Galacticism, thought-relativism, quasi-internalism

Jeff Speaks

May 27, 2020

1 The central argument

One of the great virtues of *Narrow Content* is that it provides a framework for understanding the dispute between internalists and externalists which at the same time sharpens the disagreement and makes it clear what the most defensible versions of internalism are.

Following Yli-Vakkuri and Hawthorne (Y&H), let's think of thoughts as events, and content assignments as functions from thoughts to assignments. The qualitative agential profile (QAP) of a thought is the maximal way in which that thought relates to the intrinsic qualitative properties of the agent of the thought. Then, in Y&H's terms, a content assignment is *narrow* iff the contents assigned to thoughts strongly supervene on those thoughts' QAP's. A content assignment is *truth-conditional* iff it assigns to every thought a content which has (relative to the relevant index) the truth-value which the thought has in that index.

Why the 'relative to the relevant index' qualifier? On standard views, contents will have different truth-values when evaluated at different indices. On virtually every view, contents can have different truth-values with respect to different worlds; on many views, they can have different truth-values with respect to different times. Here 'index' just means, roughly, 'whatever contents can determine different truth-values with respect to.'

One of the central questions posed in the book is: is there a narrow truth-conditional content assignment? The answer to this question, as Y&H show, depends on which parameters indices include.

The reason why can be brought out by considering one of the main styles of argument exploited throughout the book. In a *Döppelganger argument*, we consider a pair of thoughts which (i) have the same QAP, (ii) differ in truth-value, and (iii) are alike with respect to certain indices. The existence of such a pair shows that there is no truth-conditional narrow content assignment (given the relevant choice of indices). Let's say that a choice of indices is *Döppelganger-immune* iff no Döppelganger argument can be constructed for it. Then it follows from the foregoing that if there is a truth-conditional narrow content assignment, then the indices to which the assigned contents are sensitive must be Döppelganger-immune.

One way to state the conclusion of Chapter 2 of *Narrow Content* is that it turns out to be surprisingly difficult to construct Döppelganger-immune indices. The central example here is Mirror Man, a left-right symmetric agent who thinks 'Kit is human' with his left hemisphere (while pointing with his left hand to Kit_1) and simultaneously thinks 'Kit is human' with his right hemisphere (while pointing with his right hand to Kit_2). Suppose that Kit_1 is human and Kit_2 is not. Let's call thoughts of the first type L-thoughts and thoughts of the second type R-thoughts. Then a given L-thought and a simultaneous R-thought may have the same QAP and yet differ in truth value, despite having as their indices the same world, time, and agent.

2 Galacticism

How should someone with internalist sympathies respond? Here is what David Chalmers says about the central argument of *Narrow Content*:

> '. . . one can equally make an argument against galacticism, the thesis that thoughts have truth-conditional content that depends only on galactic qualitative properties, those instantiated within the galaxy of the thinker. One need only consider a Mirror Galaxy case involving a symmetrical thinker within a symmetrical galaxy in an asymmetrical universe. Say our thinker thinks *Kit_1 is being observed* and *Kit_2 is being observed*. Here Kit_1 and Kit_2 are symmetrical humans, where Kit_1 is being observed by someone outside the galaxy and Kit_2

> is being observed by no one. Then the first thought is true and the second thought is false. But the two thoughts are qualitative galactic duplicates. By the same reasoning as in the Mirror Man case, galacticism is false.
>
> Should we conclude that truth-conditional content depends on something outside the galaxy? No doubt it sometimes does, but that would be the wrong conclusion in this case. The truth-*values* of the thoughts in question depend on something extra-galactic, but the truth-*conditions* do not seem to . . . The difference between the thoughts arises from the non-qualitative numerical difference between Kit_1 and Kit_2, who may be entirely intra-galactic.
>
> Returning to the Mirror Man, what should we conclude from the argument that truth-conditional content does not depend on internal qualitative properties? . . . The reasonable conclusion is that truth-conditional content depends on non-qualitative properties, such as those involving specific objects. The argument is entirely neutral about whether those objects are internal objects, external objects, or both. So the argument as it stands does little to establish externalism over internalism.'[1]

Chalmers' discussion here points to two closely related replies to the central argument.

The first (and closer to the spirit of the above passage) is to say that the relevant theoretically interesting content assignment is sensitive, not just to the subject's intrinsic qualitative properties, but also to internal objects (like the subject's experiences and thoughts). Since the two Mirror Man thoughts are distinct token thoughts and are related to distinct token experiences, there is no reason why a view of this kind must assign the two thoughts the same truth conditions. A content assignment of this kind will not be narrow in the sense defined above, but it certainly seems that views in this family are 'internalist enough' to be genuine alternatives to familiar forms of externalism. I'll follow Y&H in calling this family of views 'quasi-internalist.'

There is a second way of responding to the Mirror Man argument which

[1] Chalmers (2018).

is suggested by some of Chalmers' other remarks.[2] This is to incorporate 'internal objects' like experiences into indices rather than into contents. Then one and the same narrow content could be assigned to L and R; it's just that that content determines different truth-values with respect to different thought-involving indices. Because on this view contents can have different truth-values with respect to indices involving different internal objects, Y&H call views of this sort 'thought-relativist' views.

With these views on the table, it is not unreasonable to wonder whether the choice between them is just a matter of bookkeeping; couldn't any theoretical role played by quasi-internalist contents also be played by thought-relativist contents + fine-grained indices (and vice versa)?

Against this, it is worth noting that in other domains the distinctions between parallel views are taken to be genuine. For example, contextualist treatments of epistemic modals (on which differences in the level of knowledge in the context of utterance make for differences in content) are typically regarded as non-trivially different from relativist treatments (on which the level of knowledge is part of the index, not the content).[3] So it is worth exploring the possibility that these two avenues for the internalist-friendly to explore differ in substantial ways.

In what follows I'll consider these two paths in turn. We'll see that, in the end, the challenges they face do have much in common. But the path to that conclusion seems to me worth tracing.

3 Thought-relativism

Let's consider the thought-relativist option first. Let's use '$[\![\text{t}]\!]_{tr}$' for the content assigned to the thought t by our thought-relativist content assignment. Then $[\![\text{L}]\!]_{tr}=[\![\text{R}]\!]_{tr}$; the two Mirror Man thoughts have the same thought-relativist content. But they would have different truth-values, since the content which they share is true relative to the index of L and false relative to the index of R.

Y&H discuss thought-relativism in §2.3. They point out that making

[2] See his discussion of the 'two tubes' example from Austin (1990) in Chalmers (2012) and Chalmers (2018).

[3] For an entryway into the now substantial literature on this, see Egan et al. (2005).

indices this fine-grained makes it trivial that there is a narrow truth-conditional content assignment of this sort, and hence makes this sort of thought-relativist view automatically immune to Döppelganger arguments. Chalmers (2018) protests: 'There are worse problems for a view to have!'

Fair enough — there certainly are. And if Y&H were arguing that its immunity to counterexample shows that there is no theoretically interesting version of thought-relativism, this would be a bad argument. Freedom from counterexample is no guarantee of theoretical pointlessness. But I take Y&H to be pointing out that the general claim that there is *some* truth-conditional thought-relativist content assignment is trivial, and so not itself theoretically interesting. This leaves unanswered the question of whether some particular thought-relativist content assignment might be theoretically interesting.

For such a content assignment to be theoretically interesting is for the content assignment to play some interesting theoretical role. Here I'll explore the question of whether this kind of content assignment can play a role familiar from Chalmers' work on these topics.

This is the role of satisfying what Chalmers has called the 'Core Thesis' of two-dimensionalism. This thesis holds that a thought is a priori iff its primary intension is necessary. Since we are now considering the possibility that primary intensions are (or are determined by) the contents assigned by a thought-relativist content assignment, we can put the relevant version of the Core Thesis as follows:

> CORE$_{tr}$ t is a priori iff $[\![\text{t}]\!]_{tr}$ is necessary

The second thesis connects the a priori with what an idealized reasoner could know. Roughly:

> IDEAL. S is a priori iff an ideal reasoner could know it to be true without justification based on experience.

An argument given by Y&H shows that CORE$_{tr}$ is false.[4] Modifying the 'Mirror Man' example slightly, let L be the type of thought one has

[4] Yli-Vakkuri and Hawthorne (2018), 134.

when looking at the object on one's left and saying 'That is human' and let R be the type of thought one has when looking at the object on one's right and saying 'That is human.' Now consider some thought token t_1 of the type

[t_1] L $\equiv$ L

and some other thought token t_2 of the type

[t_2] L $\equiv$ R

It seems that t_1 is a priori and t_2 is not. But, given that $[\![t_1]\!]_{tr}=[\![t_2]\!]_{tr}$, this falsifies CORE$_{tr}$.

It seems to me that a thought-relativist should not find this observation itself to be especially worrying. After all, the present sort of counterexample to CORE$_{tr}$ is not special to *thought*-relativism — it arises for even much more modest relativist views which relativize truth-values to times.

Let's use '$[\![.]\!]_{time}$' for a time-relativist content assignment, and consider the time-relativist version of the Core Thesis:

CORE$_{time}$ t is a priori iff $[\![t]\!]_{time}$ is necessary

Now consider someone who thinks the thought

[a_1] Amelia is happy iff Amelia is happy

where the time relevant to the evaluation of the left hand side is the same as the time relevant to the evaluation of the right hand side.[5] This is presumably a priori. But now consider someone who thinks considerably more slowly, and has the thought

[a_2] Amelia is happy iff [waits a few hours] Amelia is happy.

[5] How could they be relativized to the same time? I'll return to this below.

This is presumably not a priori; Amelia's mood could have changed during the course of the thought, and whether it did is not something which even an ideal reasoner could know independently of experience.

But note that, for all we have said, it might be that $[\![a_1]\!]_{time}=[\![a_2]\!]_{time}$. It follows that CORE$_{time}$ is false. Since this argument that the time-relativist cannot respect the Core Thesis is obviously parallel to the above argument that the thought-relativist cannot respect this principle, it is worth asking whether the responses which the time-relativist might make could also be employed by the thought-relativist.

The basic problem here is not hard to figure out. CORE$_{tr}$ and CORE$_{time}$ try to define the a priori in terms of truth at every world. But on a time-relativist or thought-relativist view, the truth-value of a thought depends on factors other than a world. It looks like our relativist core theses should be modified to take account of this fact.

How to do it? Here is one reasonably natural suggestion for the time-relativist:

> RELATIVE CORE$_{time}$ t is a priori iff $[\![t]\!]_{time}$ is, holding fixed the times to which parts of t are indexed, necessary.

Since all of the parts of a_1 are indexed to the same times, a_1 will satisfy the right hand side of RELATIVE CORE$_{time}$, and hence will come out a priori, as it should. Since the parts of a_2 are not relativized to the same times, a_2 won't satisfy the right hand side of RELATIVE CORE$_{time}$, and hence won't come out a priori — which again is what we want.

Given this, it is worth asking whether the thought-relativist can go in for a similar qualification. A natural suggestion is:

> RELATIVE CORE$_{tr}$ t is a priori iff $[\![t]\!]_{tr}$ is, holding fixed the thoughts to which parts of t are relativized, necessary.[6]

At first glance, this might seem to give us the right result about the cases discussed above, for reasons parallel to those given about the time-relativist's treatment of a_1 and a_2.

[6] Of course one might be a time-relativist *and* a thought-relativist, in which case the relevant statement of the Core Thesis should be doubly relativized; I ignore this point here for simplicity.

But in fact things are a bit more complicated than this. Just as cases like a_2 show that the time-relativist needs to be able to relativize different parts of thoughts to different times, so cases like t_2 show that the time-relativist needs to be able to relativize different parts of thoughts to distinct thought-events. But t_1, like t_2, contains distinct thoughts; it contains, after all, two different thought tokens of the thought-type L. Let's use l_1, l_2, etc. for thought tokens of the L type. Then we might represent t_1, using labels for thought tokens rather than thought types, as

> [t_1] l_1 iff l_2

We have been working under the assumption that we want t_1 to come out as a priori, and so that we want $[\![t_1]\!]_{tr}$ to satisfy the right hand side of RELATIVE CORE$_{tr}$. But it is not clear that it does.

To see why, consider the question of whether thought tokens have their causes essentially or not. This is a question which in the present context seems to call for stipulation rather than serious investigation; but either way we stipulate, we seem to get trouble.

Suppose first that we say that they do have their causes essentially. Now suppose that you know that you are playing a game in which you have to guess whether the object before you is a human or a waxwork replica of one. One first makes a guess, and then puts on a blindfold, during which time the object before one is either swapped out or remains in place. Now suppose you think the thought

> [t_3] That is human iff [puts blindfold on and waits] that is human.

It seems obvious that t_3 is not knowable priori. But suppose that in fact both 'That is human' thoughts were caused by a human, and hence are both true. On the stipulation that thought tokens have their causes essentially, this would seem to make $[\![t_3]\!]_{tr}$ necessary, and hence (by RELATIVE CORE$_{tr}$) a priori. More generally, treating thought tokens as having their causes essentially would (in the presence of RELATIVE CORE$_{tr}$) seem to wildly overgenerate instances of the a priori.[7]

[7] Notice that a similar problem seems to arise for the time-relativist. It seems possible that

So suppose that we stipulate that thought tokens do not have their causes essentially. This seems to lead to the opposite problem. Recall our earlier judgement that t_1 is a priori. But t_1 is composed of two token thoughts, l_1 and l_2. On the present proposal, we are licensed to consider worlds in which l_1 and l_2 have different causes than they actually do. In some of those worlds, t_1 will be false. So, by RELATIVE CORE$_{tr}$, it turns out not to be a priori after all.[8]

Let's pause on this second horn of the dilemma. Is this really such a bad result for the thought-relativist? After all, it does not seem like the thinker of t_1 can rule out a priori the possibility that the thing referred to by the first occurrence of 'that' $\neq$ the thing referred to by the second occurrence of 'that.' But then just for the same reason, it is not, on second thought, so implausible to say that t_1 is a posteriori.

This is fine as far as it goes. But it threatens to lead to an unacceptably restrictive view of the a priori. For surely at least sometimes we need to be able to treat distinct tokens of sentences like 'That is human' as a priori equivalent. Suppose I reason to myself as follows:

> [t_4] That is human.
>
> [t_5] If that is human then that is a mammal.
>
> So, [t_6] That is a mammal.

It seems plausible that (perhaps given some further background information about the case) it should be possible for t_6 to be an a priori consequence of t_4 and t_5. But if each these thoughts is indexed to itself, and we are licensed to consider worlds in which these thought tokens have different causes than they actually do, there will be some such worlds in which t_4 and t_5 are true but t_6 false, which (given RELATIVE CORE$_{tr}$) implies that t_6 is not an a priori consequence of t_4 and

an ideal reasoner could be in a situation in which it was not a priori knowable whether she was time traveling or not. Imagine such a reasoner entertaining the thought 'This time is earlier than [pause] this time.' Suppose that she is not in fact time traveling. If we hold fixed the times to which the two utterances of 'this time' are indexed, the thought will be necessary, and hence by RELATIVE CORE$_{time}$ a priori. But this contradicts our initial supposition that such a thought might be knowable only a posteriori. Perhaps an appeal to 'personal time' could help here.

[8] Could we say that l_1 and l_2 stand in some special relation, which does not supervene on the thought-relativist contents of the two thoughts, which provides an a priori guarantee of coreference? That would be to abandon CORE$_{tr}$ in favor of the kind of narrow relationist view Y&H explore in Ch. 6 of *Narrow Content*.

t_5 after all. This can't be right.[9]

Let's look at two ways the thought-relativist might respond to this problem.

The first response begins with the plausible thought that we should distinguish between perceptual-demonstrative uses of 'that' (like the ones in t_1, t_2, t_3, and t_4) and uses of 'that' which are meant to be parasitic on perceptual-demonstrative uses (like the ones in t_5 and t_6). Perhaps the the content of the perceptual-demonstrative token of 'that' in t_4 fixes the reference of the token as the cause of that very experience, whereas the content of the parasitic use of 'that' in the antecedent of t_5 fixes the reference of the token as the distinct thought token t_4.

To fix ideas, we might suppose that

$$[\![t_4]\!]_{tr}=[\![\text{The cause of \_\_\_ is human}]\!]$$

which is true relative to an index including t_4 itself. By contrast, perhaps

$$[\![t_6]\!]_{tr}=[\![\text{The cause of } t_4 \text{ is a mammal}]\!]$$

With a corresponding treatment of t_5, it looks like we can secure the a priority of the inference from t_4 and t_5 to t_6.

No doubt this is too simple. But even this bare sketch of a thought-relativist content assignment is enough to make clear that this is not a route which an internalist thought-relativist can travel. For we can imagine Mirror Man simultaneously carrying out the line of reasoning 'That is human; if that is human then that is a mammal; so that is a mammal' twice, starting once with an L-type thought and once with an R-type thought. By parity of reasoning the QAP's of the L-based thoughts will be the same as those of the R-based thoughts. But on the present proposal the two 'that is a mammal' thoughts will differ in content, since they involve reference to numerically distinct experiences. That shows that the present proposal is less a defense of internalist thought-relativism than a move to a kind of quasi-internalism (about which more below).

[9] There are parallels here to the problems encountered by proponents of token-reflexive accounts of indexicals in validating the rule of Repetition. For some possible moves here, see Radulescu (2018).

Let's consider a second option. The thought relativist might say that the thought in the index of some token thought t might include thoughts other than t itself. Then the thought-relativist might say that

$[\![t_6]\!]_{tr}$=[[The cause of ___ is a mammal]]

but say that the index of this thought includes, not t_6, but t_4. This is (for all we have said) consistent with the thought-assignment being internalist and (given RELATIVE CORE$_{tr}$)) secures the a priority of the inference.

The obvious question then is what determines which thought the content of a given thought is supposed to be indexed to. The simplest view says that the content of a thought is always indexed to itself; but on the present account the index relevant to a given thought-constituent may vary.

Given RELATIVE CORE$_{tr}$ and IDEAL, the facts which determine which thought figures in the index of another thought must be facts which would be knowable, without reliance on experience, by an ideal reasoner. For suppose that they were not. Then an ideal reasoner would sometimes be unable to determine the index relevant to a thought. But, since facts about the index of a thought can (given our relativized versions of the core thesis) determine whether that thought is a priori or not, there could be thoughts which are such that an ideal reasoner could not figure out whether those thoughts are a priori. And that contradicts IDEAL.

This constraint would seem to rule out certain otherwise plausible index-fixing factors. For example, one might have thought that the index relevant to a thought is fixed at least in part by causal relations between thoughts. But ideal reasoners (of the relevant sort) won't in general know these claims independently of experience.[10]

The obvious move here is to appeal to the intentions of the thinker; perhaps the index of a thought is determined (at least in some cases)

[10] Could the index instead be fixed by certain rational or inferential relations between thoughts? Perhaps. One worry here is that these relations may have a causal component. Another is that if such relations supervene on the propositional attitudes of the subject, then the problem with the appeal to intentions discussed below will arise. Another option would be to appeal to a sequence of thoughts as part of the index for a thought. These options all seem worthy of exploration. Thanks to Brian Cutter for helpful discussion here.

by the intentions of the thinker of that thought.

But this leads swiftly to the collapse into quasi-internalism just rehearsed above. For consider again the case in which Mirror Man simultaneously carries out the two lines of reasoning, culminating in his twice over drawing the conclusion 'That is a mammal.' The index relevant to those two 'That is a mammal' thought must be distinct. (In one case, it is an L-type thought and in the other an R-type thought.) That means that the intentions which fix the index of the two 'That is a mammal' thoughts must differ in content.[11] But it again looks like those two intentions could have just the same QAP. So the difference in the contents of intentions to which one wants to appeal cannot be delivered by an internalist content assignment.

The aim of the foregoing has been to present a kind of dilemma for the internalist thought-relativist. On the one hand are versions of the view which, by making every thought's truth relative to itself, yield an implausibly restrictive view of the a priori. On the other hand are versions of the view which avoid this result by building more into the contents assigned by $[\![.]\!]_{tr}$, but in so doing lapse into quasi-internalism. Does this show that the dilemma is inescapable? No.[12] But it points to some of the challenges which a more developed thought-relativist view might face in validating the conjunction of RELATIVE CORE$_{tr}$ and IDEAL.

4 Quasi-internalism

Let's turn to quasi-internalism. Y&H take up quasi-internalism in Chapter 5 of Narrow Content. The central argument they give in that chapter is that the view is 'theoretically pointless.' This charge is spelled out in the following passage:

[11] Could the thought-relativist instead say that the different index-shifting intentions themselves have different indices, and that this difference in indices is what determines the shift in the index of the two 'That is a mammal' thoughts? Maybe. But there is a worry that this just pushes the problem back a step. By reasoning parallel to that given above, it looks like the index-shifting properties of the indices of the two intentions would have to be a priori accessible to an ideal reasoner. But what properties could these be? It is not clear that this question is unanswerable; but it is also not clear to me how the thought-relativist should answer it.

[12] See note 11 for a possible escape hatch.

> 'The threat of theoretical pointlessness . . . is particularly pressing for the quasi-internalist . . . it is arguable (in fact, we think it is true) that all content assignments are quasi-narrow. If so, the quasi-internalist has failed to articulate any non-trivial condition on content assignments, and, insofar as she concedes that there is no *narrow* content assignment that can do any interesting theoretical work, her position collapses into the negation of internalism.'

Y&H then go on to give two arguments for the claim that all content assignments are quasi-narrow. Suppose that one of these arguments goes through. Why is this supposed to be bad news for the quasi-internalist?

On one way of understanding the argument, Y&H are inferring from the premise that all content assignments are quasi-narrow to the conclusion that any particular quasi-narrow content is, in some sense, theoretically pointless. This would be an obviously bad argument; the ubiquity of quasi-narrow content assignments does not imply anything about the theoretical credentials of any individual such assignment.

Fortunately, it is pretty clear that this is not the argument which Y&H have in mind. Following the discussion of theoretical pointlessness they don't consider the case closed; instead, they go on to consider various ways in which the distinctive thesis of the quasi-internalist might be strengthened. This suggests that what they take themselves to have shown to be theoretically pointless is not some specific quasi-narrow content assignment, but rather the general claim that some quasi-narrow content assignment or other is truth-conditional.

Given this, the right response to the argument is not to give up on quasi-internalism, but to consider stronger (more specific) quasi-internalist claims about content assignments and evaluate those more specific proposals for their theoretical utility. Chalmers is, it seems to me, clearly right that some specific quasi-internalist proposals, including his own, are 'internalist enough to be interesting.' The question is whether the contents they assign can do interesting theoretical work.[13] So let's turn to that question, focusing again on the Core Thesis and IDEAL.

[13] I take it that Y&H agree, which is why they focus on these views in Chapter 6 of the book.

Using '$[\![.]\!]_{qi}$' for the relevant quasi-internalist content assignment, we can state the relevant version of the core thesis as

CORE$_{qi}$ t is a priori iff $[\![\text{t}]\!]_{qi}$ is necessary

To take a toy version of the view, it might be that the quasi-internalist content of every thought-token makes reference to that very thought-token. Then $[\![\text{L}]\!]_{qi} \neq [\![\text{R}]\!]_{qi}$, which avoids one of the problems which faced the thought-relativist.

One way to pose a challenge to CORE$_{qi}$ begins with the reasons why a familiar sort of Millian-externalist content assignment $[\![.]\!]_{me}$ falsifies the relevant version of the core thesis.

On this sort of view, contents can be the same but for a difference in the external objects they are about. In particular, on many such views it will be the case that

$[\![\text{Hesperus is a planet}]\!]_{me} = [\![\text{Phosphorus is a planet}]\!]_{me}$

so that

$[\![\text{Hesperus is a planet iff Phosphorus is a planet}]\!]_{me}$

will come out necessary. But an ideal reasoner need not be in a position to know a priori that this biconditional is true, which falsifies the conjunction of IDEAL and CORE$_{me}$.

The quasi-internalist does not let external objects into contents; but she does let in internal objects. But one might worry that this leads to an argument against CORE$_{qi}$ which is parallel to the above argument against CORE$_{me}$. The problem for CORE$_{me}$ arose from the fact that there can be a pair of singular token thought-constituents x, y which are such that $[\![x]\!]_{me} = [\![y]\!]_{me}$ (because they refer to the same external object) but which are such that even an ideal reasoner would not be able to know a priori that they refer to the same external object. But couldn't there be a pair of singular token thought-constituents x, y which are such that $[\![x]\!]_{qi} = [\![y]\!]_{qi}$ (because they refer to the same internal object)

but which are such that even an ideal reasoner would not be able to know a priori that they refer to the same internal object?

This is just a schematic worry. Here is a way to sharpen it. Let's say that one token thought x *repeats* a token thought y just in case $[\![x]\!]_{qi}=[\![y]\!]_{qi}$. Then we can ask: is repetition possible, or not?

If it is not, then that seems to lead (via the line of argument involving t_4-t_6 above) to an excessively restrictive view of the a priori. After all, unless the antecedent of t_5 repeats t_4, it hard to see how t_6 could be a necessary (and hence by CORE_{qi} a priori) consequence of t_4 and t_5.

So suppose that repetition is possible, and suppose that Mirror Man, the next day, reflects upon his previous day's experiences. He thinks

> [t_7] That is human (while closing his eyes and calling one of his experiences to mind) iff that is human (while again closing his eyes and calling one of his experiences to mind).

Let's stipulate that in fact the token thoughts on either side of the biconditional both repeat some particular L-type thought. Then, Given CORE_{qi}, t7 will be a priori. But it seems that even an ideal reasoner may not be in position to know whether the constituent token thoughts repeat the same thought from the previous day. The previous day's experiences are, after all, qualitatively indistinguishable.[14]

There is an obvious parallel here to cases which make trouble for CORE_{me}. This fact might suggest a parrying reply for the quasi-internalist. If quasi-internalist theories fail to satisfy our desiderata for just the reason that Millian-externalist views fail to do so, then surely whatever response to these arguments is offered by the Millian-externalist can adopted by the quasi-internalist.

But of course the Millian-externalist is happy to concede that there is no theoretically interesting content assignment (including her own) which vindicates CORE and IDEAL. The hope that there is such a content assignment is after all one of the motivations for an alternative

[14]The case is very similar similar to the example of 'Derek Parfit is popular if and only if Derek Parfit is popular' which Y&H discuss in §6.1. The demonstrative formulation has the virtue that it avoids appeal to recurrence of tokens of the same name type. The name formulation has the virtue that it shines a light on the degree to which the quasi-internalist must appeal to problematic kinds of idealization.

to Millian-externalist views; it is not an aspiration which the Millian-externalist need share.

A more promising line for the quasi-internalist is to say that repetition is possible, but that it is not possible in cases like t_7. The question is then what principle might give us this result while still allowing the antecedent of t_5 to repeat t_4.

There are obvious candidates which one might try out here; in that sense, I agree with Chalmers (2018) that 'there is clearly a research program for internalists to be going on with' (even if the most promising versions of that research program fall on the quasi-internalist side of Y&H's taxonomy). My aim here has been less to settle the question than to follow the admirable example of Y&H in pushing the debate forward by laying out the options for, and challenges faced by, the various alternatives to externalist theories of content.[15]

References

Austin, D. F. (1990). *What's the Meaning of 'This'?: A Puzzle About Demonstrative Belief*. Cornell University Press.

Chalmers, D. (2012). *Constructing the World*. Oxford University Press.

Chalmers, D. (2018). Review of *Narrow Content*. *Notre Dame Philosophical Reviews*.

Egan, A., J. Hawthorne, and B. Weatherson (2005). Epistemic modals in context. In G. Preyer and G. Peter (Eds.), *Contextualism in Philosophy*, pp. 131–170. Oxford University Press.

Radulescu, A. (2018). Token-reflexivity and repetition. *Ergo: An Open Access Journal of Philosophy 5*, 745–763.

Yli-Vakkuri, J. and J. Hawthorne (2018). *Narrow Content*. Oxford University Press.

[15]Thanks to Brian Cutter for very helpful comments on a previous draft.