facts about Karl. Karl’s experiences of the world, which determine his history of evidence, are anchor points that help to determine, *via* the humanity constraint, the contents of his downstream beliefs and desires and ultimately language, by ruling out deviant interpretations. Thus the present local functionalist theory solves the problem evidence that plagues Lewis’s global functionalist theory.

Phenomenal functionalism should be attractive to friends of the phenomenal intentionality program for several reasons.

(i) Phenomenal functionalism avoids the problems I raise for CP-determination. First, it does not require separate “danglers” connecting the physical facts about an agent like Karl with his beliefs and desires, because it holds that once the sensory and functional facts about Karl are specified the facts about his beliefs and desires come for free. Second, it both accommodates and explains the functional constraints on belief and desire: they follow from the real definition of the belief relation and desire relation in terms of the “best interpretation”, given the functional and sensory facts about an agent. Third, it accommodates the plausible thesis of *sensory-functional supervenience*: the thought that the total intentional facts about a population supervene entirely on the total sensory and functional facts about that population, showing that there is no need to posit a distinct level of “cognitive phenomenology”.

(ii) Phenomenal functionalism is also compatible with the claim of proponents of the phenomenology intentionality program that much intentionality supervenes on phenomenology alone.

For the purposes of illustration, suppose that there really could be an accidental, life-long brain-in-a-vat that is phenomenally identical to some actual person.<sup>34</sup> Call

---

<sup>34</sup> To be clear, I am making the brain-in-the-vat supposition for the purposes of illustration. (My chief claim that much content supervenes on phenomenology alone is of course neutral on whether a brain-in-a-vat could enjoy phenomenology.) I am skeptical of the attempt by proponents of the phenomenology intentionality program such as Graham and Horgan (2004) to establish this supposition by consulting our “intuitions”. In my view, we simply do not have an *intuition* (or any pretheoretical justification) in favor of believing that *all* metaphysically possible, or even *all* nomically possible, duplicates of one’s own brain must be phenomenal duplicates of oneself (see note 4). For instance, we cannot rule out by mere *intuition* a version dualism on which there are peculiar psychophysical laws that require brains to be normally embodied if they are to support phenomenal consciousness. On such a view, phenomenology does not supervene on the brain with metaphysical or even nomic necessity. Now we might have the intuition that there are *some* metaphysically possible life-long, accidental brains-in-vats that are phenomenal duplicates of oneself (in other words, that a brain-in-a-vat *could* support one’s actual phenomenal life) because phenomenology is intuitively constitutively independent of all external physical conditions. But we also have the intuition that *complete physical duplicates* of us could lack consciousness (zombies) or have different experiences from us (spectrum inversion). Unless they are out-and-out dualists, proponents of the phenomenal intentionality program will admit that *these* intuitions about the relationship between the physical and the phenomenal are faulty. So if they admit that our intuitions about the relation-

him Karl. Karl has an experience as of a round thing, he attempts to behave as if a round thing is present, and he has experiences as of so behaving. Hence, by the humanity constraint and the behavior-rationalization constraint and the other constraints on interpretation, he counts as having the (false) occurrent belief that a round thing is present. This may fit with Horgan and Tienson's somewhat obscure claims about the relevance of "phenomenal role" to narrow content (see note 23). Indeed, maybe "round" in Karl's language get its content directly from the content of his (hallucinatory) experiences, by way of his accepting something like "round is *that* shape", while demonstrating an uninstantiated shape property presented by his experience. And maybe the sentence ' $2+2=4$ ' in Karl's language gets a certain arithmetical content, thanks to its use and "inferential role" as well as considerations of naturalness. And maybe Karl counts as occurrently believing this content, on a certain occasion, by virtue of accepting and understanding this sentence.<sup>35</sup>

Contrary to CP-determination, then, phenomenal functionalism holds that the contents of Karl's occurrent beliefs and desires are never determined merely by his "cognitive phenomenology" at the time he forms those beliefs and desires (indeed, there is, I think, no such thing). Rather, on phenomenal functionalism, Karl's beliefs and desires are always determined by a cluster of actual and dispositional facts about his sensory experiences in the past and future as well as the present (including what sentences it accepts and their patterns of use), without the need of an additional layer of "cognitive phenomenology" playing any role.<sup>36</sup>

Other cases show the importance of phenomenology to content. For instance, phenomenal functionalism allows that, owing to internal neural differences, two individuals belonging to different populations might be black-white spectrum inverted and hence have color experiences and color beliefs with radically different intentional contents, even though their sensory states play the same typical *input-output* functional roles in their populations. Lewis's global functionalism cannot handle black-white spectrum inversion, since it holds all intentional differences must be grounded in in-

---

ship between the physical and the phenomenal are generally faulty, and if their intuition concerning the possibility of brains-in-vats with phenomenal lives is simply another example of such an intuition, then they ought not put any stock in that intuition.

<sup>35</sup> On Lewis's own view (1974: 117), mental content is always prior to linguistic content. By contrast, as my examples here suggest, I think that in some cases (e. g. perceptual content) mental content is prior to linguistic content, while in other cases (e. g. arithmetical content) linguistic content is prior to mental content. However, as Lewis notes, his basic theory (and hence my modified version of his theory) is in fact compatible with a mixed view of this kind (he calls it "method 3").

<sup>36</sup> Given that phenomenal functionalism accommodates so much of what proponents of phenomenal intentionality program like Horgan and Tienson (2002) want to say, it is natural to wonder if they have provided any reason to reject such views on which intentional facts are fixed by the sensory and functional facts, with cognitive phenomenology playing no essential role. As far as I can see, they would raise two objections, but neither is successful. (i) They would object that cognitive phenomenology is needed to secure our actual level of content determinacy (Horgan and Graham 2010). But, as we saw in connection with "absent cognitive qualia cases" (in §2.1 and §3.3), they have not supported this claim, and it is not clear how they could support it. (ii) Horgan and Graham (2010) raise an important general objection to Lewis's theory of intentionality, which also applies to my Lewis-inspired "phenomenal functionalism": what *makes it the case* that naturalness enters into the determination of intentionality? For instance, what makes it the case that naturalness enters into the determination of linguistic reference (or the semantic value relation)? (Lewis holds naturalness plays a role in the determination of linguistic intentionality as well as mental intentionality, because he holds it rules out "bent" grammars.) But what are Horgan and Graham asking? Lewis holds that naturalness helps determine reference because the reference relation *just is* a relation involving naturalness among other things: call it the *naturalness-plus relation*. So in effect Horgan and Graham are asking for an explanation of an identity. But in what sense can *any* identities be explained? Of course, it is sensible to ask for an explanation of the corresponding metalinguistic fact: "By what does 'reference' refer to the naturalness-plus relation?" But Lewis has an answer which invokes his own theory: 'reference' bears the naturalness-plus relation to the naturalness-plus relation. This kind of self-subsumption is not circular in any metaphysical sense (a point repeatedly made in the literature spawned by Putnam's model-theoretic argument).

put-output functional differences.<sup>37</sup> By contrast, phenomenal functionalism easily accommodates black-white spectrum inversion, since it does not apply functionalism to the content of experience, holding instead that it is fixed by neural processing. On phenomenal functionalism, even Strawson's weather-watchers (2010) who cannot act on the world could have a rich set of beliefs and desires about their environments and own mental lives, because of facts about their experiences of the world and internal inferential roles.

(iii) Finally, phenomenal functionalism (and indeed Lewis's own functionalism) entails the following version of the chief idea of the phenomenal intentionality program that phenomenology is the *source* of all intentionality: necessarily, if a creature has never had the capacity for experiences (and does not belong to a community with such a capacity), it is not determinately belief or desire any particular contents (as I take it that we do). To see this, consider a community of zombies, who have no experiences but who otherwise resemble us as much as possible. On phenomenal functionalism, since the zombies have no experiential evidence, the humanity constraint cannot kick in to make it the case that deviant interpretations are mistaken. So, if we allow that it has beliefs at all, there is always massive indeterminacy concerning what their contents are.<sup>38</sup>

I find the source thesis to be independently plausible. Here is an argument for it. An ideal interpreter given all the fundamental physical facts about any population of zombies would not be able to deduce that they have any beliefs and desires that determinately possess some contents and not others. I think this is obvious. In our case, an ideal interpreter could rule out "deviant interpretations" provided he knew *all the facts about our sensory-phenomenal lives* in addition to the physical facts about us. He could arguably figure out what we believe and desire. But in the case of our zombie duplicates he would lack this enormous resource, and so could not rule out deviant interpretations. In my view, the fact that determinate intentional facts are not deducible from non-phenomenal facts alone is evidence that non-phenomenal facts alone are not enough to metaphysically settle determinate intentional facts.<sup>39</sup>

#### 4. Conclusion

---

<sup>37</sup> Braddon-Mitchell and Jackson (2007: 134) are Lewisian functionalists who apparently hold that in some inverted spectrum cases the phenomenal differences between the invertees are explained by neural differences, not folk-psychological functional differences. As noted in the text, I think this is the right view. But, again, it is simply not open to a global functionalist about intentionality like Jackson (Braddon-Mitchell and Jackson 2007: 190-193). Given intentionalism (which Jackson 2004 accepts), all phenomenal differences are intentional differences. And, on global functionalism, all intentional differences must be based in folk-psychological *functional* differences, as opposed to mere neural differences. Moral: black-white spectrum inversion among functional duplicates provides another reason, besides "the problem of evidence", to accept localist phenomenal functionalism over Lewis's global functionalism.

<sup>38</sup> Do our zombie twins count as having any beliefs or desires at all, despite not having the capacity for experience? Many would say they do. I find this counterintuitive. This may be semantic issue, since we agree about the fundamental facts of the case. However, even if we allow that the zombies have beliefs and desires (or beliefs\* and desires\*), the crucial point is that phenomenal functionalism entails that it cannot have the same beliefs and desires that we have (lacking acquaintance with conscious states and sensible qualities), and that it is at best it is radically indeterminate what their true contents are.

<sup>39</sup> In response to this "epistemic gap" argument, those who think that zombies might have beliefs and desires that determinately possess certain contents might reply that it has such determinate beliefs and desires even if I am right that an ideal interpreter could not deduce them *a priori* from the fundamental physical facts about its situation. In other words, they might advocate *a posteriori* physicalism about the alleged determinate beliefs and desires of zombies. Soames (1998) takes a view of this kind about our own beliefs and desires. But to make this credible one would have to least sketch physicalist theories (perhaps an *a posteriori* physicalist theories) of the belief relation and the desire relation which explains *how* the purely non-phenomenal, non-intentional physical facts about the zombies determine that it comes to determinately bear these relations to some contents rather than others. The history of failed attempts to sketch such theories provides strong grounds for doubting that one can be provided.

The CP-determination thesis is mistaken. Instead, consideration of various thought-experiments suggests that all intentionality is entirely grounded in functional facts as well as sensory and perceptual facts, which on my view are richly intentional and irreducible. The question is how this is so. I have just sketched an answer to this question - phenomenal functionalism - but it is programmatic.<sup>40</sup> Similar views have recently been proposed by Schwitzgebel and Chalmers, which also deserve to be explored.<sup>41</sup> In any case, I think one thing we have learned is that it is not enough to simply declare that “phenomenology fixes the contents of occurrent beliefs and desires”, as many proponents of CP-determination do. Those of us sympathetic to the phenomenal intentionality program need a (perhaps non-reductive) *theory* of phenomenology and of *how* it fixes content.

## References

- Alston, W. 1999. Back to the Theory of Appearing. *Philosophical Perspectives* 13: 181–203.
- Armstrong, D. 1968. *A Materialist Theory of Mind*. London: Routledge.
- Boghossian, P. 1989. The Rule Following Considerations. *Mind* 98: 507–549.
- Braddon-Mitchell, D. and F. Jackson. 2007. *Philosophy of Mind and Cognition* (2nd ed.). Oxford: Blackwell.
- Byrne, A. and D. Hilbert. 2003. Color Realism and Color Science. *Behavioral and Brain Sciences* 26: 3–21.
- Chalmers, D. 2004. The Representational Character of Experience. In *The Future for Philosophy*, ed. B. Leiter. Oxford: Oxford University Press.
- Chalmers, D. 2005. Representationalism Showdown. Online. Available: <http://fragments.consc.net/djc/2005/09/representationa.html>.
- Chalmers, D. 2006. Perception and the Fall from Eden. In *Perceptual Experience*, ed. T. Gendler and J. Hawthorne. Oxford: Oxford University Press.
- Chalmers, D. 2008 The Critique of Pure Thought. Online. Available: <http://consc.net/powerpoint.html>.
- Chalmers, D. Forthcoming. *Constructing the World*. Oxford: Oxford University Press.
- Cohen, J. 2009. *The Red and The Real*. Oxford: Oxford University Press.
- Dorr, C. 2007. There Are No Abstract Objects. In *Contemporary Debates in Metaphysics*, ed. T. Sider, J. Hawthorne and D. Zimmerman. Malden, Mass: Wiley-Blackwell.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge: MIT Press.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, J. 1994. *The Elm and the Expert*. Cambridge: MIT Press.
- Graham, G., T. Horgan and J. Tienson 2007. Consciousness and Intentionality. In *The Blackwell Companion to Consciousness*, ed. M. Velmans and S. Schneider. Oxford: Basil Blackwell.
- Graham, G., T. Horgan, and J. Tienson. 2009. Phenomenology, Intentionality, and the Unity of Mind. In *Oxford Handbook of Philosophy of Mind*, ed. A. Beckermann and B. McLaughlin. Oxford: Oxford University Press.
- Harman, G. 1990. The Intrinsic Quality of Experience. In *Philosophical Perspectives* 4, ed. J. Tomberlin. Northridge: Ridgeview Publishing Company.
- Hawthorne, J. 2004. Humeans Are Out of Their Minds. *Noûs* 38: 351–8. References are to the version reprinted in J. Hawthorne, *Metaphysical Essays* (Oxford: Oxford University Press, 2006).
- Hill, C. 2009. *Consciousness*. Cambridge: Cambridge University Press.

<sup>40</sup> There is one overlooked problem for Lewis’s theory involving the obscure notion of a *best interpretation* that phenomenal functionalism does *not* avoid (Pautz 2010a: 55). While I think phenomenal functionalism is promising, it faces many problems.

<sup>41</sup> Schwitzgebel (2002) defends “phenomenal dispositionalism” which has some similarities to my phenomenal functionalism. (In my view, Schwitzgebel’s theory is crucially incomplete, because he does not provide a general account of what it is to stand in the belief relation or the desire relation to an arbitrary proposition, an account that would explain how interpreters can go on to new cases once we grasp “belief” and “desire”.) Chalmers (2008) flirts with a “combined view” on which belief content is grounded in sensory intentionality and inferential role. This is similar to phenomenal functionalism. Likewise in his discussion of reference magnetism and the grounding of intentionality, Chalmers (forthcoming) defends a “local descriptivism”. He speculates that all content is grounded in a primitive acquaintance relation that we bear to properties in experience, inferential role, and naturalness. This is quite similar to my “local functionalism”. Indeed, his acquaintance relation seems to be just what I have called the “conscious-of relation”. We both think that the relevant relation is primitive and not grounded in anything more basic. But there are differences. Chalmers appears to favor a kind of “language of thought” approach according to which thoughts are always made up of “concepts” construed as mental representations; concepts are the basic carriers of intentional content; and the contents of complete thoughts are somehow determined compositionally, by what concepts they are made up of. In addition, Chalmers accepts “epistemic two-dimensionalism”, according to which the contents of concepts is grounded in certain more basic facts about what thoughts it is rational to have given that one has the thought that a certain complete possible scenario is actual. By contrast, my Lewis-inspired “phenomenal functionalism” is not committed to a language-of-thought-like approach or to epistemic two-dimensionalism.

- Hill, C. 2012. Locating Qualia: Do They Reside in the Brain or in the Body and the World? In *New Perspectives on Type Identity: The Mental and the Physical*, eds. S. Gozzano and C. Hill. Cambridge: Cambridge University Press.
- Horgan, T. 2010. Materialism, Minimal Emergentism, and the Hard Problem of Consciousness. In *The Waning of Materialism*, ed. G. Bealer and R. Koons. Oxford: Oxford University Press.
- Horgan, T. and G. Graham. 2010. Phenomenal Intentionality and Content Determinacy. In *Prospects for Meaning*, ed. R. Shantz. Amsterdam: de Gruyter.
- Horgan, T. and J. Tienson. 2002. The Intentionality of Phenomenology and the Phenomenology of Intentionality. In *Philosophy of Mind: Classical and Contemporary Readings*, ed. D. Chalmers. Oxford: Oxford University Press.
- Horgan, T., J. Tienson, and G. Graham. 2004. Phenomenal Intentionality and the Brain in a Vat. In *The Externalist Challenge: New Studies on Cognition and Intentionality*, ed. R. Shantz. Amsterdam: de Gruyter.
- Jackson, F. 1977. *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Jackson, F. 2004. Representation and Experience. In *Representation in Mind: New Approaches to Mental Representation*, ed. H. Clapin, P. Slezack and P. Staines. Amsterdam: Elsevier.
- Johnston, M. 2007. Objective Mind and the Objectivity of Our Minds. *Philosophy and Phenomenological Research* 75: 233-68.
- Kripke, S. 1982. *Wittgenstein on Rules and Private Language*. Oxford: Blackwell.
- Kriegel, U. 2009. *Subjective Consciousness: A Self-Representational Theory*. Oxford: Oxford University Press.
- Kriegel, U. 2011. *The Sources of Intentionality*. Oxford: Oxford University Press.
- Kriegel, U. 2011. The Phenomenal Intentionality Research Program. This volume.
- Langton, R. and D. Lewis. Defining 'Intrinsic'. *Philosophy and Phenomenological Research* 58: 333-345.
- Levine, J. 2008. Secondary Qualities: Where Consciousness and Intentionality Meet. *The Monist* 91: 215-236.
- Levine, J. 2011. On the Phenomenology of Thought. In *Cognitive Phenomenology*, ed. T. Bayne and M. Montague. Oxford: Oxford University Press.
- Lewis, D. 1972. Psychophysical and Theoretical Definitions. *Australasian Journal of Philosophy* 50: 249-258.
- Lewis, D. 1974. Radical Interpretation. *Synthese* 21: 331-344. References are to the version reprinted with a postscript in D. Lewis *Philosophical Papers: Vol. I* (Oxford: Oxford University Press, 1983).
- Lewis, D. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 63: 343-377.
- Lewis, D. 1994. Reduction of Mind. In *A Companion to the Philosophy of Mind*, ed. S. Guttenplan. Oxford: Blackwell.
- Loar, B. 1981. *Mind and Meaning*. Cambridge: Cambridge University Press.
- Lycan, W. 2001. The Case for Phenomenal Externalism. In *Philosophical Perspectives, Vol. 15: Metaphysics*, ed. J.E. Tomberlin. Atascadero, CA: Ridgeview Publishing.
- Lycan, W. 2008. Phenomenal Intentionalities. *American Philosophical Quarterly* 45: 233-252.
- Lycan, W. 2009. Giving Dualism Its Due. *Australasian Journal of Philosophy* 87: 551-563.
- Pautz, A. 2006. Sensory Awareness is not a Wide Physical Relation. *Noûs* 40: 205-240.
- Pautz, A. 2008. The Interdependence of Phenomenology and Intentionality. *The Monist* 91: 250-272.
- Pautz, A. 2010a. A Simple View of Consciousness. In *The Waning of Materialism*, ed. G. Bealer and R. Koons. Oxford: Oxford University Press.
- Pautz, A. 2010b. Do Theories of Consciousness Rest on a Mistake? *Philosophical Issues* 20: 333-367.
- Pautz, A. 2010c. Why Explain Visual Experience in terms of Content? In B. Nanay (ed.) *Perceiving the World*. Oxford: Oxford University Press.
- Pautz, A. 2011a. Can Disjunctivists Explain Our Access to the Sensible World? *Philosophical Issues* 21.
- Pautz, A. 2011b. The Phenomenal Intentionality Program: A Backhanded Defense.  
Available online at: <https://webpace.utexas.edu/arp424/www/backhanded.pdf>
- Pitt, D. 2009. Intentional Psychologism. *Philosophical Studies* 146: 117-138.
- Pitt, D. 2011. Introspection, Phenomenality and the Availability of Intentional Content. In *Cognitive Phenomenology*, ed. T. Bayne and M. Montague. Oxford: Oxford University Press.
- Price, H. 1954. *Perception*. London: Methuen.
- Quine, W. 1960. *Word and Object*. Cambridge: MIT Press.
- Russell, B. 1912. *The Problems of Philosophy*. London: Williams & Norgate.
- Schaffer, J. 2009. On What Grounds What. In *Metametaphysics*, ed. D. Chalmers, D. Manley and R. Wasserman. Oxford: Oxford University Press.
- Schwitzgebel, E. 2002. A Phenomenal, Dispositional Account of Belief. *Noûs* 36: 249-275.
- Searle, J. 1992. *The Rediscovery of the Mind*. Cambridge: MIT Press.
- Siegel, S. 2011. The Contents of Perception. In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. Online. Available: <http://plato.stanford.edu/archives/spr2011/entries/perception-contents/>.
- Siewert, C. 1998. *The Significance of Consciousness*. Princeton: Princeton University Press.
- Siewert, C. 2011. Phenomenal Thought. In *Cognitive Phenomenology*, ed. T. Bayne and M. Montague. Oxford: Oxford University Press.
- Smart, J. C. 1959. Sensations and Brain Processes. *Philosophical Review* 68: 141-156.
- Soames, S. 1998. Skepticism About Meaning, Indeterminacy, Normativity, and the Rule-Following Paradox. *Canadian Journal of Philosophy, Supp* 23: 211-250.
- Strawson, G. 2010. *Mental Reality, 2<sup>nd</sup> Edition*. Cambridge: MIT Press.
- Strawson, G. 2006. Realistic Monism: Why Physicalism Entails Panpsychism. *Journal of Consciousness Studies* 13: 3-31.
- Swoyer, C. 1982. The Nature of Natural Laws. *Australasian Journal of Philosophy* 60: 203-223.
- Tye, M. 2000. *Consciousness, Color and Content*. Cambridge: MIT Press.
- Tye, M. and B. Cutter. 2011. Tracking Representationalism and the Painfulness of Pain. *Philosophical Issues* 21.
- Williams, R. Eligibility and Inscrutability. *Philosophical Review* 116: 361-399.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.