

Phenomenal Privacy, Similarity and Communicability

Thomas Raleigh

ABSTRACT: A familiar sceptical worry about other minds is that if conscious experiences are *private* then I could never *know* whether another subject's experiences are like mine. A somewhat different sort of worry is that I could not even *understand* another subject's claims about the private features of her experience—and, conversely, that she could never understand my claims about my private features. Edward Craig (1982; 1986; 1997), taking Wittgenstein (1953) as his target, repeatedly argued that such *incommunicability* would not follow from the private nature of the features in question. The basic idea is that two subjects might each privately introduce predicates to describe their own private phenomenal features and these predicates might (albeit unknowably) have the same meaning, and that this could suffice for mutual understanding. After clarifying the notion of privacy at issue, I argue against this apparently plausible line of thought by showing how hard it is to make sense of the allegedly possible sameness of meaning.

1. Introduction

The idea that there are features of or in our conscious experience that are, in some important sense, *private* has both a long history in philosophy and a large measure of intuitive attraction. Once this idea is in place, it will be very natural to assume that one can think and judge about one's own private features. And it is then only a small step to the idea that we might communicate such thoughts and judgements about our respective private features with each other.

As students of philosophy soon learn, a traditional sort of scepticism raises the worry that we cannot ever know whether our respective private features are qualitatively the same—maybe your experience when looking at a ripe tomato is privately very different from mine; maybe your sensations when you bang your funny bone are quite different to how banging my funny bone feels to me; maybe, despite your fluent talk about experiences of various kinds, you actually have no inner conscious life whatsoever. And so on. However, even allowing that this sort of sceptic may have a point, the following thought can seem very tempting and natural: ok, perhaps I cannot be *sure* that your private stuff is the same as mine. But so long as it *is* in fact the same (and we both suppose or believe this to be so even if we may not be able to know it), then we can at least still successfully communicate and understand each other's claims about our respective private stuff. The fullest expression and development of this very natural line of thought, at least in the recent literature, occurs in a series of papers by Edward Craig (1982; 1986; 1997).

Against this very natural line of thought I will argue that, given a common and apparently highly plausible conception of phenomenal privacy, it is in fact very difficult to make sense of the required sameness of meaning for different subjects' privately coined concepts. To anticipate, I will argue that there cannot be natural similarities in the phenomenal appearance or feel of two private items given that there is no metaphysical possibility of their both appearing within a common subjective viewpoint, or 'appearance space'. But neither can there be any dispositional facts as to how I *would* conceptually or

linguistically respond to your private items, *were* I to experience them (nor vice-versa). Given that we can appeal neither to natural phenomenal similarities between the private items from our different domains, nor to any classificatory dispositions in response to experiencing each other's private items, we lack any basis in virtue of which our privately coined predicates could have the same meaning. So it turns out that on the model of a private domain in question, not only do we face the traditional sceptical worry about *knowing* whether our private features are alike, we are left unable even to *understand* each other's claims about our private features.

In arguing that the otherwise intuitively compelling model of a private domain of experience has this unwelcome consequence, I will be arguing for a thesis that occurs, as it happens, amongst Wittgenstein's famous remarks on 'private language'. But so whilst I will occasionally mention his work in what follows, the aim of this paper is not exegesis or defence of Wittgenstein. Rather, it is simply to establish an unattractive consequence for a perennially attractive way of thinking about conscious experience.

2. Privacy

So what exactly is the notion of privacy that is at issue? Craig defines the notion of a state's being private in epistemic terms:

It has been widely held that certain states of sentient creatures are private, in the technical sense that their nature cannot be known by anyone other than the subject who experiences them. (Craig 1997: 127)

One might wonder what exactly it is to 'know the nature' of something—Craig goes on to provide the following example:

For instance, the phenomenal quality of perceptual states has often been seen in this light. 'I know you call that colour by the same name, but I can't know whether you see it in the same way as I do' is a position familiar to most students of philosophy, both amateur and professional. (Craig 1997: 127)

Let's define an epistemic notion of privacy then as follows:

(EP) EPISTEMIC PRIVACY: An element/feature of a subject's conscious experience, *e*, is *epistemically* private iff it is impossible that another subject can *know* whether/how *e* is phenomenally similar/dissimilar to elements/features in her own stream of consciousness.

But rather than focusing on the impossibility of acquiring *knowledge* of the phenomenal nature of another subject's mental state, an alternative way of understanding privacy is simply in terms of the impossibility of *experiencing* the phenomenal features of another subject's mental state. I will call this alternative conception, the *metaphysical* notion of privacy:

(MP) METAPHYSICAL PRIVACY: An element/feature of a subject's conscious experience, *e*, is *metaphysically* private iff it is metaphysically impossible for *e* to likewise constitutively form part of another subject's conscious experience.

It is commonly accepted that a token experience or sensation always necessarily *belongs* to some specific subject—and hence that it is metaphysically or logically impossible for another subject to *possess* one of my experiences or my sensations. For example,

The barriers that prevent us from enjoying one another's experiences are not natural but **logical** **It is not conceivable** that there should be people who were capable of having one another's pains. (Ayer 1940: 138–139, bold type added)

Experiences ... owe their identity as particulars to the identity of the person whose states or experience they are. From this it follows immediately that if they can be identified as particular states or experiences at all they must be possessed ... in such a way that it is **logically impossible** that a particular state or experience in fact possessed by someone should have been possessed by someone else. (Strawson 1959: 97, bold type added)

... the pain you are feeling – that particular pain – is private to you. It is yours alone, and **necessarily** so. No one else could have that particular pain. Of course, conceivably someone else could have a pain that felt just like your pain, but only you could have that very pain. What is true for this one pain is true for pains generally. Indeed, it is true for all mental objects of experience. None of these items of experience can be shared. I cannot have your visual images or feel your tickles, for example. Your images and tickles **necessarily** belong to you. (Tye 2007: 24, bold type added)

We must be clear though that (MP) goes beyond the constraint that an experience or sensation necessarily *belongs* to a specific subject and no other. For it is a further claim that no other subject can *experience* my experience or sensation—a claim that does not obviously follow from the fact that no other subject can *possess* my experience or my sensation. Perhaps it is a possibility that other subjects could directly experience the very same token inner sensations as I do, but that these sensations would still necessarily belong to me (they would essentially depend for their existence on me rather than the other subjects). In order to be private in the sense defined by (MP), token instances of elements/features that occur in one subject's conscious experience cannot (with metaphysical necessity) occur in the consciousness of another subject. By way of contrast, suppose that I could somehow 'directly' experience elements/features in your stream of consciousness—for example, perhaps if our brains and nervous systems could somehow be 'fused', or perhaps, as Wittgenstein imagined in §253 of the *Philosophical Investigations* (1953),¹ if we were suitably conjoined twins, we could both have experiences constitutively featuring the one same token phenomenal element or property instance. To repeat, this would not be a matter of my merely having *qualitatively* identical experiences running in parallel to yours; the idea would be that the very same *numerically* identical phenomenal features and elements that are showing up in your stream of consciousness, and which would perhaps still essentially *belong* to you, could also constitutively feature in my stream of consciousness. If this really were a metaphysical possibility then such phenomenal features of consciousness would not be private in the sense specified by MP.

Some caution is needed here with the idea of 'directly' experiencing another subject's conscious states—for there is another important but distinct sense in which some philosophers would want to allow that one subject can *perceptually experience* or *directly observe* another's conscious mental states.

¹ Abbreviated hereafter as PI.

A number of philosophers—for example, McDowell (1982), Cassam (2007), Green (2010), Stout (2010), McNeill (2012), to name just a few—have endorsed some kind of perceptual model for our knowledge of other minds. Whilst there are many important differences between these perceptual theorists, the basic idea is that I can, in the right circumstances, simply *see* that you are angry or *see* that you are in pain when, say, I see your darkening countenance or your wincing expression. According to such theorists then, there is a sense in which I can visually experience, or ‘directly observe’, your conscious mental state of being angry or being in pain and so, perhaps, gain non-inferential knowledge that you are angry or in pain. But notice that such a perceptual theorist might well nevertheless *accept* that even though I can see that you are in pain, or see your pain embodied in your face, I still do not consciously experience the very same token instance of phenomenal painful-feeling that you, alas, are consciously aware of. In other words, even if it is claimed that your conscious mental state can be part of the content of my perceptual experience, or that I can, in one sense, ‘directly’ see and thus know that you are in pain, still one might accept that there are various phenomenal elements or aspects of this conscious mental state that I do not and cannot experience, which you alone can stand in the relation of conscious awareness to (with metaphysical necessity).

What is the relationship between EP and MP? At first glance the two conditions look like they could come apart. *Prima facie*, it might seem that an element in one subject’s stream of consciousness could satisfy MP but not EP if there were some way for another subject to gain knowledge of its phenomenal nature *other* than by ‘directly’ experiencing it for herself. For example, perhaps an omniscient angel or oracle can just tell you whether and how my metaphysically private sensation is phenomenally similar/dissimilar to your own. Or maybe you can even work out ‘what it’s like’ by hard work and inference.²

Conversely, an element, *e*, in one subject’s stream of consciousness might satisfy EP but not MP if there was some other subject who *could* directly experience *e* but who was, for whatever reason, unable to gain any *knowledge* of its phenomenal nature as a result. For example, perhaps they could not possess the conceptual capacities for such knowledge, perhaps something necessarily prevents them fully attending to *e*, perhaps they are an essentially Pyrrhonian creature who is bound always to suspend judgement rather than forming any beliefs.

However, I think that once we set aside such bizarre kinds of cases—in which a subject is somehow *necessarily* prevented from gaining knowledge of the phenomenal nature of some element/feature despite it showing up in her own stream of consciousness—there *is* an important connection from EP to MP. That is, for a normal, conceptually competent, attentive, rational (etc.) subject, if an element/feature is a constitutive component in her stream of consciousness she is *thereby* in a position to gain knowledge of the phenomenal nature of that element/feature (i.e., to gain knowledge of its phenomenal similarity/dissimilarity to other elements/features in her stream of consciousness). And so if the very same token phenomenal element/feature in my stream of consciousness could also constitutively feature in your consciousness, you would (assuming you are an otherwise normal human subject) thereby be in a position to gain knowledge of that item’s phenomenal nature. In other words, once we restrict our attention to normally rational, non-Pyrrhonian, conceptually competent (etc.) subjects, if a mental element/feature is not metaphysically private then it is not epistemically private either. For if it *is* possible for another such subject to directly experience that token element/feature, then it is thereby possible for that other subject to gain knowledge of the phenomenal nature of that element/feature. So,

² As, e.g., Dennett (1991; 2007) has claimed would be the case with Jackson’s Mary (1982; 1986).

setting aside the esoteric possibility of a subject who can directly experience another subject's stream of consciousness but who could not possibly gain any knowledge of its phenomenal nature as a result, EP effectively implies MP.

What about the other direction of entailment, from MP to EP? Above I suggested that *prima facie* it seems to be a coherent possibility that the elements/features in my stream of consciousness, and the elements/features in yours, might both satisfy MP and yet there might be some other way for us to come to know whether or how they are phenomenally similar—perhaps an omniscient demon could somehow tell you whether and how your metaphysically private features are phenomenally similar or dissimilar to my metaphysically private features. But of course the possibility of imparting such knowledge would require that there is a fact of the matter to know in the first place—concerning how our respective metaphysically private features of experience are phenomenally similar/dissimilar to each other. And one of the main contentions of this paper (see Section 5, below) will be that *there could be no fact of the matter* concerning how two metaphysically private items from different subjects' private domains are phenomenally similar/dissimilar. Given the model of a metaphysically private domain of phenomenal features/elements, phenomenal similarity is not well-defined across different private domains—or so I shall claim. And if there simply is no well-defined relation of phenomenal similarity/dissimilarity holding between metaphysically private items from different subjects, then a fortiori there is no possibility of one of these subjects gaining knowledge of how the other subject's private elements are phenomenally similar/dissimilar to her own. In other words, I will claim that MP *does* imply EP—as there could be no fact of the matter concerning how my metaphysically private items are phenomenally similar/dissimilar to someone else's private items, there could be no knowledge of any such (non-existent) relations of phenomenal similarity/dissimilarity, so the metaphysically private items would count as epistemically private also.

If I'm right then, despite the initial appearance that they can come apart, *e* satisfies EP \leftrightarrow *e* satisfies MP—at least for normal rational subjects. In what follows I will tend to focus on the metaphysical notion of privacy, MP, as it seems plausible that this is the more fundamental notion—that a conscious feature/element is epistemically private (at least for normal rational subjects) is *explained by* its being metaphysically private. I will speak of 'private domains', 'private elements', 'private features', 'private properties', 'private facts', etc., to pick out the various kinds of things that could constitutively feature in a stream of consciousness. However, I don't think it is generally very helpful to talk about a 'private language'.³ For a start, a language, qua abstract type, is not obviously the sort of thing that could be an object of conscious sensory experience. More importantly, to characterise a language, L, as 'private' could mean one of (at least) three things:

(a) L is a language whose terms are coined (have their meanings bestowed) by an act of ostensive definition involving unique experiential access to a private item/feature;

³ So far as I can tell, Wittgenstein only uses the phrase 'private language', without himself putting it in quotation marks, once in the *Investigations*, in the short and rather cryptic §259:

Are the rules of a private language *impressions* of rules? – The balance on which impressions are weighed is not the impression of a balance'. (§259)

The phrase also occurs, though now placed within quotation marks, in sections §269 and §275:

... sounds which no one else understands but which I '*appear to understand*' might be called a "private language". (§269)

... I am saying: you have not the feeling of pointing-into-yourself that often accompanies 'naming the sensation' when one is thinking about "private language". (§275)

(b) L is a language whose terms only refer to/describe distinctively private items/features;

(c) L is a language whose terms have meanings that can only be understood by a unique subject—that is, whose terms have *incommunicable* meanings.

That (a) and (b) are not equivalent is clear when we consider the possibility that a token property instance could be private, yet be an instance of a type that is not private. For example, suppose that private sense-data can literally be square-shaped. Then by attending to a square sense-datum one might inwardly coin a predicate ‘S’ which applied equally to private *and* public instances of square-ness. The term ‘S’ would then be ‘private’ in sense (a) but not sense (b). Conversely, perhaps there could be ways to form a concept that only applied to private features of one’s own experience *other* than by attending to one’s own private realm—for example, perhaps such a concept could be ‘implanted’ via ingenious neuro-surgery. In which case the concept in question would be ‘private’ in sense (b) but not sense (a).

It is not immediately obvious whether (c) follows from either (a) or (b). As we are about to see, Wittgenstein apparently held that (c) does indeed follow from (b), though this entailment was repeatedly denied by Craig.

3. Wittgenstein versus Craig

In the *Investigations*, Wittgenstein first explicitly mentions the idea of a language for describing ‘private sensations’ in the famous §243:

But could we also imagine a language in which a person could write down or give vocal expression to his inner experiences – his feelings, moods, and the rest – for his private use? – Well, can't we do so in our ordinary language? – But that is not what I mean. The individual words of this language are to refer to what can only be known to the person speaking; to his immediate private sensations. So another person cannot understand the language. (PI §243)

This passage is clearly committed to the following claim:

A language whose predicates only describe features of an ‘inner’ private domain—that is, whose meaning is solely descriptive of private properties—could only be understood by the subject in question. It would be an *incommunicative* language.

In a series of papers, Craig (1982; 1986; 1997) argued against this Wittgensteinian thesis. Craig’s core point is simple to state:

Suppose I have introduced a term, ‘F’, to describe some phenomenal feature in my private domain and you have likewise introduced a term to describe your domain, ‘F*’. The private nature of these domains may prevent us from ever *knowing* whether ‘F’ and ‘F*’ are describing the same kind of private particular/scenario. But, asks Craig, why think that *knowledge* that the two predicates mean the same is required for mutual understanding of the two predicates? It is enough that we both *believe* that ‘F’ and ‘F*’ have the same meaning and that this belief is true—‘F’ and ‘F*’ do both, in fact, pick out the same kind of private phenomenal feature. I could then understand your claims involving ‘F*’ and you could

understand my claims involving ‘F’ (though we could not know that we were in fact comprehending each other). Moreover, Craig claims, it is plausible to think that people might quite generally subscribe to the following:

BURKE’S ASSUMPTION:⁴ ‘others have, in broadly similar external circumstances, broadly similar internal states.’ (Craig 1997: 131)

If this assumption was both true of our private phenomenal states and widely accepted, then we could all communicate about these private states even though they could be neither directly experienced nor known to obtain by anyone else—that is, even though these states satisfy MP and EP.

I think that Craig’s line of thought is important to consider for a number of reasons.

Firstly, the move in PI §243 from the claim: “The individual words of this language are to refer to what can only be known to the person speaking”, to the claim “so another person cannot understand the language” could certainly do with some supporting explanation, something that is not provided by Wittgenstein, at least not in the environs of section §243.

Secondly, Craig’s argument has been and still is influential. For example, Crispin Wright, citing Craig’s 1986 paper in a footnote, states:

... [even if] the phenomenological qualia associated by each of us with the word ‘pain’ are unknowable by anyone else, and are regarded as playing a constitutive role in our respective understandings of ‘pain’, there may nevertheless *be* mutual understanding if, as it happens, our qualia are appropriately similar ...

If your pain-qualia is inaccessible to me and constitutive of your understanding of ‘pain’, what (uncontroversially) follows is not that your understanding cannot coincide with mine but only that I cannot have adequate reason to think that it does. (Wright 2001: 223–224)

And in his recent guide to the *Investigations*, Arif Ahmed writes:

Edward Craig has shown quite convincingly (Craig 1997) that a language may be epistemically privileged without being idiolectic (so Wittgenstein’s final ‘so’ [in PI 243] is out of place). (Ahmed 2010: 109)

Thirdly, and most importantly, I think that to the extent that the model of a private domain remains a tempting way of thinking of the mind, appealing to something like Burke’s assumption will remain a tempting way to explain our ability to communicate about these private features. Thus, in a recent paper, Brie Gertler endorses the idea that the meaning of one’s own phenomenal concepts can only be “conveyed” by putting one’s audience into the right kind of external circumstances so as to produce the right kind of internal private features in them:⁵

⁴ Craig chose the label ‘Burke’s Assumption’ in reference to the following passage from Edmund Burke: “We do and must suppose, that as the conformation of their organs are nearly, or altogether the same in all men, so the manner of perceiving external objects is in all men the same, or with little difference” (Burke 1757/1990: 13).

⁵ And Wittgenstein himself was clearly aware of this sort of idea:

Suppose that in my judgement *this is what it feels like to feel a pinch*, the phenomenal quality *pinching* supplies the content expressed by ‘*this*’ This means that no one can think that thought unless *pinching* is instantiated in his or her experience.

... Since understanding *this is what it’s like to feel a pinch* requires an appropriate relation to a *pinching* experience, this information cannot be conveyed in the usual way, by uttering statements that express it. Stalnaker claims that this consequence – the incommunicability of phenomenal information – is unpalatable

... But even if phenomenal information is incommunicable in Stalnaker’s sense it can nonetheless be conveyed. If you have never had a pinching experience, conveying phenomenal information about such experiences requires introducing the appropriate experience into your “internal world”—perhaps by pinching you. (Gertler 2012: 121–122)

Gertler’s suggestion then is that one can avoid ‘the incommunicability of phenomenal information’ by trading on something like ‘Burke’s Assumption’. Proper evaluation of the “private domain” model of consciousness should assess whether Burke’s assumption provides an explanation for our ability to communicate about our conscious experiences. After all, if it could be established that judgements about a private domain would be incommunicable to others, that would surely be a result of quite general interest to philosophers of mind, not just Wittgenstein scholars.

The rest of this paper will be spent criticising the very natural line of thought that Craig repeatedly gave expression to. But before proceeding I should make clear that the problems I will be raising specifically concern the sort of metaphysical model of a private phenomenal domain which satisfies (MP)—and hence also (EP).⁶ Given this specific metaphysical model for experiences it is problematic, I will argue, whether there can be the required sort of fact in virtue of which ‘F’ and ‘F*’ speak of the same kind of private phenomenal features. But I have no problem with the commonsensical idea that our conscious experiences, where these are no longer being conceived of according to the model of a metaphysically private domain, are generally similar in similar circumstances—indeed, I am extremely confident that this is true!

4. Phenomenal Similarity and Inter-Domain Natural Kinds

A natural thought, for anyone sympathetic to the notion of metaphysically private phenomenal features, is that token features from two (or more) different subjects’ private domains could be similar independently of anyone’s dispositions to *judge* or *classify* these features as similar, or to introduce predicates to apply to them. This would effectively be to maintain that the private items and features that occur in separate private domains bear *natural* similarities to one another, or fall into something like

I want to describe a feeling to someone, and I tell him “Do *this*, and then you’ll get it,” and I hold my arm or my head in a particular position. Now is that a description of a feeling? and when shall I say that he has understood what feeling I meant? – He will have to give a *further* description of the feeling afterwards. And what kind of description must it be? (*Philosophical Investigations*, Part II, Section viii, 185–186)

⁶ Again, we are setting aside the possibility of a creature that can directly *experience* the phenomenal features of my stream of consciousness but for whom it is impossible to thereby gain *knowledge* of their phenomenal natures.

shared natural kinds. *Were* there such inter-domain similarities, or natural kinds, then it could be that my predicate ‘F’ and your predicate ‘F*’ have both latched onto the same (sparse) property, a property that occurs in both mine and your private domains. There would then be a fact of the matter as to whether ‘F’ and ‘F*’ mean the same, which does not depend on how I would apply/withhold my term ‘F’ to things in your domain were I to experience them, nor on how you would apply/withhold your term ‘F*’ to features in my domain were you to experience them.

Let’s grant, for present purposes, that *if* the notion of inter-domain natural similarities or natural kinds could be made plausible, we could then tell a plausible story, perhaps appealing to something like the ‘reference magnetism’ of natural kinds, as to how our respective privately introduced predicates could both have the same meaning. However, I will argue in this section that the required notion of inter-domain natural similarities amongst private features is dubiously coherent. I will consider three different ways of understanding inter-domain natural kinds and suggest that none of them are plausible.

The first and most obvious way of thinking about natural similarities across different subjects’ private phenomenal domains is simply that: the private phenomenal items from different domains *just do* have naturally similar phenomenal characters. That is, the way that *this* feature in *my* domain feels/appears *to me* and the way that *that* feature in *your* domain feels/appears *to you* just are naturally *phenomenally* similar. However, I think that under closer scrutiny this sort of appeal to natural similarities between the subjective *feels* or *appearances* of private items/features proves difficult to make sense of, given that we are operating with a model of metaphysically private domains.

The crucial point here is that appearance properties and phenomenal properties are always essentially tied to some specific (type of) subjective viewpoint. Nothing ever looks, feels or appears a certain way *simpliciter*; any item can only have some specific appearance/feel *relative to* some particular (type of) subjective viewpoint. The petals of this flower might look one way to humans, a different way to insects, yet another way to pigeons, etc. But they don’t look any particular way at all *simpliciter*, or from no particular point of view. Nor, of course, do they look any way at all to a subject who entirely lacks vision—that is, to a subjective viewpoint or modality that cannot ‘take in’ this item/feature. Likewise, my painful sensation feels a certain way *to me*, my after-image looks a certain way *to me*; but these token phenomenal features do not feel/appear any way at all *to you*, as your conscious point of view simply cannot include them. More generally, it seems that insofar as we understand phenomenal properties via the phrase ‘what it’s like’, phenomenal property instances must be essentially tied to some subject’s conscious point of view—for we can apparently make no sense of a phenomenal property instance forming part of ‘what it’s like’ *for nobody*.

Given this essential connection of appearance/phenomenal properties to a subjective perspective or sense-modality, two items can only be *similar* in appearance or feel relative to some particular subjective perspective or perspectives. Again, these two flowers may look similar to each other *to humans*, though they look very different from each other *to pigeons*. Or, they might both *look* identical (to me), yet they might each *smell* very differently (to me). But a claim that two items or features look or feel the same way relative to no particular conscious perspective or modality would be just as ill-defined as a claim that a single item has some specific appearance or feel relative to no particular conscious perspective or modality. And when it comes to two phenomenal items or features from two different private domains, there is, by hypothesis, no possibility that both items could ever ‘show up’ within a single subjective point of view—that is, within a single subject’s conscious awareness. Thus it

seems that any claim that my private item and your private item appear/feel phenomenally similar will not be a well-defined phenomenal comparison, for there is no possible common appearance space that both items could appear in.

At this point a comparison with spatial properties might be helpful. John Divers (2014) has recently raised the following problem for Lewisian realism about possible worlds. If the truth of our modal claims is supposed to be explained by the existence of counterparts in distinct possible worlds, and those possible worlds are *not spatio-temporally related* to each other (i.e., spatio-temporal relations only obtain *within* individual possible worlds), then we face a problem with counter-factual claims like the following:

Usain might have been taller than he actually is.⁷

For the truth of such a claim is supposed to require the existence of a counterpart of Usain in some other possible world who is ‘taller than’ the actual Usain. But Usain in the actual world and Usain’s counterpart in some other possible world do not stand in any spatial relation; *a fortiori* they do not stand in the ‘taller than’ relation. In the absence of any possibility of actual-Usain and counterpart-Usain having their heights defined relative to a common spatio-temporal framework, we cannot make sense of claims that their respective heights are ‘the same’. Not only is it impossible for actual-Usain and counterpart-Usain to inhabit each other’s spatio-temporal domains, it is also impossible for any other item to function as a benchmark or ruler by which to meaningfully compare their heights—as it is metaphysically impossible for any item to figure in more than one Lewisian possible world.

Likewise, I am suggesting, assuming the metaphysical impossibility of my experiencing your private phenomenal items, or vice-versa, there is a lack of any common ‘appearance-space’ in which two items from distinct domains could have their subjective feels/appearances meaningfully compared as being phenomenally similar or different. Not only is it impossible for either of the two items in question to appear in the other’s private domain, it is also impossible for any other possible item/feature to serve as a phenomenal benchmark or standard in common to both private domains—as it is impossible for any item/feature to appear in more than one private phenomenal domain.

In being essentially tied to a conscious perspective or modality, appearance/phenomenal properties are thus analogous to spatial and temporal properties—a point that I understand Wittgenstein to be making in section §350 of the *Investigations*:

It is as if I were to say, ‘You surely know what “It’s 5 o’clock here” means; so you also know what “It’s 5 o’clock on the sun” means. It means simply that it is just the same time there as it is here when it is 5 o’clock.’ – The explanation by means of *sameness* does not work here. For I know well enough that one can call 5 o’clock here and 5 o’clock there ‘the same time’, but I do not know in what cases one is to speak of its being the same time here and there.

In exactly the same way, it is no explanation to say: the supposition that he has a pain is simply the supposition that he has the same as I. (PI §350)

⁷ In fact, Divers argues that this initial formulation of the problem still allows the Lewisian counterpart theorist some wiggle room, but that this is eliminated with the claim, ‘It is true of the tallest actual thing that it might have been taller’. However, as these refinements are not relevant for the present comparison, I have stuck with Divers’s initial example for ease of presentation.

Jennan Ismael (2007), quoting this same passage from Wittgenstein, makes a comparison with the notion of simultaneity in relativistic physics: there can be facts about temporal relations between two events *within a spatial frame of reference*. But when the two events are specified as belonging to different frames of reference, there *is no fact* as to one being earlier/later than the other. Likewise, given a model of experience which satisfies (MP), sameness or difference in the phenomenal appearance/feel of two private items/features may be well-defined relative to a particular subject's private domain. But if the experiential elements in question belong to two different metaphysically private viewpoints, then any comparison of phenomenal appearance will not be well-defined. Just as there is no such thing as simultaneity *simpliciter*, but only within a specific frame of reference, there would be no such thing as phenomenal similarity *simpliciter*, but only within a specific private viewpoint.⁸

As a last resort, a defender of the private domain model might try insisting that *even though* there is no metaphysical possibility of private items/features from two different subjects' domains figuring in a common appearance/phenomenal space, nevertheless the ways that the respective items appear/feel to the respective subjects *just do*, as a matter of brute metaphysical fact, count as 'the same'. But I think this would be a pretty desperate and ad hoc response. Part of our very grip on the notion of an item appearing or feeling some specific way, is that such appearance/phenomenal properties are always tied to a conscious perspective or modality. So the suggestion that there could be facts about two items having a certain appearance/feel in common, but where this is *not* relative to any possible common appearance space, is obscure at best. Compare: suppose the Lewisian realist about possible worlds were to insist that, despite the lack of any spatial relations between different possible worlds, it is just a brute metaphysical fact that o_1 in w_1 counts as "the same size" as o_2 in w_2 . This would clearly be a totally ad hoc manoeuvre—we would still lack any grip or insight as to *how* items from two totally unconnected spatio-temporal frameworks could count as 'the same size', we are told only that *they just do*. Or suppose that some rogue physicist were to try insisting that, as a matter of brute undetectable metaphysical fact, there is an absolute reference frame that defines absolute simultaneity for two events throughout the universe (and absolute rest and shape for objects, etc.) This extra posit would not be incoherent or inconsistent, but it would be totally gratuitous and ad hoc—no physicist would take it seriously, as there is simply no reason to think that there is any such extra fact, whose only role would be to save our pre-theoretic intuitions about simultaneity.

I have argued that metaphysical privacy would rule-out a well-defined relation of inter-subjective phenomenal similarity. As this claim is central to my case, let me now briefly make a number of clarifications:

(1) Firstly, I have denied the possibility of there being natural similarities between the *phenomenal appearances* of private items/features belonging to different subjects. This has been a point specifically about appearances—or 'what it's like' properties—which are always tied to some kind of conscious

⁸ Having presented this challenge to inter-subjective phenomenal comparisons, Ismael goes on to argue that they *can* in fact be meaningful in virtue of the possibility that two distinct subjects could undergo a *merging* of personal identity into a single subject, thus allowing for meaningful comparison of two sensations from, what were at the time anyway, two different subjects. I am not sure that I find this wholly convincing—it seems to make large assumptions about the nature of the new merged individual's 'memories' of the experiences of the two pre-merged subjects. But in any case, Ismael's strategy can be seen as saving inter-subjective comparisons of phenomenology by denying (MP)—allowing that two sensations from two different subjects can both be 'directly' experienced and compared in the *memory* of the post-merged subject.

viewpoint, perspective or sensory-modality. There is certainly not any *general* requirement that in order for two items to be naturally similar, in *any* respect, it must be metaphysically possible for a conscious subject to experience them both! But this possibility of featuring within a common ‘subjective space of appearances’ *is* an essential requirement for two items to be similar in appearance/feel.

(2) Secondly, it is perhaps worth emphasising that the issue here does not depend on tying phenomenal similarity in any way to what a conscious subject can in principle *know* or *notice* or *judge* or any other epistemic notion. The issue here is simply that items only ever appear/feel some specific way relative to some particular appearance space/phenomenal space. And so two items can only be phenomenally similar relative to some common appearance space/phenomenal space that they can both, in principle, appear in. This issue is prior to any questions about what may be possible for a subject to notice or judge about an item’s appearance/feel.⁹

Compare again the case of Lewisian possible worlds that do not stand in any spatial relation. The problem with comparing the sizes of items from different possible worlds is not due to the inability of anybody to *know* or *tell* whether the items in question are the same size. Rather, there simply *is no such fact* about their sizes being the same—the problem here is due to a metaphysical lack of the required kind of spatial relations, not the lack of any epistemic relation.

(3) Thirdly, I do not wish to dismiss as ill-defined *all* claims of the form, item *a*’s appearance/feel in subjective perspective P1 is (phenomenally) similar to the way that item *b* appears/feels in subjective perspective P2. Clearly there can be perfectly well-defined comparisons of this general form—for example, the way that *this* elliptical item looks viewed from head on is similar to the way that *that* circular item looks viewed from a 45° angle. But in this case both the items and the subjective viewpoints are *public*; it is metaphysically possible for the items to appear in more than one subjective perspective and it is metaphysically possible for more than one conscious subject to adopt/occupy the perspectives in question. Thus there can be a well-defined fact of the matter that the way *a* appears/feels in P1 is the same as the way *b* appears in P2, as there is still the possibility that both these appearances/feels could, in principle, figure in the consciousness of a single subject. And so they could, in principle, be phenomenally similar *for someone*—that is, there is a common phenomenal space in which these ways of appearing/feeling can, in principle, both figure.

Notice that for this sort of appearance comparison across different perspectives/modalities to be well-defined, all we require is just the bare metaphysical possibility that both perspectives or sensory modalities could be enjoyed by a conscious subject. So for example, it is presumably metaphysically possible for a single conscious subject to observe public object *a* via a normal human visual system and also to observe public object *b* via a normal pigeon visual system. (It is presumably a metaphysical possibility that a subject’s sensory apparatus could be radically altered/manipulated, etc.) Insofar as we treat the appearance of *a* relative to normal human vision and the appearance of *b* relative to normal pigeon vision as both in principle available to be experienced by a single subject (or indeed by indefinitely many subjects), then both of these appearances/phenomenal properties are, in principle, part of a common appearance/phenomenal space. And so there could be a fact as to their being

⁹ There may or may not be some kind of constitutive link between the phenomenal character of experience and a subject’s epistemic relation to that experience—I take no stand on that here. I am very grateful to an anonymous referee for this journal for pressing me to clarify the issues here.

phenomenally similar—for there is a possible conscious subject which they can be phenomenally similar *for*.

But when we move to the model of metaphysically private phenomenal items/features there is, by hypothesis, no possibility of any other subject occupying/enjoying my conscious perspective on my private items, nor of any other subject occupying your conscious perspective on your private items. Which is why the apparently tempting thought that we started with, that the way that *this* feature in *my* private domain feels/appears *to me* and the way that *that* feature in *your* domain feels/appears *to you* just are naturally phenomenally similar, turns out not to be a well-defined comparison.

(4) Fourthly, nothing I have said rules out that there can be *structural* similarities between two different private phenomenal domains. The relations of phenomenal similarity and dissimilarity that hold amongst items in one subject's private domain could be isomorphic to the relations of similarity and dissimilarity between items in another subject's private domain. But I take it that we do *not* think that the natures of phenomenal properties are exhausted by these kinds of purely relational/structural features. We do not think that any two appearance/phenomenal spaces which are structurally similar are thereby qualitatively identical. We think, I take it, of phenomenal properties as having intrinsic/categorical qualitative natures that go beyond their place in an abstract relational structure.

(5) Finally: I should just briefly note that the point I have been making here has some similarity to what Shoemaker (1981; 1984; 1996; 2006) has called the 'Frege-Schlick view',¹⁰ a position recently defended by Stalnaker (2003; 2006):

What I have called the Frege-Schlick view is the view that the relations of qualitative similarity, identity and difference [amongst experiences] are well defined only for the intrasubjective case ... (Shoemaker 2006: 18)

But notice that my claim has only been a conditional one: *if* experiences are conceived of on the model of a private domain, (MP), then inter-subjective comparisons of phenomenal similarity will not be well-defined.¹¹

¹⁰ Shoemaker chooses this label for the view due to some of Frege's remarks in 'The Thought' (1956) and some of Schlick's remarks in 'Positivism and Realism' (1948). However, as Shoemaker himself notes (see 2006: 18, Footnote 2), in the case of Frege it is not entirely clear whether he held that statements making inter-personal comparisons of 'sense-impressions' were strictly nonsensical or just unknowable.

¹¹ Shoemaker's 1984 paper, and also Stalnaker (2003; 2006), argue that a physicalist-functionalist account of qualitative similarity/difference between qualia would require accepting the 'Frege-Schlick view', whereas Shoemaker's 1996 and 2006 papers seek to reconcile a physicalist-functionalist theory with the 'standard view' of qualia as having intrinsic phenomenal natures which fix relations of qualitative similarity/difference between experiences of different subjects as well as those of a single subject. But even the later Shoemaker allows that the 'Frege-Schlick' view would be true *given a private, non-physical theory* of qualia:

It could be held – and this may have been Frege's view – that the qualitative character of experiences is determined by what I will call 'private properties', properties that can be instantiated only in the experiences of a single person... This would imply, of course, that experiences of different persons are qualitatively different; but they would not be different in the way experiences of red are different from experiences of green. Each person's qualia would be 'alien' relative to every other person's qualia.

But the private properties version of the Frege-Schlick view seems to be incompatible with physicalism. (Shoemaker 2006: 19–20)

I have argued that the model of a private domain does not allow us to make sense of private items from different domains having naturally similar *phenomenal appearance* properties. But might there be some *other* way of understanding how private items from different domains could be naturally similar? For example, might private items/features from different domains have a common *non-phenomenal* essence, which our predicates ‘F’ and ‘F*’ might—though unknowably so—might refer to? I will now consider two possible ways of understanding this idea and argue that neither of them work.

Let’s briefly rehearse a familiar sort of story about reference to natural kinds:

On Planet A people use the term “water” to apply to stuff that is: transparent, liquid, drinkable.

On Planet B people use “twater” to apply to stuff that is: white, solid, freezing.

Suppose further that were A-dwellers to go to Planet B they would *not* apply their term “water” to the white, hard, cold stuff there. Nor would B-dwellers apply “twater” to the transparent, liquid stuff were they to encounter it on A. But it then turns out, *on further investigation*, that these terms are, unbeknownst to the two populations, picking out the same stuff. That is, despite the different senses, and non-overlapping domains of application, the two terms turn out to be referring to the same kind of substance.

However, notice that this familiar sort of story appeals to the idea that the item/substance/kind of thing that is referred to by both terms, has a further, ‘hidden structure’ (Putnam 1975) which can be investigated and discovered, beyond the manifest, recognisable features which formed the initial senses or definitions for the terms. And it is, I suggest, very difficult to see how this could occur when we are dealing with private domains of subjective experiences.

What sense can we make of the notion of *investigating* a private phenomenal item (e.g., a sensation, sense-datum or quale) and *discovering* that in addition to its manifest phenomenal appearance/feel to me it has some further hidden private properties? For a start, the very idea of re-identifying private particulars on different occasions is, so far, undefined in a private domain. And so the usual sorts of familiar notions we associate with investigation and discovery—such as measurement, cause and effect and so on—will at least be problematic to make sense of within a private domain. But the main problem here though, I think, is simply that when it comes to a *phenomenal realm*—something that is supposed to be constituting the manifest conscious nature of our experiences—it is just very hard to make sense of ‘hidden structure’, or anything like the traditional distinction between nominal essence versus real essence¹² Anything that is hidden, that is not manifest or present to mind in having the experience, cannot be part of ‘what it’s like’ to have the experience—that is, it is simply not phenomenal. And so if a privately-introduced predicate is supposed to be picking out such purported ‘hidden structure’, it would apparently be picking out something non-phenomenal—it would not be characterising the conscious nature of experience after all.

¹² This point is familiar from the literature on phenomenal properties and physicalism—see, e.g., Kripke (1980), Chalmers (1996).

Now at this point one might be tempted by the thought that if phenomenal experiences are *identical* to some physical structures/processes in the brain, they could then have a further investigable essence, as well as their manifest phenomenal nature—for they would also have a neuro-chemical side to their nature. (Perhaps a better way to phrase this is, their manifest phenomenal nature would turn out to be richer than pure phenomenology, for it would be literally identical to some bit of the brain.) If your private phenomenal features and my private phenomenal features were both of the same natural neurological kind, then perhaps this natural similarity could be what underwrites the synonymy of our respective terms ‘F’ and ‘F*’.

However, I think it is very hard to understand how phenomenal features/properties could be *identical* to neural properties of the brain—an object in the shared physical environment—and yet still satisfy the conditions for privacy stated in MP and EP (Section 2 above).

Whilst it is perhaps not transparently contradictory to hold that a phenomenal feature of experience satisfies MP and yet is also identical to something in the shared physical environment, I think that there is very considerable theoretical pressure leading one from accepting that something satisfies MP to accepting that it is distinct from anything in the shared physical environment. For it seems that however one thinks of perceptual experience, items and features in the shared physical environment will *not* be the sorts of things that can only, with metaphysical necessity, form part of the experience of a *single* subject. Either one allows, for example, as Naïve-Realists do, that items/features in the shared physical environment can be constitutive phenomenal elements in perceptual experience—and thus can, in principle, feature in *more than one* subject’s stream of consciousness. Or one denies, for example, as ‘indirect-realists’ or sense-data theorists do, that bits of the shared physical environment can be constitutive phenomenal elements in experience—so they can constitutively feature in *nobody’s* stream of consciousness. Either way then, it seems that if something meets the metaphysical privacy condition (MP), it would be distinct from anything in the physical environment—for physical things, in particular part and processes of the brain, can either constitutively show up in the experience of more than one subject or in nobody’s. Indeed it seems plausible that it is precisely this apparently special and unique conscious access we each have to our own sensations, mental imagery, innermost thoughts, etc., that provides one of the main motivations for dualism. As John Wisdom once remarked: “The peculiarity of the soul is not that it is visible to none but that it is visible to one” (Wisdom 1950: 195).

Now perhaps there are ways that one could try to resist this line of thought pushing from accepting that a feature of experience satisfies (MP) towards accepting that it is not part of the shared physical environment. For example, perhaps it could be argued that the unique accessibility of phenomenal items/features by a single subject can be explained in terms of *indexicality*—and that indexical contents, which can only be entertained by a specific subject, can still be accounted for in physical/neural terms.¹³ But rather than pursuing this any further, let’s turn instead to the epistemic notion of privacy (EP), as here I think the incompatibility with physicalism is even clearer. If one accepts that phenomenal properties are literally identical to neurological features of the brain, then it is very hard to see how they could possibly satisfy (EP). For if the phenomenal features in our respective streams of consciousness are literally *identical* to neural features in our respective brains, then it surely *is possible* to know whether yours are phenomenally similar to mine. Phenomenal similarities/dissimilarities now will be identical to various neural similarities/dissimilarities and these are surely just the sort of thing that more

¹³ See, e.g., Perry (2001). Many thanks to an anonymous referee for suggesting indexicality in particular as a possible line of response.

than one subject could come to know. (I'm assuming here, of course, that if physicalism is true it is also knowable.)

So then, the idea that two token phenomenal elements might, in addition to their manifest phenomenal natures, share a *further* hidden essence or structure does not seem to have any application if these phenomenal elements are supposed to belong to two different *private* domains. To be clear, I certainly do not wish to suggest that *in general* the notion of a natural kind or natural similarity must always be grounded in some kind of shared inner structure—nor even that any kind of scientific analysis of the kind can always be given. The point of the foregoing discussion of hidden essences or structures was to consider *alternative* ways of making sense of natural kinds/similarities between private phenomenal items from different subjects, *given that* we have already ruled out their being naturally similar simply in virtue of their manifest phenomenal characters.

Finally, rather than a shared hidden structure, another feature of natural kinds/properties that theorists have emphasised is that they can figure in laws or inductive generalisations (e.g., Mill 1884; Quine 1969). Blackburn (1984) considers the possibility that a single subject could notice patterns in the *order* in which her private sensation properties/events occur over time, and that fitting or failing to fit these patterns could provide a degree of independent confirmation or disconfirmation on the 'private linguist's' inclination to apply the term 'F' to a particular current private item/feature.¹⁴ Might Blackburn's line of thought suggest one more way for two different subjects to have naturally similar private phenomenal items? The idea would be that instead of (or perhaps in addition to) noticing that they both are inclined to use their terms 'F' and 'F*' in similar *external* circumstances, the two subjects compare the results of their respective introspective theorising concerning the order in which private sensations/events occur—that is, concerning the *private* circumstances in which they occur. And if it turns out that these internal patterns over time are indeed the same (or similar enough), this might count as the successful investigation and discovery of a common nature shared by items/features from two different private domains. Thus, so this suggestion would go, we could after all make some sense of the idea that the predicates 'F' and 'F*' have both latched onto the same inter-domain natural kind.

Recall (to repeat): we have already ruled out understanding inter-domain natural similarities in virtue of the manifest phenomenal character of the private items in question, and we have ruled out appealing to something like a hidden essence or structure. We must be clear then about what this present suggestion would have to amount to. It is not enough that such a discovery—of isomorphic internal orderings amongst the occurrence of private items/features—provides (mere) *evidence* that the two subjects have naturally similar private phenomenal items/features. Before there can be *mere evidence* that two private domains (items/features) are the same, there must be a *fact of the matter* as to whether the two private items really are intrinsically the same/different. If there is no prior fact as to their genuine intrinsic similarity, then the discovery that these private items/features crop up in isomorphic patterns over time within their respective domains cannot be (genuine, non-misleading) evidence of similarity.

The idea then has to be not just that the isomorphic internal orderings are *an indication* that the private items classified as 'F' and 'F*' are the same natural kind; rather, the idea must be that this internal

¹⁴ Wright (1986) is pessimistic about the prospects for such introspective theorising—though his argument for this conclusion is highly technical. See Bain (2004) for an optimistic reply to Wright.

ordering (at least partially) *determines/constitutes* that the private items are of the same kind. That is, the internal order in which private sensations/items occur would have to be *essential* to the sensation being the kind of sensation it is. We could then have a genuine fact of the matter as to whether two private items are the same kind, and so we could think of the two private predicates as having latched onto this pan-private natural kind—that is, ‘F’ and ‘F*’ could count as meaning the same independently of whether the subjects have dispositions to apply these terms to the same things.

I will not attempt to give a decisive argument showing that this line of thought *cannot* be correct. I will just state that it is surely *not* part of our normal conception of sensations (or phenomenal features generally) that the internal order or patterns in which they occur over time are essential to them being the kinds of sensations (or phenomenal features) that they are. On our normal conception of sensations (or so I claim), it is simply *the feel* of a painful sensation that makes it painful—its place in any larger pattern of sensations over time is *not* part of what makes it a *pain*. Were painful sensations to occur in some distinctive pattern of occurrence over time, this might provide further *evidence* that a particular judgement: “THIS is a pain” is correct or incorrect of a particular sensation. But one balks at the idea that such a wider pattern of occurrence could be (part of) what *determines* that a particular sensation counts as a pain.

5. Classificatory Disposition and Abundant Properties

Having ruled out appeals to natural phenomenal similarities, or natural kinds, that hold across different private domains, let’s turn now to considering whether we could instead explain the alleged sameness of meaning of ‘F’ and ‘F*’ by appeal to our linguistic or classificatory dispositions. To recap, we are allowing, for the sake of argument, that there are private conscious domains of private phenomenal features. And we are also allowing, for the sake of argument, that a subject could privately introduce a predicate, which then describes some way for *her own* private features to be. So I can introduce a term ‘F’ that determinately applies to some range of items or scenarios *in my own domain*.

David Lewis (1986) influentially distinguished between two conceptions of properties: a sparse conception versus an abundant conception.¹⁵ Now, given an abundant conception of properties there is no longer any problem making sense of the idea that private items from different domains can have properties in common. Any arbitrary set of private particulars defines the extension for some abundant property,¹⁶ so two items from two different private domains are bound to be similar in respect of some such abundant property. And so it seems we should be able to make some imaginative sense of a predicate corresponding to such an abundant property—that is, whose extension included private items/features from two different private domains (or from more than two domains).

What’s the problem then in suggesting, as Craig does, that two different subjects’ privately-introduced predicates might, albeit unknowably, have the same meaning? You might think: our two private domains both contain private elements that fall in the extension of some imaginary predicate—a

¹⁵ On a sparse conception, the only genuine properties are those whose extensions are determined by some ‘natural’ or ‘joint-carving’ similarity between the members. Whereas, on an abundant conception, there will be a property corresponding to any set of particulars—including sets comprising particulars from two or more different private domains.

¹⁶ I will speak, somewhat loosely, of both predicates and properties having extensions.

predicate whose extension “spans” across multiple domains—and our respective privately introduced terms ‘F’ and ‘F*’ could both, though perhaps unknowably, be synonymous with such a predicate.

Here’s the problem: in the context of an abundant conception of properties—where no property is more genuine/natural than any other—it seems that the meaning of such a novel predicate would somehow have to be determined by the subject’s dispositions to apply the term to various possible items, and not to others, *were* she to encounter them. Given that my term ‘F’ was first introduced via experience and ostension of some item/feature in my own private domain, something that is bound to possess indefinitely many abundant properties, we must appeal to how I *would* apply/withhold the term ‘F’ in response to various other possible items *were* I to experience them in order to fix a determinate extension for ‘F’. But according to the private domain model there *is no metaphysical possibility* of my experiencing your private items and so there can be no facts about how I would apply/withhold my term ‘F’ to items in your private domain were I to experience them. So there is just nothing in my linguistic/classificatory dispositions to use ‘F’ that could determine that it has an extension that includes items in your domain as well as items in my own domain. And likewise for your predicate ‘F*’—there is, by hypothesis, simply no such possibility as your experiencing my private items, so there can be no facts about which items in my domain you would correctly apply ‘F*’ to, were you to experience them.

So whilst we can perhaps make *some* sense of a predicate whose extension includes private scenarios from both private domains, this does not mean that there is any basis in virtue of which *my privately introduced predicate*, ‘F’, applies to some specific range of items in your domain as well as to items in my own domain. In other words, I am pressing the following question: what possible basis could there be in virtue of which my privately introduced predicate “F” has an extension that spans into your private domain, rather than spanning only a range of cases in my own domain? (And what possible basis could there be for thinking that ‘F’’s extension includes *this* specific range of private elements in your domain rather than some *other* range?) As there are no facts as to how I would be disposed to apply/withhold ‘F’ in response to experiencing items in your private domain (and given an abundant conception of properties) there just is no such basis.

There is some temptation, I think, to conceive of another subject’s private domain as somewhere that is just *very difficult* to experience for oneself, like a very distant planet. Even if it is physically or nomologically impossible for me to visit some distant planet—for example, perhaps it’s located outside of my light-cone—we can still make sense of a counterfactual scenario in which, *were* I find myself on this distant planet, I *would* apply some predicate ‘F’ to various items I encounter there. And likewise for the counterfactual claim that were the alien natives of this planet to visit Earth, they would apply their predicate ‘F*’ to various items they encounter here. Thus there can be a fact of the matter as to whether ‘F’ and ‘F*’ speak of the same way for things to be. But in the case of metaphysically private domains *there is no such possibility*. There is just no such thing as your experiencing my private features of experience, or vice-versa. We cannot appeal then to merely possible scenarios concerning which range of your private items I “would” apply my predicate to “were” I to experience them, for there are no such possibilities. Hence there seems to be no fact about our dispositions to use these terms that could determine that they really do apply to some specific range of private items *outside* our own respective domains.

In his 1986 paper “Privacy and Rule-Following”, Craig appeals to just the sort of counter-factual scenario that, according to (MP), is impossible:

... if I see you standing on the beach on a summer's day facing out to sea I shall assume, so the hypothesis goes, that you see the blue of the sky, hear the ripple of the waves, smell the salt in the air much like I do. And 'like I do' means: your sensations are very like mine; **could I be transposed into your consciousness** I wouldn't find things (at any rate in these respects) much different. (Craig 1986: 175, bold-type added)

At this date, Craig apparently found nothing troublesome in appealing to "transpositions of consciousness". But by the time of his 1997 article "Meaning and Privacy", Craig seems to have become aware of how difficult this alleged "possibility" is to make sense of:

... our question should be... how would others describe my EPI [epistemically private items] if they were well-placed to do so?

Now this might seem the right moment to say that others never are or could be well-placed to describe my EPI, precisely because they are private to me. So our conditional ("If others were well-placed . . .") isn't assessable, even in principle ... (Craig 1997: 134)

In response Craig now answers:

... if Burke's Assumption is true, then we frequently are well-placed to describe other people's EPI. Perhaps we aren't *as well placed as they are*, but why should that be necessary? (Craig 1997: 134)

In other words, Craig is now granting that there is no possibility that I might directly experience items in your private domain ("be transposed into your consciousness"). But he insists that I can still have some more indirect form of "access" to your domain. I can encounter your external, environmental circumstances and so I can gain some indirect awareness of how your domain is *via inference from Burke's Assumption*. To spell this out, I notice that my private domain is some way, which I have christened 'F', whenever I am in environmental circumstance C. I then observe that you are in C, and so I judge, on the basis of Burke's assumption, that your domain is also 'F'.

This later line of thought is not an adequate response either. Appealing to Burke's assumption here is question-begging given that precisely what is at issue is whether there can be any fact/basis in virtue of which my predicate 'F' *really does* apply to your private features as well as to my own. As we had occasion to note in the previous section when discussing natural kinds, before there can be (genuine, non-misleading) evidence indicating that our respective private domains are intrinsically similar, there must be a prior fact of the matter that they really are the same.¹⁷ That C is reliably correlated with *some or other* range of (possible) items/features in your domain does not show there is any basis for thinking that these items/features can be correctly described as being 'F'.

An analogy might be helpful here:

I have a gallery that only I can enter (with strong metaphysical necessity!). So only I can see the paintings inside. The paintings change from day to day—new exhibitions replace old exhibitions. I introduce the term 'fauve' to describe how some of the paintings *look*. I can also see the external façade

¹⁷ This has nothing especially to do with privacy, it is just a truism about the relationship between (genuine) evidence and the state of affairs it is evidence for.

of my gallery and I notice that a certain poster is displayed on the outside whenever there are fauve paintings inside. You also have a gallery that is closed to me, so I cannot see the pictures inside (in a strong metaphysical sense of cannot!). But I can see your gallery's façade. And one day I notice that your gallery's façade sometimes displays the same poster. I assume—call it the 'Curator's Assumption'¹⁸—that this poster is a reliable sign that the paintings in your gallery are also fauve.

But is there any fact in virtue of which the paintings that are inside your gallery when the poster is displayed on the outside really do count as 'fauve'? We cannot appeal to how these paintings *would* look to me *were* I to see the paintings in your gallery, nor to whether I *would* use the term 'fauve' in response to seeing them, for there just *is no* such possibility. Of course, there may be some or other range of paintings inside your gallery that happens to coincide with the occurrence of the poster on the outside. And so we can imagine some possible predicate, 'P', that applies to both the fauve paintings in my gallery, which correlate with the poster on my façade, and *whatever* paintings happen to co-occur with the poster on the façade of your gallery. But this gets us no closer to having some fact/basis in virtue of which that range of paintings in your gallery really do count as *fauve*. That is, it is no help in determining whether 'fauve' has the same extension as 'P'. Unless this is determined somehow, the 'Curator's Assumption' is not an assumption that even *might* be true. If there is no fact in the first place, in virtue of which some of the paintings inside your gallery really do count as fauve, then the presence of the poster on your gallery's façade could not possibly be a reliable external sign of fauvism within.

To briefly summarise, whereas in the previous section I argued that the metaphysical privacy of our respective phenomenal items ruled out their sharing *natural* (i.e., sparse) phenomenal similarities, the argument of this section has not been that metaphysical privacy would likewise rule out their having *abundant* properties in common.¹⁹ The point is rather that, given an abundant conception of properties, we'd need some basis in my classificatory dispositions to determine what my term 'F' means—that is, to determine *which* abundant property 'F' speaks of. My actual and counter-factual usage of 'F' might be able to determine some range of actual and possible features *within my own private domain* to which it correctly applies—or so we can allow for the sake of argument. But there seems to be no fact about my actual or counter-factual²⁰ usage of 'F' in virtue of which it would apply to some determinate range of private items/features in somebody else's private domain as well as to the determinate range of items in my own domain.

Finally, I should make clear that I do *not* wish to suggest that it is a simple matter to give a non-circular account of how a subject's dispositions to apply a concept determine the concept's meaning. The main

¹⁸ Burke's original statement mentioned the similar 'conformation' of our internal organs. If we wanted to reflect this in the present analogy, we could add in to the story that I know the internal architecture, the wiring and plumbing etc., of your gallery to be identical to my gallery.

¹⁹ To repeat, any pair of items are bound to share indefinitely many abundant properties, for there are indefinitely many sets which include them both as members.

²⁰ To clarify: on the standard Lewis/Stalnaker view, counterfactuals whose antecedent is necessarily false—i.e., 'counter-possibles'—come out as trivially true (see Lewis 1973; Stalnaker 1968). So, assuming that the antecedent is indeed a necessary falsehood, a claim such as, 'IF you were to experience my sensation THEN you would apply your term 'F' to it' would come out as true. But this is surely no help for Craig's position, as we could slot *any* claim about how you would use a private predicate into the consequent and the resulting counter-factual claim would also be true. E.g., 'IF you experience my sensation THEN you would *not* apply your term 'F' to it' would likewise turn out true, as would the result of substituting any other predicate in the place of 'F'. Conversely, we could substitute any necessary falsehood in the antecedent—e.g., 'If 2+2=5 then "F" and "F*" mean the same'. That this sort of claim turns out trivially true can hardly establish that 'F' and 'F*' really do mean the same!

lesson, I take it, of Kripke's (1982) celebrated interpretation of the 'rule-following' remarks in the *Investigations* was how hard it is to get facts about a subject's dispositions to use a term/concept, *specified in non-intentional* terms, to line up with the intuitive meaning of the term/concept. So the availability of counter-factual truths about how I would apply my predicate, 'F', to as-yet-unexperienced items within my own domain does not thereby give us anything like a complete, reductive explanation of why 'F' means what it means when applied within my own domain.

However, there remains a crucial difference regarding our ability to make sense of counter-factual claims about applications of 'F' to items within the subject's own domain compared to other private domains. In the 'Kripkensteinian' context of trying to find a *non-intentional* (reductive) account of a concept's meaning, it would clearly be *cheating* to appeal to the subject's dispositions to *correctly* apply the concept in response to experiencing various private items—for, so far, our only grip on what makes some applications correct and others incorrect is in terms of the concept's meaning. Nevertheless, we can still at least make sense of claims about the subject's dispositions to correctly apply 'F' in response to experiencing items within her own domain, even if such claims are no use in providing an explanation of 'F's meaning in non-intentional terms. For there can be counter-factual scenarios concerning how a subject would correctly apply 'F' to merely possible, as-yet-unexperienced items within her own domain, *were she to experience them*. But when it comes to private items from other domains, there can be no such counterfactual scenario, for there just is no metaphysical possibility of experiencing private items from other domains. In other words, even if we 'cheat' and help ourselves to the notion of a subject's dispositions to *correctly* apply her predicate in response to experience, we *still* have no basis in the subject's actual or counter-factual use of the concept 'F' which could determine some specific range of items in other people's private domains that 'F' would also apply to.

6. Conclusion

Wittgenstein claimed that the meaning of a subject's privately-introduced predicate for describing private features of her experience would be incommunicable. Craig's objection to this claim rested on the idea that two different private domains might both, though unknowably, be the same way, and that two different subjects' privately-introduced predicates might both, though unknowably, speak of this common way for private things to be. Against this line of thought, I have argued that it is very hard to make sense of any appeal to natural kinds or natural phenomenal similarities that obtain across different private domains—and so appealing to such notions cannot provide a basis for the required sameness of meaning. And I have also argued there would be nothing in a subject's possible dispositions to use her privately-introduced predicate that could give it an extension that included private items from another subject's private domain. So given a model of private phenomenal features, as defined by MP and EP, it is very hard to make sense of different subjects' privately introduced predicates having the same meaning. So it seems that the meaning of such predicates would after all be incommunicable.

This is not yet to provide anything like a full 'private language argument'—incommunicability does not amount to outright incoherence—but it is, I think it's fair to say, a pretty unattractive consequence for an ever-tempting model of consciousness to have. Now, given how very plausible many people find both MP and the idea that there can be phenomenal similarities across different subjects, one may thus have strong motivation to reject the argument of this paper. But I hope at least to have shown that the model of a private phenomenal domain faces serious issues concerning inter-subjective similarity and communicability, and that a defence of this model would require going well beyond anything that Craig's papers provide.

Acknowledgements

Versions of this paper were presented at King's College London, The North American Wittgenstein Society in Seattle, The University of Bergen and at NTNU's 'Vitforum' in Trondheim. I am very grateful to the audiences on all those occasions. This paper has been an unusually long time in gestation, so I hope I will be forgiven for having a rather long list of people to thank. Firstly, I am immensely grateful to Michael Campbell and Michael O'Sullivan who not only read and commented on an ancestral draft but who also provided the original occasion and encouragement to begin writing this paper. Secondly, many thanks to Anders Nes and to Tim Crane, both of whom read earlier versions of this paper and provided much needed criticisms. Thirdly, many thanks also to Kevin Cahill, Salma Saab Hassen, Ole Koksvik, David Papineau, Charles Travis and Margarita Valdez for helpful discussions that very substantially improved the paper. Fourthly, I am very grateful to two anonymous referees for this journal whose comments were exceedingly generous, fair and penetrating. Fifthly, I would like to say thanks to Jane Heal who (unwittingly) planted the seed for this paper by encouraging me to read Edward Craig's papers when I was just an undergraduate, *many* years ago. And finally, my most sentimental thanks go to my father, Stephen Raleigh, for interesting me in Wittgenstein, and in philosophy, in the first place.

References

- Ahmed, Arif (2010). *Wittgenstein's "Philosophical Investigations": A Reader's Guide*. Continuum Press.
- Ayer, Alfred Jules (1940). *The Foundations of Empirical Knowledge*. Macmillan.
- Ayer, Alfred Jules (1946). 'Symposium: Other Minds'. *Aristotelian Society Supplementary Volume*, 20, 188–197.
- Bain, David (2004). Private Languages and Private Theorists. *Philosophical Quarterly*, 54(216), 427–434.
- Bar-On, Dorit (2004). *Speaking My Mind: Expression and Self-Knowledge*. Oxford University Press.
- Blackburn, Simon (1984). The Individual Strikes Back. *Synthese*, 58(3), 281–301.
- Burke, Edmund (1757/1990). *A Philosophical Enquiry into the Origins of our Ideas of the Sublime and Beautiful*. Oxford University Press.
- Canfield, John (Ed.) (1986). *The Philosophy of Wittgenstein, Volume 9: The Private Language Argument*. Garland.
- Cassam, Quassim (2007). *The Possibility of Knowledge*. Oxford University Press.
- Chalmers, David (1996). *The Conscious Mind*. Oxford University Press.
- Child, William (2015). Wittgenstein, Phenomenal Concepts, and Knowing What It's Like. In Michael Campbell and Michael O'Sullivan (Eds.), *Wittgenstein and Perception* (84–103). Routledge.
- Craig, Edward John (1982). Meaning, Use and Privacy. *Mind*, 91(364), 541–564.

- Craig, Edward John (1986). Privacy and Rule-Following. In Jeremy Butterfield (Ed.), *Language, Mind and Logic* (169–186). Cambridge University Press.
- Craig, Edward John (1997). Meaning and Privacy. In Bob Hale and Crispin Wright (Eds.), *A Companion to the Philosophy of Language* (127–145). Blackwell.
- Dennett, Daniel (1991). *Consciousness Explained*, Little, Brown.
- Dennett, Daniel (2007). What RoboMary Knows. In Torin Alter and Sven Walter (Eds.), *Phenomenal Concepts and Phenomenal Knowledge* (15–31). Oxford University Press.
- Divers, John (2014). An Inconvenient Modal Truth. *Analysis*, 74(4), 575–577.
- Frege, Gottlob (1956). The Thought: A Logical Inquiry. *Mind*, 65(259), 289–311.
- Gertler, Brie (2012). Renewed Acquaintance. In Declan Smithies and Daniel Stoljar (Eds.), *Introspection and Consciousness* (93–128). Oxford University Press.
- Green, Mitchell (2010). Perceiving Emotions. *Aristotelian Society Supplementary Volume*, 84(1), 45–61.
- Ismael, Jennan (2007). *The Situated Self*. Oxford University Press.
- Jackson, Frank (1982). Epiphenomenal Qualia. *Philosophical Quarterly*, 32(127), 127–136.
- Jackson, Frank (1986). What Mary Didn't Know. *Journal of Philosophy*, 83(5), 291–295.
- Kripke, Saul (1980). *Naming and Necessity*. Harvard University Press.
- Kripke, Saul (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Blackwell.
- Lewis, David (1973). *Counterfactuals*. Blackwell.
- Lewis, David (1986). *On the Plurality of Worlds*. Blackwell.
- Lloyd Morgan, Conwy (1894). *An Introduction to Comparative Psychology*. Walter Scott.
- McDowell, John (1982). Criteria, Defeasibility, and Knowledge. *Proceedings of the British Academy*, 68, 455–479.
- McNeill, William (2012). On Seeing That Someone Is Angry. *European Journal of Philosophy*, 20(4), 575–597.
- Mill, John Stuart (1884). *A System of Logic*. Longman.
- Quine, Willard Van Orman (1969). *Ontological Relativity and Other Essays*. Columbia University Press.
- Papineau, David (2011). Phenomenal Concepts and the Private Language Argument. *American Philosophical Quarterly*, 48(2), 175–184.
- Perry, John (2001). *Knowledge, Possibility and Consciousness*. MIT Press.
- Putnam, Hilary (1975). The Meaning of 'Meaning'. *Minnesota Studies in the Philosophy of Science*, 7, 215–271.
- Rhees, Rush (1984). The Language of Sense Data and Private Experience - I: Notes of Wittgenstein's Lectures, 1936. *Philosophical Investigations*, 7(1), 1–45
- Rosenthal, David (1999). The Colors and Shapes of Visual Experiences. In Denis Fisette (Ed.), *Consciousness and Intentionality: Models and Modalities of Attribution* (95–118). Kluwer.
- Schlick, Moritz (1948). Positivism and Realism. *Synthese*, 7(1), 478–505.
- Shoemaker, Sydney (1981). The Inverted Spectrum. *The Journal of Philosophy*, 74(7), 357–381.
- Shoemaker, Sydney (1984). Postscript to "The Inverted Spectrum". In *Identity, Cause and Mind* (327–357). Cambridge University Press
- Shoemaker, Sydney (1996). Intrasubjective/Intersubjective. In *The First Person Perspective and Other Essays* (141–156). Cambridge University Press.
- Shoemaker, Sydney (2006). The Frege-Schlick View. In Judith Jarvis Thomson and Alex Byrne

- (Eds.), *Content and Modality* (18–33). Clarendon Press
- Stalnaker, Robert (1968). A Theory of Conditionals. In Nicholas Rescher (Ed.), *Studies in Logical Theory* (98–112). Blackwell.
- Stalnaker, Robert (2003). Comparing Qualia across Persons. In *Ways a World Might Be* (219–238). Clarendon Press.
- Stalnaker, Robert (2006). Responses. In Judith Jarvis Thomson and Alex Byrne (Eds.), *Content and Modality* (251–295). Clarendon Press
- Stout, Rowland (2010). Seeing the Anger in Someone’s Face. *Aristotelian Society Supplementary Volume 84*(1), 29–43.
- Strawson, Peter Frederick (1959). *Individuals*. Methuen.
- Tye, Michael (2007). Philosophical Problems of Consciousness. In Max Velmans and Susan Schneider (Eds.), *The Blackwell Companion to Consciousness* (23–35). Blackwell.
- Wisdom, John (1950). The Concept of Mind. *Proceedings of the Aristotelian Society*, 50, 189–204.
- Wittgenstein, Ludwig (1953). *Philosophical Investigations* (G.E.M. Anscombe and R. Rhees, Eds., G.E.M. Anscombe, Trans.). Blackwell.
- Wittgenstein, Ludwig (1980). *Wittgenstein's Lectures, Cambridge 1930-1932* (D. Lee, Ed.). Blackwell.
- Wright, Crispin (1986). Does Philosophical Investigations I. 258 Suggest a Cogent Argument against Private Language? In John McDowell and Philip Pettit (Eds.), *Subject, Thought, and Context* (209–266). Oxford University Press.
- Wright, Crispin (1989). Wittgenstein’s Later Philosophy of Mind: Sensation, Privacy and Intention. *Journal of Philosophy*, 86(11), 622–634
- Wright, Crispin (2001). *Rails to Infinity*. Harvard University Press.