

1. Introduction (~2500 words; 25 minutes)

The so-called “hard problem” of consciousness¹ stems from a sense of a vast (and sometimes dreaded) “explanatory gap”² between:

1. The private, **subjective**, inexorably first-person qualitative or phenomenological character of conscious experience; and
2. The public, **objective**, detached, third-person character of empirical science (from physics and neuroscience to cognitive science to scientific psychology).

These two realms or characteristics of experience seem so unutterably dissimilar that, as Nagel put it, “ ... ”

Reactions to the gap have been varied. Scientific triumphalists, including many neuroscientists (and some scientifically-minded philosophers), believe it is just a matter of time until consciousness is explained in ordinary scientific terms—on the model of neurotransmission, or digestion.³ Whether this enthusiasm stems from a genuine belief that the gap will be scientifically bridged, or whether, rather, these people believe that the first (subjective) side of the problem will simply be washed away, is not so clear. Typically, defenders of a purely scientific view of consciousness, on a traditional conception of science—perhaps including most of those who expect its explication to be neurological—are not unduly exercised by issues of first-person subjectivity and phenomenological experience.

Those who take subjectivity seriously tend to find the gap compelling, viewing its existence as posing a challenge of enormous proportion to intellectual inquiry itself. Some (the so-called “new mysterians,”⁴ along with many phenomenologists) believe that the gap is unbridgeable: that no purely naturalistic account will ever be able to provide a satisfactory understanding of the personal, subjective character of conscious experience.⁵ To them, the public or empirical realm, and the private or transcendental realm, are virtually incommensurable. Others are more optimistic about prospects for naturalisation, but still feel that bridging the gap—providing a naturalistically satisfactory scientific account of consciousness—will require radically new resources: changes in the character of scientific knowledge, strange new types of physical mechanism (such as quantum microtubules), or some entirely new non-physical, subjective element in the world, radically different in kind from anything

¹«Reference Chalmers»

²«Reference Levine»

³«Quote Searle»

⁴The phrase is Flanagan's «ref»; examples would be McGinn ...

⁵Phenomenologists are perhaps distinct in not supposing that there is any reason to want, let alone expect, such an account

heretofore scientifically imagined.⁶

In spite of their striking differences, however, those who take the gap seriously seem to agree on one point: that nothing in our current scientific picture of the world sheds the remotest light on the essentially private, qualitative, “awake” character of what it is like to be a conscious subject.⁷ Indeed, many believe that the scientific picture of the world is somehow intrinsically inimical (ontologically and/or epistemically) to an adequate naturalistic account of the “phenomenal” character of first-person conscious experience.⁸

In this paper, I argue to approximately the opposite conclusion, claiming:

1. That the familiar first-person, phenomenological character of conscious experience, far from being independent of physics, is virtually a direct consequence of (a rather abstract property) of physical law;
2. That no new resources, either physical or metaphysical, are needed to explain it (though our meta-theoretic understanding of the nature of scientific knowledge will come in for adjustment⁹); and
3. That “bridging the gap,” from physics to consciousness, is almost simple, if only one looks at the problem correctly.

This may sound like the brief of a rabid physical reductionist. It may even suggest that I fall in the camp of those triumphalists. But that is not the case. Not only will I take the phenomenological character of experience extremely seriously; what I will defend is nothing like a traditional materialism.¹⁰ Rather, my aim is to show that an especially **rigorous physicalism**, assuming no more than (an intuitive version of) elementary field theory, gives rise, from a third-person, empirical perspective—“from

⁶Examples of those who think that the character of scientific knowledge must change include Shepherd and Hut; quantum microtubules, Penrose; and pan-informationalism, Chalmers.

⁷«get Nagel quote»

⁸Use some of the following ¶s:

... Some people (≡“optimists”) view this gap as temporary—as merely reflecting an inadequacy in the progress of science, which should be repaired. Some (≡“pessimists”—Nagel? the new mysterians?) take the conflict to be intrinsic, leading them in turn to propose disparate measures: pan-psychism, neo-dualism, etc.). Some intermediate types think the gap can be crossed, but only with radical adjustments to our conception of the physical world (Penrose, Gregg).

⁹That is: I will argue for relatively radical revisions in how we understand (all) scientific knowledge, rather than claiming that the character of scientific knowledge itself has to be adjusted in order to understand consciousness. Consciousness, I will claim, can be understood in essentially the same scientific way as we have “always already” understood other material phenomena.

¹⁰Some would say ‘materialism,’ but for reasons that emerge in §■■■, I make a distinction between these two terms.

the outside,” as it were (or “sideways-on,” to use McDowell’s phrase¹¹)—to a picture of consciousness¹² that is far more like that developed from a first-person perspective (“from the inside”), as for example in ordinary contemplative reflection, or in the phenomenological tradition, than it is like the pictures of mind one normally encounters in analytic philosophy of mind.¹³

It may seem a stretch from quantum field theory to the originary self-reference of pre-reflexive intuition. But by the time we are done I hope to make that seem like not much of a stretch at all.

1a. Argument

The argument proceeds in two phases, preceded by a clarifying step.

The clarifying step involves separating the “problem of the gap” into distinct parts. On inspection, it turns out that the alleged first-person / third-person conflict is based on something of a use-mention confusion. What is third-person is our understanding of the physical world. What needs to be first-person¹⁴ is the physical world thereby understood. Sure enough, the physical world, on current scientific understanding, isn’t first-person—but it isn’t third-person, either. In fact it isn’t “personed,” at all. The first step in closing the gap, therefore, is to develop a naturalistic psychology that takes subjectivity seriously. That is, we need an objective (third-person) scientific account of how subjective (first-person) content can arise in, or supervene on, the (seemingly impersonal) physical world.

Taking that first step—naturalising consciousness—is the aim of Part I. My attempt to do that in turn consists of two steps.

In the first step (of Part I), contrary to popular opinion, I argue that the most basic form of physically realisable representational content is (an ersatz or proto version of) first-person content. (Ersatz) first-person content is closely followed by (an inchoate version of) second-person content. Simple content has this proto-first- (or proto-second-) person character because it is permeated with a profound, intrinsic, and ultimately ineliminable indexicality. In one of the paper’s main arguments, I claim that this (semantic) indexicality stems from what—extending the ordinary use of the term—I call the “**deixis**” of the physical world: that property of physical regularities that warrants their representation in differential equations.

From a naturalistic perspective, in other words, specific facts about the nature of the physical world entail that (at least inchoate versions of) first and second person

¹¹In §■ I address the issues that, it seems to me, motivate McDowell to eschew sideways-on accounts.

¹²And for that matter all of intentionality.

¹³Let alone neuroscience.

¹⁴I.e., what needs to sustain or “subvene” first-person content, in order for consciousness to be naturalised.

content are relatively “easy.” It is third-person content, often assumed to be the default, that turns out to be hard. It is a very substantial trick, given the nature of physical law, for an embodied physical agent to achieve anything like a genuinely third-person perspective on the world around it. Developing such a perspective, I will argue, requires no less than the attainment of objectivity—not only in order to count as third-person, but also in order to count as a perspective on the world.

The second step (of Part I) involves showing, or at least suggesting, how such objectivity can be achieved. To this end, I provide a sketch how a creature that starts out with simple, inchoate, highly indexical (proto-) first- and second-person representations or intentional orientation can extend and generalise them, so as to acquire a, stable, objective—and thereby not only third-person, but also objective first- and second-person—conception of or directedness towards its environment. In brief, this requires the agent to employ a dance-like set of acrobatic skills in order to “**deconvolve the deixis**”—thereby washing out (some of) the indexical character of its primitive or primordial representations, so as to acquire a (relatively stable) picture of a (relatively stable) embedding world.

With respect to consciousness, two features of this process are particularly important. The first has to do with the content of the third-person perspective thereby achieved. For fundamental reasons (stemming once again from the character of the physical world, and from the fact that an agent must be embodied, in order to refer), the intrinsic indexicality of semantic content can never be entirely eliminated. It is this intrinsic indexicality, and the non-conceptual character of the content thereby achieved, that subvenes the originary self-referentiality of pre-reflexive content. No matter how conceptual, context-independent, or detached, third-person understanding must always remain grounded in a creature's ontologically prior indexical, non-conceptual, first-person viewpoint. In part, this is because, in order to maintain and establish reference, it is necessary that that content remain tied to practical skills and bodily movements—which, because of the deictic character of physical embodiment, require indexicality. That is not to say that scientific knowledge can be idiosyncratically private. Because objectivity in general, and science in particular, are (as will be argued) public, social phenomena, the grounding of objectivity cannot only be first-person singular; it must also be first-person plural.

The second critical feature (of objectivity) is the fact that third-person content, even of this residually indexical sort, cannot be achieved in isolation. Not only will some indexical character be retained in all third-person content, in other words, but third-person skills can only arise if based on, or grounded in, and thereby integrated with, first- and second-person capacities. Genuinely objective semantic or intentional prowess—and in a sense this can be seen as the major epistemological result of the investigation—requires an integrated suite of “personed” perspectives—crucially involving the ability to move flexibly and plastically back and forth between

and among them.

Together, these results entail the conclusion of Part I. Attaining objectivity—that is, acquiring genuine epistemic skill, not a proto or ersatz or egocentric version, but authentic, objective, intentional directedness to the world:

1. Intrinsically requires a capacity for seamless integration of first, second, and third person perspectives, both singular and plural; and
2. Must be founded on aboriginally non-conceptual, self-originary, pre-reflexive first-person stances.

In sum: objectivity (i) is personed, and (ii) requires (phenomenological) consciousness. Both results, it will have be shown, are naturalistically predicted by the character of physical law.

In Part II, given this account of naturalised subjectivity—given, that is, an (object-level) account of how an integrated representational perspective can emerge from an underlying physical substrate—I draw back one (reflective) level, to address the (meta-theoretic) question of how to meld: (i) our first-person understanding of (primordially first-person) consciousness, and (ii) our (allegedly) third-person scientific understanding of the (impersonal) physical world.

To a first approximation, the strategy is the obvious one. As depicted in figure 1, the basic idea is to interpose, between these two seemingly incommensurate positions (labeled (a) and (c), in the figure), an intermediate step (b): an (allegedly) third-person account of (first-person) consciousness. But this, of course, is exactly what was developed in Part I—or anyway is part of what was developed in Part I. In rough and ready terms, in sum, the overall aim is to break the explanatory gap into two steps, and then to argue that each step is viable, thereby allowing the whole gap to be bridged.

In reality, the explanatory solution is slightly more complex. In particular, the strategy of Part II will be to lift the object-level ontological story, developed in Part I, up to the meta level, and thereby show that scientific knowledge is not, after all, as exclusively third-person as we have been led to believe. For if the objective level account (of Part I) is correct, it follows that scientific understanding, like all other knowledge, as something like a transcendental condition for (the very possibility of) its objectivity, must rest on a suite of integrated first-, second-, and third-person perspectives and skills. The actual task of Part II, therefore, is to show, not just “from the outside,” as originally suggested, but from integrated first-, second-, and third-person points of view (both singular and plural), how the (among other things, first-person) content described in the Part I of the argument meshes with the (among other things, third-person) description of that first-person content shown or exhibited in Part I.

Then, in a final decisive move—a step that may at first appear to involve an un-

warranted leap, but on reflection can be recognised not only as a possible but requisite first-person process of recognising that one's understanding has in part been first-person all along—one is able, through something like a process of first-person disquotation, to resolve the explanatory gap. Structurally,¹⁵ I call this step **imminent induction**, both to reflect its formal similarity with, but at the same to mark its substantive difference from, transcendental deduction—to which it can be regarded as something of a first-person analogue.

In sum, the strategy involves breaking the gap into two steps: **[this needs work]**

1. A first move—effected primarily at the object level—from an (allegedly) third-person scientific understanding of the non-personed physical world to a (still allegedly third-person) scientific understanding of a personed world (integrated first-, second-, and third-person, singular and plural) of psyches and content; and
2. A second move—effected primarily at the meta level, but achieved by lifting morals from the object-level account—from the (allegedly third-person) scientific understanding of a personed world, to a more integrated first-, second-, and third-person still scientific understanding of that (same) personed world.

By arguing not only that each step is viable, but showing how to actually make it, I claim that the whole gap can be—indeed is—bridged.

Ultimately, the aim of the whole argument—substantially aided en route by some elementary use/mention clarity—is to see how an integrated, skill-based, “perspectival” conception of scientific objectivity, founded on a primordially physical, non-conceptual, egocentric (first-person) stance towards the world, based in turn on a deictic underlying physical plenum, allows us to resolve the (apparent) tension between the third-person question of what conscious is, and the first-person question of what consciousness is like.

1b. Naturalising ontology

One final preparatory comment.

Naturalism, in philosophical contexts, is often taken to be the project of showing how intentional phenomena supervene on (can be derived from, emerge on top of, be rendered intelligible in terms of, etc.) the physical world—or more generally, the world as described in science. In these naturalising contexts, intentionality is paradigmatically understood to include meaning, interpretation, reference, truth, and other familiar semantic phenomena. When discussions turn to consciousness, the intentional—or anyway what is up for naturalistic explanation—is extended to em-

¹⁵What some people (though not I) would call “formally.”

brace the phenomenological character of qualitative experience.¹⁶ Even in these generalised contexts, however, naturalism is not normally interpreted as requiring an explication of ordinary, material objects—trees, rocks, tables, chairs, people (qua physical organisms), etc. Most scientists and philosophers write as if such entities—the ordinary material furniture of the world—could be assumed, as if “for free.”

It is a primary methodological tenet of this paper that this traditional stance is wrong. Material objects, I argue, and commonsense ontology more generally, are no more part of the scientific world view¹⁷ than truth, meaning, and content. What I am calling rigorous naturalism, therefore, requires embracing a methodological commitment to **naturalise ontology**. And not only ontology, but also, critically, **abstraction**—which (I claim) figures in ontological individuation.

One might think that consciousness, physics, and the nature of science are subject matters enough; that adding ontology to the mix would only complicate the story, or generalise it beyond recognition, undermining its chance of success. On the contrary, however, I believe that recognising the reconstruction of ontology as an inherent part of naturalisation is absolutely essential. It is not just a question of simplifying the story, but of adopting the only possible perspective from which the phenomenon of consciousness can be intelligibly viewed. More strongly: I believe that the traditional view—that intentionality does, but ontology does not, need naturalising—has been the major theoretical obstacle blocking efforts to bridge the explanatory gap.

There is a certain irony here. Legions of writers, from phenomenologists to religious thinkers to new mysterians, believe that naturalisation projects can only produce sterile, desiccated, impersonal accounts of the human condition. Perversely, as I aim to show, only by enforcing a far stricter naturalism than is traditional can we overcome what anti-naturalists regard as naturalism's failure, and thereby regain a genuinely personal understanding of the world.

— end of file —

¹⁶There are other ways to do it: one could limit intentionality to non-subjective phenomena, on an argument (which I would not believe) that meaning, interpretation, truth, etc, can be understood independent of any account of subjectivity, and then take naturalisation to be a project of rendering intelligible, in terms of normal scientific accounts, of both intentional and conscious phenomena. But that variant doesn't matter to my claim here: that ontology is not normally considered to be within the scope of a successful naturalising theory.

¹⁷That is: part of the world as described by scientific theory. Obviously, scientific theories are themselves supposed to be true. They also, I believe, traffic in material objects—but only as scientific idealisations, employed (like mathematics and models) for epistemological reasons, to facilitate calculation. See OO, chapter 5.