Why I'm not a Content Pragmatist

The 2015 Minds Online Conference—the Brains Blog

Karen Neander

kneander@duke.edu

Content pragmatism is sometimes considered a viable alternative to dualism, eliminativism and naturalism, but I'm not convinced that it's even a genuine fourth alternative. Seeking a fourth alternative along the lines of content pragmatism might seem appealing, given the difficulties with the other three positions. But if we put it under pressure from one direction or another it threatens to topple over into one or another of the alternatives. I start with a quick look at dualism, eliminativism and naturalism with respect to mental content and then look at the allegedly contrasting claims of the content pragmatist.

Dualism, Eliminativism & Naturalism

The philosophical problem of intentionality is the problem of explaining how mental states (events and so on) refer to contents. If you hear a siren as coming from a certain direction, remember an absent friend, or imagine walking on Mars, these mental states refer to (are about) the direction of the siren, your absent friend and your walking on Mars. How do mental states get to be about things? *In virtue of what do they count* as referring to their contents?

When trying to understand how intentionality relates to the non-intentional facts and properties of the world, we distinguish between original and derived intentionality. *Derived intentionality* derives (at least in part) from other (independently existing) intentionality. *Original intentionality* does not. Original intentionality either derives from nothing, in which case it is a fundamental part or aspect of the furniture of the universe, or else it derives from the natural and non-intentional facts and properties of the world.

The problem of intentionality has two main parts. One is to give an account of derived intentionality. How does it derive from other intentionality? The other is to give an account of original intentionality. Is it a fundamental part or aspect of the furniture of the universe? If not, and it exists, how does it derive from (or supervene on) the other facts and properties of the world? These are the derivation and origination questions respectively. My focus here is on the origination question.

The cognitive scientist, Pylyshyn (1984, 23) describes this as the question of "what meaning is, or how it gets into my head" and he claims that it is "probably the second hardest puzzle in philosophy of mind," the hardest being the nature of consciousness (which, he adds, "probably isn't well enough defined

to even qualify as a puzzle").[1] Philosophers have pondered the nature of intentionality as far back as Aristotle and yet they've not reached anything like a general agreement on how to solve the puzzle. The puzzle has only intensified since the middle of the last century when the mind sciences began to posit mental representations and their contents in explaining cognitive capacities. In that context, reference to content is treated as an explanatory primitive.

Three main types of answer given in response to the origination question are those of the dualist, eliminativist and naturalist. The content pragmatist offers what is meant to be a fourth type of answer. Since it is easiest to state her position in terms of what she rejects in the other positions, I'll briefly sketch the other positions first.

The *dualist* with respect to intentionality holds that mental reference to content is real and one of the irreducible and, in principle, inexplicable explainers of the phenomena of the world. The dualist holds that original intentionality does not derive from (or supervene on) anything. It is a fundamental part of aspect of the furniture of the universe.[2]

(Actually, a dualist could instead claim that consciousness is such an irreducible and, in principle, inexplicable explainer, and that mental reference to content depends, constitutively, on consciousness; however, I ignore this possibility in what follows, since it will not substantially alter the argument. Either way, the dualist claims that something mental is fundamental.)

The *eliminativist* maintains that mental reference to content is not real. As well, the eliminativists claims that we should also not posit mental content in our scientific theories. Sooner or later, scientists should learn to explain behavior without positing mental reference to content, according to the eliminativist.[3]

The *naturalist*, like the dualist, believes that mental reference to content is real. We really do think about things. But the naturalist, like the eliminativist, rejects the claim that mental content is a fundamental part or aspect of the furniture of the universe. The naturalist claims that intentional phenomena derive from (supervene on) the non-intentional facts and properties of the world.

On the dualist and eliminativist's views, the puzzle of intentionality is (allegedly) dissolved. The dualist says that original intentionality is *in principle inexplicable*. The dualist says that we have thoughts about various things, but there is no *how* to it. A thought simply *is* about what it's about. The eliminativist says that we have no need for a theory because intentionality *does not exist*. So, again, there is

---

[1] See Pylyshyn's *Computation and Cognition: Toward a Foundation for Cognitive Science* (MIT Press, 1984). I don't think that Pylyshyn was drawing a sense/meaning versus reference distinction when he said this.
[2] See (e.g.) Plantinga, "Against Materialism" in *Faith and Philosophy* 23 (1): 3-32 (2006)
[3] It's hard to find thoroughgoing eliminativists with respect to all intentionality. Usually, what is argued is that a particular type of intentional phenomena (such as folk psychological intentional phenomena) is alleged not to exist. This is the case, for example, in Churchland's classic paper. See Churchland, P.M. (1999). "Eliminative Materialism and the Propositional Attitudes." In Lycan, W.G., (Ed.), *Mind and Cognition: An Anthology*, 2nd Edition. Malden, Mass: Blackwell Publishers, Inc.

no *how* to it, since it isn't real. Only the naturalist is stuck with the second-hardest puzzle. She can't shrug it off because she claims that, in principle, there is a story to be told concerning how intentional phenomena derive from non-intentional phenomena. That buys the naturalist a whole lot of trouble.

The naturalist might or might not be optimistic about success. The pessimistic naturalist might think that we're still far from understanding intentionality, or that the relevant understanding will always remain beyond our reach, because intentional mental phenomena are too complicated or we lack the ability to develop the right conceptual repertoire, or something along these lines. But every naturalist will probably support the naturalized semantics project to some extent, even if only in the form of moral support from the sidelines. The naturalized semantics project is an attempt to solve rather than dissolve the problem of intentionality. According to the dualist and the instrumentalist, anyone involved in this project (in other than a critical capacity) is on a fool's errand. According to the naturalist, there is a solution, in principle, and we certainly will not succeed, nor even thoroughly understand the difficulties standing in our way, if we do not try.

Content Pragmatism

Where does the content pragmatist stand? In her view, there are no immaterial souls simply thinking about things, and no fundamental psychophysical laws relating complex brain states to fundamental intentional properties. The content pragmatist denies the need for the ontological extravagance of dualism. She also rejects the eliminativist's methodological claim. According to the content pragmatist, positing mental reference to content is useful in explaining cognitive capacities. Yet she agrees with the eliminativist that mental content is not real. Since the content pragmatist agrees with the eliminativist that mental content is not real, the content pragmatist believes that the naturalist is on a fool's errand. Like the dualist and the eliminativist, her claim is that there is, in principle, no true story to be told concerning how intentional phenomena derive from the natural and non-intentional facts and properties of the world. She seeks a fourth way.

The content pragmatist sees no need to eliminate talk of mental reference to content from the mind sciences. Despite the fact that mental content is not real, we should, as one content pragmatist puts it, 'keep content as a gloss'.[4] The content pragmatist believes that it is useful to ascribe mental content in explaining cognitive capacities but that the appropriateness and usefulness of specific content ascriptions

---

[4] This paper is springs from a comment that I delivered to the 2015 SPP at Duke University in response to Frances Egan's paper, which was in defense of her version of content pragmatism. It benefits greatly from that session and I am especially indebted to Egan's paper, on which I rely in sketching the content pragmatist's motivation and view. The phrase 'keep content as a gloss' is Egan's. But I am not attempting to faithfully reproduce Egan's view at the SPP (and she may anyway have altered her view since the SPP). Those interested in Egan's position might read Egan (forthcoming), "Pragmatic Aspects of Content Determination," in Consciousness and Intentionality, Models and Modalities of Attribution, (Ed.) by Denis Fisette.

depends on the explanatory aims of the researchers and other pragmatic aspects of the explanatory context.

The rejection of naturalism is generally based on the argument from despair. The argument from despair starts with the claim that no available naturalistic theory of mental content has yet succeeded, nor seems near succeeding, and ends with the conclusion that the devil is not in the details (as the naturalist might want to claim) but in the very idea that a naturalistic theory of mental content is possible. The content pragmatist's negative argument relies on a large claim concerning the state of play of a whole literature devoted to developing and assessing naturalistic theories, and so it isn't possible to fully explain it, let alone assess it, here.[5]

But perhaps I can give a quick impression of some of the difficulties that the naturalist faces, all or some of which a content pragmatist might think insuperable. A central concern is that there are certain content determinacy challenges that have not (the content pragmatist claims) been met. Essentially, a content determinacy challenge is a challenge to explain why a given intentional mental state (or a given mental representation) counts as representing C, and not something else, Q (when C is not the same as Q).

To illustrate, consider the relatively simple case of toad prey-capture behavior. This has been much studied by cognitive neuroethologists, who speak of it as involving 'object recognition'. The sign stimulus for a toad's prey-capture behavior (i.e., what normally triggers it) is a stimulus that is (roughly), within certain size parameters, elongated and moving parallel to its longest axis. That is, it is a visual target that is moving in worm-like way. Something relevantly worm-like need not be a worm—it could be a cricket, a millipede or some other kind of creature, including (in the case of large toads) a small frog or mouse. Nor need a worm be relevantly worm-like. For example, a worm (cricket, etcetera) that is stunned, strung up, and moved perpendicular to its longest axis, is not relevantly worm-like. The relevant internal state, by which stage 'object recognition' has occurred, is thought to be a high frequency of firings in cells in a layer of the toad's optic tectum (T5-2 cells). Call the high frequency of firings in this layer of cells a 'M-state'.

What is the content of these M-states? For example, are such states representing the visual target as prey, as food, or as something worm-like (i.e., as something elongated and moving parallel to its longest axis)? Those offering naturalistic theories of content have, in effect, defended a variety of answers. Different answers license different evaluations with respect to whether there is correct or incorrect representation on a given occasion. For example, suppose that M-states represent visual targets as food. In that case, if a cardboard rectangle is moved past the toad in a worm-like way and this triggers a M-state, the cardboard will be misrepresented. If, in contrast, M-states represent visual targets as worm-

---

[5] There are a number of online introductions available. For example, see Adams & Azaiwa's entry, 'Causal Theories of Mental Content' and Neander's entry 'Teleological Theories of Mental Content' in the Stanford Encyclopedia of Philosophy.

like, the cardboard rectangle will be represented correctly. But if toad food that is not worm-like (e.g., a stunned cricket moved perpendicular to its longest axis) triggers an M-state (as can occur if there is neurological damage) it would on that interpretation be misrepresented.

This content determinacy challenge is often posed as a special challenge for teleosemantic theories of mental content, which appeal to the biological functions of the cognitive mechanisms involved. For this reason it is sometimes called 'the functional indeterminacy problem'. The idea is that mere appeal to such functions does not settle which content ascription is correct because several different descriptions of the relevant functions seem equally plausible. We could, with equal plausibility, claim that M-states have the function to carry information about (or promote the capture and ingestion of) members of a prey species, *or* toad food, *or* items displaying the relevant configuration of visible features. (Thus teleosemantic theories are also said to suffer from a fatal dose of 'disjunctivitis'.)

There are numerous content determinacy challenges, but I'll mention just two more. Another one concerns the distal nature of content. Roughly, the challenge in this case is to explain why a thought refers to (say) a cow rather than something more proximal, such as the cow's hide or the light reflected from the cow or the retinal impressions produced when one looks at a cow. For instance, suppose that it is suggested that the content of a perceptual state is what causes it (or what is "supposed" to cause it). The further suggestion might be that the contents of thoughts derive from the contents of perceptual states—that is, cows cause cow-perceptions and then these cause thoughts about cows, and that's why a thought about a cow counts as being about a cow. Or the suggestion might just be that at least the contents of perceptual states are determined by their causes. One problem with this simple starting idea is that, when a cow causes a cow-perception, so do other things. There are, for example, other links in the causal chain between the cow and the resulting mental state, such as the cow's hide or the light reflected from the cow or the retinal impressions produced when the light impacts the eyes.

A third content determinacy challenge concerns the Quinean gavagai problem. Quine posed this as a problem for the possibility of 'radical translation'. A radical translation is a translation of a language that begins from scratch, and relies on the observation of behaviors (including speech utterances), as well as on observations of the contexts in which those behaviors occur. Quine argued that the totality of such clues leave the interpretation indeterminate. For example, suppose that the speakers of the language tend to utter 'gavagai' in the presence of a rabbit. Quine argued that there will be no principled way to determine that the reference of 'gavagai' is *rabbit*, as opposed to *time slice of a rabbit*, *a rabbity appearance* or an *undetached rabbit part*.

The content pragmatist will argue that, even if we are permitted to investigate the inner workings of people (and can take into account their evolutionary and learning histories, and so on) such content indeterminacy will remain. According to the content pragmatist, there is no determinate mental content, although there are more or less appropriate and useful *ascriptions* of determinate mental content.

Whether they are more or less appropriate and useful depends on explanatory aims and other pragmatic aspects of the explanatory context.

Perhaps the naturalists have made more progress in addressing these kinds of content determinacy challenges than the content pragmatist thinks. But the answers that naturalists have offered certainly remain controversial, even among the naturalists themselves. In any event, for the purpose of this discussion, let's assume that no fully satisfactory naturalistic theory of mental content has yet been given, if only for the sake of the argument. My claim here is not that the naturalized semantics project has succeeded and that the content pragmatist should appreciate its success. I am wondering whether content pragmatism even provides a clearly distinct fourth alternative.

In contrast to the eliminativist, the content pragmatist thinks that it is useful to ascribe content in certain explanatory contexts. The content pragmatist's rejection of eliminativism is based on the recognition that mind scientists have continued to posit mental representations when explaining cognitive capacities, despite the problems with naturalism. Why have they continued to do so? One reason that might be given is that we cannot otherwise explain how a given neurological process 'amounts to' the exercise of a cognitive capacity. For example, if we want to understand how a process amounts to decision-making, we need to interpret it as one in which one decision is made as opposed to another, or if we want to understand how a process amounts to recognition, we need to interpret it as one in which something is recognized as one thing as opposed to another. But, according to the content pragmatist, the intentional interpretation of the inner causal process is 'just a gloss'. In her opinion, there is in reality no content to which this inner causal process refers. Instead, determinate content ascriptions are more or less appropriate, or more or less useful, depending on our explanatory aims and other pragmatic aspects of the explanatory context.

Let me sum up so far. The dualist dismisses the content determinacy challenges that create so many headaches for the beleaguered naturalist by telling us that the intentionality of mental states is real but that it is, in principle, inexplicable and irreducible. The eliminativist dismisses the content determinacy challenges too, instead telling us that mental states do not really refer to contents and that there is no need to ascribe such contents to mental states. The poor naturalist takes on the burden of explaining how mental content supervenes on the non-intentional facts and properties of the world, but as yet (we are assuming) has not yet succeeded in doing so. The content pragmatist rejects all three of these positions and seeks a fourth one, which involves a negative and a positive claim. The negative claim is that nothing non-intentional grounds intentionality because reference to content is not real. The positive claim is that, nevertheless, reference to content is appropriately and usefully ascribed in explaining cognitive capacities.

The pragmatist & the dualist

Now I want to put some pressure on content pragmatism from several different directions. It seems to me that, when we put pressure on it in one way or another, it threatens to topple over into one or the other of the three main alternatives. Let's start with content pragmatism in contrast to dualism.

Recall that a dualist claims that intentionality is irreducible and an inexplicable explainer. It is, according to the dualist, a fundamental part or aspect of the furniture of the universe. This is an ontological claim. There really is mental content and it really is basic—it just is. The distinction between dualism and physicalism is, notoriously, hard to draw in a satisfactory fashion. In part this is because we do not have an ideal, complete physics on hand, nor even one that is internally consistent. We do not yet know what an ideal, complete physics will posit. But, as a working assumption, the physicalist assumes that nothing mental is fundamental. The dualist denies this. The content pragmatist rejects the dualist's rejection of physicalism. (Or, at any rate, the content pragmatist, like the physicalist, denies that we need to posit fundamental mental phenomena to explain intentional phenomena.)

Yet the content pragmatist claims that determinate content ascriptions are appropriate or inappropriate depending on explanatory aims. The problem for the content pragmatist is that explanatory aims are intentional phenomena. If I aim to explain how human vision works then I am in a mental state that has one content as opposed to another. My aim is to explain how human vision works, and not how geese fly or how economic recessions occur. So, the content pragmatist seemingly claims that mental content is not real and yet also claims that what makes content ascriptions more or less appropriate and useful in a given explanatory context is mental content. This further content seems to be treated as an unexplained explanatory primitive in the content pragmatist's account. That looks a lot like dualism.

Of course, one might be a content pragmatist about content ascriptions in cognitive science, while holding a different view about the intentional states of whole people. Since the middle of the last century, cognitive scientists have posited inner mental representations in explaining cognitive capacities. That is, they have explained cognition in terms of internal processes that involve causally efficacious elements, some of which are given semantic interpretations. One might consistently hold the view that the referential content ascribed to these interpreted causally efficacious elements involved in cognition is not real, and yet also hold the view that whole people really have intentional mental states. But then content pragmatism would owe us an account of the real thing—the intentionality of the mental states of people.

I pass quickly from this worry, since it can be quickly expressed and other worries take longer to explain. But the fact that the content pragmatist seems to be appealing to intentional mental states while denying the reality of intentionality is a whopping big worry. On its own, it seems sufficient to scuttle content pragmatism's standing as a viable alternative.

The pragmatist & the eliminativist

The content eliminativist says that we should give up content. The content pragmatist says that we should not give up content but that we should, instead, keep the intentional interpretation of cognitive processes 'as a gloss'. What is meant by keep it 'as a gloss'? This is a key question when comparing content pragmatism with eliminativism.

It helps to mention a view that I shall call 'instrumentalism' here. The instrumentalist makes two claims: one ontological and the other methodological. The ontological claim is that we should eliminate mental content from our ontology—it does not exist. The methodological claim is that we should nevertheless posit mental reference to content on the part of mental states because it is useful to do so. It is 'instrumental' to do so. Here the instrumentalist disagrees with the eliminativist. The eliminativist says that we should not invoke mental reference to content in explaining cognitive capacities. The instrumentalist says that we should.

Is the content pragmatist an instrumentalist? As discussed in the previous section, the content pragmatist seems not to deny the reality of *all* mental content. If no mental content exists, there are no aims to explain anything. But let's set that elephantine problem for the content pragmatist aside. Let's instead ask whether the content pragmatist is an instrumentalist concerning the intentional interpretation of mental representations in a scientific context?

Traditionally, a *thoroughgoing* pragmatist is someone who says that ontological questions only make sense up to a point, beyond which they are mere mumbo jumbo. According to the thoroughgoing pragmatist it makes sense to ask if it is useful to posit something for explanatory purposes, but once this question has been asked and adequately answered there is no sense in further pressing the question of whether it really exists. If it is useful to posit it, that's as real as it gets, according to the thoroughgoing pragmatist.

The thoroughgoing pragmatist will want to qualify her position, in order to avoid cheap shots. She will want to make it clear that she is *not* endorsing the reality of everything that it could be useful to posit at any time or under any circumstances. (She is not endorsing the reality of the tooth fairy, even though it might be useful to posit a tooth fairy if one wants to comfort a small child.) She will want to adopt a general principle along the following lines: a phenomenon is real if and only if it is useful to posit it for explanatory purposes in the long run and all things considered.

Now, suppose for a moment that the content pragmatist is a thoroughgoing pragmatist. And suppose that she endorses this aforementioned general principle. If she endorses the prediction that it will be useful to posit content for explanatory purposes in the long run and all things considered, as the instrumentalist does, then such a content pragmatist is committed to the claim that mental content *is* as real as it gets.

But the content pragmatist is claiming that content is *not* as real as it gets. She is recommending that we keep mental content *just* as a gloss. So, presumably, she is not a thoroughgoing pragmatist. Presumably, she is a selective pragmatist, a pragmatist about mental content but not about everything. Presumably, the claim is that some phenomena (for example, electrons, neural spiking patterns, cows and so on) are more robustly real than mental content is, and yet it is nevertheless useful to posit mental content for explanatory purposes.

But does the selective content pragmatist think that it will be useful to posit mental content for explanatory purposes in the long run and all things considered? Again, it depends on how the view is further developed. Maybe the content pragmatist *rejects* the prediction that it will be useful to posit mental content in the long run and all things considered. But then the content pragmatist is an eliminativist, even if one who is less impatient than some. She might differ from other eliminativists on the question of how soon we should eliminate the positing of mental content from scientific explanations of cognitive capacities, but the content pragmatist's view is not an alternative to eliminativism on this understanding of it. It is a version of eliminativism, albeit one that recommends patience.

It is only if the content pragmatist is a selective pragmatist who *accepts* the prediction that ascriptions of mental content will be useful for explaining cognitive capacities in the long run and all things considered that her position seems genuinely distinct from eliminativism. Now, however (assuming that it is not a version of dualism in disguise) I wonder if it is a genuine alternative to naturalism.

The pragmatist & the naturalist

At this point, we're supposing that the content pragmatist is a selective pragmatist, not a thoroughgoing pragmatist. The content pragmatist says that it will be useful to posit mental content in the long run and all things considered, but mental content is neither real, nor as real as it gets. If mental content is not real, it does not really supervene on anything, and so it does not supervene on the non-intentional natural facts and properties of the world. Thus the view is overtly opposed to naturalism. But now we must ask, why would content ascriptions be useful and appropriate for explanatory purposes in the long run and all things considered if mental content is not real or as real as it gets? Isn't a phenomenon really explanatory in the long run and all things considered only if it is real, or at least as real as it gets?

Maybe this is too strong but there's a tension here that needs to be addressed. I take it that the content pragmatist might respond that determinate mental content is *relative* to explanatory aims.[6] However, the most plausible ways to develop this claim do not seem to me to support irrealism with respect to mental content.

The content pragmatist maintains that we should abandon hope of a naturalistic solution to the content determinacy challenges. Instead, she claims, different determinate content ascriptions are more or less appropriate or useful, depending on explanatory aims and other pragmatic aspects of the explanatory context. But the naturalist can agree that different determinate content *ascriptions* are more or less appropriate, depending on explanatory aims and so on. The content pragmatist needs a stronger claim. The content pragmatist needs to maintain that the way in which the appropriateness or usefulness of a content *ascription* is relative supports content irrealism.

---

[6] An alternative claim (proposed by Egan in her response at the SPP) is that content ascriptions are a form of idealization. Generally, scientific idealizations are thought to involve falsification or fictionalization and so this could support a distinctive content pragmatism position, but this needs more development and then more discussion than I could manage here.

Recall the functional indeterminacy problem again. In the case of the toad, three possible alternative content ascriptions are *prey*, *toad food* and *worm-like* (i.e., within certain size parameters, elongated and moving in the direction of the longest axis). The content pragmatist might be (wrongly) accused of, so to speak, 'resolving' the content determinacy challenges pragmatically by saying something like this. If one wants to explain how the toad recognizes visual targets as prey then one posits a representation of prey. If one wants to explain how the toad recognizes visual targets as toad food then one posits a representation of toad food. If one wants to explain how the toad recognizes visual targets as worm-like then one posits a representation of something worm-like, and so on. This would be an uninteresting pragmatic 'resolution' of this type of content determinacy challenge. It would not address the question of whether one of these explanatory aims is more useful or appropriate than the others. Note that, if the toad is not representing anything as X, then it is not appropriate to seek an explanation of how it does so.

However, this is not what the content pragmatist would want to say. Instead, I think that the content pragmatist wants to say something more like this (in relation to this type of case). If one wants to explain the toad's role in its ecological niche then one best ascribes the content *prey*, and if one wants to explain why the use of the representation is adaptive then one best ascribes the content *toad food*, and if one wants to explain how the toad's visual system works then one best ascribes the content *worm-like*.

To provide some structure to the discussion, let's distinguish between three different kinds of ways in which explanatory aims and so on might be relevantly different. (This is not meant to be an exhaustive list.) (1) Some people might want to explain cognitive capacities (e.g., human or anuran vision), whereas other people (or the same people on different occasions) might want to explain something else (e.g., a creature's ecological role or its fitness relative to its environment). This is the kind of difference adverted to in the previous paragraph. But (2) different people (or the same people on different occasions) might also be interested in explaining different cognitive capacities. That is, some might be interested in explaining vision, whereas others (or the same people on different occasions) might be interested in explaining (say) motor control, decision-making or learning. And (3) different people might also choose to adopt different methodological approaches to explaining psychological phenomena. For example, some people might try to explain behavior by means of a folk-psychological approach, whereas others (or the same people on other occasions) might try to explain psychological processes by means of an information processing or a Gibsonian approach. In my view, all such differences might lead to different content *ascriptions* being made.

Let's consider them in order. First off, it seems to me that explaining cognitive capacities takes precedence for determining the appropriateness and usefulness of content ascriptions. This is because it matters what contents are ascribed in that context, but it does not matter what content is ascribed to a creature if one is (say) explaining its ecological role or its fitness to its environment. For example, the toad need not represent its prey *as prey* in order for the ecologist to refer to the toad's prey as prey. The ecologist, not the toad, needs a concept of prey in order to do that. Similarly, in order for the biologists to explain how the toad manages to get fed and is in that respect fit relative to its habitat, there is no need

for the biologist to assume that the toad represents its food *as food*. It would, for instance, be sufficient to explain that the toad recognizes a certain configuration of visible features (if that is what it does), which are often enough displayed by edible and nourishing creatures in the toad's habitat, and that the recognition of this certain configuration of visible features triggers frequently successful hunting behavior in the toad.

The second potential source for the relativist claim is interest in different cognitive capacities. Plainly, if one is interested in explaining anuran vision, and is working within an information processing approach, one should consider how light is reflected from a distal source and how information about the distal source is extracted from the play of light on the retina and so on. I agree that the *vision* theorist should assign *visual* contents, since the vision theorist is interested in explaining vision.[7] I also agree that, in many cases, the same inner state is involved in multiple cognitive capacities. For example, a high level of activation in the toad's T5-2 cells is also involved in *locating* a worm-like stimulus. Different T5-2 cells respond to moving stimuli in different parts of the toad's visual field. This localization is not precise on any dimension: up/down, left/right or near/far. More precise localization of the stimulus is performed elsewhere in the toad's brain. Still, if we want to explain how enough localization occurs to allow the toad's initial orientation toward a moving worm-like stimulus, it is appropriate to mention this further aspect of the content, this localization content.

But this isn't a *competing* content ascription. This isn't a case where we have two equally plausible interpretations licensing different evaluations with respect to misrepresentation. If the content concerns the presence of something worm-like in such and such a location, a representational state could be correct with respect to one aspect of the content and incorrect with respect to another. That is, a single representational state could be right with respect to *what* the visual target is and wrong with respect to *where* it is, or *vice versa*. Thus, this is not the same as saying that there are two equally good interpretations of *what* is represented, or two equally good interpretations of *where* it is represented as being.

If we go on in this vein, to consider other cognitive capacities in which the same spiking pattern is involved, we will find the same kind of thing. For instance, the same cells in the optic tectum of the toad have a role with respect to the motor system. They send neural signals to the motor system that initiate orientation toward the stimulus. Again, this motivates a more complicated content ascription, but not a

---

[7] See Neander, "Content for Cognitive Science", in *Teleosemantics* (2006) edited by Macdonald & Papineau. I argue that the other content ascriptions favored in the philosophical literature seem to fly in the face of an information-processing explanation of anuran vision. Unfortunately, most of the philosophical literature has ignored all of the work in cognitive neuroethology on anuran vision that has been published since the seminal study conducted by Lettvin *et al* (1959). According to the Lettvin *et al* study, the relevant 'object-recognition' occurred in anuran retinal ganglion cells but, in the decades that followed, it was discovered that the relevant information processing is considerably more complicated and a number of mid-brain structures are involved. If one takes a careful look at the cognitive neuroethologists' explanations of anuran vision, one can soon see (or so I maintain) that only a content ascription that captures the relevant visible features of the stimulus is appropriate, for that explanatory context.

competing one. The same representational event can have visual content as well as motor content. It could represent *worm-like motion at such and such a location, orient toward such and such a location.* Again, the visual content could be correct even if the motor instruction is not executed properly (i.e., even if the toad fails to orient, or orients inaccurately).

The content ascription might become even more complicated as we consider more cognitive capacities. The information processing that produces the relevant state is really quite complicated. For instance, these high frequency firings in the T5-2 cells are not normally produced during the mating season and are inhibited by the presence of large looming (predator-like) shapes. But still, the point remains that an interest in explaining how the toad reduces the risk of predation or how the toad balances its need to hunt with its need to mate does not lead to *competing* standards of evaluation. Adding further cognitive capacities that we want to explain can complicate the content without justifying competing content ascriptions.

Turning now to the third source of differences in explanatory aims, there are definitely different methodological sympathies in play in the debate over which content ascription is correct in a given case. Some philosophers appear to have folk psychological explanations in mind. They are explicitly interested in which content ascriptions rationalize behavior.[8] (They presumably think of the toad as a toy example.) Others seem more sympathetic to a Gibsonian or neo-Gibsonian approach.[9] On a Gibsonian theory of perception, there is no representation prior to the representation of affordances (such as the edibility of an apple, which affords the opportunity for eating). Still others seek content ascriptions that can play a role in information processing explanations of the relevant cognitive capacities. On this kind of approach, visual processing is understood as extracting information from retinal impressions and it is a golden rule that, in vision, the visible features of the stimulus must be represented before any invisible ones (such as nutritive potential) are.

I find the information processing approach the most illuminating of these three. Some readers might favor an alternative. Perhaps the best approach is still to be discovered. In any event, the content pragmatist will need to maintain that two or more approaches, which support different content ascriptions, will be equally well able to explain the relevant cognitive capacities in order to maintain that content relativism follows. Whether or not that is plausible is not something that I could adequately discuss here.

But, even if it were true, a further question is whether content irrealism would follow. Suppose that Method 1 supports one content ascription (Content C) and Method 2 supports another (Content Q) and that Method 1 and Method 2 are equally explanatory in the long run and all things considered. If so, one might conclude that there really is Content C and Content Q. That is, one might conclude that the content really is C relative to Method 1 and Q relative to Method 2. (E.g., one might say that it has the

---

[8] E.g., Price (2001) *Functions in Mind, a Theory of Intentional Content.* (OUP).
[9] E.g., Millikan (2000) *On Clear and Confused Ideas: An Essay About Substance Concepts.* (CUP)

folk psychological content *food*, and the information-processing content *worm-like*, or something of this sort.) If Content C is really the most useful and appropriate content ascription, conditional on use of Method 1, and Content Q is really the most useful and appropriate content ascriptions, conditional on use of Method 2, would this not have some basis in reality? While the content pragmatist apparently needs to answer in the negative, in order to sustain a distinct position from the naturalist's, it is hard to understand how this answer can be sustained.

Concluding remarks

I am questioning whether the content pragmatist can stake out a middle ground, which is not already held by one of the other three—the dualist, the eliminativist or the naturalist. To begin with, she appeals to the background explanatory aims of researchers. These are what make determinate content ascriptions more or less appropriate and useful, she maintains. She thereby appeals to the intentional mental states of people—their intentions to explain this or that. The intentionality of these mental states is implicitly left as an explanatory primitive, on the content pragmatist's view. Thus the content pragmatist is either left without any account of original intentionality, or is positing original intentionality as a part of the fundamental furniture of the universe. If the content pragmatist does *not* expect content ascriptions in cognitive science to be useful in the long run and all things considered, then her view collapses into a version of eliminativism. If she *does* expect content ascriptions to be useful in the long run and all things considered, then it seems reasonable to think that mental content must be real. This is not undermined, so far as I can see, even if various different content ascriptions relative to different explanatory aims will be useful in the long run and all things considered.