Consciousness Semanticism: A Precise Eliminativist Theory of Consciousness

Jacy Reese Anthis[1,2][0000-0002-4684-348X]

[1] University of Chicago, Chicago, IL 60637, USA
jacy@uchicago.edu
[2] Sentience Institute, New York, NY 10006, USA

Abstract. Many philosophers and scientists claim that there is a 'hard problem of consciousness', that qualia, phenomenology, or subjective experience cannot be fully understood with reductive methods of neuroscience and psychology, and that there is a fact of the matter as to 'what it is like' to be conscious and which entities are conscious (Chalmers, 1995). Eliminativism and related views such as illusionism argue against this; they claim that consciousness does not exist in the ways implied by everyday or scholarly language. However, this debate has largely consisted of each side jousting analogies and intuitions against the other. Both sides remain unconvinced. To break through this impasse, I present *consciousness semanticism*, a novel eliminativist theory that sidesteps analogy and intuition. Instead, it is based on a direct, formal argument drawing from the tension between the vague semantics in definitions of consciousness such as 'what it is like' to be an entity (Nagel, 1975) and the precise meaning implied by questions such as, 'Is this entity conscious?' I argue that semanticism naturally extends to erode realist notions of other philosophical concepts, such as morality and free will. Formal argumentation from precise semantics exposes these as pseudo-problems and removes their apparent mysteriousness and intractability.

Keywords: philosophy of mind, consciousness, hard problem of consciousness, artificial intelligence, neuroscience, eliminativism, materialism, semantics, semanticism

1 Introduction

This paper attempts to add a new perspective to the debate on 'What is consciousness?' I sidestep the conventional approaches in an effort to revitalize intellectual progress. Much of the debate on the fundamental nature of consciousness takes the form of intuition jousting, in which the different parties each report their own strong intuitions and joust them against each other in the form of intuition pumps (Dennett, 1980), gestures, thought experiments, poetic descriptions, and analogies. Consider, for example, the 'deflationary critiques' of Chalmers' argument for the 'hard problem of consciousness'. Chalmers asserts that there are some 'easy problems of consciousness' that could eventually be explained with reductive methods of scientific inquiry, but that 'the problem of experience' seems to remain even with a full behavioural and neuroscientific understanding of the

Please cite as: Anthis, Jacy Reese. 2022. "Consciousness Semanticism: A Precise Eliminativist Theory of Consciousness." In *Biologically Inspired Cognitive Architectures 2021*, edited by Valentin Klimov and David Kelley, 1032:20–41. Studies in Computational Intelligence. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-96993-6_3.

human brain (Chalmers, 1995). Critics argue that it clearly does not make sense to speak of a 'hard problem' for other concepts, such as life, perception, cuteness, light, and heat (Churchland, P. M., 1996; Churchland, P. S., 1996; Dennett, 1996). They argue that because there is no good reason to think consciousness is relevantly different from these concepts, there is no 'hard problem of consciousness'. Chalmers replies by insisting that phenomena related to consciousness 'cry out for explanation' beyond their material function while these other phenomena do not (Chalmers, 1997). This intuition jousting leads to an impasse that continues to encumber progress in the field of consciousness studies. In this paper, I take an alternative approach. I readily yield that the claim, 'Consciousness doesn't exist', is counterintuitive, and I try to critique that intuition with neither intuitions nor analogy. I make use of analogy to explain my theory, but unlike most arguments on the nature of consciousness, analogy and intuition do no work in the theory itself.

The current literature is also plagued by crosstalk. Consider Strawson's (2018) claim that eliminativism is 'the silliest claim ever made'. This seems to result from Strawson's focus on a strawman of consciousness eliminativism—portraying it as the claim that there is no mental phenomenon that we can directly gesture at, whatever it corresponds to after third-person analysis. This directly accessible phenomenon is not the sense of consciousness that I deny, and I have not seen it explicitly denied elsewhere. Indeed, in Dennett's (2018) reply and Strawson's (2019) subsequent reply, both acknowledge they are using different definitions of consciousness. The cornerstone of my approach is thus to clarify a distinction between 'consciousness-as-self-reference', to denote it as the mere contentless self-reference, and a different phenomenon, manifest in a statement such as, 'Is this entity conscious?', which I call 'consciousness-as-property' because it is the assignment of a property rather than an ostensive self-reference. Both of these phenomena appear to fall under definitions of 'phenomenal consciousness'[1] (Block, 1995), 'qualia' (Lewis, 1929), and other references to subjective mental phenomena, though I will discuss at length the issue of vagueness in these definitions. This distinction between the direct (arguably undeniable[2]) datum of one's consciousness-as-self-reference and the ambitious thesis of a real property of consciousness applicable across entities seems to be a neglected analytical perspective.[3]

The theory argued for in the following pages, *consciousness semanticism*, asserts that consciousness-as-property does not exist in the way commonly implied by everyday and scholarly discourse. This theory is closely related to other

---

[1] This article is oriented towards phenomenal consciousness, rather than Block's (1995) related concept of access consciousness, because the former seems to be the bigger challenge for intellectual progress. I thank Keith Frankish for helping me distinguish each of these notions.

[2] Debates on consciousness fallibilism are beyond the scope of this paper, but one argument is that the move from 'I directly observe my experiences' to 'My experiences exist' still hinges on logic, such as modus ponens from the conditional, 'If I directly observe something, it exists', and we cannot have absolute certain in even basic logic. Nonetheless, debates on eliminativism do not hinge on fallibilist or infallibilist claims.

[3] McGinn (1989) can be read as gesturing at such a distinction in his argument that the 'property P, instantiated by the brain, in virtue of which the brain is the basis of consciousness' is 'closed to perception', though the distinction is not developed in the sort of detail required here. I thank David Chalmers for raising this point. There is also a similar point made by Gloor (2020, Footnote 18).

formulations of eliminativism, type-A materialism,[4] and illusionism, but it is not intended to perfectly align with any particular one. The main difference is that I avoid argument by intuition and analogy—rather, the work in the argument is done by highlighting the semantic vagueness of the definitions of 'consciousness' used in everyday and scholarly discourse, arguing that this is inconsistent with the precision implied in everyday and scholarly questions we ask about consciousness, such as 'Is this computer program conscious?', and thus, consciousness does not exist as we like to assume. In other words, I argue that the way in which 'consciousness' is used by consciousness realists to imply that there is a 'hard problem' hinges on its vague definition, which is incoherent. Once we have established the fundamental nature of consciousness, there are many important conceptual and empirical research questions that naturally follow, such as the extent to which structures and processes such as 'global workspace' (Baars, 1988) and 'integrated information' (Tononi, 2008) obtain in humans, nonhuman animals, and artificial intelligences, and the extent to which these phenomena correlate with and cause self-reports of human and machine consciousness. These are deeply important questions, but the arguments herein suggest that there is no fact of the matter—no discovery that could be made—about which of these phenomena are the correct description of consciousness. If we add precision to the consciousness debates, as is more obviously necessary to evaluate other properties such as 'life' or 'brightness', then notions of the hard problem and consciousness-as-property evaporate, clearing the intellectual quagmire and making way for intellectual progress in a rigorous understanding of biological and artificial minds.

2 Terminology and Concepts

Semanticism is intended to be a precise, semantics-driven theory, so I begin with an extensive articulation of the terms and concepts involved in its argument before presenting the formal argument in the following section. When we force definitions to be precise, we must smooth out the scatterplot of semantic intuitions just as we would with statistical data points in a mathematical regression, retaining some valuable intuitions while discarding contrary ones to achieve a parsimonious description of reality. As argued by Wittgenstein (1922), we must select definitions that best 'picture' the world. My position here is especially vulnerable given how explicit I have made its substance, including those definitions, but I think that is how we will make intellectual progress and, perhaps optimistically, how we will achieve far greater agreement among philosophers on the nature of consciousness.[5]

---

[4] Eliminativism is not always materialist or physicalist. Eliminativists tend to think dualist explanations are incorrect, but it is possible that the best explanations of conscious-seeming mental phenomena will invoke phenomena outside physical reality as currently understood. Under semanticism, even if we were to find such explanations, we would still need to decide whether to categorize those phenomena as conscious, non-conscious, or indeterminate.

[5] It may be that even staunch defenders of consciousness' existence would be eliminativists under my view. I take this as confirmation that the argument is sound and that intellectual consensus is more tractable than commonly assumed, but with other approaches to philosophical inquiry, one can take this as rendering the eliminativist position trivial and thus uninteresting. I thank David Chalmers and Jake Browning for developing this point.

2.1 Eliminativism

The specific claim I will articulate and defend is that the property of consciousness implied in our everyday language and scholarly discourse does not exist. I take 'property' and 'exist' as the key terms herein. Semanticism is one example of the eliminativist view, defined by *The Stanford Encyclopedia of Philosophy* as, 'the radical claim that our ordinary, common-sense understanding of the mind is deeply wrong and that some or all of the mental states posited by common-sense do not actually exist and have no role to play in a mature science of the mind' (Ramsey, W., 2020). Eliminativism is a general view, in that it does not specify exactly what mental states do not exist; I call the specific view 'consciousness semanticism' because it is grounded in a criticism of the inconsistent semantics of consciousness in everyday and scholarly discourse.[6]

2.2 Property

Here, I adopt the conventional usage of the term 'property' to mean 'those entities that can be predicated of things or, in other words, attributed to them' (Orilia and Swoyer, 2020). I expect this usage to be uncontroversial or at least easily translated into different framework of definitions (e.g., some formulations of structuralism).

2.3 Existence

The definition of 'exist' is much more fraught. Depending on your definitional approach, the definition can vary based on context or other factors (see, for example, van Inwagen, 2005). What exactly does it mean for a property such as consciousness, life, or wetness to exist? There seems to be no consensus in the extant literature on how to approach this. Moreover, there is no established definition that is sufficiently precise to resolve debates about consciousness' existence.

How then do I approach establishing a definition? I take it that there is no right or wrong definition without preestablishing some criteria for our definitions.[7] In this case, I aim for the criteria of simplicity, precision, and approximation of the usage of the word in everyday language. Thus, I propose this working definition:

---

[6] Labeling philosophical concepts is challenging, given almost every plausible English word already has an established meaning, especially the most meaningful words. 'Semanticism' fortunately only has one significant established usage, as far as I can tell. According to Akiba (2015), 'Semanticism about vagueness (or the semantic theory of vagueness) holds that vagueness exists only in language and other forms of representation and not in the world itself'. While this internalist view is related to my argument, it is not isomorphic, and given a large majority of potential philosophical terms (e.g., 'relevant word' + 'ism') have already been used somewhere in the field, 'semanticism' still seems to be the best descriptor.

[7] As Putnam (1987) says, 'the logical primitives themselves, and in particular the notions of object and existence, have a multitude of different uses rather than one absolute "meaning".' I am not claiming that there is a univocal or unambiguously correct definition of existence. In fact, much of my criticism of the current consciousness debate can be read as a rejection of that claim, at least in the sense that for a definition to be correct, we need to specify criteria. My choice of meaning is thus only motivated as the best operationalization of the term that I know of for the purpose of resolving debate over the nature of consciousness. Without such operationalization, we could not make progress one way or the other.

> *Existence*: A property exists if and only if, given all relevant knowledge and power, we could categorize the vast majority of entities[8] in terms of whether and to what extent, if any, they possess that property.

While it is counterintuitive to be so specific about the definition of such a fundamental term, this definition corresponds with common-sense questions such as 'Is this computer program conscious?' or 'Is there something it is like to be a rock?'[9] By asking these questions, we act as if there is an objective, potentially discoverable answer as to whether any given entity in the universe has conscious-as-property. This is ubiquitous in everyday and scholarly discourse on consciousness, but the semanticist view entails that, even if we could observe everything in the universe with unlimited time and intellectual capacity to analyse those observations, we could never answer these questions or even assign coherent probabilities to possible answers given the common definitions of consciousness.[10]

Despite its appeal, this operationalization of consciousness will inevitably be contentious. The 'potential discoverability' definition seems to best fit with the criteria of simplicity, precision, and approximation of the usage of the word in everyday language. However, if one disagrees with this definitional approach and wants to keep the definition of 'existence' vague, that entails more than the nature of consciousness being a challenging or confusing topic, but that it is a fundamentally impossible debate to resolve. In other words, I cannot tell you if I am a consciousness eliminativist or not unless we start the discussion with a precise definition of what it means for consciousness to 'exist'. You might have a different definition you want to use for existence, such as that all possible properties exist or all properties we can intelligently discuss, even if we do so vaguely, exist. Again, I think there are good arguments for various operationalizations, and I think we should engage all of them, in order to avoid what Chalmers calls a 'verbal dispute' (Chalmers, 2011).

Under a different definition or an insistence to avoid precisification of 'existence', the rest of this paper may be unpersuasive. However, the question addressed by this paper can simply be read instead as 'Is consciousness potentially discoverable?' or 'Is consciousness vaguer than we make it out to be?'[11] The view outlined in this

---

[8] We can define existence more strongly, requiring the categorization of all entities. The strong and weak versions of existence both seem worth consideration to me, and the argument for consciousness semanticism works with both. I thank Keith Frankish for raising this point.

[9] Sometimes the referent of 'conscious' is not an entity per se, but a mental state itself, such as, 'Is anger conscious?' or 'Is the vestibular sense conscious?' I take these claims to be similarly vague in most cases. The exception is a deictic definition of consciousness, in which they may be true by definition, but as I discuss below, the one example does not constitute a definition that can be extended to the world at large.

[10] With definitions that are at least somewhat precise, such as those defining consciousness as 'brain or no brain' or 'integrated information', we can of course assign coherent probabilities in some cases (e.g. an animal with no brain or a computer with no integration has zero probability of consciousness). While these definitions are interesting to discuss, they have not been widely endorsed in the literature.

[11] I am not the first to suggest that consciousness is a vague concept and that this leads to problems in conventional approaches to understanding consciousness. For example, Papineau (2002) writes, 'My thesis will not be that there is anything vague about how it is for the octopus itself. Rather, the vagueness lies in our concepts, and in particular whether

paper may be called ‘consciousness incoherentism’. This would not change the significant practical implications. The map is not the territory (Korzybski, 1933).

2.4 Omniscience and Potential Discoverability

One could also argue that this definition is simply a restatement of the ‘hard problem’ itself, an impassable gap between our scientific knowledge of the physical world and the ineffable state of consciousness (Chalmers, 1995), or a restatement of the ‘problem of other minds’, that we cannot peer into the minds of other beings to know if they are conscious (Chalmers, 1996). Prima facie it seems that these claims also render the nature of consciousness impossible to discover.

My claim is different than these because I specify omniscience. That something exists does not mean we will ever in practice have the tools to discover it. Rather, it means that in principle, that something exists means that it could be discovered by some sufficiently powerful entity—powerful not just in observing everything, but with the intellectual capacity to analyse and fully understand all observations. An omniscient entity therefore would not have trouble overcoming the ‘hard problem’ or the ‘problem of other minds’. In other words, one can assert consciousness exists in terms of ‘potential discoverability’ while still maintaining that the ‘hard problem’ is still an impassable gap for mere humans lacking omniscience.

Semanticism implies that the ‘hard problem’ and ‘problem of other minds’ do not represent meaningful issues in consciousness studies, except perhaps as meta-meaningful facets of our imperfect scholarly understanding. On the far side of these proposed gaps—in the realm of subjective experience or other ineffable or impermeable properties—there is nothing that matches our common-sense definition of consciousness. Gaps that have nothing on the other side are merely borders of reality. In the language of Chalmers (1997), semanticism implies that all that needs explaining about consciousness is its functions and any other similarly effable features.

There are many different definitions of omniscience we can use here, but none have major implications. We can try to imagine a category of entities that exist but are not potentially discoverable. This seems intuitively reasonable, but I have not read a coherent operationalization of this possibility that jibes with common sense. These distinctions seem relatively arbitrary. Take a very large number, 3^^^^^^3 in up-arrow notation (Knuth, 1976) with however many exponents you need to make it so large that the decimal representation cannot be computed with all the computing power made with all the atoms in the universe. The decimal representation of this number seems undiscoverable, but does it exist? Is it real? It seems that the everyday and intuitive definitions of ‘exist’ and ‘real’ are unclear in this kind of edge case, so it becomes an arbitrary semantic choice. Moreover, nobody is claiming that consciousness exists but is hidden from us by computational limitations.

Another class of entities are those are too far away in time or space from us to know about given the laws of physics, even with all knowledge and power up to but not including breaking those laws. For example, ‘Is there currently a blue moon around any planet in the Andromeda galaxy?’ which is unknowable in the sense that it would take us millions of years to see that amount of light from Andromeda, even if we can build equipment sensitive enough to detect it. Similarly, assuming time travel is impossible, we cannot examine the past, such as, ‘Were any atoms in the

---

such phenomenal concepts as pain draw a precise enough boundary to decide whether octopuses lie inside or outside.’

White House today also in the eggs George Washington ate for breakfast on March 17, 1760?' which is unknowable in the sense that we cannot go back and check what Washington had for breakfast, and it currently does not seem like a computer could ever reverse-engineer the physical processes of the universe with this level of fidelity. This speaks to a broader observation that there is a wide gradient of possibilities when one specifies a property such as omniscience: Exactly which laws of the universe does omniscience allow you to break?[12]

Some logical positivists also wrestled with these possibilities, such as Blumberg and Feigl (1931), countering the claims of critics who say some scientific assertions are not verifiable. Blumberg and Feigl say those critics are confused about the meaning of 'possibility of verification', which assumes we can surpass the limitations of scientific instruments and laws of nature. In the present work on consciousness, it is sufficient to leave these as ambiguous edge cases because nobody is claiming that consciousness exists but is hidden from us by natural laws such as the light barrier or arrow of time.

Of course, some people use very loose definitions of 'exist' and 'real', such as if something is real if you are simply able to talk about it—this means that dragons and other imaginary creatures are real (Bowles, 2018)—or the even stronger definition of modal realism, that arguably views mere possibility as sufficient for realness (Lewis, 1986). With those definitions, I agree that consciousness is real and exists. But the most common scholarly and everyday usages of 'consciousness' seem to imply a much stronger sense of reality or existence. The discussion of whether consciousness exists also seems to evade debates of idealism, whether the physical world exists out-side of mental phenomena, since consciousness is very clearly a mental phenomenon (Russell, 1912).

2.5 Related Views

A number of philosophical views can be interpreted as claiming that consciousness does not exist in some fashion: consciousness eliminativism (Tomasik, 2015), eliminative materialism (Ramsey, W., 2020), type-A materialism (Chalmers, 2003), some definitions of materialism itself (e.g., Papineau, 2002), illusionism (Frankish, 2016), and the view of consciousness as a 'strange loop' (Hofstadter, 2007, Chapter 22). I do not claim that the view I propose here is perfectly aligned with any existing versions of eliminativism. Developing a novel view allows me to position my arguments with more freedom and clarity. The overlap with extant discussions will be discussed throughout the paper as relevant.

For example, while I am sympathetic to the broad gesture that consciousness (or at least qualia) is an illusion, or specific assertions that experiences do not have 'what it is like' properties, consciousness is analogous to a trick of stage magic, and so on, none of the available formulations of illusionism seem to capture my precise claim regarding imprecision and potential discoverability. In other words, I consider

[12] Another category is dualist or non-physical phenomena, which are commonly posed as answers to, 'What is consciousness?' As noted above, the semanticist argument does not rely on physicalism. If interactionist dualist phenomena exist or if non-interactionist dualist phenomena exist and we have dualist means of knowledge production, then they could be part of a precise definition of consciousness, and semanticism applies. If non-interactionist dualist phenomena exist and we are limited to physical means of knowledge production, then they could not be a part of a precise definition of consciousness, and semanticism applies.

myself an illusionist, but my claim is also more precise than that. Similarly, while I largely agree with previous discussions of eliminativism, I worry that the ontological claim that consciousness does not exist has been conflated with the pragmatic claim that we should eliminate the term ‘consciousness’ from consciousness studies, which is a difficult epistemic and sociological question about the best way to make scholarly progress as a field.

3 The Semanticism Argument

Now that terminology is established, the semanticism argument is brief and straightforward.

1. Consider the common definitions of the property of consciousness (e.g., ‘what it is like to be’ an entity) and the standard usage of the term (e.g., ‘Is this entity conscious?’).
2. Notice, on one hand, each common definition of ‘consciousness’ is imprecise.
3. Notice, on the other hand, standard usage of the term ‘consciousness’ implies precision.
4. Therefore, definitions and standard usage of consciousness are inconsistent.
5. Consider the definition of exist as proposed earlier: Existence of a property requires that, given all relevant knowledge and power, we could precisely categorize all entities in terms of whether and to what extent, if any, they possess that property.
6. Therefore, consciousness does not exist.

First, define consciousness using any of the available definitions: most commonly, ‘something it is like to be that organism’, ‘subjective experience’, or the homage to jazz, ‘If you got to ask, you ain't never gonna get to know’ (Nagel, 1974; Strawson, 2016). Also consider the common deictic and ostensive definitions, which simply gesture at consciousness by referring to personal examples, such as saying it is the common feature between you seeing the colour red, imagining the shape of a triangle, and feeling the emotion of joy. Also consider standard usage of ‘consciousness’ in scholarly and everyday discourse, such as the question, ‘Is this computer program conscious?’

Second, notice that all of these definitions are imprecise. They do not clearly differentiate all possible entities into conscious or non-conscious, even if we know all there is to know about such beings. It might seem clear via introspection that there is something that it is like for you to see red and feel joy. In fact, that is loaded into the deictic definition. But one example does not constitute a definition that can be extended to the world at large. Consider not just extensions to the mental lives of other beings (e.g., ‘Is this computer program conscious?’), but also borderline cases in your own mental life, such as whether the vestibular sense (i.e., sensing the spatial orientation of your body) is a conscious experience.

Third, notice that scholarly and everyday usages of the term ‘consciousness’, such as the question, ‘Is this computer program conscious?’ imply precision, that there is in principle a correct answer as to whether any particular being is conscious.

Fourth, because all standard definitions of consciousness are imprecise (again, except for precision regarding one individual if using the deictic definition), yet common usage implies precision across individuals, there is inconsistency. You could not, even as an omniscient and omnipotent being, categorize any entities (except yourself, if you use the deictic definition) in terms of whether and to what

extent they are conscious. Consider a hypothetical example where humanity builds a sophisticated AI and understands every single detail of its inner workings. In this case, what exactly would you check for to determine if the being is conscious or non-conscious? You would have no reasonable basis for claiming it is or is not. This implies there is some issue with common definitions and standard usage.

Fifth, as discussed at length above, define 'existence' such that: 'A property exists if and only if, given all relevant knowledge and power, we could categorize all entities in terms of whether and to what extent, if any, they possess that property'. This allows us to test whether a property exists.

Sixth, notice that an imprecisely defined property cannot be used to categorize all entities as having to a full extent, having to some extent, or not having that property. As such, the property of consciousness fails the test for the existence.[13]

4 Objections

As with previous literature on the subject, it seems that the best way to respond to objections and thus illustrate my view is to analogize 'consciousness' to properties that lack the intellectual and intuitive baggage burdening 'consciousness' itself. But, importantly, the analogies are explanatory and not part of the formal argument.

> *Since I am familiar with, and in fact certain of, my first-person experience, I can reasonably guess that other entities like me have their own first-person experiences.*

This is a useful heuristic for most contexts, but not for consciousness. I contend that the following situation is analogous to the case of consciousness: I show you a few distinct shapes on a piece of paper: circles, triangles, squiggly lines, etc. Then I point to one shape and say, 'That is a baltad', which is a word I just made up. Can you now categorize all the shapes on the paper as baltads or non-baltads? Of course not. Since all we know about baltads is that one example, baltadness does not exist. Even if I give you an arbitrarily large number of examples, you still would not be able to take a new shape and tell me whether it is a baltad or not. Of course, you could guess the definition I have in mind based on human psychology, such as using a machine learning classifier and coding the shapes as a matrix of pixels, but if I have no hidden definition behind the curtain—no definition I came up with for baltad that I have not yet shared—then there is no empirical strategy that can estimate it because it does not exist. That is the case with the purported property of consciousness, where proponents do not seem to endorse any behind-the-curtain definition. This is one benefit of defining existence the way I do in this paper, since we can say that the property 'baltadness' does not exist despite knowing for a fact that a single 'baltad' exists.

[13] Currently this semanticism argument seems conclusive, but to consider another argument, the existence of consciousness would make the world more complicated than a world without it because it adds an extra feature, which adds the weight of parsimony in favor of semanticism relative to most non-eliminativist views. Absent the semanticism argument, since parsimony is not conclusive (i.e., it is a heuristic, not a proof), it would then need to be weighed, rather subjectively, against the intuition that consciousness exists as a property. Moreover, some realist views would consider consciousness realism to be more parsimonious.

This raises a fatal issue for questions such as, ‘Is this computer program conscious?’ in the same way we would struggle to answer, ‘Is a virus alive?’ (about the property of life) or ‘Is Mount Davidson, at 282 meters above sea level, a mountain?’ (about the property of mountainhood). We cannot hope to find an answer, or even give a probability of an answer, to these questions without creating a more exact definition of the term.

‘Life’ is a particularly interesting comparison to ‘consciousness’ because it has endured some of the same discursive challenges. It is useful to talk about living beings versus non-living beings and get a sense of the properties that match our intuitions regarding what is alive and what is not. But do we need academic journals filled with papers asking what is alive and what is not? Do we assign probabilities to whether each different entity is alive? Do we need to spend precious research resources trying to figure out whether a virus, which cannot reproduce on its own but can do so within a living host, is alive? No, we simply accept that it depends on how we define life, and then move on to more important research questions such as, how exactly do viruses work? As philosopher Carol Cleland said about her seminal 2002 paper on life’s definition, ‘I argue that it is a mistake to try to define 'life'. Such efforts reflect fundamental misunderstandings about the nature and power of definitions’ (Cleland and Chyba, 2002; NASA, 2003). Personally, I expect the same evolution of thought will happen to consciousness research, though it might be much harder to reach that resolution.

Scientists and philosophers regularly ask meaningless questions about consciousness, assuming that they have the potentially discoverable answers that normal empirical questions do, such as, ‘Will this coin land on heads?’ In that case, it is very clear to everyone involved what it would mean for the coin to land on heads. It might be very clear to you that you are conscious—if we are using a deictic definition of self-reference—but that does not speak to whether a virus, a computer program, or even another human is conscious. In other words, there will never be a ‘conscious-o-meter’ as some have imagined, even if we have a perfect understanding of what is happening inside every mind (The Brain Dialogue, 2016).

This ties into another objection:

> *But if we accept the semanticist argument against consciousness realism, and we’re extending this argument to baltads, life, mountainhood, and coin-flipping, doesn’t that imply that the vague everyday properties we refer to like wetness or brightness don’t exist in this way either?*

Yes, it does, and it is a very important realisation about our everyday language. We do not spend decades of research on questions like, ‘Is this lightbulb bright?’ or ‘Is my raincoat wet?’ or ‘Is my uncle bald?’ because it is so clear that the answers simply hinge on how we choose to define the terms. We instead use these terms only in a much more limited sense of existence, a cluster of things in the multidimensional definition space; we never presume a fact of the matter. For most terms, the ambiguity is an entirely reasonable part of discourse. For example, we all know the sun is very bright and an ember is only slightly bright, and we can ask perfectly fruitful questions such as, ‘Is this lamp bright?’, even if we have not

precisely laid out lumens as a unit of measurement.[14] In this sense, wetness and brightness do not exist in the sense of 'exist' this paper rests on, but wetness and brightness do exist in the sense of 'exist' implied by everyday discourse. To put it another way, if our discourse referred to consciousness only as 'consciousness-as-cluster' and not as consciousness-as-property, there would be no issue.

Moreover, we do not have obfuscating intuitions on brightness and wetness, so we have never cascaded into a 'hard problem of brightness', and we probably never will. Yet, because of the different intuitional landscape for consciousness, much of scholarly discourse on consciousness is hopelessly wrapped up in questions of, 'Is this entity conscious?' which conflates two very different questions:

1. What mental capacities does the entity have?
2. Which mental capacities do we want to categorize as components of consciousness?

The first is an important and substantive scientific question that deserves much attention and detailed neuroscientific research. The second is a relatively arbitrary decision, which we might make based on practical considerations of scholarly or public communication. Neither derives an objective answer as to whether the entity is conscious, and the only reason left to believe in an objective answer is the fallible intuition that consciousness exists.

Using current definitions, we will never discover which beings are conscious, and claims like, 'There is a 10% chance that this computer program is conscious', are incoherent. But we can decide on a more precise definition and thus know which beings are conscious, and this can be based on genuinely interesting discoveries about the mental lives of other beings, such as whether honeybees possess mood-like states (in fact, current evidence in Bateson et al. (2011) suggests they do).

This is the same step we have taken for other proto-scientific concepts such as 'star'. Before we could see the sky with lenses and telescopes, it made perfect sense to use the term 'star' to refer to bright little things in the sky. But as we discovered more about the physical properties of different celestial objects, we decided to narrow the definition to mean masses of plasma held together by gravity, as distinct from planets or comets that are also bright little things in the sky. But this was not finding an answer to, 'Is Polaris a star?' It was simply an evolution of language as we gathered more detailed empirical information about celestial objects, and our language about consciousness can evolve the same way as we are getting more and more detail on the workings of brains and AI. For another example, consider metals (e.g., gold, iron). It is potentially discoverable whether a certain hunk of metal ore is gold or iron, but only after humanity has decided what exactly gold and iron are—chemical elements with precise atomic numbers.[15] This is related to debates on

[14] Some terms are ambiguous simply because of indexicality, where the term is ambiguous until placed in a certain context, such as 'me' (which depends on who is using the word) or 'the first item in the list' (which depends on which list is being referred to). An example of an indexical property is 'in our group' (which classifies entities based on whether they are in the group of the speaker). This is not the sort of ambiguity I'm referring to here.

[15] One could argue that even atomic elements may not be entirely precise. What if scientists encounter an atom identical to a gold atom but with a new, undiscovered subatomic particle included? This seems impossible based on current physics, but most epistemic views imply we cannot completely eliminate its possibility. In this case as in others, we

semantic ‘internalism’ and ‘externalism’, whether linguistic meaning is internal to the speaker or an external aspect of the speaker’s environment (Chomsky and Smith, 2000; Speaks, 2010). An example of the internalist view is that ‘meanings are instructions for how to build concepts of a special sort’ (Pietroski, 2018), while externalism is famously exemplified in Hilary Putnam’s (1973) ‘Twin Earth’ thought experiment. The mapping of these semantic positions to consciousness realism and eliminativism is not straightforward. For example, the Twin Earth comparison between water and Twin water entails a singular category while the proto-star comparison entails a plural category, given that bright little things in the sky includes at least (i) masses of plasma and (ii) masses of rock or gas. The claim herein is only that such meanings can change over time based on new empirical knowledge, and this claim may be fit into an internal or external conceptualization, though the internalism fit is more natural in the sense that internal meaning seems more malleable. This is also related to the recent literature on ‘topic continuity’ and ‘conceptual engineering’ (Cappelen, 2018; Knoll, 2020), alongside older work such as Quine (1951) that describes definitions and observations as co-evolving over time. For example, if one sees philosophical inquiry as grounded not only in analysing the current usage of a concept but in engineering future usage (Jackman, 2020), then the vagueness entailed in current usage of ‘consciousness’ may be more comfortable. Nonetheless, the vagueness would still entail that precisifying the concept is a part of answering the central question of consciousness, ‘Is this entity conscious?’

> *You admit that there is something—you call it ‘consciousness-as-self-reference’—exists and is such that I have 100% confidence in it. Is this first-person knowledge not ineffable? Even if we solve all the ‘easy’ problems of consciousness, how could someone else ever gain that knowledge?*

It is true that self-reference is a unique phenomenon in that it seems to remain even if you are in any imaginable brain-in-a-vat or Laplace’s demon scenario. We need to be careful not to conflate this with the broader claim that the indexicality or positionality of the observer implies an existing property. No additional knowledge is accrued by indexicality: If I am the first person in a single-file queue at a marketplace, I am unique in my indexicality in the sense that no other person could ever be first in line at the same time, but what knowledge does that endow me with that others necessarily lack? Other people in line could assess what my visual field is from that position, how happy I must be to be first in line, and so on. Ineffable knowledge does not follow from unique indexicality.

> *You say that the common definitions, such as ‘something it is like to be that organ-ism’, are not precise. But first-person experience offers exactly the precision you’re missing! Observe, in yourself, the ‘what it is like’-ness, such as the redness of red or the feeling of your hand on a hot stove. This is a very specific thing that an omniscient being could look for and use to categorize the vast majority of (or, more strongly, all) entities.*

---

can accommodate much vagueness by speaking of atomic elements given our current physical models but not by presuming a discoverable answer as to which entities constitute the element in scenarios when those physical models no longer apply.

When I observe my most personal of mental features, I share this intuition about the nature of those features. And I agree that the core semanticism argument (1–6 above) rests on whether we accept or reject the intuition that consciousness exists. I agree that the property of consciousness, then, seems to exist in a very real, undeniable way; explaining this intuition is known as the 'meta-problem of consciousness' (Chalmers, 2018). However, there are two issues with the reliability of this intuition.

First, as explained above, we must differentiate the two uses of the word 'consciousness': the act of self-reference and the broader property. Reasonable people may consider the former a brute fact, but the latter is not self-evidenced through introspection—it transcends that datum. In other words, we may have strong intuitions about consciousness-as-property, but we cannot have direct introspective evidence about consciousness-as-property in the way we do for consciousness-as-self-reference. In this paper, I exclusively take aim at consciousness-as-property's nonexistence. There is no 'hard problem' of consciousness-as-self-reference because the self-reference is simply a datum; it has no extension across individuals as a property. There is no category that we know my consciousness-as-self-reference and your consciousness-as-self-reference both belong to other than those following from the similar processes by which we made those self-references.

Second, if we properly disentangle self-reference from property, it seems that humans do not have reliable intuitions about the category of deep question to which questions of consciousness belong.[16] Humans did not evolve making judgments of and getting feedback on our answers to deep questions, such as the nature of sentience, quantum physics, molecular biology, or any other field that was not closely related to the day-to-day decision-making of our distant ancestors. Such intuitions are unrefined extrapolation from our intuition-building, evolutionary environment. Moreover, there are strong, specific reasons to expect humans to have an intuition that consciousness exists even if it does not. The idea of an objective property of consciousness is in line with a variety of intuitions humans have about their own superiority and special place in the universe. We tend to underestimate the mental capacities of nonhuman animals; we struggle to accept our own inevitable deaths; and even with respect to other humans, most of us suffer from biases a la the Dunning–Kruger effect. Consciousness realism is the same sort of phenomenon: it places our mental lives in a distinct, special category, which is a quite enticing prospect. Indeed, another objection I often hear is that consciousness eliminativism cannot be accepted because it would not allow us to give consciousness the moral consideration it deserves, but the moral impetus yielded by a belief is not evidence of that belief's validity.

As Elizabeth Anscombe said of an interaction with Ludwig Wittgenstein, 'He once greeted me with the question: "Why do people say that it was natural to think that the sun went round the earth rather than that the earth turned on its axis?" I replied: "I suppose, because it looked as if the sun went round the earth." "Well," he asked, "what would it have looked like if it had looked as if the earth turned on its axis?"' (Anscombe, 1959). In the case of consciousness, we should reflect on how it would seem if consciousness does not exist as a property; it would seem[17] no

[16] An unreliable intuition is still interesting and worth discussion, but less reliability should correspond to proportionally less evidential weight in our beliefs.

[17] One may respond that any 'seeming' is itself consciousness. I take this to be an uncommon and almost always dismissed definition upon reflection, but if one's definition

different than the current situation, and thus our intuition provides no net evidence against (or in favour of) eliminativism. The intuition that consciousness exists may appear more definitive than the intuition in favour of geocentrism, but again, I insist that we are discussing consciousness-as-property, not consciousness-as-self-reference—and thus there is no reason to give it special weight in our belief system.[18]

If one continues to testify that they have direct introspective evidence that consciousness-as-property exists, this creates an impasse. Once we have carved out all the discursive space around an individual's testimony, there is no argument I can offer on this or any other subject that will defeat brute insistence. Moreover, if intuition weighs heavily in this analysis, as seems to be the preference of most philosophers of mind, then we should account for the fact that new survey research suggests most people do not even agree there is a 'hard problem' (Díaz, 2021). And if the response to this evidence is that most people have not engaged in the proper reflection on those intuitions, then we should also consider the new experimental research suggesting that lay-person judgments about philosophical cases tend to stay the same after such reflection (Kneer et al., Forthcoming).

> *Ah, but now you have trapped yourself. Can't the semanticist argument now be applied to your own claims about the property of 'existence', and thus the entire debate between eliminativism and realism is meaningless?*

This is the objection I am most sympathetic to: Rather than saying consciousness does not exist, I could say that there is no meaningful or determinate answer as to whether consciousness exists or does not. The upshot of my argument would remain. This approach, removing the superficial layer of a philosophical question, could be seen as a version of logical positivism, or more precisely verificationism (Misak, 1995): Where I say a property 'does not exist', you can replace that with, 'is cognitively meaningless'. However, verificationism seems to not determinately resolve the debate on consciousness because realists could simply assert that the existence of consciousness is empirically verifiable or discoverable through introspection. In the language of Carnap (1928), I am arguing that the 'hard problem' is a 'scheinproblem' or 'pseudoproblem', a philosophical problem that is worded as if it has meaningful content, yet it cannot 'be translated into the formal mode or into any other unambiguous and clear mode'. Or, in Ryle's (1949) language, I am arguing that consciousness realism is a 'category error', mistakenly putting consciousness-as-property in the category of precise, real, or meaningful whereas it is actually in the category of vague, unreal, or meaningless. Each of these formulations is a reasonable translation of the view laid out in this paper.[19]

---

of consciousness extends that widely across human mental activity, then of course it exists. It simply does not get us anywhere in our understanding of the mind.

[18] Because of how strong religious doctrine was circa 1500, the geocentrism intuition may have felt even stronger than the consciousness intuition at that time.

[19] I view semanticism as a precisification of logical positivism alongside eliminativism, anti-realism, and so on, but a full development of such ideas is beyond the scope of this paper.

5 Implications and concluding remarks

The semanticist claim, if correct, is deeply important. Not only would it mean that a deeply seated intuition about an intimate component of the human experience is wrong, but it would force us to revaluate how we assess the consciousness of other beings. What does it mean to conduct neuroscientific research on consciousness if that property does not exist? How do we make legal and ethical decisions about brain-dead patients, foetuses, or nonhuman animals if there is no fact-of-the-matter regarding who is conscious and who is not? What properties of the mind will we imbue with normative value if we can no longer rest on a vague gesture towards consciousness or sentience?

Perhaps even more importantly, humanity seems to be rapidly developing the capacity to create vastly more intelligent beings than currently exist. Scientists and engineers have already built artificial intelligences from chess bots to sex bots. Some projects are already aimed at the organic creation of intelligence, growing increasingly large sections of human brains in the laboratory. Such minds could have something we want to call consciousness, and they could exist in astronomically large numbers. Consider if creating a new conscious being becomes as easy as copying and pasting a computer program or building a new robot in a factory. How will we determine when these creations become conscious or sentient? When do they deserve legal protection or rights? These are important motivators for the study of consciousness, particularly for the attempt to escape the intellectual quagmire that may have grown from notions such as the ‘hard problem’ and ‘problem of other minds’. Andreotta (2020) argues that the project of ‘AI rights’, including artificial intelligences in the moral circle, is ‘beset by an epistemic problem that threatens to impede its progress—namely, a lack of a solution to the “Hard Problem” of consciousness’. While the extent of the impediment is unclear, a resolution of the ‘hard problem’ such as the one I have presented could make it easier to extend moral concern to artificial intelligences.

So, how should our discussions move forward if we accept semanticism? Let me first clarify what semanticism is not. It is not just a view on semantics, such as ‘This is what “consciousness” means, or a view on epistemology, such as, ‘This is what we can know about consciousness’. These are both deeply involved in the analysis but only insofar as they are components of the ontological question, ‘Does consciousness exist?’. I take that to be a primarily ontological question, but answering that question requires semantics (i.e., meaning) of its three words, and the meaning of ‘existence’ involves epistemological facts about what we can know, as argued above. Moreover, even if you choose to define the word ‘exist’ differently than how I use it here and thus have a different answer to the question, ‘Does consciousness exist?’, the empirical upshot remains.

Let me also clarify that semanticism is distinct from the popular question of whether we have consciousness or only seem to have consciousness (an operationalization of the claim that ‘consciousness is an illusion’), which simply depends on whether we choose to define our individual conscious experience in a way that allows for a distinction between the seeming and the conscious experience itself. If we choose this definition, then this operationalization of illusionism can be correct. If we do not, then it is obviously wrong. I think this amounts to a verbal dispute, or specifically what I would call a definitional trap, in which the two sides are talking past each other with different definitions of ‘illusion’, and the question is trivially resolved if we just pick a definition. Moreover, this relatively arbitrary

choice of definitions seems much less important than the substantive question of potential discoverability.

For now, I suggest we continue to use the word ‘consciousness’. While vague, the term still fills an important social niche that no other term is currently poised to fill. With certain definitions of ‘eliminativist’ or ‘reductionist’, that empirical view makes me no longer qualify for either identity because I am not saying we should eliminate our use of the term, but I suggest that we disentangle ontological views on what exists from strategic views on how we should use language in intellectual discourse. Deciding on the best words to use relies on empirical investigation into what makes for effective communication, and it hinges on a variety of psychological and sociological variables, so I have much less confidence in my view on what term we should use than I do the ontological question of whether consciousness exists.

While we may continue using the term ‘consciousness’, I suggest that we no longer approach consciousness as if it is some potentially discoverable property and that we avoid assumptions that there is a ‘hard problem’, a ‘problem of other minds’, ‘neural correlates of consciousness’, or any other sort of monumental gap between scientific understanding of the mind and the ‘mystery’ of conscious experience. Research projects resting on those assumptions are wild goose chases. We should merely use our scientific knowledge to precisify the discourse. As ‘life’ has been broken down into reproduction, growth, homeostasis, and other characteristics, we may break consciousness down into more precise characteristics. Personally, I break it down as reinforcement learning, goal-oriented behaviour, mood-like states, complexity, and integration. There are various empirical descriptions we can give for these characteristics, usually found through neuroscience or behavioural tests, and those are all the explanations we need for a full account of consciousness. We will never discover what consciousness is, except that it is a vague gesture towards certain interesting, important mental phenomena. Theories of consciousness can only succeed in describing such phenomena, perhaps in a relatively unified way such as ‘global workspace theory’ (Baars, 1988) or ‘integrated information theory’ (Tononi, 2008). However, the success of a particular theory consciousness will have been a semantic decision, not an objective discovery—a feat of engineering, rather than a feat of analysis.[20]

There is much new ground to be broken in a new line of consciousness research with a more concrete framework that avoids being caught up in claims of ineffable mystery. We can talk about properties even if they do not exist in this sense, as long as our talk does not imply that they do. In neuroscience, we can still figure out

[20] It is possible that the precise phenomena associated with consciousness may be tightly clustered in feature space. For example, with advanced brain imaging and thalamic bridging, we may notice that all adult, non-vegetative humans share a specific information processing system, and when we turn that circuit off (e.g., through transcranial magnetic stimulation), subjects consistently report, ‘Wow, everything is exactly the same, except now it doesn’t feel like anything to be me’, a more generalized version of pain asymbolia. Then we notice that if and only if we place this circuit into artificial intelligence (e.g., the ‘emotion chip’ in *Star Trek*) does the AI report a ‘what it is like’ to be them. No other circuits have this effect. In this hypothetical scenario, while semanticism would still be correct, it would not matter much in practice because the vagueness could be somewhat resolved by empirical experimentation. Of course, this sort of scenario seems extremely unlikely, especially the consensus of consciousness evaluations of dissimilar entities, such as simple computer programs or alien species, where introspective beliefs about consciousness have little application.

neural correlates of what we intuitively think of as consciousness, and we can still figure out all the wondrous machinery that causes our reports of conscious experience. In fact, because semanticism allows us to disentangle these two phenomena, it seems we can make neuroscientific discoveries more efficiently, casting off the metaphysical and semantic baggage.

There is an unfortunate cyclical effect in consciousness studies: our misguided intuition fuels vague terminology and makes philosophers and scientists work hard to justify that intuition—as they have for centuries—which then perpetuates that intuition.[21] I believe that if we can get past this mental roadblock, accepting the imprecision of our current terminology and that there is no objective truth to consciousness as it's currently defined, then we can make meaningful progress on the two questions that are actually very real and important: What exactly are the features of various organisms and artificial intelligences, and which exact features do we morally care about?

There are also two important moral implications of eliminativism, particularly semanticism, outside of consciousness research. First, it reduces the likelihood of moral convergence (i.e., human descendants settling on a specific moral framework in the future). This is because one-way moral convergence could happen is if humanity discovers which beings are conscious or sentient and uses that as a criterion for moral consideration. This reduced likelihood should make us more pessimistic about the expected value of the far future (in terms of goodness versus badness) given humanity's continued existence, which then makes reducing extinction risk a relatively less promising strategy for doing good (Anthis and Paez, 2021).

Second, eliminativism tends to increase the moral weight people place on small and weird minds, such as insects and simple artificial intelligences, which is an important topic in the burgeoning field of research on the moral consideration of artificial entities (Harris and Anthis, 2021).[22] This is not a necessary consequence of the view, but it tends to happen for the following reason: when you view consciousness as a discoverable, real property in the world, you tend to care about all features of various beings (e.g., neurophysiology, behaviour, but also physical appearance, evolutionary distance from humans, substrate, etc.) because these are all analogical evidence of a real property. However, if you instead view consciousness

---

[21] I sympathize greatly with physicists defending the Everett or Many-Worlds Interpretation of quantum mechanics. It is easy for eliminativists like me to imagine an alternate history of theoretical physics where the notion of wavefunction collapse was never assumed (analogous to never assuming consciousness realism) and Many-Worlds took off as the default interpretation in the early 1900s instead of taking until at least the 1980s to catch on among quantum field theorists. I also hear woes from theoretical physicists who see a morass keeping string theory in place despite its challenges.

Similar dynamics obtain in theology. Consciousness realists ask, 'Without the reality of consciousness, can we still have compassion for and seek to protect other beings? How can we prevent suffering if suffering does not exist?' Religious people ask their nonreligious alters, 'If God doesn't exist, why don't you just steal and murder like a selfish hedonist?' The nonreligious person replies, 'how scary it would be if my belief in God were the only compelling reason I had to not steal and murder'.

[22] Eliminativism and illusionism are gateway drugs to panpsychism, in the sense that they encourage us to focus on specific mental features such as nociception that exist in a wide range of entities. However, discussion of panpsychism is beyond the scope of this paper and, predictably, hinges on exactly how we define panpsychism.

as a vague property, you tend to care less about the features that seem less morally relevant in themselves (e.g., physical appearance, evolutionary distance). Those features may still be indicators, since neither eliminativists nor realists have full knowledge of mental capacities of different entities. However, the indication for an eliminativist is more direct, such as from evolutionary distance to capacity for reinforcement learning rather than the realist evidential pathway from evolutionary distance to capacity for reinforcement learning to ineffable qualia. In other words, if an insect has the capacity for reinforcement learning, mood-like states, and integration of mental processes, then the eliminativist seems to have more freedom to say, 'That's a being I care about. My moral evaluation could change based on more empirical evidence, but those are mental features I want to consider'. Eliminativism places an evidential burden on those who deny that animals and other nonhuman entities lack consciousness: they are compelled to point to at least one specific, testable mental feature that those entities lack.

When we take a close look at the arguments for or against the existence of consciousness, our common-sense understanding evaporates, and that is okay. In fact, morality, free will, the meaning of life, the purpose of life, and thick concepts that serve as 'useful fictions' (Rosen, 1990), such as justice, evaporate in analogous ways. The modern stall of intellectual progress on philosophical questions—if the field decides that it sincerely wants to make progress—may be overcome by generalizing semanticism or other versions of eliminativism beyond consciousness. Developing such a view is beyond the scope of the current work, but it seems that the arguments for consciousness eliminativism, moral anti-realism,[23] free will reductionism, personal identity reductionism, as well as empiricism, positivism, and verificationism, which also sweep away much of philosophical discourse, are valuable starting points for generalization.[24] The thrust of semanticism is that a great philosophical clarity comes from accepting the vagueness of most philosophical terms. Questions such as 'Does free will exist?' evaporate, leaving only tractable questions such as, 'What notions of free will are most useful in society today?'

For consciousness studies, my hope is that semanticism, illusionism, or another view under the eliminativist umbrella—broadly construed—will take off and reach escape velocity, driving a new, scientific, and unadulterated understanding of consciousness. Which exact perspective, along with its respective jargon and

[23] This kind of moral anti-realism also seems to constitute a counterargument to moral uncertainty (MacAskill et al., 2020), the idea that we should account for being factually wrong about morality, analogous to empirical uncertainty, though a version of moral uncertainty could persist where the moral agent simply decides to care about their future moral preferences and account for new occurrences changing those. There may also be a sort of Pascal's wager for moral realism where anti-realists should account for the expected realist value of their actions, but I am not persuaded by such argumentation because it hinges on the plausibility of moral realism, whereas to me it seems semantically mistaken and thus cannot be assigned even a tiny probability.

[24] There are differences in the current discourses on these topics. For example, while consciousness is mostly referred to as a fact-of-the-matter of which discovery is at least theoretically possible (if not practical), some other concepts are more often properly acknowledged as useful fictions, where we are merely smoothing out a scatterplot of intuitions, which would remove the force of my argument. Currently it seems to me that none of these discourses fully acknowledge that fictitious, subjective nature of their object of study. Instead of trimming the literature with Occam's razor, it may be so distended that we need to launch Occam's nuke.

conceptual Lego, takes off is much less important than the overarching objective of clearing out the intellectual quagmire. I think that clarity could be reached through a variety of intellectual campaigns. For example, it could be that we say, 'Consciousness exists, but qualia do not', which may not be eliminativism per se but could have the same underlying claims about the world. Again, the map is not the territory. In any case, I expect that we will cast off the baggage of the 'hard problem' and similarly confused concepts, entering a new era of clarity in consciousness studies.

The nonexistence of consciousness can be one of the most challenging claims to accept: it pushes against a deeply held intuition about the nature of human experience. We crave a unique, unsolvable mystery at the core of our being. We want something to hang onto in this perilous territory, and due to academic happenstance, the terms carved out as handholds have been 'hard problem' and 'qualia' and other words that mistakenly gesture at ineffability and grandiosity. If we want to advance our species' understanding of who we are, we need to let go of these unsubstantiated intuitions. There is no insurmountably hard problem of consciousness, only the exciting and tractable problems that call out for empirical and theoretical study. The deepest mysteries of the mind are within our reach.

Acknowledgments. I am grateful for insight from Kelly Anthis, Peter Brietbart, Liam Bright, Jake Browning, Rachel Carbonara, David Chalmers, Patricia Churchland, Rodrigo Diaz, Kynan Eng, Keith Frankish, Douglas Hofstadter, Peter Hurford, Tyler John, Ali Ladak, Tom McClelland, Kelly McNamara, Seán O'Neill McPartlin, Caleb Ontiveros, Jay Quigley, Jose Luis Ricon, Ilana Rudaizky, Atle Ottesen Søvik, and Brian Tomasik.

References

1. Akiba, K.: How Barnes and Williams have Failed to Present an Intelligible Ontic Theory of Vagueness, Analysis (2015), https://doi.org/10.1093/analys/anv074.
2. Andreotta, A. J.: The Hard Problem of AI Rights, AI & Society (2020), https://doi.org/10.1007/s00146-020-00997-x.
3. Anscombe, G. E. M.: An Introduction to Wittgenstein's Tractatus. Hutchinson University Library, London (1959).
4. Anthis, J.R., Paez, E.: Moral circle expansion: A promising strategy to impact the far future, Futures, 130, 102756 (2021), https://doi.org/10.1016/j.futures.2021.102756.
5. Baars, B. J.: A cognitive Theory of Consciousness. Cambridge University Press, Cambridge (1988).
6. Bateson, M., Desire, S., Gartside, S. E., Wright, G. A.: Agitated Honeybees Exhibit Pessimistic Cognitive Biases, Current Biology, 21 (12), 1070–1073 (2011), https://doi.org/10.1016/j.cub.2011.05.017.
7. Block, N.: On a confusion about a function of consciousness, Behavioral and Brain Sciences, 18 (2), 227–247 (1995), https://doi.org/10.1017/S0140525X00038188.
8. Blumberg, A. E., Feigl, H.: Logical Positivism, The Journal of Philosophy, 28 (11), 281–296 (1931).
9. Bowles, N.: Jordan Peterson, Custodian of the Patriarchy, The New York Times (2018).
10. Cappelen, H.: Fixing Language: An Essay on Conceptual Engineering. Oxford University Press, Oxford (2018).
11. Carnap, R. Scheinprobleme in der Philosophie (1928).
12. Chalmers, D. J.: Consciousness and its Place in Nature. In: S. P. Stich, T. A. Warfield (eds.), Blackwell Guide to the Philosophy of Mind. Blackwell Publishing, Oxford (2003).

13. Chalmers, D. J.: Facing Up to the Problem of Consciousness, Journal of Consciousness Studies, 2 (3), 200–19 (1995).
14. Chalmers, D. J.: Moving Forward on the Problem of Consciousness, Journal of Consciousness Studies, 4 (1), 3–46 (1997).
15. Chalmers, D. J.: The Conscious Mind: In Search of a Fundamental Theory. Oxford University Press, Oxford (1996).
16. Chalmers, D. J.: The Meta-Problem of Consciousness, Journal of Consciousness Studies, 25 (9–10), 6–61 (2018).
17. Chalmers, D. J.: Verbal Disputes, Philosophical Review, 120 (4), 515–66 (2011).
18. Chomsky, N., Smith, N.: New Horizons in the Study of Language and Mind. Cambridge University Press, Cambridge (2000).
19. Churchland, P. S.: The Hornswoggle Problem, Journal of Consciousness Studies, 3 (5–6), 402–8 (1996).
20. Churchland, P. M.: The Rediscovery of Light, The Journal of Philosophy, 93 (5), 211 (1996).
21. Cleland, C. E., Chyba, C. F.: Defining 'Life', Origins of Life and Evolution of the Biosphere, 32 (4), 387–393 (2002), https://doi.org/10.1023/A:1020503324273.
22. Dennett, D.: Facing Backwards on the Problem of Consciousness, Journal of Consciousness Studies, 3 (1), 4–6 (1996).
23. Dennett, D.: Magic, Illusions, and Zombies: An Exchange, The New York Review (2018).
24. Dennett, D.: The milk of human intentionality, Behavioral and Brain Sciences, 3 (3), 428–430 (1980), https://doi.org/10.1017/S0140525X0000580X.
25. Díaz, R.: Do people think consciousness poses a hard problem? Empirical evidence on the meta-problem of consciousness, Journal of Consciousness Studies (2021).
26. Frankish, K.: Illusionism as a Theory of Consciousness, Journal of Consciousness Studies, 23 (11–12), 11–39 (2016).
27. Gloor, L.: Moral Anti-Realism Sequence #2: Why Realists and Anti-Realists Disagree, Effective Altruism Forum (2020), https://forum.effectivealtruism.org/s/R8vKwpMtFQ9kDvkJQ/p/6nPnqXCaYsmXCtjTk.
28. Harris, J., Anthis, J. R.: The Moral Consideration of Artificial Entities: A Literature Review, Science and Engineering Ethics, 27 (4), 53 (2021), https://doi.org/10.1007/s11948-021-00331-8.
29. Hofstadter, D. R.: I Am a Strange Loop. Basic Books, New York (2007).
30. van Inwagen, P.: Existence, Ontological Commitment, and Fictional Entities. In: M. J. Loux, D. W. Zimmerman (eds.), The Oxford Handbook of Metaphysics. Oxford University Press, Oxford (2005).
31. Jackman, H.: Construction and Continuity: Conceptual Engineering without Conceptual Change, Inquiry, 63 (9-10), 909-918 (2020), https://doi.org/10.1080/0020174X.2020.1805703.
32. Kneer, M., Colaço, D., Alexander, J., Machery, E.: On Second Thought: Reflections on the Reflection Defense. In: Lombrozo, T., Knobe, J., Nichols, S. (eds.), Oxford Studies in Experimental Philosophy. Vol 4. Oxford University Press, Oxford (forthcoming).
33. Knoll, V.: Verbal Disputes and Topic Continuity, Inquiry, 1-22 (2020), https://doi.org/10.1080/0020174X.2020.1850340.
34. Knuth, D. E.: Mathematics and Computer Science: Coping with Finiteness, Science, 194 (4271), 1976, 1235–1242 (1976), https://doi.org/10.1126/science.194.4271.1235.
35. Korzybski, A.: Science and Sanity: An Introduction to Non-Aristotelian systems and General Semantics. International Non-Aristotelian Library Publishing Company, Lancaster (1933).
36. Lewis, C. I.: Mind and the World-Order, C. Scribner's Sons, New York (1929).
37. Lewis, D. K.: On the Plurality of Worlds. Basil Blackwell, Oxford (1986).
38. MacAskill, W., Bykvist, K., Ord, T.: Moral Uncertainty. Oxford University Press, New York (2020).

39. McGinn, C.: Can We Solve the Mind–Body Problem?, Mind, XCVIII (391), 349–366 (1989), https://doi.org/10.1093/mind/XCVIII.391.349.
40. Misak, C. J.: Verificationism: Its History and Prospects. Routledge, London (1995).
41. Nagel, T.: What Is It Like to Be a Bat?, The Philosophical Review, 83 (4) 435–450 (1974), https://doi.org/10.2307/2183914.
42. NASA, Life's Working Definition: Does It Work?: Interview with Carol Cleland (2003), https://www.nasa.gov/vision/universe/starsgalaxies/life's_working_definition.html.
43. Orilia, F., Swoyer, C.: Properties. In: Zalta, E. N. (ed.), The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University (2020).
44. Papineau, D.: Thinking about Consciousness. Oxford University Press, Oxford (2002).
45. Pietroski, P. M.: Conjoining Meanings. Oxford University Press, Oxford (2018).
46. Putnam, H.: Meaning and Reference, The Journal of Philosophy, 70 (19), 699 (1973), https://doi.org/ 10.2307/2025079.
47. Putnam, H.: The Many Faces of Realism. Open Court, Chicago (1987).
48. Quine, W. V.: Two Dogmas of Empiricism, The Philosophical Review, 60 (1), 20 (1951), https://doi.org/10.2307/2181906.
49. Ramsey, W.: Eliminative Materialism. In: E. N. Zalta (ed.), The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University (2020).
50. Rosen, G.: Modal Fictionalism, Mind, 99 (395) (1990).
51. Russell, B.: Idealism. In: The Problems of Philosophy (1912).
52. Ryle, G. The Concept of Mind. Hutschinson's University Library, London (1949).
53. Speaks, J.: Theories of Meaning. In: Zalta, E. N. (ed.), The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University (2021).
54. Strawson, G.: Appendix: Dunking Dennett, Estudios de Filosofia (59) (2019).
55. Strawson, G.: Consciousness Isn't a Mystery. It's Matter., The New York Times (2016).
56. Strawson, G.: The Consciousness Deniers (2018).
57. The Brain Dialogue: Towards a Conscious-O-Meter (2016), https://www.cibf.edu.au/towards-a-conscious-o-meter.
58. Tomasik, B.: The Eliminativist Approach to Consciousness, Center on Long-Term Risk (2015), https://longtermrisk.org/the-eliminativist-approach-to-consciousness.
59. Tononi, G.: Consciousness as Integrated Information: A Provisional Manifesto, The Biological Bulletin, 215 (3), 216-242 (2008), https://doi.org/10.2307/25470707.
60. Wittgenstein, L.: Tractatus Logico-Philosophicus. Harcourt, Brace & Company, New York (1922).